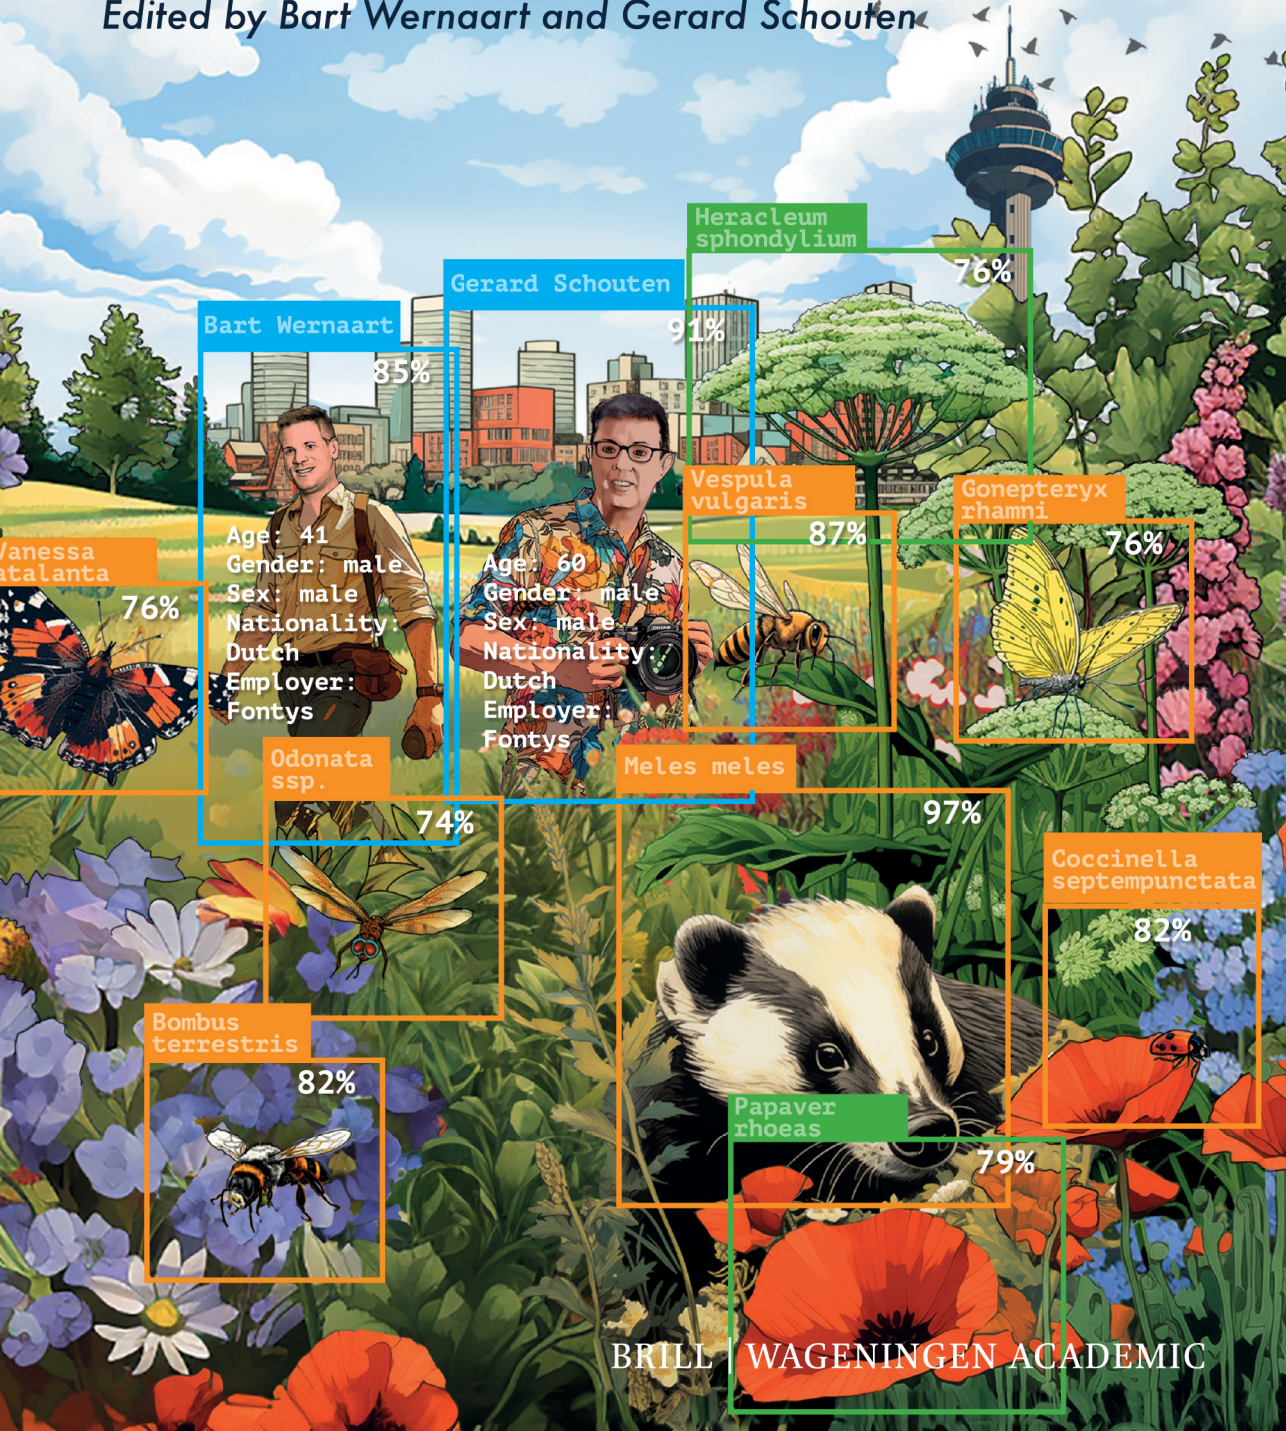


Moral design and green technology

Edited by Bart Wernaart and Gerard Schouten



Bart Wernaart

85%

Age: 41
Gender: male
Sex: male
Nationality:
Dutch
Employer:
Fontys

Odonata
ssp.

Gerard Schouten

91%

Age: 60
Gender: male
Sex: male
Nationality:
Dutch
Employer:
Fontys

Meles meles

Heracleum
sphondylium

76%

Vespula
vulgaris


87%

Gonepteryx
rhamni

76%

Vanessa
atalanta

76%



74%

Bombus
terrestris

82%




97%

Papaver
rhoeas

79%

Coccinella
septempunctata

82%



Moral design and green technology

Moral design and green technology

Edited by

Bart Wernaart
Gerard Schouten



BRILL | WAGENINGEN ACADEMIC



This is an open access title distributed under the terms of the [CC BY-NC-ND 4.0](https://creativecommons.org/licenses/by-nc-nd/4.0/) license, which permits any non-commercial use, distribution, and reproduction in any medium, provided no alterations are made and the original author(s) and source are credited. Further information and the complete license text can be found at <https://creativecommons.org/licenses/by-nc-nd/4.0/>

The terms of the CC license apply only to the original material. The use of material from other sources (indicated by a reference) such as diagrams, illustrations, photos and text samples may require further permission from the respective copyright holder.

Cover illustration: Exploring city nature. The cover (as well as the book's illustrations) was created by Berry Sanders with the assistance of the generative AI tool Midjourney.

The Library of Congress Cataloging-in-Publication Data is available online at <https://catalog.loc.gov>
LC record available at <https://lccn.loc.gov/2025010888>

Typeface for the Latin, Greek, and Cyrillic scripts: "Brill". See and download: brill.com/brill-typeface.

ISBN 978-90-04-71167-9 (hardback)

ISBN 978-90-04-73077-9 (e-book)

DOI 10.3920/9789004730779

Copyright 2025 by Bart Wernaart and Gerard Schouten. Published by Koninklijke Brill BV, Plantijnstraat 2, 2321 JC Leiden, The Netherlands.

Koninklijke Brill BV incorporates the imprints Brill, Brill Nijhoff, Brill Schöningh, Brill Fink, Brill mentis, Brill Wageningen Academic, Vandenhoeck & Ruprecht, Böhlau and V&R unipress.

Koninklijke Brill BV reserves the right to protect this publication against unauthorized use.

For more information: info@brill.com.

This book is printed on acid-free paper and produced in a sustainable manner.

Contents

- List of figures and tables VII
Contributors XI
- 1 Moral design and green technology 1
Bart Wernaart and Gerard Schouten
 - 2 Sustainability struggle: economics, business, and technology
A brief history and future challenges 5
Bart Wernaart
 - 3 Democratizing green technology with the public stack 27
Max Kortlander, Anne-Marie Sweep and Imme Ruarus
 - 4 Behavioural insights for moral design and green technology 48
Jeske Nederstigt
 - 5 The moral programming of XR, and what we can learn from the
AI experience 68
Leon Kester, Bart Wernaart and Nadisha Marie Aliman
 - 6 Citizen science for nature 82
Simona Orzan and Gerard Schouten
 - 7 The role of technology in human–nature connectedness
Case studies on citizen participation 97
Derk Jan Stobbelaar and Jetske G. de Boer
 - 8 Food ethics and technology
Towards food innovation with crowdsourced ethics 120
*Bart Wernaart, Sonja Floto-Stammen, Marieke van Vliet, Anika Kok and
Natalia Naranjo Guevara*
 - 9 How natural is our food?
How relevant is that word for the design of future food? 136
Niels Louwaars
 - 10 How to apply green AI in practice?
Moving from FLOPs to CO₂ footprint 153
Qin Zhao and Gerard Schouten

- 11 Lessons learned from developing green software 168
Luís Cruz and Petra Heck
- 12 A daily data workout!
Being in correspondence for a green data revolution 186
Danielle Arets and Jessie Harms
- 13 Added value of AI for studying urban plants 196
Barbara Gravendeel and Yannick Woudstra
- 14 Adding contextual information to object detection models: a wildflower
monitoring case 208
Georgiana Manolache and Gerard Schouten
- 15 ARISE: a Dutch dataspace connecting nature and people 233
*Elaine van Ommen Kloeke, W. Daniel Kissling, Julian Evans,
Chantal Huijbers, Jacob Kamminga and Gerard Schouten*
- Index 253

Figures and tables

Figures

- 3.1 Private stack, state stack, and public stack illustrations 29
- 3.2 Public, private, and state stack continuum 29
- 4.1 Two systems in the brain that control behaviour (Kahneman, 2011) 54
- 4.2 Original Technology Acceptance Model (Davis, 1986) 57
- 4.3 The diffusion process (Rogers, 1995) 58
- 4.4 The Strategies and Motives for Resistance to Persuasion (SMRP) Framework (Fransen *et al.*, 2015) 59
- 4.5 Self-Determination Theory (Ryan and Deci, 2000) 62
- 5.1 Illustration of the Socio-Technological Feedback Loop from a human-machine interaction perspective (Aliman and Kester 2022a) 75
- 5.2 Illustration of the Socio-Technological Feedback Loop from a system life-cycle perspective 76
- 6.1 Equipped with a torch and landing net, a group of nature enthusiasts visit pools in and around the city of Eindhoven 83
- 6.2 The number of active environmental citizen science projects through the years, according to the inventory maintained by the European Commission (JRC, 2018) 85
- 6.3 Impact of all citizen science projects on SDGs 88
- 7.1 Technology tends to reduce our connection with nature 99
- 7.2 Theoretical model of the mutual influence between system level and daily life level 100
- 7.3 Case study 1, walking apps may be designed to stimulate people to connect more strongly to nature, e.g. by providing information and stories 108
- 7.4 Case study 2, citizen science organized by the Dutch Butterfly Conservation foundation. Buckets with LED-light are used by citizen scientists to record moths in the Netherlands, while butterflies may be recorded along transects 111
- 7.5 Case study 3, IJsselstein is a municipality in the Netherlands with a large collection of fruit trees in the public green space. Technology, including a website and GIS application, provides insight into the uniqueness of this collection and helps to coordinate activities associated with it 113
- 8.1 The value profiles of two people that participated in the moral food lab. Note that the size of the words is in no way correlated with the number of times the words were mentioned 129
- 8.2 The average percentage of values people spoke in, grouped by how they responded to question B 130
- 10.1 Typical AI lifecycle model 155

- 10.2 The computational effort of AI models increases according to Moore's Law of the AI era 157
- 10.3 Carbon intensity for an assortment of locations 159
- 10.4 Benchmark of CNN models for image classification 160
- 10.5 Clean data improves prediction accuracy 161
- 10.6 Clustering the E-waste dataset into device groups 164
- 10.7 Training process of a neural network 166
- 11.1 The five dimensions of sustainable software engineering 171
- 11.2 ISO 25000 quality model 173
- 11.3 Energy monitoring setup for mobile app development 175
- 11.4 Development process for AI-enabled systems, including a data and model (ML) loop 182
- 11.5 Green AI at the root of Trustworthy AI 182
- 12.1 Leaflet Media Gym, Studio Cream on Chrome 187
- 12.2 Screenshot of taste workshop by E-missions 191
- 13.1 Examples of plants in an urban environment. a) Bird's-foot Trefoil (*Lotus corniculatus*), b) Ivy-leaved Toadflax (*Cymbalaria muralis*), c) Kidney Vetch (*Anthyllis vulneraria*), d) White Clover (*Trifolium repens*), e) Dandelion (*Taraxacum officinale*), f) Yarrow (*Achillea millefolium*), g) Common Poppy (*Papaver rhoeas*), h) Ground Elder (*Aegopodium podagraria*), i) Wallflower (*Erysimum cheiri*) 198
- 13.2 Impact of green design of private urban gardens on quality of living environment. Left: tiled backyard with overheated owner with irritated respiratory tract due to allergenic pollen released by ornamental olive shrubs. Right: backyard filled with non-allergenic trees, shrubs and herbs, providing shade, water retention, food to wild animals and a general feeling of well-being to owners 201
- 13.3 Example of urban trees encouraging bird safaris. Migrating Bohemian Waxwings (*Bombycilla garrulus*) foraging for berries in a tree planted along the canal of a typical Dutch historical urban center with bird watchers enjoying the scene, while keeping a respectful distance 202
- 14.1 Challenges of flower identification 'in the wild': 1) viewpoint variations (*Papaver rhoeas*); 2) occlusion (*Ranunculus repens*); 3) clutter (*Achillea millefolium*); 4) light variation (*Leucanthemum vulgare*); 5) deformations (*Bellis perennis*); 6) intra-class variation (*Ficaria verna*); 7) inter-class similarity (*Bellis perennis*, *Leucanthemum vulgare*, *Matricaria chamomilla*) 209
- 14.2 Image recognition and object detection comparison. Image recognition (first and third) labels the entire image, while object detection (second and forth) localizes objects in an image by drawing bounding boxes around them and then labels them accordingly. Photos are crops from EWD images (Schouten *et al.*, 2024) 212
- 14.3 Diagram of the Faster R-CNN architecture 213
- 14.4 Data fusion techniques: (left) early fusion, and (right) late fusion 215
- 14.5 An overview of our multimodal object detection solution 216

- 14.6 An example of an EWD image sliced into tiles, taken from Schouten *et al.* (2024). The dashed lines show equally sized tiles. Note the difference in the number of cut wildflowers (*Calyta palustris* in this case) between the two tiling schemes 218
- 14.7 Selected flower species grouped by visual similarity. Group 1: Buttercup (aggregate), *Caltha palustris*, *Ficaria verna*; Group 2: *Bellis perennis*, Chamomile (aggregate), *Leucanthemum vulgare*. Photos are crops randomly sampled from EWD 219
- 14.8 Data alignment overview for both groups: (top) flowering phenology estimates from NDFF for the selected species and (bottom) histogram of objects counts from EWD for the selected species. The horizontal axis is the day of year ranging from 1 to 365, while the vertical axis is (top) the phenological index, normalized from 0 to 1, and (bottom) an object count 220
- 14.9 Confusion matrix with confidence threshold over 0.75 and IoU threshold over 0.50 for image-only and learned feature-level fusion elementwise addition models 227
- 15.1 Overview of the end-to-end architecture in ARISE 235
- 15.2 Manual data collection and identification (top) versus a fully automatic AI-powered solution for monitoring biodiversity (bottom) 237
- 15.3 Diversity of digital biodiversity sensors tested in ARISE. (a) Location of the three ARISE monitoring demonstration sites in the Netherlands and the deployed sensors to monitor biodiversity non-invasively and remotely. (b) Different sensors and their data volumes 239
- 15.4 Overview of the Biocloud architecture with the different layers of processing the original data sources from raw to enriched and curated data for future use and access 243
- 15.5 The active learning cycle of advanced species identification in ARISE 247

Tables

- 7.1 Examples of changes through technology in various aspects of daily life 103
- 7.2 Types of Human Nature Connectedness (after Ives *et al.*, 2018) and examples of the role of technology 103
- 7.3 Ways in which the different types of HNC may be influenced in the three case studies 112
- 7.4 Future areas of life enhancing HNC 116
- 7.5 The potential role of technology in the symbiocene 116
- 10.1 Metrics of green AI 156
- 14.1 Selected flower species dataset description. Subspecies visually indistinguishable in the field are merged in the EWD dataset: *Ranunculus acris* and *Ranunculus repens* are labelled as Buttercup (aggregate), while *Matricaria chamomilla* and *Matricaria maritima* are labelled as Chamomile (aggregate) 219

- 14.2 Test results on Group 1 and Group 2 for all models averaged over five seeds. Best results are shown in bold 224
- 14.3 Test results per species class for the image-only baseline and the best performing feature fusion models averaged over five seeds 225
- 14.4 Test results on Group 1 and Group 2 for all classification models averaged over five seeds, hence average precision (AP) 226
- 15.1 Mapping of processes and scenarios to ARISE components 247

Contributors

Nadisha-Marie Aliman

Postdoctoral visiting scholar at Utrecht University

Danielle Arets

Professor Journalism and Innovation at Fontys University of Applied Sciences

Jetske de Boer

Professor Biodiversity and Nature-inclusive Society at Aeres University of Applied Sciences, and honorary fellow at NIOO-KNAW

Luís Cruz

Assistant professor Green AI and Sustainable Software Engineering at Delft University of Technology

Julian Evans

Researcher at Institute for Biodiversity and Ecosystem Dynamics (IBED), University of Amsterdam

Sonja Floto-Stammen

Senior researcher and lecturer in the field of business innovation, food marketing and consumer science at Fontys University of Applied Sciences

Barbara Gravendeel

Special professor in Plant Evolution at Radboud University and team lead of Evolutionary Ecology at Naturalis Biodiversity Center

Jessie Harms

Researcher, project manager at the Dutch National Climate Platform

Petra Heck

Lecturer, researcher and coordinator at Fontys University of Applied Sciences

Chantal Huijbers

Senior project manager at Naturalis Biodiversity Center

Jacob Kamminga

Assistant professor at the University of Twente

Leon Kester

Senior research scientist Responsible AI at TNO

W. Daniel Kissling

Associate professor of Quantitative Biodiversity, Institute for Biodiversity and Ecosystem Dynamics (IBED), University of Amsterdam

Max Kortlander

Writer and researcher at Waag Futurelab

Niels Louwaars

Law group at Wageningen University, and CEO at Plantum

Georgiana Manolache

Lecturer and researcher Data Science and AI at Fontys University of Applied Sciences and PhD candidate at Eindhoven University of Technology

Jeske Nederstigt

Consumer psychologist and a senior researcher Sustainable Behavioral Change at Fontys University of Applied Science

Natalia Naranjo Guevara

Researcher edible insects and senior project manager at New Generation Nutrition (NGN) [with specialisms in entomology, agribusiness, insect biology and behaviour]

Elaine van Ommen Kloeke

Program manager ARISE at Naturalis Biodiversity Center

Simona Orzan

Teacher and researcher at Fontys University of Applied Science

Imme Ruarus

Lab lead Smart Citizens Lab at Waag Futurelab

Derk-Jan Stobbelaar

Teacher and researcher at Van Hall Larenstein University of Applied Sciences

Gerard Schouten

Professor AI and Data and chair Centre of Expertise AI for Society at Fontys University of Applied Sciences

Anne-Marie Sweep

Researcher Collective Creativity and Innovation at Fontys University of Applied Sciences

Marieke van Vliet – van der Graaff

Senior researcher Moral Design Strategy at Fontys University of Applied Sciences, owner of Breinpaleis

Bart Wernaart

Professor Moral Design Strategy and chair Centre of Expertise Sustainable and Circular Transitions at Fontys University of Applied Sciences. Head of Research Fontys Business and Communication.

Yannick Woudstra

Postdoctoral researcher at Netherlands Institute of Ecology (NIOO-KNAW) and evolutionary botanist at Stockholm University

Qin Zhao

Teacher and researcher Green AI and Sustainable ICT at Fontys University of Applied Sciences

Moral design and green technology

Bart Wernaart and Gerard Schouten

1.1 Welcome to the Anthropocene

How can we involve citizens in designing green smart cities? What can we learn from fungi networks in developing new technologies? How can we design insect-farms for more sustainable food-products? And how can we use AI to measure biodiversity?

Welcome in the Anthropocene, the era where humans have put a tangible mark on our planet. But also, the era where humans and technology have the potential to shape the necessary transitions towards a sustainable world. Technology and nature are often considered as two opposing phenomena. However, they are increasingly intertwined, for better or worse. In this book, we explore how technology and nature relate to one another in the moral design of new, green, technology.

This book is relevant for IT- and engineering-professionals, business leaders and policy makers with (green) innovation in their portfolios and students of (applied) science who are interested in either sustainable and green design of technology or in the application of technology – with an emphasis on AI and IT – to create a greener, more sustainable world. The chapters are written by experts and leading researchers in an attractive, accessible, and practical writing style. Each chapter offers colourful examples and challenges the reader – without becoming paternalistic – to critically think through moral decision making and the design of innovations considering our planet’s perspective. This is a conceptual change in values. Nature should not be considered as a resource; it is the fabric of life that makes our own existence possible.

1.2 Chapter overview

In two parts, we discuss crucial relationships in shaping a world based on the morale that technology is green and sustainable. In the first part, we focus on the role of humans. How do humans relate to green tech, and what human processes are needed or should be optimized to help green technology flourish?

first part

Humans, business organizations and governments have been struggling to shape the economy in a sustainable way since the Industrial Revolution. The emergence of new technologies has always influenced these efforts both positively and negatively. In chapter 2, Bart Wernaart reflects on this by discussing groundbreaking

books – even including the leading voices in science fiction – that set the agenda on sustainability struggles and challenges since – roughly – 1760.

To make sure green technology is aligned with democratic principles around citizen participation and agency, the democratizing of technology becomes an urgent topic. To this end, in chapter 3, Max Kortlander, Anne-Marie Sweep and Imme Ruarus introduce the Public Stack model as a framework within which to root technological development in the public interest via participatory design practices and open technology.

Since human behaviour is often the bottleneck in the diffusion of green smart technology, Jeske Nederstigt reflects in chapter 4 on the acceptance of green and smart technological innovations, and what we can learn from behavioural sciences to stimulate and normalize sustainable attitudes.

In chapter 5, Leon Kester, Bart Wernaart and Nadisha Marie Aliman reflect on XR and how citizen can participate in its moral programming. XR technology is becoming more complex, is applied in countless contexts (as a professional instrument, a research tool or in the field of leisure), and its use is becoming more widespread amongst citizens. It can be considered as the ‘next big thing’ after Artificial Intelligence (AI). They propose a moral procedure called ‘moral programming’ that should enable technology developers and tech-policy makers to recognize and identify moral issues that need to be answered by those who are to be affected by the new technology, and equip them with a meta-method to establish a bottom-up crowd sourced ethics approach (including feedback loops) in their design process, to align the values of those whom it concerns and those who develop and/or govern.

In Chapter 6 Simona Orzan and Gerard Schouten present inspiring examples of ‘citizen science for nature’ and discuss the impact of these initiatives for society. They argue that humans are curious beings, intelligent and (mostly) social – the precise core ingredients that make a good scientist. They reflect on the challenges for science and the opportunities for citizens.

Technology in general has alienated us from nature. However, technology can also play a role in restoring or even reinventing human-nature connectedness. This duality is discussed by Jetske de Boer and Derk Stobbelaar in chapter 7. In this chapter they present three citizen participation case studies and reflect on the use of technology in the context of their Social-Technical-Ecological System framework.

Food technology in particular is at the core of a new moral revolution. It can alter the way we think about cultural acceptance of certain foods, our perception of what is healthy, and our perception of what it is fair (or sustainable) to consume. Despite the consequences that new technologies bring, the end consumer is rarely involved in the design process. In chapter 8, Bart Wernaart, Sonja Floto-Stammen, Marieke van Vliet and Anika Kok and Natalia Naranjo Guevara propose how to use

crowdsourced ethics as input for reshaping our food system by building a food tech moral lab.

Interestingly, ‘technology’ often has a negative connotation when connected to food, as it is considered ‘unnatural’. In chapter 9, Niels Louwaars reflects on the world ‘natural’ and rethinks the importance of this word for the design of future food and food technologies.

In the second part, our lens shifts towards green technology. What does green technology look like, and how can it help humans to become more sustainable, obtain a healthier livelihood and take better care of our planet?

second part

AI is revolutionizing industry and transforming society. However, the rapid growth of AI has also raised concerns about the environmental impact of this technology, as the energy consumption required to train and use AI models is increasing exponentially. To address this issue, a new research field called Green AI has emerged. In chapter 10, Qin Zhao and Gerard Schouten explore how to apply green AI in practice by discussing various case studies.

A similar approach can be found in chapter 11, written by Luís Cruz and Petra Heck, with a focus on best practices for sustainable software development.

And to conclude this triptych Danielle Arets and Jessie Harms argue in chapter 12, that engaging with artistic research is essential to raise awareness and inspire the public at large to stay ‘in correspondence’ with green data and AI.

Urban flora is increasingly appreciated for improving the quality of our living environment. Sadly, many do not notice the diversity of their green surroundings and miss out on how urban flora enriches their lives. In chapter 13, Barbara Gravendeel and Yannick Woudstra discuss how AI might cure ‘plant blindness’, validate plant-based ecosystem services, and change perceptions of the added value of urban green spaces with multiple stakeholders.

Chapter 14, written by Georgiana Manolache and Gerard Schouten, shows in detail how this promise can be fulfilled. They present an advanced AI model that can automatically identify and count wildflowers. They demonstrate that a multimodal approach – in which vision is combined with phenology data – significantly improves model performance and is able to solve the typical wildflower puzzle of high intra-class variation and inter-class similarity.

Worldwide, biodiversity is in rapid decline. It is important to help restore the connection of people and nature. One of the ways of doing this is to provide a data platform that supports large-scale monitoring and provides (near) real-time biodiversity services. Therefore, Elaine van Ommen Kloeke, Daniel Kissling, Julian Evans, Chantal Huijbers, Jacob Kamminga and Gerard Schouten introduce in chapter 15 the ARISE infrastructure, that is currently being built and that targets semi-automatically identification of all multicellular species in the Netherlands.

1.3 Building bridges and inclusive authorship

In this book we try to bridge many worlds, and this bridging is done on various levels. First, there is the disciplinary level: the chapters represent a wide range of academic disciplines and professions.

Herein lies the second gap we tried to close: that between the sciences and applied sciences. The moral design of green technology cannot be a merely academic discussion. Instead, it is something actionable. In this book we offer both scientific reflections and applied science perspectives.

Third, we tried to find a proper gender balance in the team of authors. Not only in the authorship of the contributing authors, but also in our references to literature and how we write about individuals. We tried either to use gender-neutral language, or make sure to use he/she references in equal amounts. We also realise that this does not fully cover the diverse gender types that exist and would welcome suggestions to improve that. In this book, the geographic scope of authors is – with some exceptions – predominantly European and North American, with a professional focus on the Netherlands. We are aware that this could be seen as a shortcoming. However, throughout the chapters the reader may recognise a global mind-set.

1.4 On a more personal note

The editors want to express their gratitude to the following people and organizations. First and foremost, we would like to thank the contributing authors who were so kind to share their knowledge and talents with us. We are extremely proud to have been able to work with this author team with such remarkable, creative and bright minds. We would also like to express our gratitude for the cover artwork and the book's illustrations, created by Berry Sanders with the assistance of the generative AI tool Midjourney.

Second, we would like to thank Fontys University of Applied Sciences for their ongoing support, and their financial contribution to make this publication an open access work. In particular, we would like to thank our colleagues from the Fontys Centre of Expertise AI for Society and Centre of Expertise Sustainable and Circular Transitions for sharing our enthusiasm for this work.

Third we would like to acknowledge our publisher Brill-de Gruyter for placing their trust in us by publishing our work, and for simply being an outstanding publishing company.

Sustainability struggle: economics, business, and technology

A brief history and future challenges

Bart Wernaart

Abstract

In this chapter we will reflect on how the struggle for sustainability as a goal for business and technology developers ended up in literature, science fiction and academic works. We first discuss some groundbreaking novelists and science fiction writers (1760–1987), then we will highlight the emerging role of business and how this affected the sustainability debate (1969–2002). Then, we will reflect on the interplay between economic policies, business models, technology and sustainability (2002–2024). Last, we will discuss some future challenges.

2.1 Books and the struggle for sustainability

Since 1760, the application of the steam engine has led to the industrialization of Western societies and the emergence of mass production in factories. Besides the enormous technological progress, production efficiency and upscaling of industrial business, this has led to several societal challenges, including urbanization, pollution on a large scale and social inequality (Trinder, 2013; Dimanov, 2024). Also referred to as ‘the social question’ (Vercherand 2014), this was a pressing issue in the 19th and early 20th century. Thereafter, industrialization led to – more or less – similar effects elsewhere in the world (e.g. Angotti, 1996; Baloch *et al.*, 2018; Liang and Wang, 2019). In various groundbreaking works, critical novelists and science fiction authors foreshadowed the devastating effects of large-scale industrialization.

2.2 Some groundbreaking novelists and science fiction writers (1760–1987)

Charles Dickens was one of the first British writers to highlight the social and environmental challenges that came with industrialization. In novels such as *Oliver Twist* (1838), *A Christmas Carol* (1843) and *Hard Times* (1854) he famously criticized the unequal distribution of the means of production and the poor and polluted living conditions of workers.

Charles Dickens

Upton Sinclair

Another critic of the social question in the US was Upton Sinclair. In his book *The Jungle* (1906) he focusses on the corrupt behaviour of those who are in power in both government and the (mass production) food business industry. In preparation for his novel (first published in the socialist newspaper 'Appeal to Reason'), Sinclair worked undercover in a meatpacking plant in Chicago. In his novel, he describes the cruel treatment of animals, the poor (and even dangerous) food safety standards, and the exploitation of workers in the industrialized meat industry.

In the industrialized countries, the so-called 'social question' was slowly resolved during the 20th century by increasing protection of workers through labour laws, democratization of the industrialized countries, and the founding of labour unions. In the long term, the positive effects of industrialization also became visible, with technological developments and scientific progress that would – amongst other things – lead to the reduction of pollution effects, and better healthcare and mobility (Trinder, 2013).

One way or another, the technological shockwave stirred by the steam engine, and its resulting mass-manufacturing, was a warning: technology had side effects that were not only positive, but could also be dangerous or counterproductive. Not only for people, but also for our (living) environment. It is perhaps not surprising that the science fiction genre thrived alongside large-scale industrialization. Authors explored the possible future effects of a world with advanced technology. Most notably, the 'godfathers' of the genre – H.G. Wells and Jules Verne – already fantasized about what the world would look like if technology worked in a destructive or unethical way. H.G. Wells warns us about bioengineering, and shows us a world in which we can artificially alter human or animal bodies. In the character of Dr Moreau we find the typical 'mad scientist' who changes animals into men (and thereby acts as some sort of God), with disastrous, unethical and inhumane effects (*The Island of Doctor Moreau*, Wells, 1896). In his famous novel *The Time Machine* (1895), Wells shows us a utopian future made possible by advanced technologies, and a raw-dystopian beneath-the-surface necessary to maintain these techniques. And in *The World Set Free* (1914), Wells explores the possibility of weapons of mass destruction (including the atomic bomb), and warns us about the desire use massive amounts of energy. And consider also Jules Verne's works. In all his books, he warns us of the environmental consequences of new technologies (Evans, 1988).

H.G. Wells

In *The Child of the Cavern* (*Les Indes noires*, 1877), Verne raises the question of what would happen if all the coal were used up and exhausted by humankind. In both *The Begum's Fortune* (*Les Cinq cents millions de la Bégum*, 1879) and *Paris in the Twentieth Century* (*Paris au xxe siècle*, 1863), Verne sketches a grim image of the future, including dystopian and heated places, all caused by technologies that turned out to have a destructive long-term effect. The latter work was even rejected by Verne's publisher due to its overly dystopian tone of voice, and subsequently put aside by the author. More than a century later, in 1994, it was published as a 'lost manuscript'.

Jules Verne

Interestingly, both Wells and Verne also imagine a future in which human beings use technology for good and are able to master our planet with sustainable and innovative technologies. Whereas, in the book *The Purchase of the North Pole* (*Sans dessus dessous*, 1889), Verne warns about the risks of terra-forming technologies, he is more positive about ‘cultivating the environment’ in *The Mysterious Island* (*L’Île mystérieuse*, 1875), where the main characters transform from shipwrecklings to colonists, deploying their technological skills to survive and cultivate most of the island, and using water as a power source (the coal of the future). And then we have the electrical submarine in *Twenty Thousand Leagues Under the Sea* (*Vingt mille lieues sous les mers*, 1870), where the Nautilus is used for marine research. Wells also shows us different future utopias in which humankind is able to control other species and the planet’s environment through technology. In books such as *A Modern Utopia* (1905), *Men Like Gods* (1923), and *The Shape of Things to Come* (1933), Wells shows us an optimistic view of technology, suggesting that, hypothetically, men should be able to overcome the threat of extinction by mastering ‘mother nature’. His ultimate idea and solution for governing our planet in a meaningful way is to strive for a world government. In his historic work *The Outline of History* (1919–1920) he explores the idea of the ‘United States of the World’ (a term that he also coined in his earlier work).

Such viewpoints reflect Enlightenment thinking, in which the idea prevailed that mankind could control the environment through technology. It needs to be noted here, however, that both Verne and Wells developed a darker view on technology when they grew older. Wells in particular was deeply moved by the events of the two world wars, which changed his perception of humankind’s relationship with technology. This resulted in a grim societal reflection in his last book *Mind at the End of its Tether* (1945) in which he predicts climate change, amongst other things. With despair, he describes our natural habit of ‘race suicide by gigantism’ through aggression and greed (Chapter V). He then continues (in Chapter VIII):

The writer sees the world as a jaded world devoid of recuperative power. In the past he has liked to think that Man could pull out of his entanglements and start a new creative phase of human living. In the face of our universal inadequacy, that optimism has given place to a stoical cynicism. The old men behave for the most part meanly and disgustingly, and the young are spasmodic, foolish and all too easily misled.

In later science fiction, technological pessimism was explored in more detail, in particular in the so-called ‘new wave’ of science fiction from the 1960s onwards. So we have seen plenty of examples in which authors warn us about technology ‘gone wrong’, resulting in biohazards and environmental destruction (Roberts, 2016). Perhaps Frank Herbert’s *Dune* series is one of the first works to widely explore the theme (although the desert-like planet Arrakis was not a result of a manmade

later science
fiction

disaster). Since the '70s in particular, destructive pollution as a result of manmade technologies, such as nuclear wars, failed terraforming and corporate greed, are central themes. Books like *The Sheep Looked Up* (John Brunner, 1972), *The Lathe of Heaven* and *Always Coming Home* (Ursula K. le Guin, 1971 and 1985), *Parable of the Sower* (Octavia E. Butler, 1993) are classics in this respect.

2.3 From science fiction to science and politics (1962–1987)

While in science fiction, authors have warned us about the environmental effects of technological progress, red flags were also being raised gradually in both science and politics.

Rachel Carson

One of the first was Rachel Carson, a marine biologist, whose groundbreaking book *Silent Spring* (1962) addressed her grave concerns regarding the widespread use of dichlorodiphenyltrichloroethane (DDT) and other pesticides in the agricultural sector. The use of these pesticides led to a maximization in efficiency in agriculture. Carson postulated that the effect on biodiversity loss and human health were being ignored by the industry (and politics). She was able to combine hard-core scientific insights with an eloquent and activist writing style (using chapter names such as 'elixirs of death', 'rivers of death', 'the human price', etcetera). Her work concluded as follows:

the 'control of nature' is a phrase conceived in arrogance, born of the Neanderthal age of biology and philosophy, when it was supposed that nature exists for the convenience of man. The concepts and practices of applied entomology for the most part date from the Stone Age of science. It is our alarming misfortune that so primitive a science has armed itself with the most modern and terrible weapons, and that in turning them against the insects it has also turned them against the earth.

Carson's ideas led to quite a lot of resistance in both the industry and politics (Beijl, 1992). She was attributed all kinds of labels ('bird lover', 'cat lover', 'fish lover', 'nun of nature', and 'priestess of nature') and was criticized for favouring the interest of insects above the well-being of thousands of malnourished human beings who would have a better life as a result of more efficient agriculture. The United States Department of Agriculture (USDA)'s Agricultural Research Service even published a film (*Fire Ants on Trial*) in support of the use of pesticides to convince people of its usefulness and necessity. In Carson's view, this was 'flagrant propaganda'.

Her work is often considered a starting point for the environmental movement in the U.S.A. and elsewhere in the world, and created a widespread awareness around the safety of food production and food consumption as well as environmental well-being (Lytle, 2007). In the U.S.A., eventually, environmental laws were

significantly revised, and a ban on the use of DDT was introduced. The effects of the book still resonate in the debate around environmental well-being and agricultural technologies (Epstein, 2014; Arp, 2023).

At the global level, after the Second World War and most certainly after the Cold War, the prevailing macro-economic model was that of neo-liberalism, with a strong focus on free trade (i.e. lowering trade barriers and creating a level playing field for business) and economic growth. Amongst others, the United Nations (UN) aligned Bretton Woods institutions (World Bank Group, the International Monetary Fund and the World Trade Organization, formerly known as the General Agreement on Tariffs and Trade (GATT)) would put the neo-liberal mindset at the core of its work, to stimulate a global economy that would raise countries from the ashes of WWII. Increasingly, however, the effects of industrialization and (global) trade on the environment became more visible and problematic. At the UN level, this led to various landmark gatherings to discuss environmental issues as a result of human interference. One of the first to be held was the United Nations Conference on the Human Environment, in Stockholm in 1972. Despite being boycotted by the Warsaw Pact countries and not welcomed by the informal 'Brussels Group' (including Britain, the US, Italy, Belgium, the Netherlands and France), the conference resulted in the Declaration of the United Nations Conference on the Human Environment (the Stockholm Declaration). In this declaration some pressing issues were addressed, and countries would attempt to solve them in line with the 26 Principles therein. These included the rapid exhaustion of natural resources, the unequal distribution of non-renewable resources, the serious consequences of pollution (including ocean pollution), the weak position of developing countries in these challenges, and the threat of weapons of mass destruction.

Stockholm
Declaration

In the same year, the Club of Rome published its computer-simulated analysis, using the World3 computer model, in a book called *The Limits to Growth* (1972). The analysis involved a forecast of the development and interaction between the variables of population, food production, industrialization, pollution, and consumption of non-renewable natural resources. The conclusion was alarming: if humankind continued in the way they were evolving, the limitations of growth would be reached within at least a century.

Club of Rome

More than a decade later, the principles set forth in the Stockholm Declaration had not been met, which led to the setting up of the UN World Commission on Environment and Development, later known as the Brundtland Commission (named after its chairperson, the Norwegian Gro Harlem Brundtland). The famous Brundtland report, *Our Common Future* (UN, 1987), was a milestone in global sustainability thinking. In this report, the authors emphasize the risks that come with exhausting natural resources, and that natural resources are limited. They also identify poverty reduction, gender equality and redistribution of wealth as key features in strategies for environmental preservation. In general, the idea is that poverty stands in the way of sustainability, and need to be tackled alongside. This also

Brundtland
report

means that economic growth must be balanced with ecological sustainability. The Brundtland report has been a continuous source of inspiration for governmental policies in the field of sustainability, and is at the core of sustainable business strategies (Schubert and Láng, 2005).

2.4 The emerging role of business (1969–2002)

In the post-World War period, it was not only governments and global institutions such as the UN that were confronted with the harmful effects of technology on the planet's environment. Increasingly, the role of the private sector was put on the agenda, and businesses were called on to take their share of responsibility in achieving sustainable development.

John Ellington is one of the pioneers in sustainable business and corporate social responsibility. In his famous book *Cannibals with Forks* (1997) he coins his main ideas on the role of business in this discourse. In Ellington's analysis, the sense of urgency concerning sustainability did not gradually evolve but was rather triggered by so-called 'shockwaves'. He labels them as 'limits', 'green', and 'globalization'.

limits The first shockwave, 'limits', were triggered by the notion that there are limits to growth. The idea was greatly inspired by Rachel Carson's *Silent Spring*, and coincided with the Stockholm Declaration and the Club of Rome report *The Limits to Growth*. The peak of this shockwave was around 1969–1973 and was met with strong resistance and strident lobbying by the private sector against environmental laws.

green A second shockwave was labelled 'green' and was triggered by various ecological disasters, including the Bhopal gas tragedy (1964), the nuclear disaster of Chernobyl (1986) and the Exxon Valdez oil spill (1989). In the Exxon and Bhopal cases in particular, the slow and responsibility-avoiding responses of the respective companies involved – Exxon and Union Carbide – attracted negative media attention, which led to a wider debate about the responsibility of the corporate sector regarding the affected people and their environment. This triggered a response in environmental law-making in which the notion was felt that reducing pollution and the use of natural resources was not enough. The invention and use of more sustainable technologies and production processes needed to be encouraged. Also, amongst businesses, there was a change of focus. Instead of merely obeying the law and ticking the 'ethics' box, a discourse evolved in the private sector in which the urge was felt to take a broader responsibility than before, and the idea emerged that sustainable entrepreneurship could be a unique selling point and therefore a way to compete. The peak of this shockwave was around 1988–1990.

globalization A third shockwave was 'globalization'. Increasingly, the role of international financial institutions was critically assessed. The lowering of trade barriers led to large-scale global trade, but also to large 'regulatory competition'. In other words, states (in particular those with vulnerable economies and/or instable state

structures) could compete with more flexible rules in the field of environmental protection or human rights. Product chains had become complex and long distanced. This led to the optimization of production costs, but at the expense of the environment or human well-being. Multinational corporations were called out for pollution and poor labour conditions within their product chains in developing countries. Notorious examples are the Shell scandal in the Niger Delta (in the noughties), or the collapse of the Rena Plaza Building in Bangladesh involving many multinational garment brands (2013). This shockwave peaked around 2002.

Ellington's central idea built on the idea that a business should not merely focus on making profit as a single bottom line, but should also report on social wealth creation and environmental responsibility. This notion was introduced by Freer Speckley in 1981 (Spreckley, 1981), and later rephrased by John Elkington as the 'triple bottom line': a business should act according to three bottom lines: People, Planet and Profit (3Ps).

One of the key ideas of Triple Bottom Line thinking is that profit is not the only goal. This also has an effect on how a business defines and relates to its stakeholders. From a strategic business perspective, R. Edward Freeman evolved his theory on strategic management in the early '80s, for the most part in tandem with Elkington's development of the Triple Bottom Line (although both authors are in completely different academic circles and do not actively refer to one another's work). In essence, Freeman argues that concepts such as 'corporate social responsibility' or 'triple bottom line' are not per se useful concepts. Instead, he has a rather pragmatic, managerial approach: in the preface to the 2010 edition of his core work (1984) he argues:

I believe that the central characters (...) must be companies and their customers, employees, suppliers, communities and financiers. Other groups such as NGOs, governments, unions etc. may also be important to particular businesses. The idea that one of these groups (financiers, for instance) always has priority over others, simply misses the main contribution of business and capitalism. Business works because the interests of all of the stakeholders can be satisfied over time. It is the intersection of these interests which is central to effective and sustainable stakeholder management.

In other words, Freeman's central point is that from a strategic point of view a business should move away from a shareholder approach (with a sole focus on the financial interest of the financiers of a business), and instead be driven by a stakeholder approach in order to achieve sustainable corporate success. A central question then is how to define business stakeholders and their interests. In general, stakeholders could be defined as those individuals or groups that can affect or are affected by the company's decision-making (see Freeman, 1984; see also Crane *et al.*, 2019: 59–65). In earlier models, shareholders were presented as relatively strictly separate groups

that included mostly shareholders, suppliers, customers and employees. Later, governments, competitors and civil society were added to the list (Freeman, 1984). However, such ‘lists’ do not necessarily do justice to the business and societal reality a company finds itself in. First of all, stakeholders amongst themselves interact and have fluid interests that are context dependent. Second, stakeholders are mostly translated to ‘groups of people’. The environment is therefore barely a direct stakeholder (at most it is represented through NGOs or governments) while it is the one and only supplier of natural resources. To solve this, stakeholder approaches are increasingly being transformed into a full network model that not only maps out the various stakeholders, but also shows how these stakeholders in turn are related to their own stakeholders and how their interests interact (Rowley, 1997). Increasingly, we see models in which the environment is recognized as one of the primary stakeholders of a business (Jacobs, 1997; Phillips and Reichart, 2000; Madsen and Ulhøi, 2001; Haigh and Griffiths, 2009).

network model

2.5 Economic policies, business models, technology and sustainability (2002–2024)

Alongside the increased urgency of environmental issues we see new ideas from various academic disciplines on (interrelated, but too often separately discussed) economic policies, business responsibility and the role of technology/innovation.

2.5.1 *Grim, grimmer, grimmest: crises, global warming and disruptive technologies*

Since the start of this century, several crises have triggered an increasingly heated debate about the desirable shape and form of our economies, the responsibility of the business sector and the role of technology. Some of these crises are clearly identifiable and can be pinpointed on a timeline, such as the financial crisis around 2008 and the COVID-19 pandemic in 2019–2022. Others are ongoing, such as man-made global warming, and the adversarial effects of disruptive technologies.

financial crises

2.5.1.1 The financial crises (around 2008)

At the end of 2007, global financial markets collapsed. The disaster started in the United States of America, with – among other things – the large-scale (and risky) lending of subprime mortgages and the sharp decline of real estate value. This resulted in a chain reaction across the world that brought entire countries to the brink of bankruptcy and had a negative impact on huge numbers of people. The financial crisis was attributed to the widespread immoral behaviour of financial institutions and companies and the failure of governmental supervision. This resulted in an ethical discussion about responsible banking, responsible investing, responsible accounting and responsible supervision. In addition, the

way in which governments and financial institutions must (be able to) supervise the financial markets, and in particular the way in which companies should take responsibility in this regard, shot to the top of the agenda (Foster and Magdoff, 2009; Westbrook, 2009).

2.5.1.2 The COVID-19 pandemic (2019–2022)

COVID-19
pandemic

In late 2019, the Corona virus (COVID-19) raged throughout the entire world, resulting in many deaths, lockdowns, and social and economic disruption. The pandemic affected the business world and the business world affected the way the pandemic could be ended, which led to new unexplored ethical challenges. Technology played a central role here on multiple levels. On the one hand, the pandemic had a disruptive effect on (especially) SMEs in certain branches (Belitski *et al.*, 2022), e.g. tourism, restaurants, the arts. On the other hand, it led to a sudden growth of other industries (e.g. ICT and multimedia businesses) and above all the flourishing of big tech (Fernandez *et al.*, 2020), whose inventions were crucial in the new ‘working from home’ reality.

And last but not least, the pharmaceutical industry played a key role in developing vaccines that could ‘open up’ the world again. This also led to ethical discussions on various levels. How safe are quickly produced vaccines? Can companies make use of their patent rights when they lack the capacity to produce enough vaccines quickly? How powerful/how much influence does the pharmaceutical industry have in deciding how vaccines are distributed across the globe, and to what extent can profit be a priority in these decisions? (Mucchielli, 2020; Binagwaho *et al.*, 2021)

2.5.1.3 Ongoing global warming as a result of manmade climate change and its consequences

global warming

Since 1988, the UN Intergovernmental Panel on Climate Change (IPCC) has been assessing the worldwide (peer-reviewed) science related to climate change, with written reports on the climate status of our planet, and discussions on future scenarios. Since the first report in 1990, these future scenarios become grimmer and more undesirable by the year. The Earth is gradually heating up as a result of manmade global warming. Manmade global warming is for the most part caused by the use of fossil fuels/greenhouse emissions, especially in industrialized regions. These alarming reports were core to the adoption of the Kyoto Protocol (1997) and the Paris Climate Agreement (2015) with the aim of reducing greenhouse gas emissions, as well as development agendas such as Agenda 21 (1992), the Millennium Goals (2000) and the Sustainable Development Goals (SDGs) in 2015. While these initiatives are essentially addressed to States, and UN Member States are its signatories, the corporate sector increasingly uses these goals as yardsticks for their own business performance. For instance, the Sustainable Development Goals address multiple challenges that directly relate to business performance, such as SDG 12 (responsible consumption and production). The SDGs are widely used, and not

only by governments (Abhayawansa *et al.*, 2021; Bexell and Jönsson, 2019) and companies (Rashed and Shah, 2021; Kumi and Kumi, 2021); and the 2030 Agenda for Sustainable Development emphasizes that governments, businesses and civil society have a joint responsibility to contribute to the realization of the SDGs.

disruptive
technologies

2.5.1.4 Disruptive technologies that led to sudden societal changes

In the post-WWII era, we have seen the societal and economic transformation from industrialization to digitalization. At its centre is the rise of information technology, and so this transformation is referred to by some as the information age, or digital revolution. These technologies greatly influenced the way we communicate, do business and look at the world. Inventions such as the home computer (1969–1989), the world wide web (internet) (1989–2005), social media and smartphones (2005–2020), profoundly changed multiple aspects of daily life. Increasingly, the power of tech companies became a major concern, with only a handful of corporate organizations owning (and deciding on) the world's digital infrastructure. Since 2010, but especially since 2020, new major technologies have once again been driving fundamental changes, and even further this time round. At the core is Artificial Intelligence (AI) and – perhaps to a lesser extent – Extended Reality (an umbrella term for Augmented and Virtual Reality technology), and in the near future, the effects of Quantum Technology are to be expected. When considering AI, the wide and accessible use of Large Language Models (LLMs) has had a particularly tangible effect on people's lives and has become a subject of moral concern in its own right (e.g. Felten *et al.*, 2023).

When we consider these digital developments, matters about privacy, consumer autonomy, intellectual property and equality are high on the agenda (McAfee and Brynjolfsson, 2017; Van Dijk *et al.*, 2018; Eubanks, 2019; Wernaart, 2022). Such issues relate to social challenges and citizens' rights. However, these digital developments also affect environmental issues in many ways, as we can see elsewhere in this book. Digital innovation is considered both a solution (e.g. Biswas, 2022) and a problem (e.g. Tamburrini, 2022) where environmental sustainability is concerned. However, this has long been an underrepresented theme in the general discourse on new technology and responsible innovation (Al-Khouri, 2013; Lichtenthaler, 2021)

2.5.2 *Economic models that challenge unlimited growth*

To address the concerns that relate to the effects of unlimited growth and unsustainable industries, various macro-economic ideas and models were created to rethink the (neo)-liberalist mindset that coincided with industrialization.

The purpose of economic growth in particular is challenged from multiple perspectives. A first aspect is that of capital and income, as advocated by the famous economist Thomas Piketty. In short, he concludes that interest on capital usually grows faster than the average economy can grow. This will always lead to a growing

income gap in the longer term (with some historical exceptions), leading to economic instability with disproportionate wealth in the hands of the few. The idea that economic growth is beneficial to all is therefore an erroneous reasoning if governments do not interfere. He emphasizes the necessity to redistribute wealth, e.g. through higher and worldwide taxes on capital and progressive income taxes (Piketty, 2014; 2021). With these conclusions, he argues against some of the important concepts we find in neo-liberalism, which is usually characterized by free markets and limited governmental interference.

Piketty's work has a strong focus on social justice, and not so much on environmental sustainability. Kate Raworth's so-called 'donut economy' has both. In this model, the donut is used as a physical metaphor for the desirable size of our economies. The outer side of the donut represent the limits of our planet (the ecological ceiling), and the inside represents the economic minimum that is required to provide all people with an adequate living standard (the social foundation). Somewhere between these two extremes, our economies can develop. This is called a safe and just space for the economy. Economies that are below the social foundation represent a shortfall, whereas economies that are beyond the environmental ceiling represent an overshoot (Raworth, 2012; 2017).

Kate Raworth

Whereas Piketty and Raworth are not per se against economic growth as such (as long as it leads to a fair distribution of wealth, or fits within the donut), some economists argue that growth is no longer necessary. In 'The Economics of Arrival', Katherine Trebeck and Jeremy Williams (2019) hold that some economies are simply 'grown up' and people need to learn how to live in these grown-up economies rather than continuously striving for even more growth.

Katherine
Trebeck and
Jeremy
Williams

Some authors even introduce the concept of 'degrowth', the main idea being that economic growth is no longer a means to an end, but rather social well-being and ecological sustainability, for which, usually, economic degrowth is required.

Serge Latouche advocates degrowth in his famous book 'Farewell to Growth' (2009). He proposes the 8R model that can help in rethinking our economy: re-evaluate, relocalize, reduce, reconceptualise, restructure, redistribute, reuse and recycle. In short, it means we need to re-evaluate our perception of wealth: this is not the same as money, but instead includes well-being and happiness. Also, we need to produce locally (and not globally), and product chains need to be designed in such a way that they naturally lead to the reuse and recycling of these products. Also, we need to strive for a fair redistribution of resources.

Serge Latouche

Jason Hickel (2020) also advocates degrowth. In a nutshell, he argues that capitalism always leads to the exploitation of people and our planet, and is therefore not a useful economic model. He proposes a radical change in the way we think about our economies. Amongst other things, he proposes valuing things based on their usefulness rather than a price that is a result of supply and demand. This would require some systemic changes in our economies. For example, he suggests

that a company should not focus on sales amounts but instead on duration of use of their products; we should shift our focus from ownership to usership, make sure that destructive industries no longer have a right to exist, redistribute resources and wealth in a fair way, stop advertisements, and stop food waste.

2.5.3 *Business models and responsible technology*

As mentioned in Section 2, there was a growing notion after WWII that the private sector was a problem for sustainable development, and could also be part of the solution. We have seen the Triple Bottom Line model of John Elkington, and the Stakeholder or Network Approach triggered by the works of R. Edward Freeman. But how do we put such ideas into practice? How do we organize companies not only around profit-making, but also around positive social and ecological impact? Countless models have been introduced ever since it was suggested to organize businesses in a more sustainable manner, and use and develop technologies with a positive social and ecological impact.

These models, however, have been developed from different academic corners, and as such result in a rather scattered approach. Also, the warning voiced by macroeconomist Milton Friedman in the late sixties resonates to this day: ‘the social responsibility of business is to increase its profits (Friedman, 2007, originally published in the New York Times Magazine, 1970). Or, in more popular wording: ‘the business of business is business’. In his work, Friedman advocated a no-nonsense shareholder approach whereby he argued that if businesses were to do something beyond the making of profit it would lead to spending someone else’s money (that of the shareholders) to a goal in the general interest. This would lead to a hidden and undemocratic tax, and, above all, it is not in the nature of a business to do something beyond making profit. So, with a scattered academic approach and the profit-driven nature of a business, one of the greater risks here is that of ‘greenwashing’ (de Freitas Netto *et al.*, 2020).

greenwashing

2.5.3.1 Accountancy and finance: integrated reporting

Each business must give an account of its financial achievements. In most legal systems, there are fixed rules and regulations for annual financial statements for private and public limited companies, and slightly more flexible rules for other business forms, such as (general) partnerships, sole proprietorships and foundations. In 2010, various ways of so-called ‘external cost accounting’ or ‘integrated reporting’ were introduced. The central idea is to broaden the scope of the materiality principle. This principle in accounting embodies the idea that an annual report should contain all necessary information so that its stakeholders can properly assess the performance of a company. Increasingly, this is understood more broadly than financial performance only, and company performance regarding social and ecological impact are incorporated in the annual reporting.

CASE 1.1

Integrated reporting and sportswear

Puma, a German sportswear company, was one of the frontrunners in integrated reporting. In their 2010 annual report, Puma not only provided accounts for their financial achievements, but also calculated the social and environmental costs of their production processes. In this report, Puma calculated that for each € 100 of sales, the societal and environmental costs were € 6.70. This way of calculating costs could lead to more – and product chain wide – insights into societal and environmental impact. It was concluded that most of this (negative) impact was caused in the supplying countries, and not per se in the sales environment. This could lead to specific policies to improve and mitigate this negative impact.

One of the landmark initiatives was driven by the German-based company Puma, as we can see in Case 1.1. Increasingly, other businesses undertook similar initiatives, and various approaches and methods were introduced, including the review of business achievements against the Sustainable Development Goals (SDGs). Also, on a global and regional scale, various initiatives and methods were developed to stimulate integrated reporting of businesses, such as the Carbon Disclosure Project and the Global Reporting Initiative for companies, and the UN System of Environmental Economic Accounting for national industries.

In Europe, integrated reporting is now considered obligatory for so-called public-interest entities. These are usually listed companies, credit institutions and insurance undertakings, and some other (sometimes nationally designated) organizations, such as – in some countries – pension funds or investment companies. Since the adoption of the Corporate Sustainability Reporting Directive,¹ these organizations are obliged to report according to European Sustainability Reporting Standards.² This means that they have to report on risks and opportunities that relate to social and environmental aspects of the business, as well as their impact on these social and environmental issues.

In the literature, there is a feisty debate about the actual impact of integrated reporting (Stubbs and Higgins, 2014; Brown and Dillard, 2014), and some parties raise the concern that methods of integrated reporting (or the data that is used)

1 Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 amending Regulation (EU) no. 537/2014, Directive 2004/109/EC, Directive 2006/43/EC and Directive 2013/34/EU, as regards corporate sustainability reporting

2 Commission Delegated Regulation (EU) 2023/2772 of 31 July 2023 supplementing Directive 2013/34/EU of the European Parliament and of the Council as regards sustainability reporting standards

are usually unclear or random, which easily leads to companies drawing desirable conclusions, rather than having a serious and measurable impact on social and environmental issues (Beske *et al.*, 2020). In other words, the chosen method and dataset can heavily influence the 'greenness' of a business, and therefore lead to ambiguous conclusions (e.g. Fay, 2012; Murtha and Hammlton, 2012).

2.5.3.2 Business transition and sustainability

Earlier in this chapter, we introduced Elkington's triple bottom line. He defined the necessary prerequisites and transitions to become a truly sustainable private organization. Elkington observed various revolutionary dimensions that would be indispensable for a sustainable capitalism transition, and that could lead the private sector away from unlimited and unsustainable growth. The first dimension is a shift from compliance to competition. A sustainable and successful business organization can no longer sustain itself by merely complying with legislation. Instead, as a result of consumer demand and green investment policies, social and environmental responsibility are part of a business case to create unique selling points that lead to more competitiveness. A second dimension relates to values. Elkington observes that societal values changes and keep changing. This also means that what society expects from business is not based on rigid and fixed values but rather changes over time. In his observation, businesses can no longer take values for granted, but need to put effort in their societal value alignment. This means, for instance, the permanent consultation of NGOs and other social organizations. A third dimension relates to transparency. Businesses need to open up and can no longer operate as separate, closed units in society. Driven by new information technologies, the communication channels of businesses have changed dramatically, and/or demand – as a result of the previously discussed changing values – more open insights into the intentions, strategy, mission and vision of a business. A fourth dimension relates to life-cycle technology. Whereas companies would normally focus production processes on products, a shift is required from product to function. A business does not produce gasoline, but is rather 'in' the mobility industry. This would mean that innovation is focussed on functionality, which opens the possibility to innovate sustainably, and aim for products and services that have a long or endless life cycle through recycling and reuse. In other words: production from cradle to cradle instead of cradle to grave. A fifth dimension relates to partnerships. The network and alliances of business organizations are changing. Unusual or unexpected collaborations emerge and are necessary for a networked sustainability revolution: NGOs, governments, business and education (quadruple helix) that work together and tackle societal challenges. A sixth dimension relates to time. On the one hand, (information) technologies allow us to communicate and respond within seconds to things happening anywhere in the world. This applies to e.g. social media, the way money is invested and circulated in a global financial market, the speed of online marketing and the way news and media respond to actualities. So, time

is – as Elkington puts it – ‘wider’. At the same time, our actions and our technologies have an increasing long-term impact that stretches far beyond a couple of years, but extends to generations ahead. One of the main issues here is global warming as a result of manmade climate change. This means that time has another dimension: time is not only wide, but also ‘long’. Companies need to be able to deal with both concepts of time, but have the tendency to focus on the immediate, wide, concept of time. A shift is necessary to also be able to oversee and understand the (very) long-term consequences of business decisions (think of the use of microplastics or PFAS). The last dimension Elkington observes relates to the way business organizations are governed. Corporate governance needs to shift from exclusive to inclusive governance. By this, he means that triple bottom line thinking should be incorporated in the business DNA and its governance. When sustainable issues are considered a separate issue from standard business processes, they remain an exclusive ‘extra’ on top of business as usual. As a result, a business organization will not be able to make the necessary transition towards a sustainable organization. Instead, sustainability should be at the core of decision-making processes in the boardroom, and the driving value in the product and supply chain organization.

corporate
governance

Eventually, such dimensions should encourage business transition to so-called ‘corporate honeybees’: these are regenerative companies with a high societal impact. In contrast, ‘corporate caterpillars’ are degenerative (unsustainable) organizations with a low impact. ‘Corporate locusts’ are degenerative companies with a high impact, and ‘corporate butterflies’ are regenerative companies with a low impact. Businesses need to transform into regenerative businesses, and preferably organizations with a high impact. These are characterized by a sustainable business model, a clear set of ethical guidelines, strategic sustainable management of natural resources, a capacity for sustained heavy lifting, powerful symbiotic partnerships, sustainable production of natural, human, social, institutional and cultural capital, and the capacity to help transform other businesses in their supply chain into regenerative organizations (Elkington, 2004).

corporate
honeybees

2.5.3.3 Business processes and technology development

The revolutions around industrialization and digitalization have led to major challenges in production processes. When production capacity, driven by technology, is scaled up massively, linear production processes lead to the exhaustion of basic materials, and pollution caused by waste and fossil fuels. Digital services (especially AI-based services) and processes become more widespread and require an unprecedented amount of energy, which needs to be produced and stored. In both cases, production processes that are part of a linear economy are unsustainable and will have to come to an end one way or another: either because the basic resources that are required are exhausted because they are non-renewable, or because the waste that is caused by the production and consumption of the product has nowhere to go. Various models are used to encourage the move towards a circular rather than

a linear economy. In short, a circular economy is driven by the assumption that waste no longer exists. The more eco-efficient production and consumption processes are, the more circular the economy becomes. Back in 1979 already, a model was proposed by the Dutch politician Ad Lansink (later referred to as ‘the ladder of Lansink’ or ‘the waste pyramid’) that could form a blueprint for more sustainable waste management. With some variations, the waste pyramid is now a worldwide model used for classifying sustainability levels in production and consumption processes. The pyramid is composed of three parts. At the top is ‘prevention’. The prevention of producing is always better than producing something. Avoid the use of products, and if you use them, use them only when really needed. Smart technology in smart houses that lead to energy saving are examples. The second part deals with the reuse or recycling of products. In many cases, the product lifetime is not the same as the actual period of use of a product. It therefore makes sense to organize our businesses in such a way that a product is reused with the same purpose until the economic lifespan is ended. Models around a sharing economy are good examples in which reuse forms a central part of the collective sharing and consumption of products and services (Cheng, 2016; Puschmann and Alt, 2016; Hossein, 2020). Recycling differs from reuse insofar as in recycling, waste equals basic materials. In other words, energy or product elements are used again to form new products with a different purpose. Recycling can be upcycling, when the process leads to an increase of value, and downcycling when it leads to a decrease of value. These approaches resonate in the famous work *Cradle to Cradle* (McDonough and Braungart, 2002), in which the authors suggest an economy without waste. Production processes should be design (or redesigned) in such a way that waste per definition forms the new basic materials for new products, so that there is an endless circular production cycle. This would prevent the exhaustion of basic materials, and pollution as a result of incarnation or disposal. These latter are the last stage of the waste pyramid (the bottom) which focusses on waste. Sometimes waste is incarnated for energy recovery, adding an extra loop in the process before waste is disposed of (mostly in the form of CO₂ emissions). Incarnation and waste disposal are the last two steps on the waste pyramid.

Elements of the waste pyramid are now used in the R-hierarchy, R-ladder or 10Rs model. This model is slightly more detailed and categorizes sustainability strategies in the design phase, consumption phase, end-of-life or return phase, and loss of resources phase. In the design phase, there is a focus on refuse, rethink and reduce (in the waste pyramid this would be similar to the prevention phase). The focus is on smarter product use and smarter manufacturing processes, leading to less or no use of basic materials. In the consumption phase, there is a focus on life extension strategies through reuse, repair, refurbish, remanufacture and repurpose. In the end-of-life or return phase, there is a focus on creative material application through recycling and recovering. The last phase is the loss of resources through incarnation or disposal.

The approaches proposed around ‘cradle to cradle’, the waste pyramid and the 10Rs model have a strong focus on environmental issues and how to innovate more sustainably. Other approaches have a broader reach and integrate multiple values in the design processes of new products and services. Such approaches are often labelled value-sensitive design (Friedman and Hendry, 2019), human-centred design (Kurosu, 2009; Shneiderman, 2022) or moral design (Wernaart, 2022). What they have in common is that they involve methods to internalize external (societal) values in design processes. Design processes in organizations tend to have an inward focus (Isil and Hernke, 2017); the integration of societal values at the start of design processes can help in aligning societal values and organizational values. This is a particularly useful approach when technology is complex, and its societal impact is difficult to predict.

2.5.3.4 MNOs and chain responsibility

chain responsibility

Multinational organizations (MNOs) typically operate in a multitude of countries. These countries may have different trade rules, making it worthwhile to strategically organize the business activities in such a way that the MNO maximally benefits from the most favourable rules that are available. The globalization movement after the Second World War (as discussed in section 2) helped on the one hand to create global markets with countless possibilities and positive effects on economies. At the same time it opened the doors for questionable practices among MNOs. A typical example is that of trade deflection, where the MNO enters (with its products and services) a free trade zone in the country with the lowest tariffs to trade. Another example is that of tax avoidance, where MNOs are registered in countries with the most favourable rules on taxation, optimizing their financial results by paying the lowest possible tax rate on profits.

trade deflection

tax avoidance

When states deliberately compete with the flexibility of their rules and regulations to attract more businesses, this is called ‘regulatory competition’. Some countries are famous for their mild tax regimes, and are labelled ‘tax havens’ (e.g. Wier, 2022). Such practices are usually met with criticism, where the negative consequences of not paying enough tax are pointed out. Regulatory competition does not only focus on tax and tariffs, but can also be applied on social and environmental issues (Stark, 2019). For instance, in the garment industry there is a long tradition of MNOs outsourcing the production of clothes to countries with flexible rules on employee rights and environmental laws (e.g. Gahlot and Singh, 2024). A similar tradition exists in the IT industry (e.g. Cai *et al.*, 2022).

regulatory competition

MNOs are confronted with the challenge to take chain responsibility. This usually leads to industry self-regulation, where MNOs adopt private rules that are imposed on the entire product chain to guarantee a certain level of environmental and human rights protection. The SDGs are frequently a key concept used to review business practices in the entire product chain. On the one hand, this leads to the private regulation of industries beyond the scope of individual lawmakers.

chain responsibility

On the other hand, it sometimes gives rise to chain regulation that is not always feasible to all the businesses that are part of it, especially if they are small and lack the resources to comply with the rules imposed on them by the large multinational organization that is at the top of the product chain.

2.6 Future challenges

The role of business organizations in green transitions is substantive, both as a hindering factor, and as potential part of the solution. From various academic corners and applied practices we see the emergence of models and approaches that may help business organizations to live up to societal expectations and contribute to the shift towards sustainable economies.

There are some main challenges that we need to overcome to make these models and approaches truly meaningful and to optimize their functioning.

moral
decoration

First, there is the risk of ‘moral decoration’. Popular frameworks such as the Sustainable Development Goals, or the triple bottom line, are deeply incorporated in managerial language. There is sometimes a thin line between sustainable claims and marketing. It is not without reason that in 2021, John Elkington withdrew his 3P model. He was disappointed by its misuse, and held that the systemic change he envisioned triggering with his model had scarcely been accomplished. Instead, his model was being misused as a means to green- and social washing rather than improving business performance in the field of social and ecological values.

language

Another challenge lies in the language that is used around sustainable issues. When we consider the language of law, we speak of fundamental rights and principles, or processes. When we speak in a managerial language, we speak the language of business models and accountancy. When we speak the language of technology and programming, we speak the language of maths or IT. For instance, it is hard to translate fundamental rights to AI programming, or make sure a business model is in alignment with public values. The bridging of these language gaps is deeply underexplored, and deserves much more attention if we want businesses to contribute to a more sustainable world.

multinational
companies

Last but not least, technology is in the hands of a limited number of large multinational companies that are at the top of large supply chains in both physical and digital products and services. Due to their global or near global business activities, they are beyond the reach of national legislatures, and due to the lack of serious competition, it is unlikely that other businesses, consumers or even governments can use their buying power to enforce moral or sustainable behaviour. This leads to a power gap in our globalized society that can stand in the way of sustainable development, since the choice to become a sustainable and responsible business

organization has to come from within that organization, which might feel like an unnatural response in a profit-driven entity.

References

- Abhayawansa, S., Adams, C.A., and Neesham, C., 2021. Accountability and governance in pursuit of Sustainable Development Goals: conceptualising how governments create value. *Accounting, auditing & accountability journal*, 34:923–945.
- Al-Khouri, A.M., 2013. Environment sustainability in the age of digital revolution: a review of the field. *American journal of humanities and social sciences*, 1:202–211.
- Angotti, T., 1996. Latin American urbanization and planning: inequality and unsustainability in North and South. *Latin American perspectives*, 23:12–34. <https://doi.org/10.1177/0094582X9602300403>.
- Arp, H.P.H., Aurich, D., Schymanski, E.L., Sims, K., and Hale, S.E., 2023. Avoiding the next silent spring: our chemical past, present, and future. *Environmental science & technology*, 57:6355–6359.
- Baloch, A., Shah, S.Z., Noor, Z.M., and Magsi, H.B., 2018. The nexus between income inequality, economic growth and environmental degradation in Pakistan. *GeoJournal*, 83:207–222.
- Bexell, M., and Jönsson, K., 2019. Country reporting on the Sustainable Development Goals – the politics of performance review at the global-national nexus. *Journal of human development and capabilities*, 20:403–417.
- Belitski, M., Guenther, C., Kritikos, A.S., and Thurik, R., 2022. Economic effects of the COVID-19 pandemic on entrepreneurship and small businesses. *Small business economics*, 1–17.
- Beske, F., Haustein, E., and Lorson, P.C., 2020. Materiality analysis in sustainability and integrated reports. *Sustainability accounting, management and policy journal*, 11:162–186.
- Beyl, C.A., 1992. Rachel Carson, Silent Spring, and the environmental movement. *HortTechnology*, 2:272–275.
- Binagwaho, A., Mathewos, K., and Davis, S., 2021. Time for the ethical management of COVID-19 vaccines. *The lancet global health*, 9:e1169–e1171.
- Biswas, S.S., 2023. Potential use of Chat GPT in global warming. *Annals of biomedical engineering*, 51:1126–1127.
- Brown, J., and Dillard, J., 2014. Integrated reporting: on the need for broadening out and opening up. *Accounting, auditing & accountability journal*, 27:1120–1156.
- Cai, X., Zhu, B., Zhang, H., Li, L., and Xie, M., 2020. Can direct environmental regulation promote green technology innovation in heavily polluting industries? Evidence from Chinese listed companies. *Science of the total environment*, 746:140810.
- Carson, R., 1962. *Silent Spring*. Boston: Houghton Mifflin Harcourt, Boston, MA, USA.

- Crane, A., Matten, D., Glozer, S. and Spence, L., 2019. *Business ethics* (5th ed.). Oxford University Press, Oxford, UK.
- Cheng, M., 2016. Sharing economy: a review and agenda for future research. *International journal of hospitality management*, 57:60–70.
- de Freitas Netto, S.V., Sobral, M.F.F., Ribeiro, A.R.B. *et al.*, 2020. Concepts and forms of greenwashing: a systematic review. *Environmental sciences Europe* 32:19 <https://doi.org/10.1186/s12302-020-0300-3>.
- Van Dijck, J., Poell, T., and De Waal, M., 2018. *The platform society: public values in a connective world*. Oxford university press, Oxford, UK.
- Kolev Dimanov, D., 2024. Social and human costs of industrialization: a review of some negative effects on society during the first industrial revolution. *KNOWLEDGE – international journal*, 62:123–127.
- Elkington, J., 1997. *Cannibals with forks*. Capstone Publishers, Oxford, UK.
- Elkington, J., 2004. Enter the triple bottom line. In: Henriques, A. and Richardson, J. (editors). *The triple bottom line, does it all add up?* Earthscan, London, UK.
- Epstein, L., 2014. Fifty years since silent spring. *Annual review of phytopathology*, 52:377–402.
- Eubanks, V., 2019. *Automating inequality*. Picador, New York, NY, USA.
- Evans, A.B., 1988. *Jules Verne rediscovered: didacticism and the scientific novel*. Greenwood Press, Westport, CT, USA.
- Fay, C., 2012. Domini and BP. In: Krosinsky, C. (ed.), *Evolutions in sustainable investing*. John Wiley & sons Inc, Hoboken, NJ, USA.
- Felten, E., Raj, M., and Seamans, R., 2023. How will language modelers like ChatGPT affect occupations and industries?. *arXiv preprint arXiv:2303.01157*.
- Fernandez, R., Adriaans, I., Klinge, T.J., and Hendrikse, R., 2020. *Engineering digital monopolies. The financialisation of Big Tech*. Amsterdam: Somo.
- Foster, J.B., and Magdoff, F., 2009. *The great financial crisis: Causes and consequences*. NYU Press.
- Freeman, E., 2010 (original: 1984). *Strategic management*. Cambridge University Press, Cambridge, UK.
- Freeman, E., 2010. *Stakeholder theory*. Cambridge University Press, Cambridge, UK.
- Friedman, M., 2007. The social responsibility of business is to increase its profits. In: Zimmerli, W.C., Holzinger, M., Richter, K. (eds) *Corporate ethics and corporate governance*. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-70818-6_14.
- Friedman, B., and Hendry, D.G., 2019. *Value sensitive design, shaping technology with moral imagination*. MIT press, Cambridge, UK.
- Gahlot, S. and Singh, K.D., 2024. Post-pandemic challenges of textile industry workers in India: analysis of social security laws of select Asian countries. *Law and development review*, 17:331–365. <https://doi.org/10.1515/ldr-2024-0036>.
- Haight, N., and Griffiths, A., 2009. The natural environment as a primary stakeholder: the case of climate change. *Business strategy and the environment*, 18:347–359.

- Hickel, J., 2020. *Less is more: how degrowth will save the world*. Random House, New York, NY, USA.
- Ho, O., Iyer-Raniga, U., Sadykova, C., Balasooriya, M., Sylva, K., Dissanayaka, M., Sukwanchai, K., Pal, I., Bhatia, A., Jain, D. and Sivapalan, S., 2024. A conceptual model for integrating circular economy in the built environment: an analysis of literature and local-based case studies. *Journal of cleaner production*, 449:141516.
- Hossain, M., 2020. Sharing economy: a comprehensive literature review. *International journal of hospitality management*, 87:102470.
- Isil, O., and Hernke, M.T., 2017. The triple bottom line: a critical review from a transdisciplinary perspective. *Business strategy and the environment*, 26:1235–1251.
- Jacobs, M., 1997. The environment as stakeholder. *Business strategy review*, 8:25–28.
- Khan, S.A.R., Yu, Z., Golpira, H., Sharif, A., and Mardani, A., 2021. A state-of-the-art review and meta-analysis on sustainable supply chain management: future research directions. *Journal of cleaner production*, 278:123357.
- Kumi, E., Yeboah, T., and Kumi, Y.A., 2020. Private sector participation in advancing the Sustainable Development Goals (SDGs) in Ghana: experiences from the mining and telecommunications sectors. *The extractive industries and society*, 7:181–190.
- Kurosu, M. (Editor), 2009. *Human centered design* (pp. 94–103). Springer, Berlin Heidelberg.
- Latouche, S., 2009. *Farewell to growth*. Polity, Cambridge, UK.
- Liang, W., and Yang, M., 2019. Urbanization, economic growth and environmental pollution: evidence from China. *Sustainable computing: informatics and systems*, 21, 1–9.
- Lichtenthaler, U., 2021. Digitainability: the combined effects of the megatrends digitalization and sustainability. *Journal of innovation management*, 9:64–80.
- Lytle, M.H., 2007. *The gentle subversive: Rachel Carson, Silent Spring, and the rise of the environmental movement*. Oxford University Press, Oxford, UK.
- Madsen, H., and Ulhøi, J.P., 2001. Integrating environmental and stakeholder management. *Business strategy and the environment*, 10:77–88.
- McAfee, A., and Brynjolfsson, E., 2017. *Machine, platform, crowd: harnessing our digital future*. W.W. Norton & Company, New York, NY, USA.
- McDonough, W., and Braungart, M., 2002. *Cradle to Cradle, remaking the way we make things*. North Point Press, New York, NY, USA.
- Meadows, D.H., Meadows, D.L., Randers, J., and Behrens, W.W., 1972. *The limits to growth*. Club of Rome.
- Mucchielli, L., 2020. Behind the French controversy over the medical treatment of Covid-19: the role of the drug industry. *Journal of sociology*, 56:736–744.
- Murtha, T.O., and Hamilton, A., 2012. Sustainable asset management. In: Krosinsky, C. (Ed.), *Evolutions in sustainable investing*. John Wiley & sons Inc, Hoboken, NJ, USA.
- Phillips, R.A., and Reichart, J., 2000. The environment as a stakeholder? A fairness-based approach. *Journal of business ethics*, 23, 185–197.
- Piketty, T., 2014. *Capital in the twenty-first century*. Harvard University Press, Cambridge, MA, USA.

- Piketty, T., 2021. Capital and ideology. In: Capital and ideology. Harvard University Press, MA, USA.
- Puschmann, T., and Alt, R., 2016. Sharing economy. *Business & information systems engineering*, 58, 93–99.
- Rashed, A.H., and Shah, A., 2021. The role of private sector in the implementation of sustainable development goals. *Environment development and sustainability*, 23, 2931–2948.
- Raworth, K., 2012. A safe and just space for humanity: can we live within the doughnut? Oxfam, Oxford, UK.
- Raworth, K., 2017. A Doughnut for the Anthropocene: humanity's compass in the 21st century. *The lancet planetary health*, 1:e48–e49.
- Roberts, A., 2016. The history of science fiction. Palgrave histories of literature, London, UK.
- Schubert, A., and Láng, I., 2005. The literature aftermath of the Brundtland report 'Our common future'. A scientometric study based on citations in science and social science journals. *Environment, development and sustainability*, 7:1–8.
- Shneiderman, B., 2022. Human-centered AI. Oxford University Press, Oxford, UK.
- Spreckley, F., 1981. Social audit – a management tool for co-operative working 1981. Beechwood College Ltd, Leeds, UK.
- Stark, J., 2019. 'Regulatory competition', law for sale: a philosophical critique of regulatory competition, Oxford Academic, Oxford, UK. <https://doi.org/10.1093/oso/9780198839491.003.0002>.
- Stubbs, W., and Higgins, C., 2014. Integrated reporting and internal mechanisms of change. *Accounting, auditing & accountability journal*, 27:1068–1089.
- Tamburrini, G., 2022. The AI carbon footprint and responsibilities of AI scientists. *Philosophies*, 7:4.
- Trebeck, K., and Williams, J., 2019. The economics of arrival: ideas for a grown-up economy (1st ed.). Bristol University Press, Bristol, UK.
- Trinder, B., 2013. Great cities, old and new. In: Britain's Industrial Revolution: the making of a manufacturing people, 1700–1870. Carnegie Publishing, Lancaster, UK.
- UN, General Assembly. (1987, 4 August). Report of the World Commission on Environment and Development, UN doc. A/42/427.
- Vercherand, J., 2014. The 'social question' since the 19th century. In: Labour. London: Palgrave Macmillan.
- Wells, H.G., 1896. The island of Dr Moreau. Heinemann, London, UK.
- Wernaart, B. (editor), 2022. Moral design and technology. Wageningen Academic Publishers, Wageningen, the Netherlands.
- Westbrook, D., 2009. Out of crisis: rethinking our financial markets. Paradigm Publishers, Boulder, CO, USA.
- Wier, L., 2022. 'The Dutch damage done: how tax haven activities in the Netherlands affect the world' in: Cnossen, S., and Jacobs, B. (eds), Tax by design for the Netherlands. Oxford: Oxford Academic.

Democratizing green technology with the public stack

Max Kortlander, Anne-Marie Sweep and Imme Ruarus

Abstract

In this chapter we introduce the public stack as a framework within which to root technological development in the public interest via participatory design practices and open technology. We consider existing models of citizen participation, participatory development, and design for the common good and identify the need for practical approaches to building technology in the public interest. We explore two case studies to illustrate how the public stack may be further applied in the development of green technology – particularly as a resource to involve non-technical stakeholders in technical design and governance decisions – and we provide a ‘public stack reflection tool’ to help others adopt a public stack approach in their own technical development. The public stack is proposed as a viable framework for centering development of green technology around shared public values but has not yet been widely applied as such.

3.1 Introduction

Democratic societies ought to approach the development and implementation of green technology with caution: the values underpinning technology are manifested not only in the technology itself (Haraway), but also in the physical (*the ville*) and societal spaces (*the cité*) (Sennett) in which that technology operates. In this view, where the values embedded in technology have huge impacts upon people’s lived experiences, everyday lives, and interactions with broader society, green technology is not enough. Considering democratic values and the shortcomings of private development models, and in contrast to top-down smart city development models controlled by private and/or state power, we argue that not only does the development of green technology need to be motivated by the well-being of people and the planet, but it also needs to provide people (citizens, communities, and the general public) with agency over technological development and deployment. Green technology ought to be open, public, and promote the shared public interest rather than be controlled by private or state power.

3.2 Overview of the public, private, and state stack

There is an ongoing question about how to operationalize values that we find in democratic societies – like openness, transparency, participation – into technology that is appropriately reflective of those values. The public stack considers the position of human beings and planetary boundaries in relation to technology (Kortlander and van Duijnen, 2022), and approaches technological development and implementation through four general layers:

- the foundation (those values and assumptions which underpin technological development);
- the design process (especially as relating to relative ownership, transparency, and participation);
- the technology layer itself (including hardware, software, networks, supply chains, etc.);
- people and the planet (those relating to and affected by technology).

A core concept behind this model implies that the values underpinning a technical development process – whether based on for-profit values, state power values, or shared public values – are central determinants of the impact that the technology will have on people and the planet. This happens by way of a process in which each layer informs the layer above it: the foundation informs the design process, which influences concrete technical choices; in turn, technical choices (around dilemmas like decentralization vs. centralization; anonymity vs. identifiability; open tech vs. closed tech; and user control vs. automatic processes) impact society, people, and the planet.

The model defines and contrasts the public stack in relation to two other archetypes: the private stack and the state stack. Each of these three stacks consists of the same four layers listed above. The private stack archetype has a foundation of market- and profit-based values. Its design processes and technology are closed and proprietary, and it positions human beings as consumers. The state stack archetype is based on state values. Its technology is closed and confidential, and it positions human beings as subjects. In today's reality, most technical development is closely aligned with the private stack archetype, and to a lesser (but nonetheless prominent) extent, with the state stack.

The public stack model, in response, advocates for the uptake of a public stack archetype to the extent possible. The public stack archetype has a foundation which is based on shared, public values; utilizes open and participatory (co-creative) design processes; uses and develops open-source technology; and positions people as human beings with fundamental rights and democratic agency, while positioning the planet as a more-than-human commons (see figure 3.1 and 3.2).

In practice, no piece of technology or technological initiative fits neatly into any one of these archetypes. Instead, technology is often a mix of all models – any given piece of technology likely embodies aspects from each archetype throughout its

private stack
archetype

state stack
archetype

public stack
archetype




FIGURE 3.1 Private stack, state stack, and public stack illustrations

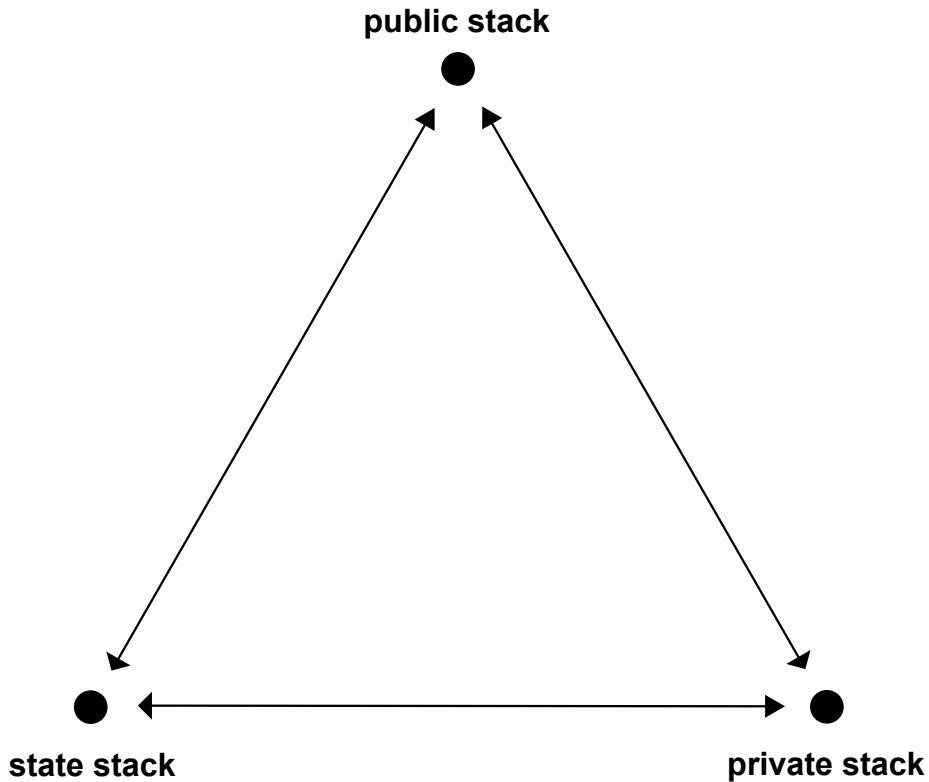


FIGURE 3.2 Public, private, and state stack continuum

various layers. For example, the Fairphone is arguably based on public, shared values (like openness and sustainability), was initially funded through public (state) funds, and currently operates with a business model in a market of other phone providers. In this way, the Fairphone is neither fully in the public stack, nor fully in the state or private stack. It *is*, however, more closely aligned with the public stack than the iPhone, which is more closely aligned with the private stack.

The public stack, private stack, and state stack archetypes are displayed on a continuum. No single project or technology fits neatly into any single archetype, and most would be positioned somewhere in the middle area of this continuum. The public stack model encourages development to move closer to the top of the triangle through design processes and technology that are aligned with public values.

How can developers and facilitators of participatory processes utilize these different archetypal lenses (public stack, private stack, and state stack) to operationalize the development of green technology? Most immediately, the public stack lens raises questions at each 'layer', which can assist developers and other stakeholders (ideally also policymakers and citizens/humans) to reflect upon the myriad components that affect the way people are positioned in relation to that technology. We consider how the public stack was developed and utilized in two projects contributed to by Waag Futurelab to illustrate how the model works in practice.

3.3 Literature overview

3.3.1 *Embedding public values and planetary boundaries in digital infrastructures: a call for participatory design*

Smart cities often implement technology with large-scale investments in ambitious digital infrastructure projects that raise concerns about the lack of transparency and privacy around the technologies that shape public services and urban life (Calzada, 2020a). Private, proprietary, and closed models tend to prioritize profit over social and/or environmental impact, leading to a misalignment with the common good. The inherent shortcomings of these models prompt a critical examination of alternative approaches that can better serve the broader community.

Living in democratic societies arguably necessitates a design approach for new (green) technologies that reflects and embeds public values. The challenges of climate change and environmental degradation further underscore the importance of aligning technological developments with societal needs to stay within planetary boundaries (Rockström *et al.*, 2009; Richards K *et al.* 2023). The need arises for an approach that not only acknowledges public values but also integrates them into the very fabric of technical design processes.

To better ensure this in the future, there is a call for a more participatory approach to technology development that engages both policymakers and citizens alike in the development of the 'smart city'. There is a need to level the playing field between citizens, policymakers, and tech: non-technical stakeholders like citizens may lack technical expertise but possess crucial insights into the societal impacts of technology in everyday life. Engaging the broader public in the technology design and implementation trajectory becomes not only a smart step in the design process, but ultimately a shared, value-based responsibility. In this light, if we fully recognize the complexity of challenges such as climate change and biodiversity loss,

we must ensure that the development and implementation of technology are both inclusive and considerate of these diverse perspectives.

In order to move forward, there is a need to further develop practical frameworks that empower all stakeholders to actively participate in shaping value-based, smart city technology for the common good.

3.3.2 *Citizen participation and democracy: Navigating the dynamics of direct democracy*

Citizen participation in democracy extends beyond the act of voting, encompassing a spectrum of active involvement of citizens in decision-making processes. In doing so, it unravels the dynamics surrounding citizen engagement in relation to democracy, to provide context for participatory design principles.

The broader context of democracy plays a vital role in understanding how individuals can actively contribute to shaping technology with the potential for participatory design to transcend its traditional boundaries. A common critique regarding technological developments and implementations in smart cities is that the driving force behind them is often not citizenship or democracy but rather developers seeking self-interested acknowledgement of the advantages of user-centred design to effectively execute digital technology projects. This raises concerns regarding control, representation, participation, and democracy when *'Citizens are to be steered, nudged, controlled; they can browse, consume, and act. If there is civic engagement, it is in the form of a participant, tester, or player who provides feedback or suggestions, rather than being a proposer, co-creator, decision-maker, or leader'* (Kitchin et al., 2019: 11).

Roberts (2003) provides a comprehensive theoretical framework for how this unwanted situation can be transformed into direct citizen participation, highlighting its vital role in community life and public accountability. However, scepticism exists, with representative democracy seen as a safeguard against potential pitfalls such as uninformed public opinion and the tyranny of the majority. This tension is rooted in the need for a balance between citizens' right to engagement and the complexities of modern, highly technological societies where decision-making often requires a high level of expertise. Barber (1984) and Box (1998), as cited by Roberts, emphasize the ambivalence surrounding citizen participation. They argue that while citizenship participation is the cornerstone of democracy, direct involvement is often viewed with scepticism. Representative democracy is in this light presented as a protective mechanism, buffering citizens from the potential dangers associated with direct engagement, such as corruption and uninformed decision-making. It should also be noted that citizen involvement is not monolithic, and the capacity and luxury to be involved in public decision-making processes varies greatly between individuals and communities.

Desario and Langton (1984) trace the evolution of citizen demands for greater participation – particularly in complex technological issues – back to public unease

direct citizen
participation

about technology and expert power. They note the shift in public sentiment following incidents like Love Canal, Three Mile Island, and Agent Orange. This shift led to an increased demand amongst active citizens for deeper involvement in issues such as air and water management, hazardous waste control, nuclear power, and DNA research. Evaluating the effectiveness of direct citizen participation in coping with wicked problems, and drawing on Robert's insights, we argue that participatory processes offer a unique opportunity for meaningful dialogue and deliberation and are valuable despite their time-intensive nature. Criticisms about efficiency in crisis situations may not hold when considering the long-term costs and consequences. In the end, according to Smith and Martin (2022), 'technology, built through participatory social relations, is key to smart city platform success, because digital platforms for urban democracy are – like democracy itself – a perpetual work in progress' (Smith, A., and Martín, P.P., 2022).

3.3.3 *Citizen participation in smart cities: navigating from top-down towards community-up approaches*

Increasingly, cities like to consider themselves 'smart cities', defined by David and Benson as a city that uses a '*citizen-centric urban development and governance strategy facilitated by technology to improve, overall, sustainability*' (David, N.P. and Benson, T.S., 2021).

Many smart cities have to bridge the gap between being both technology driven and simultaneously embedding public values. Smart cities provide new solutions in the domains of mobility, environment, economy, governance, quality of life, and education thanks to the innovative use of Information and Communication Technologies (ICT). Simonofski *et al.* (2017), however, argue that smart cities often do not optimally reach their objectives if the citizens, the end users, are not involved in their design (Simonofski *et al.*, 2017).

Smart city development has historically favoured top-down corporate models, prompting a growing acknowledgment of people's 'right to the city' and the necessity to transition towards bottom-up approaches. In the increasingly complex landscape of digital infrastructures in smart cities, it is crucial to embed public values of the people living in those cities into the design and implementation process (Kresin *et al.*, 2020).

In the case of operationalizing citizen participation in smart cities, Lim *et al.* (2018) assert that despite the rhetoric of citizen-centric smart cities, citizens often remain passive users, with their needs merely addressed rhetorically. Criticisms have been directed at technology-driven smart cities, emphasizing the neglect of citizens. Questions surrounding how to practically involve citizens in democratic processes persist, necessitating a clarification of their roles in participatory processes (Lim *et al.*, 2018).

Smart city development is a domain where technology design intimately intersects with urban life. It demands infrastructures that enable citizens to actively

citizen
participation

smart city
development

contribute to the development of ‘their cities’, again emphasizing the need for inclusive practices that consider the diverse values and perspectives present in the living and evolving community of a (smart) city.

Alizadeh (2018) critically compares crowdsourced smart cities with corporate smart cities, highlighting the deficiencies in existing smart city theories and frameworks. The crowdsourced approach champions a democratic bottom-up strategy, dynamic public-private partnerships, and an emphasis on participatory governance. It challenges the corporate vision, promoting social inclusion and a shift in power from corporations to communities. However, concerns arise regarding the lack of attention to implementation in alternative smart city visions (Alizadeh, 2018).

Addressing the need for a more level playing field between citizens and experts, the smart citizens’ approach prioritizes long-term participation ‘in which local residents, organizations, and public administrations co-create bottom-up solutions to ensure that the use of technology in the city is guided and owned by the citizens themselves’ (Kresin *et al.*, 2020). This involves a community of citizens, experts, and administrators physically building technology together over an extended period. While the prolonged co-creation process may be perceived as a limitation, it aligns with Roberts’ assertion that meaningful participatory processes are particularly useful for tackling complex and wicked problems.

The term ‘bottom-up’ to describe community involvement in fields like democratic decision-making and participatory design might not be the preferred term as it fails to capture society’s multi-layered social fabric, and implies an unequal hierarchy. The term ‘community-up’ may be preferable when referring to citizen participation/initiatives. This semantic adjustment better aligns with the ‘smart citizens’ approach to technological development in cities, as described by Kresin *et al.* While promising strategies such as the smart citizens’ approach, crowdsourcing and other grounded theory approaches have emerged, challenges in operationalization and implementation of these community-up approaches to citizen participation persist, especially amidst rapid technological advances.

3.3.4 *Existing frameworks, theories and models for participatory design based on shared values*

Key Enabling Methodologies (KEMs), Value Sensitive Design, and Human-Centred Design, are existing models for participatory design. KEMs have a role in breaking down barriers and enabling diverse participation. Value Sensitive Design integrates human values into the design process. And Human-Centred Design emphasizes user needs and incorporates them into the technology development cycle. These existing frameworks, theories, and models contextualize the public stack’s role and the added value of the public stack within the ecosystem of participatory design. Key Enabling Methodologies are a ‘toolbox, or set of instruments, consisting of methods, models, strategies, processes and tools. This includes ways of

smart citizens

bottom-up

community-up

Key Enabling
Methodologies

working (together), dealing with problems and creating interventions; tools with which ‘change’ professionals, such as designers, policymakers or public administrators, are able to structure their work, give direction and realize impact’ (ClickNL). KEMs may be considered a collection of tools that can help designers in a process – specifically, in processes to address complex and wicked societal challenges. As their website describes, ‘Social challenges are without exception complicated: they affect people, society, the environment and cover several disciplines at the same time.’¹

The origins of KEMs lie in (mission-driven) Innovation Policy, through which the Dutch government aims to further develop the innovative power of the Netherlands in order to tackle social challenges jointly and thoroughly, and emphasize social inclusion in *technical* design processes.

As with many participatory frameworks when put into practice, the context in which KEMs are used determines their impact. A top-down application of KEMs risks placing the implicit leading role on designers to engage society. Designers using KEMs ought therefore to heed the methodologies’ open and participatory approaches to ensure that societal needs drive design and development. Nonetheless, KEMs have the potential to address complex societal challenges if mindfully applied to place societal and community needs at the centre.

Many of the KEMs fit into the public stack layer of ‘design processes’ as part of the open, participatory, and co-creative design process that it advocates. While not officially classified as a KEM, and arguably not a methodology, the public stack shares many similarities with KEMs and may one day develop into one.

The public stack also shares commonalities with human-centred design and value-sensitive design.

Human-centred design, as its name implies, puts the needs and wants of human beings at the centre of a design process. As van der Bijl-Brouwer and Dorst describe:

Human-centred design (HCD) is a group of methods and principles aimed at supporting the design of useful, usable, pleasurable and meaningful products or services for people. The main principle of these methods is that they describe how to gain and apply knowledge about human beings and their interaction with the environment, to design products or services that meet their needs and aspirations. (van der Bijl-Brouwer and Dorst, 2017)

Value-sensitive design can be considered a subset of human-centred design, in that value-sensitive design ‘is a theoretically grounded approach to the design of

human-centred
design

value-sensitive
design

¹ <https://kems-en.clicknl.nl/introduction-to-the-agenda/1-introduction/1.2-key-enabling-methodologies-or-kems>.

technology that accounts for human values in a principled and comprehensive manner throughout the design process' (Friedman *et al.*, 2006). As this definition notes, value-sensitive design places human beings at the centre of a design process, but goes beyond aspects like user preferences, needs, and wants and instead places an emphasis on values: value-sensitive design has an inherently ethical component.

A potential limitation of both human-centred design and value-sensitive design is that they do not necessarily prioritize environmental concerns prescriptively (although these models may address environmental concerns if they are a priority of the participatory community at hand). These models may also be used in conjunction with life-centred design (Borthwick *et al.*) or a more-than-human perspective (Tsing *et al.*). The public stack takes such a combinatorial perspective that considers the impact of technology on both people and the planet.

Value-sensitive design draws from participatory and co-creative methods in human-centred design to identify what values ought to drive a development process. Friedman *et al.* describe this identification of values as including periods of 'conceptual, empirical, and technical investigations'. Generally, the conceptual phase scopes the values of direct and indirect stakeholders. The empirical phase utilizes 'the entire range of quantitative and qualitative methods used in social science research including observations, interviews, surveys, experimental manipulations, collection of relevant documents, and measurements of user behavior and human physiology' to answer questions like 'How do stakeholders apprehend individual values in the interactive context? How do they prioritize competing values in design trade-offs?' Finally, a technical investigation phase 'focus[es] on how existing technological properties and underlying mechanisms support or hinder human values' (Friedman *et al.*, 2006).

Value-sensitive design embodies a participatory approach to design, which seeks to be proactive by 'influenc[ing] the design of technology early in and throughout the design process' (Friedman *et al.*, 2006). Notably, the basis for defining values in value-sensitive design is relatively open-ended, and dependent on the outcomes of a participatory design process.

As we will describe in the following chapter, the public stack approach shares commonality with value-sensitive design, and could perhaps be viewed as a form of value-sensitive design. Importantly, however, the public stack is more prescriptive – values like openness, fairness, and inclusivity are prescriptive to the public stack. This includes an embedded preference for open source technology and privacy. Value-sensitive design does not have such rigid starting points; while co-creation and participatory approaches are used to define values from scratch in value-sensitive design, the public stack approach uses such methods to further define prescriptive values and additional shared values, and understand the ways in which both identified and prescriptive values will be embedded in both the design process and the technology itself.

3.3.5 *Conclusions of literature review: towards technical development for the common good*

These participatory approaches reviewed thus far acknowledge the dominance of private, proprietary, and closed models in technology development, and raise concerns about their efficacy in delivering for the common good. Acknowledging this, the technology development landscape has seen a growing recognition of the need for frameworks that prioritize the common good.

Waag's Public Stack model acknowledges the limitations of private development models and proposes a transformative approach. It defines a foundation for technical development processes based on public values and shared goals. This foundation can then become a guiding force throughout the development process, influencing decisions about design and technology for smart cities. The emphasis on revisiting and adhering to public values ensures a continuous alignment with the common good, creating a robust framework for sustainable development. Both the platform for the common good and Waag Futurelab's public stack model share a common theme: the maximization of social impact. This unifying goal stands out as a much-needed alternative to profit-centric approaches, emphasizing the importance of technology development in contributing positively to society.

In summary, there is a necessity to transition from private development models to those prioritizing the common good. The platform for the common good and the public stack serve as examples of this type of transformative way of approaching technology, offering strategies that maximize social impact and address the deficiencies of technocratic government or profit-centric approaches. By adopting these models, the technology development landscape can shift towards a more inclusive, sustainable, and socially responsible future.

3.4 Case Studies: Hollandse Luchten and ACROSS

The public stack has been developed iteratively through the contributions of several projects at Waag Futurelab. Illustrations and general descriptions of the public stack were developed as part of the Digital European Public Spaces project (Waag Futurelab, Digital European Public Spaces). These illustrations and descriptions were found to be useful tools in subsequent projects when communicating to technical and non-technical experts alike about the need and potential to develop technology based on public values. The two case studies below demonstrate how the public stack was developed and used in subsequent projects at Waag Futurelab.

The Shared Cities Smart Citizens project utilized the public stack as a reflection tool to reflect upon and influence the Hollandse Luchten citizen science project. The ACROSS project developed the public stack further by implementing it as a design approach and formalizing its use as a reflection tool via the development of public stack reflection cards.

reflection tool

design approach

3.4.1 *Application of public stack in Hollandse Luchten*

'Hollandse Luchten' (Dutch Skies) is a citizen science project where citizens measure air quality in their environment in the Province of North Holland through the use of open source and affordable sensor technologies (Waag Futurelab, Hollandse Luchten). In this project, Waag Futurelab, in collaboration with the Province of North Holland, local municipalities, and the RIVM (the Dutch national environmental agency), facilitates a process to promote citizens' insight into and agency over the air quality in areas where the quality of life is under pressure.

In 2022 the Hollandse Luchten project was evaluated through the research project Shared Cities Smart Citizens using the public stack as a reflection tool in internal (Waag team) and external (with community and governmental stakeholders) collaborative sessions. The goal was to co-creatively evaluate the extent to which choices made in the project were in line with public values and the values of the community. At this stage in the project, the public stack model was utilized as a reflection tool to provide a basis for further decision-making in Hollandse Luchten and to level the playing field amongst the project's various stakeholders.

Hollandse Luchten had been underway for roughly three years at the time this evaluation took place. Core attributes of the project's direction – values, participatory communities, the development and deployment of open-source sensors – had already been well established. The layers of the stack were formed into specific questions for the Hollandse Luchten initiative:

- Foundation: What values and goals are most widely shared by participants? By facilitators?
- Design: Does Hollandse Luchten's design process promote the values and goals in the project's foundation? How might specific aspects of Hollandse Luchten's participatory process change to better further shared aims?
- Technology: Does the technology help to further the aims of participants? Is the technology in line with shared values like openness and sustainability?
- People and Planet: What impact does Hollandse Luchten have? How does this align with or deviate from the values and goals established in the foundation? How could the impact be better aligned with the values and goals identified in the project's foundation?²

The reflection process consisted of three steps: citizen participant interviews, an internal co-creation session, and an external co-creation session. The goal of the citizen participant interviews was to identify use cases and goals the participants had for using the data gathered through the Hollandse Luchten project. These so-called 'data goals' are:

2 <https://shared-cities-smart-citizens.nl/hollandse-luchten-mapping-use-cases-with-the-public-stack/>.

- Community knowledge sharing – Hollandse Luchten data can be used to help inform people in pilot communities about air quality and how it is measured.
- Policy change – Hollandse Luchten data can be used to apply pressure to policy-makers and advocate for policy change.
- Personal insight – Hollandse Luchten data can help an individual to be more informed about the air quality and make decisions accordingly in their own living environment.
- Transparency and accountability – Hollandse Luchten data can be used to check and contribute to existing air quality measurements (for example from Tata Steel and RIVM) towards an accepted shared reality about air quality in the region and keep the polluters and/or the government accountable. (Kortlander and van Duijnen, 2022).

Following the interviews with citizen participants, facilitators held an internal session to uncover potential conflicts in reaching values and ‘data goals’. This internal reflection session highlighted several areas to be addressed with other stakeholders in the project, including (bold emphasis added):

- the need to build common ground between project stakeholders – The facilitators write: “... stakeholders tend to refer to different data, accept different facts, and thus lack a common ground and shared reality from which to begin collaborating towards positive change. Most generally, the problem will be solved when all stakeholders – including people in the communities, governmental organisations, and Tata Steel – have a shared reality (in the form of data, analysis, and experience) to which they can refer. Success in this regard ought to be primarily defined by community participants.”
- the question of how to practically “place local people on a more equal playing field along with those in government and industry”, specifically via official governmental decision making processes.
- the dilemma that “participants would like to draw definitive conclusions based on the data, while experts in other stakeholder groups emphasise the uncertainties and vagueness inherent to this and other data analysis” (Kortlander and van Duijnen, 2022).

The exercise focused the facilitation team’s attention on specific steps to prioritize in the immediate and long term. In the short term, the team identified the need to host a similar reflection session along with citizen participants, to understand what processes (*design layer*) and technology (*tech layer*) they would need to make their data goals a reality.

Following the identified short-term action, a reflection session was organized with citizens participants, experts, and policymakers. The goal of the session was to identify which processes (*design layer*) and technology (*tech layer*) they would need to make each of the four data goals a reality and what impact the data goal would have on people and the planet.

The session led to numerous findings, notably that there is a dilemma between openness and accuracy in the technology used in the project. Stemming from the transparency and accountability goal, the following was concluded: ‘Everyone – from Tata Steel to RIVM to local residents – should be able to draw shared conclusions from the available data and the underlying motivations of all involved should be public knowledge.’³ However, at this time, the data collected by the open source sensors used in the project were only indicative; they contributed to knowledge development and formed the basis for a dialogue, but had no legal value. This uncovered a dilemma about which value to prioritize in this instance – whether to utilize a sensor that was less accurate but more open source, or to utilize a sensor that was more accurate but less open source. In this instance, citizens voiced their opinion to prioritize a more reliable sensor over an open sensor, with the goal of making *Hollandse Luchten* more robust for drawing conclusions.

3.4.1.1 Outcomes of the public stack reflection in *Hollandse Luchten*

Through the use of the public stack as a reflection tool, *Hollandse Luchten* could be thoroughly evaluated with the input and perspective of stakeholders. A stated goal of *Hollandse Luchten* is to level the playing field between citizens, industry, and the government. The public stack reflection process helped non-technical stakeholders take part in technical decision-making in the process. These non-technical citizen stakeholders are also non-political stakeholders; by pulling apart various layers of policy at play, this process also helped citizens to take part in political decision-making discussions around air quality and the environment. Furthermore, the reflection uncovered insights into the strengths and weaknesses of the project in relation to the data goals in its foundation and consequently the further development of the project.

Three concrete steps that were carried out based on the evaluation are:

- The development of a data-analysis platform based on community needs;
- The organization of workshops about how citizens can influence local policies;
- The quality of the sensor and its data is continuously researched and calibration of the sensor was carried out at the end of 2023 by RIVM to improve the quality of the data.

The reflection process also made decisions clear, such as the decision to move to a less open but more reliable sensor, which was based on the community need to deliver higher data quality. This part of the process also helped facilitators and citizens alike to be aware of conflicts to be addressed in the future. The need for a sensor that is both relatively open *and* relatively reliable was identified and can now be communicated to experts and developers. Note that at the time of writing,

3 <https://shared-cities-smart-citizens.nl/hollandse-luchten-mapping-use-cases-with-the-public-stack/>.

even with the more accurate sensor, community-gathered data in the Hollandse Luchten project is still not considered a suitable set of data on which to base decisions regarding air quality policies and actions by governments, because juridically significant data must meet strict requirements, for example measurement uncertainty may not exceed 25% (European Parliament and European Council). This forms a challenge for citizens wanting to improve their living environments, as they are limited to using the data for certain purposes like raising awareness of issues and gaining community insights.

A limitation of the public stack in the way it was used together with the community is its potential inaccessibility for citizens and other non-technical stakeholders. Some of the concepts of the public stack are quite technical and do not align with how citizens experience their interaction with technology. This issue was mitigated by framing the activity in concrete (non-academic) terms that related directly to the needs of participants. In order to minimize confusion and jargon, facilitators did not explain or present the public stack as a model; rather, participants were asked the straightforward questions listed above in response to their respective data goal. The exercise identified outstanding issues and steps to be addressed for each data goal.

In summary, using the public stack as a reflection tool in an ongoing project around technology and citizen agency has helped to provide clarity on the different goals of the community and level the playing field by including technical and non-technical stakeholders more equally in design decisions. With the insight into these goals, future choices could be made that are in line with the needs of the community, while raising awareness of conflicting values and goals. Furthermore, the public stack as a reflection tool enabled an inclusive dialogue on the direction of the project, and thus functioned as a tool for participatory decision-making.

3.4.2 *Application of public stack in ACROSS*

ACROSS is a project funded by the European Commission to develop workflows that support specific user journeys in cross-border mobility (e.g. a student moving from Greece to Germany). ACROSS developed an open source and largely decentralized platform that connects people who move across borders with necessary service providers. Waag Futurelab is a project partner, whose role involves the development of a governance approach and methodology.

ACROSS developed and employed a public stack approach from the project's outset to shape the project's governance approach and methodology. Project partners refined the public stack as an analysis tool (building from its application in Hollandse Luchten/Shared Cities Smart Citizens), and also applied the public stack more intentionally as a design approach throughout the project's duration. Project facilitators describe it as follows:

analysis tool

design approach

ACROSS has adopted and further developed a public stack governance model and service design approach. Most broadly, this entails a process that begins with identifying a values-based foundation for technical development; implementing an open and participatory design process; developing technology that is itself open, fair, and inclusive; and culminates in technology and digital services that are aligned with public values, respect planetary boundaries, and position people (users and non-users alike) as citizens with democratic agency (rather than as consumers or subjects). (Kortlander and Hoefsloot, 2024)

The public stack approach in ACROSS involved first defining the foundation through a series of co-creation sessions with the core project team and external stakeholders. ACROSS partners summarized the values in its foundation as:

- Interoperability, Functionality, and Technical Completeness
- Citizen Control and Privacy
- Trust and Openness
- European values, laws, and ethical guidelines as expressed through the ECHR and relevant legislation such as GDPR (Kortlander and Hoefsloot, 2024).

Once every six months (on average), partners revisited this foundation during each general assembly of the project consortium. Partners iterated and further defined the project's foundation, reflected on the ways in which development over the subsequent six months would be in line with the foundation's values and goals. Kortlander and Hoefsloot document the steps in this process:

- Identified values (with the project consortium).
- Identified gaps in cross-border services (through interviews and research) and validated those gaps (with external stakeholders from industry, academia, government, and the general public through co-creation sessions).
- Refined the public stack reflection process with the consortium and set legacy goals based upon the gaps and values which formed the foundation.
- Refined the public stack reflection process with external stakeholders from industry, academia, government, and the general public through co-creation sessions.
- Identified key design dilemmas where technical decisions had to balance various requirements and values.
- Conducted a mid-project reflection exercise with the consortium.
- Iterated and finalized this reflection exercise into the 'Public Stack Reflection Cards' which are intended for use by other development teams, especially those working on public technology (Kortlander and Hoefsloot, 2024).

The process was more cyclical than linear, as the foundation was unfixed and refined throughout the processes. Nonetheless, these cycles generally moved from

the ‘inner core’ outwards: co-creation began with the consortium and development team first, then expanded to include external stakeholders, before repeating this cycle again as needed.

3.4.2.1 Outcomes of the public stack design approach in ACROSS

ACROSS partners identified the following outcomes of using a public stack design approach:

- Alignment around shared and public values with an emphasis on ensuring that ‘values identified earlier in the project (protection of digital identity, control and privacy over personal data, data minimisation) are adhered to via ACROSS’s design process and technical development’.
- Identification of contextual limitations for cross border services in utilising certain types of PETs (privacy enhancing technologies) in particular cases which are ‘reliant on service providers’ willingness and capacity to adopt certain practices or technologies’.
- Identification of challenges to address including the goal to facilitate ‘revocable consent for the use of personal data’.
- Provision of a guiding light to follow when faced with design dilemmas. Examples of such design dilemmas include centralisation vs decentralisation, as ‘[w]orkflows are centralised while much of the system architecture is decentralised’; and identifiability vs anonymity, as the ‘[u]se of credentials in a wallet helps to identify necessary aspects of a user while allowing them to maintain as much anonymity as is possible in a given information exchange’. (Kortlander and Hoefsloot, 2024).

As was the case in *Hollandse Luchten*, the public stack as a governance and service design approach in ACROSS allowed non-technical stakeholders and non-technical project partners to contribute to technical design decisions. By separating the project into several layers, and by considering the impact that design decisions could have on people and the planet, this approach allowed technical and non-technical partners and stakeholders to equally identify design choices, discuss their reasoning, consider their impact, and ultimately converge upon design decisions with consensus. The ‘identification of challenges’ similarly took place as a result of the public stack approach in both ACROSS and *Hollandse Luchten*. However, the provision of a ‘guiding light’ was unique to the public stack’s use in ACROSS and provided continuity and uniformity over the course of several years. Such a guiding light developed from the establishment of a clear foundation, which was then referred back to when facing design challenges and dilemmas.

Project partners in ACROSS view the public stack as a worthwhile lens for guiding a development process, as evidenced by their development of the Public Stack Reflection Cards, which make this process accessible for developers of other technology. The public stack’s use and development into a design process in ACROSS was iterative, experimental, and co-creative: its particular form was not defined at

the beginning of the project, but rather took shape over the course of three years. A few of these cards are available to be cut and used from the following pages of this book. The full set is open access and can be downloaded and printed from <https://publicstack.net/cards/>.

How to use the public stack as a design tool

- In the early stages of a project establish a clear foundation that identifies the shared values and goals of project partners and stakeholders.
 - co-create from the core outwards – first with your core team, then with the full development team, and ultimately with other stakeholders (citizens, public administrators, academia, and industry).
 - At each stage, identify shared requirements, goals, or values that already bind your development team. For example, teams developing public technology in Europe can look to values in the ECHR and European Digital Strategy as a starting point.
 - Utilize co-creation methods to identify further shared goals and values of your development. <https://ccn.waag.org/navigator/theme/direction> offers a few practical methods, along with instructibles on how to use them in your own context.
- Iteratively and repeatedly return to those values and goals. Do they still hold true? Has the project's recent progress furthered those values and goals? Do our plans for upcoming progress further those values and goals?
 - ACROSS partners returned to the reflection questions we asked our consortium over the project's duration and curated them into a set of questions that are applicable for other development teams, in particular (but not only) those developing public technology.
 - Use the Public Stack Reflection Cards to prompt this discussion during a team meeting. Rinse and repeat as needed throughout development.
- Use the shared values and goals you've identified as a guiding light when faced with design dilemmas.

3.5 Conclusions: lessons learned, recommendations, and the application of the public stack to green technology

The public stack model emphasizes the human side of technology and views the world from the perspective that technology is neither good nor bad, but operationalized. Since technology can be used for 'good ends' or 'bad ends', in the 'right way' or the 'wrong way', society needs to have a role in developing it, scrutinizing it, and making decisions about when and how it is used.

Society is not primarily composed of technical experts. The public stack and its layers help people – technical and non-technical stakeholders alike – to

become aware of various (often unseen) factors that determine the impact that technology has on their community and the planet. Its organization of complex socio-technological processes into graspable components helps people consider which values are embedded throughout the technology and design process layers – and reflect on the extent to which the technology and design processes are (or are not) in line with foundational shared values and intentions. From there, people may consider the extent to which a technology is emblematic of the public stack, private stack, or state stack.

In comparison to other models, theories, and frameworks around citizen participation in development processes, the public stack may be considered as normative or prescriptive – it is a somewhat aspirational model in that it implies what technology should be, rather than assessing what technology often is. It posits that technology should be based on public values, developed through co-creation, open source, and improve the well-being of people and the planet (to the fullest extent possible). In doing so, the public stack does not prescribe one ‘right way’, while challenging developers and practitioners to ‘move closer’ to the public stack archetype by developing based on these values in addition to those identified by citizens and communities through a co-creative design process.

The public stack has not yet been widely applied to green technology. It remains to be seen how green values around sustainability, carbon emissions reduction, and resource protection will resonate with development teams in balance with democratic values like openness, equity, and inclusive citizen participation. This is a much-needed debate that shines a light on two valuable perspectives of the development of technology. As demonstrated by the conflict between openness and accuracy of sensor data quality in relation to environmental enforcement in the *Hollandse Luchten* case study, the public stack can help communities to identify and address such conflicts with awareness and intention. This can ultimately support them in making decisions in line with their needs and values.

The public stack embraces certain conundrums. It is aspirational, yet open ended; normative and prescriptive in its championing of open technology and design processes; yet it embraces the unknown which arises through citizen participation. In the development of democratic green technology, the public stack leaves it to the discretion of developers and participants as to where they will draw the line when balancing difficult decisions about conflicting values. As such, the public stack implicitly acknowledges the necessity of wiggle room in values-driven technical development processes.

The public stack’s embrace of nuance and imperfection is a nod to its humanness. The model helps human beings to put ourselves (along with the planet and its ecosystems of which we are a part) at the centre of technical design. It helps people to organize and parse the complexities of technical development to make them understandable by and relevant for citizens and policymakers who are not technical experts but nonetheless have a vested interest in the social and environmental

impacts of technology. By placing shared values in the driver's seat, it pushes developers to relinquish control and cede power to citizens and communities. This is not only a matter of political ethics, but also one of practicality, as it encourages people to design technology that is relevant for society.

As demonstrated by the case studies, the public stack is useful both as a reflection tool and design approach that helps developers and communities to better understand why certain design choices are made and to prioritize which issues to address next. Building on these findings, we see the potential and value of the public stack best represented when used as a design approach. This application ensures the use of the public stack from the outset and throughout a project's course, to help consistently root design decisions in shared values. As we increasingly develop and use technology to address shared and complex problems like climate change and resource depletion, the public stack provides a necessary design perspective on top of greenness and sustainability to also allow room for considerations around human, civil, and democratic rights and values.

Further research can help to clarify the public stack's operationalization at the intersection of the physical, digital, and social city. In this case, the question of funding remains of primary importance: How ought public stack development to be funded? Ideally, all technical development, regardless of funding, ought to move towards alignment with public values, but the question of how to incentivize this movement towards the public interest remains a complex topic for further deliberation. Furthermore, the public stack relies upon shared values, but society is not a monolith. Different individuals will be affected in different ways, and undoubtedly have different priorities – again, there is often no single 'right answer' to be identified, and consensus will not be found around every case. It remains to be seen how a public stack approach might scale to work for larger and more diverse communities of citizen participants and non-technical stakeholders. Such questions will continue to be addressed and posed as part of the ongoing pursuit to include society in its own decision-making processes.

References

- Alizadeh, T., 2018. IOP conference series: earth and environmental science, Volume 158, The 1st ITB Centennial and 4th PlanoCosmo International Conference 3–5 April 2018, Bandung, Indonesia. <https://iopscience.iop.org/article/10.1088/1755-1315/158/1/012046/meta>.
- Van der Bijl-Brouwer, M., and Dorst, K., 2017. Advancing the strategic impact of human-centred design. *Design studies*, 53:1–23.
- Borthwick, M., Tomitsch, M., and Gaughwin, M., 2022. From human-centred to life-centred design: considering environmental and ethical concerns in the design of interactive products. *Journal of responsible technology*, 10:100032.

- Calzada, I., 2020. Democratising smart cities? Penta-helix multistakeholder social innovation framework. *Smart cities*, 3:1145–1172. <https://doi.org/10.3390/smartcities3040057>.
- Cardullo, P. and Kitchin, R., 2019. Smart urbanism and smart citizenship: the neoliberal logic of 'citizen-focused' smart cities in Europe, *Environment and planning C: politics and space* 37: 813–830.
- David, N.P., and Benson, T.S., 2021. Citizen-centrism in smart cities: reality or rhetoric? Holistic approach for decision making towards designing smart cities, 245–266.
- DeSario, J., and Langton, S., 1984. Citizen participation and technocracy. *Policy studies review*, 3(2). <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1541-1338.1984.tb00116.x>.
- European Parliament and European Council, 2008. Guidelines: regarding air quality and cleaner air in Europe.
- Friedman, B., Kahn, P., and Borning, A., 2002. Value sensitive design: theory and methods. University of Washington technical report, 2:1–8.
- Haraway, D., in Kunzru, H., 1997. You are cyborg. *Wired Magazine*, 5:1–7.
- Kortlander, M., and van Duijnen, S., 2022. A public stack approach to governance in the Hollandse Luchten citizen science initiative. *Shared cities, smart citizens*. Waag Future Lab, Amsterdam, the Netherlands.
- Kortlander, M., and Hoefsloot, M., 2024. ACROSS Governance framework including service design approach, Horizon 2020 Project ACROSS. H2020-SC6-GOVERNANCE-2018-2019-2020.
- Lim, C., Kim, K.-J., and Maglio, P.P., 2018. Smart cities with big data: reference models, challenges, and considerations. *Cities*, 82:86–99. <https://doi.org/10.1016/j.cities.2018.04.003>.
- Malek, J.A., Lim, S.B., and Yigitcanlar, T., 2021. Social inclusion indicators for building citizen-centric smart cities: a systematic literature review. *Sustainability*, 13:376.
- Mansoux, A., Howell, B., Barok, D., and Heikkilä, V.-M., 2023. Permacomputing aesthetics: potential and limits of constraints in computational art, design and culture, *Ninth Computing within Limits*.
- Richardson, K., Steffen, W., Lucht, W., Bendtsen, J., Cornell, S.E., Donges, J.F., ... and Rockström, J., 2023. Earth beyond six of nine planetary boundaries. *Science advances*, 9:eadh2458.
- Roberts, N., 2003. Direct citizen participation: building a theory. Paper presented at the 7th National Public Management Research Conference, October 9–11, Georgetown University, Washington, D.C., USA.
- Rockström, J., Steffen, W., Noone, K., Persson, Å., Chapin III, F.S., Lambin, E., ... and Foley, J., 2009. Planetary boundaries: exploring the safe operating space for humanity. *Ecology and society*, 14.
- Sennett, R. (2021). *The uses of disorder: personal identity and city life*. Verso Books. New York, NY, USA.
- Simonofski, A., Asensio, E.S., De Smedt, J., and Snoeck, M., 2017. Citizen participation in smart cities: evaluation framework proposal. In: 2017 IEEE 19th conference on business informatics (CBI), 1:227–236.

- Smith, A., and Martín, P.P., 2022. Going beyond the smart city? Implementing technopolitical platforms for urban democracy in Madrid and Barcelona. In *Sustainable smart city transitions*. Routledge, Oxfordshire, UK, pp. 280–299.
- Tsing, A.L., Deger, J., Keleman Saxena, A., and Zhou, F., 2020. *Feral atlas: the more-than-human Anthropocene*. Stanford University Press, Redwood City, CA, USA.
- Veenkamp, J., Kresin, F., and Kortlander, M., 2020. Smart citizens in Amsterdam: an alternative to the smart city. In A. Karvonen, F. Cugurullo, and F. Caprotti (Eds.), *The Routledge companion to smart cities*. Routledge, Oxfordshire, UK, pp. 144–146.
- Waag Futurelab. (n.d.). Digital European public spaces. <https://publicstack.net/deps/>.
- Waag Futurelab. (n.d.). Hollandse Luchten <https://hollandse-luchten.org/>.
- Waag Futurelab. (n.d.). Public stack. <https://publicstack.net/>.
- Waag Futurelab. (n.d.). Public stack puzzle [Co-creation tool]. <https://ccn.waag.org/navigator/tool/public-stack-puzzle>.
- Waag Futurelab. (2023). Public stack reflection cards. ACROSS project. <https://publicstack.net/cards/>.

Behavioural insights for moral design and green technology

Jeske Nederstigt

Abstract

Human behaviour is often the bottleneck in the diffusion of smart green technological innovations. Behaviour, in this case embracing green technology, is not only a matter of willingness (motivation) but also of ability (capacity) and facilities (opportunity). We often see a mismatch between attitudes and behaviour. This can partly be explained by a lack of knowledge and financial resources (capacity) and facilities (opportunity) to act according to the attitude. An additional explanation is the 'laziness' of consumers and their tendency to rely on their autopilot instead of making conscious decisions.

The acceptance of green, smart technological innovations only partially depends on the – partly controllable – characteristics of the innovation. It mainly relies on the perceived ease of use and the perceived usefulness of the innovation. Co-creation can help to safeguard these conditions. Furthermore, personal differences, such as the extent to which a person is open to innovations, play a role. Many consumers feel resistance to change and influence. In order to achieve lasting behavioural change, intrinsic motivation and a related sense of autonomy are essential.

Because of differences between groups of consumers it is important to align communication about innovations to the target group.

Finally, the conflict between individual and collective interest is important in consumer decisions about whether to make sustainable choices. To tackle this 'social dilemma' it is recommended to make the positive effects of sustainable behaviour visible and to emphasize that sustainable behaviour is the 'social norm'.

4.1 Introduction

The Intergovernmental Panel on Climate Change (2022) is clear on this topic: behavioural change can make a major contribution to reducing climate change and achieving a 40 to 70% reduction in greenhouse gas emissions by 2050. Many consumers find climate change important and say they are willing to change their behaviour. In practice, however, sustainable choices still appear to be very limited. There seems to be a gap between what people say and think, and what they do.

In the literature (Curtis and McConnel, 2002) a distinction is sometimes made between ‘citizens’ (with sustainable attitudes) and ‘consumers’ (who make less sustainable choices). In this chapter we view consumers as citizens and vice versa. We try to understand the gap between sustainable behaviour and sustainable attitudes and the possibilities (and limitations) to influence behaviour to go in a smart, green, innovative direction.

The success of green technological innovations depends largely on the people who do (or do not) embrace and use these innovations. What stops people from embracing these innovations? And what encourages them to actively participate in the development of green technology? To answer these questions, we first need to understand a bit about human behaviour. That’s what the first paragraphs are about. Next, we’ll take a closer look at the psychological aspects of acceptance of and resistance to innovations. After that, we will discuss the role of autonomy in involving citizens in designing green smart cities. Finally, we look at sustainable behaviour as a social dilemma; a situation in which individual interest and collective interest are conflicting.

4.2 Understanding human behaviour

In order to influence behaviour, we must understand why people behave (or don’t behave) in a certain way. We discuss two basic models: the Triad model (Poiesz, 1999 and Poiesz, 2014) and Kahneman’s theory about thinking fast and slow (Kahneman, 2011).

4.2.1 *Three conditions for behaviour*

The Triad model (Poiesz, 1999 and Poiesz 2014) is a simple practice-oriented model that helps to understand and explain behaviour and to develop a strategy for influencing behaviour. When it comes to influencing behaviour, we often think of tempting, convincing or motivating people. But motivation alone does not guarantee a certain behaviour (see CASE 4.1).

CASE 4.1

Waiting for gas-free solutions

Peter is genuinely concerned about the condition of our planet. He wants to make his house gas-free as soon as possible and has investigated the possibilities of heating his house with a heat pump and solar panels. Fortunately, he has the technical insight that allowed him to make a well-considered choice based on an enormous amount of information he has processed about heat

pumps and solar panels. Moreover, he has sufficient financial resources to pay for everything. Unfortunately for Peter, there is a huge waiting time for the delivery and installation of the products he has chosen. Two years later he still uses gas to heat his house.

According to the Triad model, there are three conditions for behaviour:

- Motivation: you must *want* it; you find the behaviour (or the consequence of the behaviour) attractive or important.
- Capacity: you must *be able* to do it; you have the necessary knowledge, skills and financial resources.
- Opportunity: the *circumstances* must allow the behaviour; availability, accessibility, facilities.

Only if all three conditions are met will the behaviour occur. In the case (1) of Peter and his heat pump, it is clear that Peter's motivation and capacity are sufficient, but because there is no opportunity in the short term (there is a waiting period), Peter will not have a heat pump installed in the short term.

According to the Triad model, the probability that certain behaviour will occur can be 'estimated', based on the product of the motivation, capacity and opportunity to perform the intended behaviour. Each of the factors can be expressed as an estimated percentage. This leads to the following formula:

$$\text{Probability of the behaviour} = \text{motivation} \times \text{capacity} \times \text{opportunity}$$

It is not necessary to calculate the scores very accurately (that is not feasible). In Peter's case we could estimate that motivation = 1 (100%), capacity = 1 (100%) and opportunity = 0.05 (5%; perhaps there are ways to get higher on the waiting list). The probability that Peter will soon install a heat pump = $1 \times 1 \times 0.05 = 0.05$, or 5%.

It is an approximative estimation that shows which of the factors needs to be 'edited' to increase the probability of the behaviour.

It is important when applying the model that the intended behaviour is described as specifically as possible. Only then is it possible to make a somewhat reliable estimation of the three factors. The opportunity to install a specific type of heat pump in the short term can be estimated more accurately than the opportunity to 'make your house more energy efficient'.

The three Triad factors are not independent; they can influence each other. We call these balance effects. If Peter is very motivated to install a heat pump at short notice, he will do everything he can to shorten the waiting time. He may be able to pick up the pump himself somewhere; or he bribes someone high on the waiting list. In other words: a high motivation can stimulate people to create opportunity or capacity. Because Peter was so motivated to install a heat pump, he started to look

into it; initially he had no knowledge, but under the influence of high motivation he acquired that knowledge. The influence of high motivation on capacity and opportunity is a positive balance effect.

positive balance
effect

motivation (+) → (capacity, opportunity) (+)

Please note, it does not work the other way around: a high capacity and/or opportunity does not mean that motivation increases. The fact that you can afford an expensive Electric Vehicle (EV) (capacity) and that such EVs are available (opportunity) does not automatically mean that you also *want* to buy an EV (motivation). Maybe you just love the sound of a gasoline engine.

Balance effects can also be negative. Suppose Peter had had no technical insight and had been completely stuck in the information about heat pumps and solar panels, he might have dropped out. His motivation decreased due to the lack of knowledge (capacity). And the long waiting time (low opportunity) could also lead to Peter's motivation decreasing. A negative balance effect therefore means that motivation decreases as a result of low capacity and/or opportunity.

negative balance
effect

(capacity, opportunity) (-) → motivation (-)

Here too, it does not work the other way around: low motivation does not mean that capacity and opportunity will decrease.

The Triad model makes a distinction between objective and subjective Triad scores.

CASE 4.2

The jump

Tom considers jumping over a wide ditch (which is filled with water). His motivation is high (by making the jump he can save a half hour of walking; and he is in a hurry) There is also opportunity: there is room for a long run-up. Tom also rates his capacity as being high (he used to be good at athletics at school). Unfortunately for Tom, his jump ends in the middle of the ditch.

As we can see in CASE 4.2, Tom misjudged his own capacity. As a result, Tom has tried the behaviour, but without successful results. The same applies to the other factors. Regarding motivation, it may seem strange to distinguish between objective and subjective as well (as clarified in CASE 4.3).

CASE 4.3

Fitness

Emma really wants to work on her fitness! She thinks it would be great to cycle on a fancy home trainer bike for 45 minutes every morning before work from now on. Three weeks after the home trainer bike was delivered to her home, Emma concludes that she is actually less motivated to get up early every morning to cycle than she had estimated.

The values of motivation, capacity and opportunity as estimated by the person him/herself (subjective scores) are decisive for the behavioural intention. However, these values are not always accurate and reliable. The actual (objective) values determine the chance of a successful outcome.

It may seem obvious to point to opportunity and capacity as conditions for behaviour, in addition to motivation. But in practice, it turns out all too often that these factors are overlooked. This has happened for example in the energy transition in the Netherlands: citizens and companies are being encouraged to switch to sustainable electricity but the electricity grid can't handle it. The opportunity factor is the bottleneck.

Of the three Triad factors, capacity and opportunity are often more practical in nature. Psychological insight relates in particular to the motivation factor. Besides the distinction between objective and subjective motivation, the Triad model distinguishes between motivation to perform the behaviour itself and motivation to achieve certain consequences with the behaviour. Or satisfaction in the short term (e.g. eating a tasty but unhealthy hamburger) versus satisfaction in the longer term (not having that hamburger to feel healthier in the long term.) This explains why we sometimes do things that we actually don't like to do at that moment. (Or do *not* do things, although we know it would be better for the long term.) Another important distinction is the difference between intrinsic and extrinsic motivation. We will get back to this in Section 4. Motivation can also have substantively differing sources. Social motivation is about 'wanting to belong' (or wanting to distinguish yourself from the crowd). We will get back to this in Section 3. Motivation can also come from fear (a person who is afraid of dogs is motivated to take a long detour if it means avoiding an encounter with a dog); or from hope (for example, the motivation to eat less meat in the hope that animal welfare will improve).

The acceptance of green technology requires new behaviour from citizens and consumers. This will only work if there is sufficient motivation (in addition to capacity and opportunity) to change habits and make sustainable choices. The next section explains why it is not always easy to encourage citizens and consumers to change their habits.

4.2.2 *Thinking Fast and Slow*

Nobel Prize winner and psychologist Daniel Kahneman, stated (2011):

Thinking is to humans as swimming is to cats; they can do it, but they'd prefer not to.

According to Kahneman, behaviour comes about by two thinking systems in our brain (see Figure 4.1).

System 1 concerns unconscious, automatic processes. They are not very flexible or accurate, but they run quickly and require little energy. A kind of autopilot. Think of jumping aside when a car comes towards you, or automatically moving when the traffic light turns green. System 1 always responds, even if you don't consciously think about it. But the quick thinking in System 1 is sometimes sloppy and the preference for instant gratification can lead to decisions that are illogical in the long term. Think of that tasty hamburger that beckons while you know that it is better for your health and for the planet to have an apple. Resisting the temptation to eat the hamburger requires System 2 activity. System 2 concerns conscious, deliberate processes. They are accurate and flexible but run slower than processes in System 1 and require a lot more energy. System 2 only kicks into action when conscious attention is paid. While System 1 is responsible for emotional and affective responses, System 2 mainly exhibits responses.

autopilot

System 1

System 2

New behaviour is often learned under the control of System 2 and later automated via System 1. Just think of your first driving lesson: you have to consciously think about every action. As you gain more driving experience, driving requires less effort and the action becomes more automatic.

System 2 also ensures that we sometimes do things that we find unattractive at that moment, but still do because we like the consequences (for example, a visit to the dentist). We need System 2 to resist temptations and to attentively process information or to make future plans and act accordingly. But because people are naturally lazy and prefer not to invest too much energy in System 2, we operate on autopilot as much as possible. This is unfavourable for the acceptance of innovations as it usually requires new, not yet automated behaviour and therefore activity from System 2.

The fact that people don't like to think much also implies that, if we want to influence behaviour, we should not necessarily do so by providing information and substantive arguments. The chances are high that people won't pay any attention to it. The easiest way to influence behaviour fast is to use the properties of the autopilot. In other words, the probability of the desired behaviour increases because:

- the behaviour is easier (requires less energy from System 2). The ease of use of Apple devices in the noughties (from 2000 to 2010) has significantly boosted their use. And/or because:
- the behaviour itself is more attractive. Think of gamification (see CASE 4.4).

System 1	System 2
Automatic, unconscious; associative	Conscious, reasoned; deductive
Fast	Slow
Inflexible	Flexible
Sometimes irrational	Reliable/accurate
Requires little energy	Takes a lot of energy

FIGURE 4.1 Two systems in the brain that control behaviour (Kahneman, 2011)

CASE 4.4

Gamification

A piano staircase was installed at the Central Stadium in the Dutch city of Rotterdam in 2016. The steps, which look like piano keys, contain sensors that emit a tone when you step on them. This intervention was supposed to make it more appealing for people to take the stairs instead of the elevator or escalator.

But there is more to the autopilot than 'making it easy' or 'making it fun'. Reactions from the autopilot often follow certain 'rules of thumb' or heuristics. If you understand these, it is easier to elicit responses from consumers. Commonly used heuristics are, for example: 'I do what I always do because I am satisfied with that.' Or: 'I do what most people do, because if many people do it, it must be good.' Or: 'I choose the most expensive option; then it is certain to be good.' Cialdini (2016 and 2021) has incorporated some of these rules of thumb into seven 'persuasion principles' that are often applied in marketing.

Authority: According to this principle, people are more likely to be convinced by an authority, someone with expertise. When the dentist says that electric brushing

rules of thumb
heuristics

authority

is better than manual brushing, you are more likely to accept it than when the neighbour says so.

Sympathy: People are more likely to accept a person they find sympathetic or attractive than an unlikable person. Interesting detail: the more people resemble us (in appearance and/or opinions), the more sympathetic we find them.

Social proof: If many people make a certain choice, 'then it will be right'. If one restaurant is (almost) full, while the other has no customers, a passing person is more likely to choose the busy restaurant than the empty one.

Reciprocity: Those who receive something from someone else are more likely to give something back. If your friend helps you paint your house, you are more likely to lend him a hand when he needs help.

Commitment and consistency: Once you have committed to something (for example, you answered affirmatively to the question about whether animal welfare is important to you), you want to adjust your behaviour accordingly. For example, you will be more likely to buy free-range chicken instead of farmed chicken. People like to be consistent in their behaviour.

Scarcity: The impression of scarcity makes a product or service more attractive. Consider offers under 'sold out' conditions or concerts for which 'only a few tickets are still available'.

Unity: This principle is about the need to belong. By appealing to people about shared values, group feeling (unity) is stimulated and the tendency to adapt to the group increases.

A limitation of the route to behavioural change via the autopilot is that the behaviour does not come from conviction and therefore quickly disappears as soon as the triggers disappear. Yet it is also conceivable that more or less automatically provoked reactions, i.e. habit formation, in the end will also lead to attention from System 2.

Although people naturally use System 1 more often than System 2, often System 2 is needed for lasting behaviour change. As mentioned before, we need System 2 to learn more complex new behaviours, as well as to resist temptations. However, it is much more difficult to trigger System 2 than to trigger System 1. The use of System 2 requires attention and energy, and these are scarce.

The probability of gaining attention from the target group increases as the trigger – and if aimed at System 2, this is often information – is better aligned with the values and lifestyle of the target group. Both in terms of content and tone of voice, as well as in terms of source and media. It is therefore important to empathize with the target group you want to influence: What is their initial attitude towards sustainability? Do they have sufficient financial resources and knowledge to make sustainable choices? What media do they use? And most of all: What is their initial opinion on the subject and which underlying values is that opinion based on? By aligning with the initial value pattern, you encourage the target group to feel understood.

According to Social Judgment Theory (Sherif *et al.*, 1965), consumers compare influencing information with their own prevailing ideas about the subject in question. The more strongly a consumer's existing attitude about a subject is confirmed, the smaller the latitude of acceptance (Taylor, 1981). That is, the range within which the arguments that are acceptable to him fall is smaller. A consumer who does not (yet) have a strongly confirmed attitude on the subject is still open to more arguments and has a broader 'latitude of acceptance'. In an attempt to influence behaviour through substantive arguments, you can use the assimilation effect. This means that information that falls within the latitude of acceptance is received extra positively by the consumer, while information that falls outside the latitude of acceptance is received extra negatively. (This second phenomenon is the contrast effect.) By putting yourself in the position of the target group and initially reasoning from their point of view, a certain degree of trust is built up, so that the consumer adopts a more positive attitude and is more likely to pay attention to subsequent arguments. When arguments that fall outside the latitude of acceptance are immediately started, the consumer – due to the contrast effect – actually takes a negative attitude towards the influence and pays no attention to the following arguments. If you want to convince a die-hard carnivore to eat less animal protein, it is better to first propose that they try a vegetarian dish once than to immediately propose that they always eat vegan from now on. A consumer with a less strong opinion about eating meat could perhaps immediately be persuaded to eat vegetarian for a few days a week.

But either way, it is still a pitfall to overestimate the extent to which people are willing to pay attention to your information.

4.3 Acceptance of and resistance to innovations

4.3.1 *Acceptance and diffusion of innovations*

The acceptance of green technological innovations is not only about objective characteristics of the innovation, but just as much about user perception. The Technology Acceptance Model (TAM; Davis, 1986) argues that the intention to use an innovation is determined by the attitude towards the innovation. And that attitude is determined by the perceived usefulness and the perceived ease of use, where the weight of perceived usefulness is approximately one and a half times the weight of perceived ease of use (see Figure 4.2).

- Perceived usefulness (PU) was defined by Davis as 'the extent to which a person believes that using a particular system would improve his/her performance'. So, it is about the extent to which the user thinks that the technology is useful or not for what he/she wants to do.
- Perceived ease of use (PEOU) was defined by Davis (in 1989) as 'the extent to which a person believes that using a particular system would be effortless'.

User Motivation

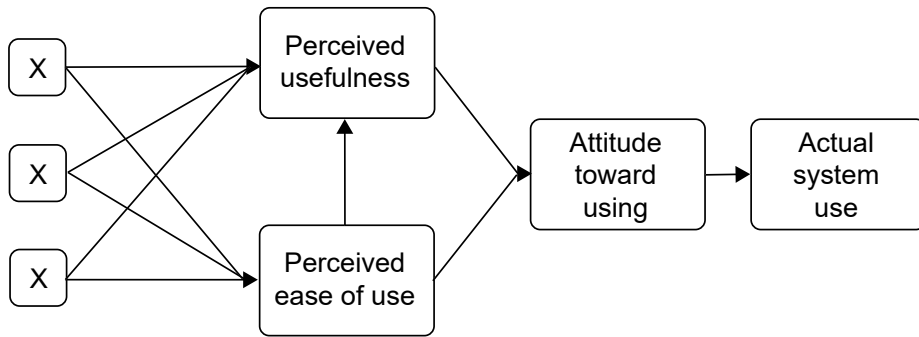


FIGURE 4.2 Original Technology Acceptance Model (Davis, 1986)

TAM is in line with Kahneman's (2011) idea that 'convenience' plays an important role in behavioural change. In terms of the Triad model, perceived usefulness is comparable to motivation. And although TAM views perceived ease of use as a part of motivation, it is also comparable with (subjective) capacity. The third condition, opportunity, can be seen as the availability and/or trialability of the innovation; however, this is not included in the TAM.

The perceived usefulness and the perceived ease of use are also determined by characteristics of the innovation itself, and personal factors such as the social environment, personality traits and experience.

The original TAM was created in 1986, by Davis (see Figure 2). Since then, several adjustments have been made (Lai, 2007). The most important upgrades are the TAM 2 (Venkatesh and Davis, 2000) and the Unified Theory of Acceptance and Use of Technology (UTAUT, Venkatesh *et al.*, 2003). The upgrades mainly consist of a further elaboration of the independent variables, making the model more accurate but also more complex.

With regard to the diffusion of innovations across the target group, Rogers (1995) distinguishes different groups, depending on the moment at which they accept the innovation (see Figure 4.3; the percentages mentioned are arbitrary and have to do with the definition of the subgroups):

- innovators (the first 2.5% of the target group)
- early adopters (the next 13.5%)
- early majority (the next 34%)
- late majority (the next 34%)
- laggards (the last 16%)

The rate of diffusion varies with the nature of the target audience. Innovation spreads faster in a target group of mainly young, highly educated consumers than in a target group of older, low-educated consumers.

diffusion of
innovations




FIGURE 4.3 The diffusion process (Rogers, 1995)

The first group to use an innovation are the innovators. They are best convinced by objective information and their own personal experiences. They play an important role in convincing the following groups, who (from left to right) are increasingly sensitive to the opinions and experiences of other users and less open to innovations. Each group (and therefore each phase of the process of introducing innovations) requires its own communication approach.

co-creation

To stimulate acceptance, it is important to increase perceived usefulness and perceived ease of use. Therefore, we should involve the end user in the development process as early as possible, for example by co-creation. Co-creation increases the perceived benefit and support for the innovation (Prahalad and Ramaswamy, 2004). The role of the end user in the co-creation process can range from evaluating different prototypes to brainstorming for innovations. In the most intensive form of co-creation, the end user is part of the development team throughout the entire development process. Co-creation can create ambassadors, which accelerates the diffusion across the target group. It is therefore important that consumers involved in co-creation not only belong to the target group but also have the characteristics of opinion leaders. Opinion leaders are often regarded as experts by people around them, and their influence on the consumer behaviour of others is great. Opinion leaders are curious, open to innovation and often extrovert. They get a lot of information from the media to base their consumer choices on; they are less sensitive to the opinions of other consumers. The amount of opinion leaders is relatively high among the group of innovators (Rogers, 1995).

4.3.2 *Resistance to innovations*

People are naturally inclined to stick to their own attitudes and habits. Attempts to influence behaviour or attitudes (for example, through innovations) may therefore encounter resistance. Vonk (2017) mentions the following possible causes of resistance:

- Reluctance: for example, if the acceptance of an innovation costs time, money or effort.
- Inertia: the tendency to stick to the status quo; resistance to change in general.
- Scepticism: doubts about the substantive arguments.
- Self-protection: the feeling that you ‘must’ change can imply that you have been doing something wrong up to now. This damages the consumer’s ego.
- Autonomy: also called Reactance (Brehm and Brehm, 1981); people do not like to be influenced by others and like to make their own choices.

In a review, Fransen *et al.* (2015) place the motives for resistance and the strategies that people use in response to resistance in an integrative framework (see Figure 4.4). This model distinguishes three motives for resistance that correspond to the reasons that Vonk (2017) mentions. Freedom threats covers Autonomy and Self-protection, Concerns of deception is similar to Scepticism, and Reluctance to change covers Reluctance and Inertia. The Avoidance strategy – which can occur with any of the motives – is an anticipatory strategy in which people avoid attempts to influence in advance. The other strategies are a response to attempts to influence.

freedom threats /
concerns of
deception /
reluctance to
change /
avoidance

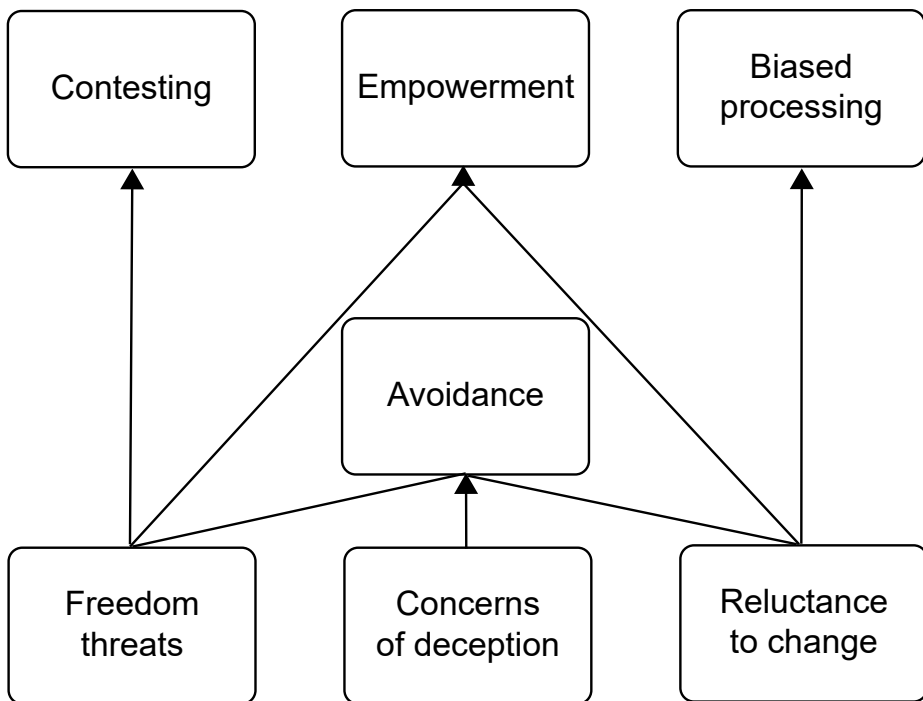


FIGURE 4.4 The Strategies and Motives for Resistance to Persuasion (SMRP) Framework (Fransen *et al.*, 2015)

- Empowerment When people are reluctant to change, they are likely to use Empowerment (i.e. reinforcing assertions of confidence or reinforcing the particular attitude or belief that is challenged) or Biased processing (processing information in such way that it aligns with existing attitudes and behaviour) as strategies to resist persuasion.
- Biased processing
- Contesting When concerns of deception are present, people are likely to use Contesting strategies (i.e. counter arguments) to resist persuasion. In response to persuasive messages that are perceived to threaten one's freedom, people are likely to use both contesting and empowerment strategies to resist persuasion.

CASE 4.5

Fake meat

Fransen gives a striking example of Empowerment in response to 'perceived Freedom Threats' in 'Het Grote Duurzaam Gedrag magazine' (2021). A consumer posted on Twitter a photo of an enormous amount of meat that he had just bought. He wrote: *'I saw all kinds of leaflets passing by with "fake meat" in which I was told that I should no longer eat meat ... Well, that makes me recalcitrant!'*

- Entitlement The reaction in CASE 4.5 is also an example of Entitlement, people's feeling that they have a 'right' to something; in this case, eating meat. This is a difficult phenomenon, because how do you convince people that there is nothing obvious about having the right to eat meat (or having a car, or flying) without making the consumer feel that their freedom being threatened? The feeling of Entitlement can be strengthened by Moral Self Licencing (Blanken *et al.*, 2015). This means that a consumer allows himself immoral behaviour to compensate for other moral behaviour ('I don't eat meat, so I can take a plane to go on holiday).
- Moral Self Licencing

Motives for resistance differ between consumers. Innovators are less reluctant to change, but critical of deception. Among the late majority and the laggards there are expected to be more consumers with perceived freedom threats or reluctance to change than among the innovators, the early adopters and the early majority.

For the acceptance and diffusion of innovation, it is important to prevent resistance as much as possible. The less moralizing a message is perceived, the less people feel that their freedom of choice is threatened. Factual, objective information from independent sources can also reduce concerns or deception. An objective, independent approach is particularly important for acceptance by innovators and early adopters. For later groups, independent reviews from earlier users can be effective. The core of the message is that the chance of acceptance of an innovation increases, the more consumers feel that it is their own voluntary choice.

4.4 Citizens and autonomy

People by nature like to stick to existing habits (Kahneman, 2011) and are often resistant to external influence (Vonk, 2017; Franssen *et al.*, 2015). Citizens will therefore only take on an ambassadorial role in spreading smart, green, technological innovations if they are truly convinced of the usefulness and perceived ease of use of the innovation and, moreover, feel that it is their own voluntary decision to adopt the smart green technology.

This sense of autonomy not only prevents resistance to influence, but is also closely related to intrinsic motivation (which is a prerequisite for lasting acceptance of innovations). Autonomous, intrinsically motivated citizens or consumers are indispensable for the broad-based, bottom-up diffusion of innovative green technology.

If people start using an innovation because they are utterly convinced of its usefulness, acceptance is intrinsically motivated. If acceptance comes from an expected reward, the expected approval of others or the expectation of being able to avoid negative outcomes, then the acceptance is extrinsically motivated. An intrinsically motivated consumer purchases solar panels based on sustainability ideals, while an extrinsically motivated consumer purchases them to save on energy costs or to impress the neighbours.

Extrinsic motivation is easier to influence by third parties through rewards than intrinsic motivation. Unfortunately, extrinsically motivated behaviour quickly disappears when the reward disappears. When gas prices are sky high, many households significantly reduce their gas consumption, not for sustainability reasons (intrinsic motivation), but to save on costs (extrinsic motivation). If gas prices fall, consumption will increase again.

For the lasting acceptance of green technology, intrinsic motivation is therefore preferable to extrinsic motivation. Ryan and Deci (2000) distinguish in their Self-Determination Theory between autonomous and controlled motivation.

Controlled motivation is driven by extrinsic motives, such as material or social rewards (discount on your energy bill or 'esteem'). When the consequences of behaviour are internalized – i.e. when the consumer believes that his/her behaviour matches his/her self-image and his/her personal values – the behaviour is autonomously motivated (or based on intrinsic motivation). The distinction between autonomous and controlled motivation is not black and white, rather there is a gradual continuum (see Figure 4.5).

The less controlled (extrinsic) and more autonomous (intrinsic) a consumer is motivated, the more willing he/she is to spend time, money and effort on the behaviour (i.e. accepting and using an innovation). According to the Self-Determination Theory of Ryan and Deci (2000), intrinsic or autonomous motivation is related to three factors:

1. **Autonomy:** the feeling that you are making the decision by yourself, voluntarily.
2. **Competence:** the feeling of being able to utilize your own capabilities.
3. **Connectedness:** the feeling of being connected to others in the same situation.

The feeling of autonomy can be promoted by allowing consumers to make the choice for customization (or the option to tailor the innovation to your specific situation) instead of offering one-size-fits-all solutions only. A consumer may want to eat insects as a source of protein but prefers it in the form of powder rather than as a hamburger.

The feeling of competence can be stimulated by really listening to the consumer and taking his/her opinion seriously. This can be done in the design phase (by co-creation), or by jointly looking for a suitable solution that meets the wishes of this consumer.

Finally, encouraging connectedness can be done by showing who the other users of the innovation are or by creating a community. For example, an online platform where consumers can share experiences with the innovation and exchange tips.

According to the Self-Determination Theory, two things are important when aiming at lasting, intrinsically motivated acceptance of innovations. The first is to offer 'customization' (to stimulate the feeling of autonomy and competence). The second is to facilitate the exchange of experiences between 'innovators' and 'early adopters'.

Given the resistance to external influence, it is important that the acceptance and diffusion of green technological innovations occur bottom-up. Autonomy and intrinsic motivation are crucial in this regard.

A-motivation	Extrinsic motivation				Intrinsic motivation
No motivation	Controlled motivation	Controlled motivation	Autonomous motivation	Autonomous motivation	Autonomous motivation
No behavioral intention	Avoiding punishment / getting a reward	Avoiding guilt/shame; maintain self-esteem	Assigning personal value to the behavior	The behavior fits perfectly with your own values	The behavior itself is satisfying
No regulation	External regulation	Regulation by introjection	Regulation by identification	Regulation by integration	Intrinsic regulation
Low quality motivation ← → High quality motivation					

FIGURE 4.5 Self-Determination Theory (Ryan and Deci, 2000)

4.5 Sustainable choices: a social dilemma

One of the main reasons why sustainable attitudes do not always lead to sustainable behaviour has to do with the contradiction between individual interest and collective interest (see for an example, CASE 4.6). Sustainable consumer behaviour is behaviour in which the consumer gives greater weight to the positive effects of his/her behaviour on society than the negative individual effects.

CASE 4.6

Recycling

Suppose Jack offers all his discarded items for repair or recycling and buys only (expensive) recyclable products. This will cost him time, effort and money. The benefits – less climate change – are uncertain (because Jack does not know whether many other people make sustainable choices) and will only be noticeable in the longer term. If Jack himself chooses ‘cheap and easy’ opportunities while many other people make sustainable choices, he will benefit from a better climate in the long term. But if everyone were to reason and act like this, global warming would continue and even accelerate beyond a possible tipping point.

A situation in which the individual interest conflicts with the collective, social interest and every individual also depends on what others do for a positive effect, is called a social dilemma (Dawes, 1980).

social dilemma

Hardin described this phenomenon in 1968 as the ‘tragedy of the commons’: if everyone is committed to the common interest except one person, then that one ‘free rider’ is best off. However, if everyone thinks this way, the common interest is not achieved and ultimately that is worse for everyone.

Objectively speaking, there is a direct relationship between individual and social interests. The consumer is part of society, so what is bad for society is ultimately also bad for him. However, a number of psychological effects – which can be directly linked to the autopilot from Section 4.2.2 – ensure that consumers do not always choose in the interests of society.

4.5.1 *Self and others*

People consider themselves more important than others. Therefore, a consumer will give more weight to self-interest in consumption decisions than to the interests of unknown others. Moreover, a sense of injustice plays a role here: ‘If I am the only one who takes the social effect into account and others do not, I have a double disadvantage. I bear the ‘costs’ and – because others do nothing – I still don’t

sense of injustice plays

experience the desired effect.' Many consumers are quite willing to contribute, but only if others do so too. But if there is a suspicion that others are not contributing to the (implicit?) agreement, they stop contributing themselves. (Mary is happy to eat vegetarian food two days a week, but not when she hears that her neighbour continues to eat meat every day.) The need for justice is deeply rooted in human nature: sometimes we would rather have nothing than 'less than the other'. In the ultimatum game (Harsanyi, 1961), subjects (the proposers) were instructed to divide an amount of money between themselves and another person (the responder). The responder could refuse or accept the proposal. If he/she refuses, both get nothing. Rationally speaking, it makes sense to accept any offer; even if the proposer keeps 99% of the amount and the responder only gets 1%. After all, 'something' is always better than nothing. Yet it appears that proposals in which the responder gets less than 20 to 30% are often rejected: the responder would rather get nothing than (much) less than the other (it also turned out that many 'dividers' themselves propose a 50–50 division).

4.5.2 *Nearby and at a distance*

Socially responsible consumption requires that remote effects be taken into account. Consider the use of margarine that contains palm oil (with a negative impact on tree felling in South East Asia) or the poor working conditions of workers in some Asian countries. The mental distance also plays a role. Social effects of consumption are often invisible or difficult to link to consumer choices. Consumers do not always realize that eating pork necessitates keeping pigs, that pig farms produce manure that is spread on fields, that this ultimately ends up in the air and damages forests and buildings. The 'mental leap' from pork chop to forest degradation is large. Moreover, the autopilot (Section 2.2) gives more weight to what is close.

4.5.3 *The short term and the long term*

Short-term consequences of choice are generally more decisive than long-term consequences. Taking a flight to go on holiday fulfils the need for relaxation in the short term, but it does not benefit the long-term need for 'a liveable planet'. The autopilot, again, gives more weight to the short term than to the long term.

4.5.4 *Certain and uncertain*

Socially responsible consumption is more obvious if the effects are certain. But if we buy an electric car to save the environment, we are not sure whether it will have an effect. There is still uncertainty about the effects of battery production and way in which electricity is generated. And there is uncertainty about the behaviour of others: only together can we make a difference.

In order to stimulate sustainable behaviour, we must tackle the social dilemma in a targeted manner. This can be done in two ways (Nederstigt and Poiesz, 2022).

With a structural approach, the dilemma itself is removed by creating individual benefits to sustainable behaviour. If travelling by train is cheaper and faster than flying, there is no longer a reason for consumers to fly. Other examples are subsidies on electric cars or solar panels. Legislation can also help: if the government bans single-use plastic, consumers no longer have to choose between (cheap) plastic straws or more expensive sustainable options. Structural solutions are effective in the short term but do not lead to intrinsic conviction among consumers.

structural
approach

A psychological approach does not remove the dilemma but is aimed at reducing the pursuit of self-interest. An obvious strategy is to provide information about the severity and urgency of the problem and the need for sustainable behaviour. For example, a documentary about the effects of the loss of biodiversity. This approach requires people's attention and thus energy. Therefore, you mainly reach people who are already consciously engaged in sustainable behaviour; others will probably pay little attention to this type of information. It is primarily a method to reach the pioneers, the innovators, who in turn may gain followers (Section 3 and Figure 3). Another psychological strategy is to increase the perceived manageability of the problem. By making visible what the effect of the desired behaviour is (preferably at the individual level), uncertainty about the usefulness of sustainable behaviour decreases. In other words: make clear that it really works! In this context, it is important to provide feedback on results already achieved! The media often pays more attention to bad news than to good news. More items about, for example, declining CO₂ emissions can stimulate sustainable behaviour. A third psychological strategy is addressing the social norm. Communicating that many people exhibit sustainable behaviour not only increases the idea that it is useful ('together we make a difference'), it also reduces the sense of injustice and the fear of 'free riders'. The psychological approach requires more time than the structural approach, but ultimately leads to intrinsically motivated behaviour more often.

psychological
approach

severity and
urgency of the
problem

perceived
manageability

social norm

4.6 Wrap up

Changing the behaviour of consumers and citizens so that it moves in a sustainable, smart, innovative direction is quite a challenge.

The acceptance of green, smart technological innovations only partially depends on the characteristics of the innovation itself. Much more important for the degree of success of innovations is the perceived ease of use, the perceived usefulness and the extent to which the consumer is open to innovations.

There seems to be a gap between sustainable attitudes and sustainable behaviour. This can be partly explained by a lack of capacity and opportunity but also by the 'laziness' of consumers and their tendency to rely on their autopilot instead of making conscious decisions. Finally, the contradiction between individual interest and collective interest plays a role in the lack of sustainable behaviour.

Promising interventions to bridge the gap include focusing on facilitating behaviour, aligning interventions with the functioning of the autopilot, making the positive effects of sustainable behaviour visible, and last but not least, emphasizing that sustainable behaviour is the social norm. Ambassadors or influencers can help to create such a social norm. Ambassadors can in turn be created by involving consumers in a process of co-creation at an early stage of innovation development.

Consumers naturally feel resistance to change and influence. In order to create support for smart green technology and convince consumers to adopt lasting behavioural change, intrinsic motivation and a related sense of autonomy are essential. True conviction requires conscious attention (and thus energy) from consumers. Attention is scarce, and to increase the probability of attention it is important to tailor persuasive communication in terms of content, tone of voice, source and media to the values and lifestyle of the target group.

References

- Blanken, I., Van De Ven, N. and Zeelenberg, M., 2015. A meta-analytic review of moral licensing. *Personality and social psychology bulletin*, 41:540–558. <https://doi.org/10.1177/0146167215572134>.
- Brehm, S.S. and Brehm, J.W., 1981. *Psychological reactance: a theory of freedom and control*. Academic Press, Lake Charles, LA, USA.
- Cialdini, R.B., 2021. *Influence, new and expanded: the psychology of persuasion (Revised.)*. HarperCollins, New York, NY, USA.
- Cialdini, R., 2016. *Pre-suasion: a revolutionary way to influence and persuade*. Simon and Schuster, New York, NY, USA.
- Curtis, J.A. and McConnell, K.E., 2002. The citizen versus consumer hypothesis: evidence from a contingent valuation survey. *Australian journal of agricultural and resource economics*, 46:69–83. <https://doi.org/10.1111/1467-8489.00167>.
- Davis, F.D., 1986. *A technology acceptance model for empirically testing new end-user information systems: theory and results*. Massachusetts, United States: Sloan School of Management, Massachusetts Institute of Technology, MA, USA.
- Davis, F.D., 1989. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS quarterly*, 13:319–340 <https://doi.org/10.2307/249008>.
- Davis, F.D., Bagozzi, R.P., and Warshaw, P.R., 1989. User acceptance of computer technology: a comparison of two theoretical models. *Management science*, 35:982–1003. <https://doi.org/10.1287/mnsc.35.8.982>.
- Dawes, R.M., 1980. Social dilemmas. *Annual review of psychology*, 31:169–193. <https://doi.org/10.1146/annurev.ps.31.020180.001125>.
- Fransen, M.L., Smit, E.G., and Verlegh, P.W., 2015. Strategies and motives for resistance to persuasion: an integrative framework. *Frontiers in psychology*, 6:146377. <https://doi.org/10.3389/fpsyg.2015.01201>.

- Hardin, G., 1968. The tragedy of the commons. *Science*, 162:1243–1248.
- Harsanyi, John C., 1961. On the rationality postulates underlying the theory of cooperative games, *The journal of conflict resolution*. 5:179–196. doi:10.1177/002200276100500205.S2CID 220642229.
- IPCC, 2022. *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [H.-O. Pörtner, D.C. Roberts, M. Tignor, E.S. Poloczanska, K. Mintenbeck, A. Alegria, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, B. Rama (eds.)]. Cambridge University Press, Cambridge, UK.
- Kahneman, D., 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York, NY, USA.
- Lai, P.C., 2017. The literature review of technology adoption models and theories for the novelty technology. *JISTEM – Journal of information systems and technology management*, 14:21–38. <https://doi.org/10.4301/S1807-17752017000100002>.
- Nederstigt, J., Poiesz, T., 2022. *Consumentengedrag*. Noordhoff, Groningen, the Netherlands.
- Poiesz, T. (Ed.), 2014. *Redesigning psychology. In search of the DNA of behaviour*. Eleven International Publishing, the Hague, the Netherlands.
- Poiesz, T.B.C., 1999. *Gedragmanagement: waarom mensen zich (niet) gedragen*. Inmerc, Wormer, the Netherlands.
- Prahalad, C.K., and Ramaswamy, V., 2004. *The future of competition: co-creating unique value with customers*. Harvard Business Press, Brighton, MA, USA.
- Rogers, E.M., 1995. *Diffusion of innovations*. 4th ed., The Free Press, New York, NY, USA.
- Ryan, R.M., and Deci, E.L., 2000. Intrinsic and extrinsic motivations: classic definitions and new directions. *Contemporary educational psychology*, 25:54–67. <https://doi.org/http://dx.doi.org/10.1006/ceps.1999.1020>.
- Sherif, C.W., Sherif, M., and Nebergall, R.E., 1965. *Attitude and attitude change: the social judgment-involvement approach*. Saunders, Philadelphia, PA, USA. pp. 127–167.
- Venkatesh, V., and Davis, F.D., 2000. A theoretical extension of the technology acceptance model: four longitudinal field studies. *Management science*, 46:186–204. <https://doi.org/10.1287/mnsc.46.2.186.1926>.
- Venkatesh, V., Morris, M.G., Davis, F.D., and Davis, G.B., 2003. User acceptance of information technology: toward a unified view. *MIS quarterly*, 27:425–478. <https://doi.org/10.2307/30036540>.
- Vonk, R., 2017. *Sociale psychologie*. Boom, Meppel, the Netherlands.

The moral programming of XR, and what we can learn from the AI experience

Leon Kester, Bart Wernaart and Nadisha Marie Aliman

Abstract

Extended Reality (XR) technology is on the rise. The technology becomes more complex, is applied in countless contexts (as a professional instrument, a research tool or in the field of leisure), and its use becomes more widespread amongst citizens. It can be considered as the ‘next big thing’ to Artificial Intelligence. As with each new technology, its development comes with new moral problems, and new moral problems urge new moral solutions. Usually, these solutions are found in the short term in ethical frameworks or codes of conducts and on the long term in (enforceable) laws. In this contribution, we propose to not step in the pitfalls of the past and learn from the AI experience in the context of setting moral boundaries for technology development. We propose a moral procedure called ‘moral programming’ that should enable technology developers and tech-policy makers to recognize and identify moral issues that need to be answered by those who are to be affected by the new technology, equip them with a meta-method to establish a bottom-up crowd sourced ethics approach (including feedback loops) in their design process, to align values of those whom it concerns and those who develop/govern.

To this end we reflect on 1) when are moral issues high stake moral issues that need the ‘special treatment’ of a moral programming approach? 2) how can we decide who should be the moral authority on the moral programming in itself? And 3) what does moral programming look like, and how can it be done?

In conclusion, we will focus on several particular challenges that are exemplary for XR (and in particular AIXR) technology.

5.1 Introduction¹

Extended Reality (XR) is usually considered the umbrella term for Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR) (Vasarainen *et al.*, 2021). Extended reality could be defined as ‘the tool that sits between the digital twin and

¹ This chapter was written as a part of XR4HUMAN: a three-year EU-funded project. Its mission is to co-create living guidance documents on ethical and related policy, regulatory, governance, and interoperability issues of XR technologies whilst building public trust and acceptance and a strong and competitive European XR ecosystem.

real-world information' (Casini, 2022). Virtual Reality encompasses virtually produced worlds, replacing the real world, where AR supplements reality through digital products. Mixed Reality is a term that is used for AR and forms of VR that are complemented with real world elements (also referred to as Augmented Virtuality, AV). Extended Reality is increasingly applied in many different fields, e.g. in education (Alnagrat *et al.*, 2022; Sharma, 2021), the business world (Marr, 2021), health care (Andrews *et al.*, 2019; Morimoto *et al.*, 2022; Zhang *et al.*, 2023), urban planning (Çöltekin *et al.*, 2020), design (Schneider; 2021; Kharvari and Kaiser, 2022), military training (Boyce *et al.*, 2022; Garcia Estrada *et al.*, 2024), tourism (Santoso *et al.*, 2022) and the Internet of Things (IoT) (Andrade and Bastos, 2019). The broad application scope touches many professions and individuals in society, including the major theme of this book (green technology), and the overall impact is considered to be significant (see all previous sources). At the same time, serious concerns around ethics issues are identified, which has led to an extensive academic debate (Tromp *et al.*, 2018; Carter and Egliston, 2020; Burns, *et al.*, 2022; Ramirez *et al.*, 2023; Indradevi *et al.*, 2024). Amongst others, the question is raised on how to regulate XR developments in the near future, and how to solve moral dilemmas that come with the wide use of XR applications.

Regulation is done by law and assessment of conformity to that law supported by standardisation documents. It usually does not focus on how to resolve moral dilemmas calculating harm/risk/benefits but on principles how to avoid harm. The language of 'weighing harm' is far away from the principled language of 'fundamental rights'.

regulation

The EU AI act, that also applies to XR systems that make use of AI components, was recently adopted and breaks with this principles approach. This law has a risk-based product safety approach for the design and development of AI systems where risk is understood as the probability of harm and the severity of that harm where severity can be both qualitative and quantitative. Harm is further detailed in relation to safety, health and fundamental rights. Therefore for high-risk AI systems we have to shift from the principle of avoiding all harm to the concept of acceptable level of risk since for high-risk AI systems not all possible harm can be excluded. The consequence of that is that AI systems engineers need a more explicit formulation of harm. In addition, to the list of AI systems to which this law applies, special attention is given to 'General Purpose AI' (GPAI) models, also known as foundation models or Generative AI models. Special 'systemic risk' management requirements are formulated where systemic risk relates to the whole of (EU) society.

EU AI act

System applications are classified into categories of low, medium, high, and unacceptable risk. Most attention is on high-risk AI system applications since there are potential high risks, but the benefits are considered higher than the (potential) costs in terms of harm to citizens or society at large.

In the development of high-risk AI systems, developers typically encounter ethical and moral dilemmas pertaining to safety, health, and fundamental rights. It is imperative for them to mitigate these risks to an acceptable level.

AIXR system
risks

Dealing with moral dilemmas in system design is not specific for AI but also holds for XR systems. In addition, most XR systems make extensive use of AI in these systems and will therefore also have to comply to the AI act. Also, a wide range of specific AIXR system risks and subsequent ethical/moral dilemmas might arise particularly in relation to GPAI where multimodal generative AI can be harnessed to mimic or forge more and more aspects of human behavior. What is unique about introducing AI in XR systems is that these systems can adapt during operation in order to avoid or reduce risk which necessitates a more formal (mathematical) specification of how to deal with moral dilemmas.

5.2 Moral problems

To discuss moral problems that come with AI or XR (or both), we first need to be clear of the nature of moral problems. This has been a matter of (Western) academic debate since the ancient Greek (the so-called Platonic conflict) and continued to be discussed ever since (e.g. Sartre, 1957). A moral problem would encompass a situation in which a moral actor feels the moral obligation or has moral reason to two or more things (moral actions) that rule out (but not override) the other moral actions (Sinnott-Armstrong, 1987) and/or both do something and refrain from doing so at the same time (Statman, 1990).

moral truth

Whether or not a moral problem can exist in the first place depends on how this is approached in theories on normative ethics. One crucial element is how we consider the existence of moral truth (Dubbink, 2018). In absolutist normative theories, there are moral truths. A moral problem then exists when there is a conflict between moral truths, and one needs to prioritize one over the other. In some absolute normative theories, it is assumed that a normative theory should not allow the existence of (genuine) moral problems e.g. Immanuel Kant's categorical imperative (Timmermann, 2013), and some approaches in utilitarianism (Gowans, 1987), where moral dilemmas cannot exist within the normative theory in itself, since there is always one objectively justifiable and reasonable way to act in situations with moral challenges. In contrast, there are relativist theories that assume that moral truth does not exist, and moral righteousness is context and culture dependent (De Haan, 2001). In relativist theories therefore, moral problems are widely accepted phenomena, and there is a focus on handling differences between and prioritizing multiple moral solutions of whom each is morally justified/desirable according to the dominant moral viewpoints in one moral context, but not (per se) the other.

moral solutions
are comparable

In the second place, we need to establish that if moral problems exist, its possible moral solutions are comparable. When we accept that moral solutions are incomparable, the existence of moral problems becomes problematic since they cannot be solved (Chang, 2012). This implies that we need a method that allows the comparability of multiple moral solutions to a moral problem. In so called 'moral hard

cases' (borrowed from the legal domain, e.g. (Statman, 1990) there is more than one solution to a moral problem that can be defended by moral theory, but these moral solutions contradict one another (Messerli *et al.*, 2017). Therefore, it is necessary to find ways to meaningfully compare these moral solutions in order to choose one. Existing approaches in meta-ethics to compare moral solutions are comparative methods, such as the Agents, Deeds and Consequences (ADC)-model (Dubljević, and Racine; Dubljević, 2020), or the Expected Moral Value (EMV)-approach (Lockart, 2000) as we will discuss further in section 5.4

However, in the reality of technological innovation, tech-companies and tech regulators, it would be rather far-fetched to assume there is an absolute moral truth or universal moral procedure that drives their techno moral decision-making, or that the involved moral actors are capable of finding/using such a moral truth or procedure. We therefore have to assume that moral problems exist, regardless whether or not this is something that ought to be. Also, scenarios of human-crafted algorithms in unsupervised settings without human interference complicates matters in the application of meta-models, that are initially designed for human moral action (Wernaart, 2021). We therefore have to conclude there is a need for a moral procedure that does justice to and helps solve the practical challenges of techno-moral problems on a meta level.

5.3 Moral authority

When we accept that there are techno-moral problems, and there is a need to compare multiple solutions to a techno-moral problem, we need to address the matter of moral authority: who decides on the techno-moral problem (Wernaart, 2022)? In general, we could say that those who are affected by the techno-moral problem ought to have the moral authority to decide on these problems. With the design of technology, it is mostly the designers of the new technology that are able to make the most important design decisions, and strongly influence the way techno-moral problems are decided on. Considering that it is mostly not the designers who bear the moral consequences of their design choices, they usually do not have the moral authority to take such decisions (Millar *et al.*, 2017). Ideally, designers are capable of sensing and incubating the values of those who bear the moral consequences of their design. This is however not as straightforward as it may seem. We need to find ways to establish *who* should have moral authority, and then map *what* type of moral solutions that are supported by those who have moral authority.

There are multiple ways to establish who has moral authority. Millar (2017) suggests making a distinction between high and low stake moral problems, and individual or societal consequences of these moral problems. When moral problems represent low stakes, it is more practical to assume that the designer of technology can take the moral decision and assume moral authority. When they are high

high and low
stake moral
problems

stake, it would be irresponsible if the designer would decide on behalf of the moral authority, and therefore would have an obligation to incorporate the moral views of the moral authority in their design. If the moral consequences are centralized around a specific group of people or individuals (e.g. users) then they should decide. Sometimes, moral consequences are felt society-wide, and therefore would require the involvement of society. This is a rough sketch of how to distribute moral authority. A more detailed approach is offered by Wernaart (2021) through a moral dashboard for engineers, using the issue-contingent model that was introduced by Jones (1991). The moral intensity is measured according to six dimensions: magnitude of consequences, level of social consensus, probability of the moral effect, temporal immediacy, proximity and the concentration of the effects. These dimensions can help to establish into more detail whether moral issues are to be considered high or low stake, and to what degree moral consequences are individual or societal. This dashboard can help individual engineers or organizations to better understand the stake of their moral programming, and reflect on the allocation of moral authority.

So, considering that with most technological innovations the moral authority is not the same as the moral programmer, it is important to be able to separate those, and find useful manners to enable a moral programmer to involve the moral authority when designing/engineering new technologies (Wernaart, 2022). It is also important to find useful approaches to decide on how to establish the moral preferences of the moral authority, and in case of conflicting preferences, how to settle these. To this end, theories in moral programming may be of help, involving meta-ethical calculation methods.

5.4 Moral programming

As discussed earlier, the EU AI act has risk-based product safety approach for the design and development of AI systems where risk is understood as the probability of harm and the severity of that harm. It is therefore required that during product design and development there should be a risk management procedure where someone in the company is responsible for risk management. Some confusion has come up with recent developments in AI regarding the misconception that an AI system would be capable of fully autonomous risk management. However, an AI system cannot understand morality and its underlying explanations and therefore cannot construct a new run-time-adaptive moral theory itself so it is up to humans to craft machine processable (and therefore mathematical) heuristic moral models (Aliman and Kester 2022a) which needs continuous human supervision and feedback loops to keep evaluating if the moral model still functions as intended.

It is also for this reason that the EU AI act states that the objectives of the AI system should be human specified. For high-risk AI systems these objectives should be updatable by-design and machine readable (i.e. a mathematical objective function)

and contain this moral model. Thereby, it is possible to select a strategy wherein the responsibilities of designer/developer and moral authority are clearly separable (Aliman *et al.*, 2019) – a procedure termed “orthogonality-based disentanglement” – where the moral authority is engaged in the crafting of the objective function (also known as ethical goal function) which encodes the moral model while the designer is responsible for designing a system that optimizes on this updatable objective function/ethical goal function.

But how are we to heuristically represent morality in a mathematical model? For this we first need to address the meta-ethical question: “what would a proper mathematical representation of the moral model be”? In academia, various meta-ethical models have been proposed that can help make a decision between various normative ethical frameworks. One example is a comparative approach (Chang, 2012), in which the extent of which a value is realized per normative ethical theory is compared and expressed in percentages. Another related approach is the Expected Moral Value Theory (Lockhart, 2000) which adds the subjective assessment of the credibility of normative theories to the equation. Also here, multiple actions are reviewed against multiple moral theories, comparing to what extent the moral action that is proposed fulfils the values that are supposed to be fulfilled by each moral theory. The result is multiplied by the credibility of that moral theory according to the moral authority in a scale of 0.0 to 1.0. Yet another example of a meta-ethics model is the Agent-Deed-Consequence (ADC) Model (Dubljević, and Racine; Dubljević, 2020), interlacing consequentialism, deontology and virtue ethics through a score system.

meta-ethical
models

However, these methods are still based on predefined (mostly Western-oriented) normative theories that are both abstract and not translatable to a fine-grained moral programming code. Instead, a science based non-normative approach is required that allows for broader cultural perspectives, that involves the moral authority and other stakeholders to give feedback as they perceive the harm done and a method to continuously provide updates as society evolves.

So, instead of putting forward a normative theory, and instead of purely basing it on empirical ethical enquiries, Aliman and Kester (2022a) have developed a descriptive and explanatory framework which follows the scientific method to resolve the issues discussed above which is coined Augmented Utilitarianism (AU).

Augmented Utilitarianism (AU) is a generic non-normative framework insofar as it does not prescribe any values that should be set as ethical goals. Instead, it leaves it open to the moral authority to fill in the blanks. In short, it is a mathematical model that does not specify what should be optimized (as the conventional utilitarianism does). It does not apply one abstract, rigid theory, nor does it deliberately mimic cognitive dissonances or mirror moral contradictions. This also allows the framework to be applicable cross-culturally as the attributes and values can be adjusted according to the relevant society’s worldview and values. Insofar as it prescribes a method, the idea of optimization inevitably imposes its own ontological

Augmented
Utilitarianism

and epistemological assumptions, that is, respectively, its own theory of being and theory of knowledge. To ensure that the values of a pluralistic society can be encoded into moral models for AI systems, “AU targets what one could conceive of as a possible smallest heuristic moral superset (SHMS) capturing the plurality of candidate ethical frameworks available in practice for moral programming – with the pre-condition that AU must flag known formal inconsistencies to forestall predictable practical safety problems in AI deployment”. Such a SHMS aims at accommodating for the breadth and variety of widespread moral frameworks by covering the following parameters: the agent (as e.g. in the focus of virtue ethics), the action (as e.g. key to deontology), the user (as e.g. in the focus of consequentialism) and the experienter of the moral situation (as e.g. foregrounded in care ethics and many non-Western ethical frameworks).

On the whole, AU has its scientific basis in affective and cognitive neuroscience, and in moral psychology literature. Regarding the latter, AU is compatible with Gray and Schein’s theory of dyadic morality or dyadic completion. They argue that moral judgement and the nature of harm are not an objective matter of reason and should instead ‘be redefined as an intuitively perceived continuum’ (Gray and Schein, 2018). Dyadic Morality postulates that an act is assigned moral ‘value’ according to ‘norm violation’, ‘negative affect’ and crucially, ‘perceived harm’. In sum, AU allows the mathematical combination of consequentialism, deontology, virtue ethics and other frameworks and in addition perceived moral intuitions. It positions itself as a descriptive and explanatory framework that aims to capture the nuances in human morality’s functioning and moral pluralism rather than as a normative framework. Indeed, it recognizes that values can be attributed differently cross-culturally, and that harm is often perceived differently.

mathematical
scaffold

As such, AU is a ‘mathematical scaffold’ of how humans view scenarios and attribute value and is mathematically operationalizable in terms of ethical goal functions. This methodological framework follows the tenets of the scientific method. The functions are composed of attributes and values, which can be adjusted and determined by the relevant society. What matters is that this specification matches the world view of this society. Once specified, these functions can be tested in simulation environments and adjusted based on feedback from the relevant society. As such, AU adopts a dynamic process which is referred to as the ‘socio-technological feedback (SOTEF) loop’ (Aliman *et al.*, 2019). If a function did not capture what was originally intended by programmers, it can be amended to better match the relevant conception. The ethical goal functions are the requirements we have on the system – with such settings, a system may be able to deduce how to adapt at run-time. To the extent that these functions match with the world view of (the) relevant society, such functions should be able to resolve impossibility theorem scenarios, as described in more depth elsewhere (Gros *et al.*, 2024) and the engagement enables that ethical considerations are grounded in the lived experiences of individuals and adapt to shifts in societal attitudes over time.

socio-
technological
feedback
(SOTEF)
loop

5.5 Implications for moral programming in XR contexts

A first concrete step in constructing the ethical goal function for a specific application is to identify the relevant attributes in the ‘mathematical scaffold’.

The identification of relevant attributes is just a first step. The next step is how the different attributes are mathematically related to each other. Do they just add up like in a cost benefit analysis or are the different mathematical components related to the attributes correlated in a more complicated way? A final step is how the different attribute specific utility functions are shaped (e.g. concave, convex or linear) within the mathematical scaffold of AU. For all these different steps specific designed elicitation process will be needed.

The SOTEF loop describes how the different actors are engaged in the various processes. In figure 5.1 an illustration, presented in Aliman and Kester 2022a, of such a loop is shown.

Another angle (Heijnen *et al.* 2024) on interpreting the SOTEF loop is the life cycle of the AIXR system as depicted in figure 2. In this view four layers can be distinguished, the governance loop, the design loop, the development loop and the operational loop. The idea is that the actors identified in figure 1 are grouped in governance, design, development and operational teams. It is clear from both these pictures that already at an early stage of the AIXR system life cycle the process of moral programming should be centre stage. After all, the ethical goal function can

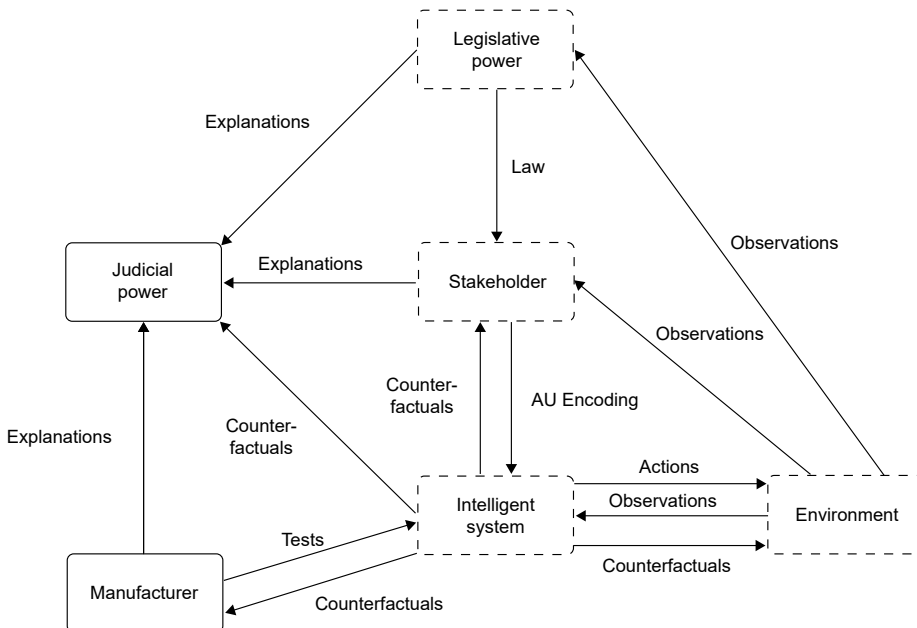


FIGURE 5.1 Illustration of the Socio-Technological Feedback Loop from a human-machine interaction perspective (Aliman and Kester 2022a)

also be regarded as a requirement specification that is needed not only to make run-time decisions but also to make design choices for the AIXR system. It can also be regarded as the function on which the system is optimising run-time, through which the system can be controlled by the operational team, by which the functioning of the system can be explained, and against which the system can be tested. It is therefore crucial that the relevant actors in the SOTEF loop are familiar with what the ethical goal function stands for and how it affects the functioning of an AIXR system. Relevant actors in the early stages of the AIXR system development are, referring to figure 1, the manufacturer and the stakeholders. Being familiar with AU encoding, they will identify the relevant attributes and, after consulting the law, come up with a first version of the ethical goal function. Thereafter, manufacturer, stakeholders and legislative power can evaluate the functioning of the AIXR system. In figure 1, we can also identify the inner feedback loop between stakeholder and intelligent system as a constant interaction where the AU function can be adapted and updated during operation that relates to the control loop in figure 5.2.

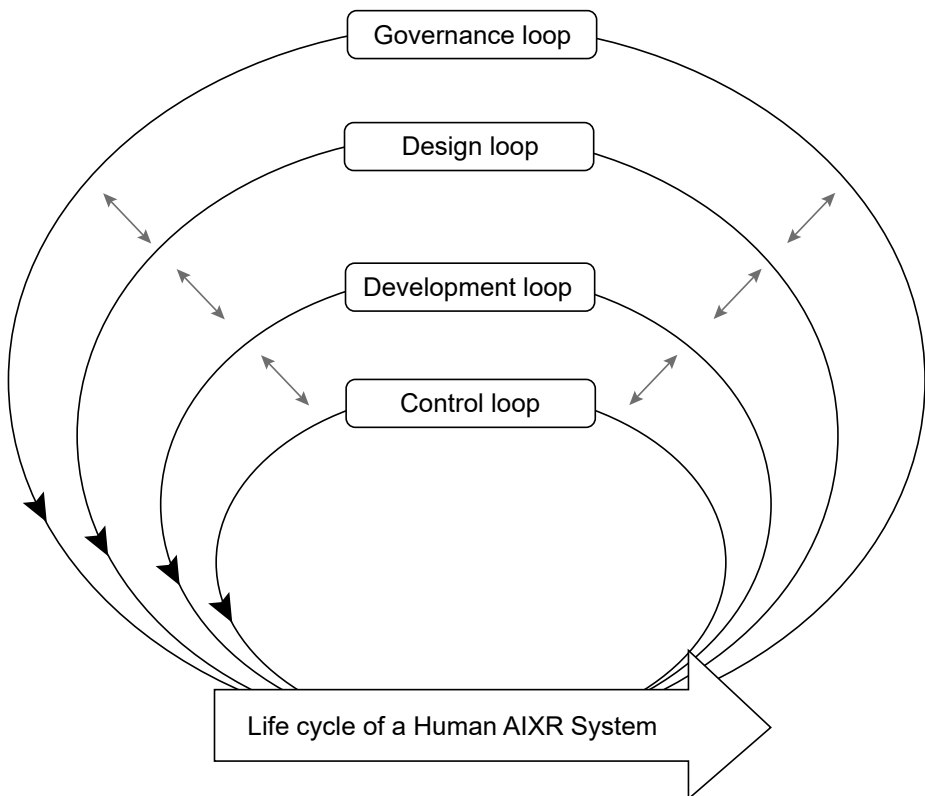


FIGURE 5.2 Illustration of the Socio-Technological Feedback Loop from a system life-cycle perspective

Also, a wider societal feedback loop with environment, legislative power and judicial power can be identified that acts as moral authority at the longer timescale of the governance loop in figure 2.

For high-risk AIXR systems it may often be too dangerous to test the system in real world scenarios. In that case, at early stages of the development, the test environment could e.g be a VR environment where certain aspects of the real world are simulated, also known as sandbox approach. The moral programming procedure with AU and the SOTEF loop also suggests that in such a simulated environment all the different people involved in construction of the ethical goal function could already have rich experiences able to augment their moral understanding and judgement. There is already experience with moral decision making by people in VR environments (Aliman and Kester, 2019), however, in many of these experiments the question was what people would decide themselves or how to augment their own moral decision making, sometimes even under time constraints. In the case of moral programming the question is more: “what would the relevant actors in particular the moral authority as experiencer want the AIXR system to do?”. For this question it is important to realize that the time pressure on decisions does not play a role since whatever the moral authority decides to value, the AIXR system could execute that instantly. As Lisa Barrett explains (Barrett, 2017) human core affect and rationality are inseparable. Interoception and core affect are part of the construction of moral judgments. Moral authorities should take as much time as needed for a more detailed construction of particularly the moral model of societal-level ethical goal functions.

high-risk AIXR
systems

5.6 Looking ahead

We have outlined the concept of a new approach of moral programming for AIXR systems dealing with many scientific and philosophical issues. One might wonder why such an approach is not already a well-established procedure? Many issues may play a role:

1. The procedure of moral programming is a transdisciplinary approach and involves scientists from many different domains. This can easily lead to misunderstanding, inconsistent concepts from the different disciplines, incompatible doctrines and lack of trans-disciplinary oversight.
2. Some people feel uncomfortable talking or analysing themselves about their values particularly if it concerns severe harm. They rather talk about what choice they would make in particular circumstances, but this is much harder to generalize.
3. Developers of systems do not always want to be transparent about the moral choices they make during system design, particularly if the objective of the

- system is not in the interest of users/society or it would unveil malfunctioning of the system, or when transparency would potentially damage Intellectual Property benefits or stands in the way of the overall business model.
4. The moral authority is not always keen to exercise that authority particularly if it is publicly transparent, also it is not always clear to the moral authority where he/she takes responsibility for.
 5. Policy makers who are key players in regulating procedures that enable the moral authority to put forth its preferred moral programming are hesitant in regulating this. Sometimes this is due to a lack of technological expertise, sometimes this is due to a powerful tech-lobby who opposes such regulation.
 6. A new type of problem arises with Large Language Models where model providers could use their model to spread misinformation and disinformation about AI that is in the favour of that company. Particularly misinformation and disinformation about misconceived existential risks (see e.g. (Duarte *et al.*, 2024)) can complicate value agreements and regulation activities.

It is for these reasons that for moral programming an integral approach is needed that also takes the issues into account that are not directly related to moral theory, value elicitation and moral programming.

In addition, the possibility of so-called AI agents in XR (Wan *et al.*, 2024) is a new challenge that needs to be addressed. Advances in Generative AI facilitate the development of these agents both for beneficial and malicious use. Unintentional substantial risks can arise in the domain of privacy and concerning the overestimation of AI agent capabilities. Defences against malicious use are even more challenging. Detection of artificial agents deployed by malicious actors will be very difficult because every detection mechanism can in principle be countered unless the AI agent is tested on something AI agents cannot do: the construction of new scientific or moral theories. In Aliman and Kester 2022b and 2023 possible defences in XR environments are discussed. Moral programming can also help here to determine possible risk mitigation strategies.

References

- Aliman, N.M., and Kester, L., 2019. Extending socio-technological reality for ethics in artificial intelligent systems. In: IEEE international conference on artificial intelligence and virtual reality (AIVR), pp. 275–2757.
- Aliman, N.M., Kester, L., Werkhoven, P. and Yampolskiy, R., 2019. Orthogonality-based disentanglement of responsibilities for ethical intelligent systems. In: Hammer, P., Agrawal, P., Goertzel, B., Iklé, M. (Eds.) Artificial General Intelligence. AGI 2019. Lecture Notes in Computer Science, pp. 22–31. Springer Verlag. https://doi.org/10.1007/978-3-030-27005-6_3.

- Aliman, N.M. and Kester, L., 2021. Epistemic defenses against scientific and empirical adversarial AI attacks, *CEUR Workshop Proceedings*, vol. 2916.
- Aliman, N.M. and Kester, L., 2022. Moral programming: crafting a flexible heuristic moral meta-model for meaningful AI control in pluralistic societies, in: B. Wernaart (Ed.), *Moral design and technology*. Wageningen Academic Publishers, Wageningen, the Netherlands, pp. 63–80 https://doi.org/https://doi.org/10.3920/978-90-8686-922-0_4.
- Aliman, N.M. and Kester, L., 2022. VR, deepfakes and epistemic security. In: *IEEE international conference on artificial intelligence and virtual reality (AIVR)*. IEEE, pp. 93–98.
- Aliman, N.M., Kester, L., 2023. AI-related misdirection awareness in AIVR. *PhilPapers*.
- Alnagrat, A., Ismail, R.C., Idrus, S.Z.S., and Alfaqi, R.M.A., 2022. A review of extended reality (XR) technologies in the future of human education: current trend and future opportunity. *Journal of human centered technology*, 1:81–96.
- Andrade, T., and Bastos, D., 2019. Extended reality in iot scenarios: concepts, applications and future trends. In: *5th Experiment International Conference (exp. at'19)*. IEEE, pp. 107–112.
- Andrews, C., Southworth, M.K., Silva, J.N., and Silva, J.R., 2019. Extended reality in medical practice. *Current treatment options in cardiovascular medicine*, 21:1–12.
- Barrett, L.F., 2017. The theory of constructed emotion: an active inference account of interoception and categorization. *Social cognitive and affective neuroscience*, 12. <https://doi.org/10.1093/scan/nsw154>.
- Boyce, M.W., Thomson, R.H., Cartwright, J.K., Feltner, D.T., Stainrod, C.R., Flynn, J., ... and Rovira, E., 2022. Enhancing military training using extended reality: a study of military tactics comprehension. *Frontiers in virtual reality*, 3:754627.
- Burns, M.B., Lebkuecher, G., Rahman, S., Roytman, M., Samoska, S., and Vukov, J., 2022. Extended frameworks for extended reality: ethical considerations. *AJOB neuroscience*, 13:171–173.
- Carter, M., and Egliston, B., 2020. *Ethical implications of emerging mixed reality technologies*, the University of Sydney.
- Casini M., 2022. Extended reality for smart building operation and maintenance: a review. *Energies*, 15:3785. <https://doi.org/10.3390/en15103785>.
- Chang, R., 2012. Are hard choices cases of incomparability? *Philosophical issues*, 22:106–126.
- Çöltekin, A., Lochhead, I., Madden, M., Christophe, S., Devaux, A., Pettit, C., ... and Hedley, N., 2020. Extended reality in spatial sciences: a review of research challenges and future directions. *ISPRS international journal of geo-information*, 9:439.
- De Haan, J., 2001. The definition of moral dilemmas: a logical problem. *Ethical theory and moral practice* 4:267–284. <https://doi.org/10.1023/A:1011895415846>.
- Duarte, T., Barrow, N., Bakayeva, M., and Smith, P., 2024. The ethical implications of AI hype. *AI and ethics*, 1–3.
- Dubbink, W., 2018. A typology of ethical problems. *Ethical perspectives*, 25: 683–714.
- Dubljević, V., and Racine, E., 2014. The ADC of moral judgment: opening the black box of moral intuitions with heuristics about agents, deeds, and consequences. *AJOB neuroscience*, 5:3–20. <https://doi.org/10.1080/21507740.2014.939381>.

- Dubljević, V., 2020. Toward implementing the ADC model of moral judgment in autonomous vehicles. *Science and engineering ethics* 26:2461–2472. <https://doi.org/10.1007/s11948-020-00242-0>.
- Garcia Estrada, J., Prasolova-Førland, E., Kjeksrud, S., Themelis, C., Lindqvist, P., Kvam, K., and Ricci, S., 2024. Military education in extended reality (XR): learning troublesome knowledge through immersive experiential application. *The visual computer*, 1–30.
- Gros, C., Kester, L., Martens, M., and Werkhoven, P., 2024. Addressing ethical challenges in automated vehicles: bridging the gap with hybrid AI and augmented utilitarianism. *AI and ethics*. <https://doi.org/10.1007/s43681-024-00592-6>.
- Gowans, C.W. (Ed.), 1987. *Moral dilemmas*. Oxford University Press, USA.
- Heijnen, M., Schoonderwoerd, T., Neerincx, M., van der Waa, J., Kester, L., van Diggelen, J., and Elands, P.A., 2024. Socio-technical feedback loop for responsible military AI life-cycles from governance to operation. In: *Responsible use of AI in military systems*. Chapman and Hall/CRC, Oxford, UK, pp. 17–36. <https://doi.org/10.1201/9781003410379-3>.
- Indradevi, R., Venkataswamy, K.P., Saravanan, S.A., and Mittal, A., 2024. Ethical considerations as limitations of extended reality: an empirical study through the lens of retail consumers. In: *Multidisciplinary applications of extended reality for human experience*. IGI Global, Hershey, PA, USA, pp. 44–63.
- Jones, T.M., 1991. Ethical decision making by individuals in organizations: an issue-contingent model. *Academy of management review*, 16:366–395.
- Kharvari, F., and Kaiser, L.E., 2022. Impact of extended reality on architectural education and the design process. *Automation in construction*, 141:104393.
- Lockhart, T., 2000. *Moral uncertainty and its consequences*. Oxford University Press, Oxford, UK.
- Marr, B., 2021. *Extended reality in practice: 100+ amazing ways virtual, augmented and mixed reality are changing business and society*. John Wiley & Sons, Hoboken, NJ, USA.
- Messerli, M. and Reuter, K., 2017. Hard cases of comparison. *Philosophical studies*, 174:2227–2250. <https://doi.org/10.1007/s11098-016-0796-y>.
- Millar, J., 2017. Ethics settings for autonomous vehicles. In: Patrick Lin, P., Abney, K. and Jenkin, R. (editors). *Robot ethics 2.0: from autonomous cars to artificial intelligence*, Oxford Academic, Oxford, UK, pp. 20–34.
- Morimoto, T., Kobayashi, T., Hirata, H., Otani, K., Sugimoto, M., Tsukamoto, M., ... and Mawatari, M., 2022. XR (extended reality: virtual reality, augmented reality, mixed reality) technology in spine medicine: status quo and quo vadis. *Journal of clinical medicine*, 11:470.
- Ramirez, E.J., Jennett, S., Clay, D., and Gandhi, M., 2023. Extended reality, control, and problems of the self: an ethical analysis. In: *Exploring extended realities*. Routledge, Oxfordshire, UK, pp. 104–129.
- Santoso, H.B., Wang, J.C., and Windasari, N.A., 2022. Impact of multisensory extended reality on tourism experience journey. *Journal of hospitality and tourism technology*, 13:356–385.

- Sartre, Jean-Paul, 1957. Existentialism is a humanism, Trans, Philip Mairet, in Walter Kaufmann (ed.), *Existentialism from Dostoevsky to Sartre*. Meridian, New York, NY, USA, pp. 287–311.
- Schein, C., and Gray, K., 2018. The theory of dyadic morality: reinventing moral judgment by redefining harm. *Personality and social psychology review*, 22:32–70. <https://doi.org/10.1177/1088868317698288>.
- Schneider, R.O., 2021. Extended reality: the next frontier of design. *The international journal of design management and professional practice*, 15:43.
- Sharma, R., 2021. Extended reality: It's impact on education. *International journal of scientific & engineering research*, 12:247–251.
- Sinnott-Armstrong, W., 1987. Moral dilemmas and 'ought and ought not'. *Canadian journal of philosophy*, 17:127–139. <https://doi.org/10.1080/00455091.1987.10715904>.
- Statman, D., 1990. The debate over the so-called reality of moral dilemmas. *Philosophical papers*, 19:191–211. <https://doi.org/10.1080/05568649009506337>.
- Timmermann, J., 2013. Kantian dilemmas? Moral conflict in Kant's ethical theory. *Archiv für geschichte der philosophie*, 95:36–64. <https://doi.org/10.1515/agph-2013-0002>.
- Tromp, J., Le, C., Le, B., and Le, D.N., 2018. Massively multi-user online social virtual reality systems: ethical issues and risks for long-term use. *Social networks science: design, implementation, security, and challenges: from social networks analysis to social networks intelligence*, pp. 131–149.
- Van de Poel, I. and Kroes, P., 2014. Can technology embody values? In: Kroes, P. and Verbeek, P.P. (eds.). *Moral agency and technical artefacts*. Springer, Dordrecht, the Netherlands, pp. 103–124.
- Vasarainen, M., Paavola, S., and Vetoshkina, L., 2021. A systematic literature review on extended reality: virtual, augmented and mixed reality in working life. *International journal of virtual reality*, 21:1–28.
- Wan, H., Zhang, J., Suria, A.A., Yao, B., Wang, D., Coady, Y., and Prpa, M., 2024. Building LLM-based AI agents in social virtual reality. In: *Extended abstracts of the CHI Conference on Human Factors in Computing Systems*, pp. 1–7.
- Wernaart, B., 2021. Developing a roadmap for the moral programming of smart technology. *Technology in society*, 64:101466.
- Wernaart, B., 2022. Moral authority. In: *Moral design and technology*. Wageningen Academic Publishers, Wageningen, the Netherlands, pp. 29–46.
- Zhang, J., Lu, V., and Khanduja, V., 2023. The impact of extended reality on surgery: a scoping review. *International orthopaedics*, 47:611–621.

Citizen science for nature

Simona Orzan and Gerard Schouten

Abstract

Public engagement in research activities through citizen science is growing in popularity and significance, especially in the context of increasing concerns about the impacts of climate change, biodiversity loss and degradation of our natural environment. This chapter explores the citizen science movement, mainly using real-life examples, and aims to provide a short overview of its history, status, opportunities and challenges. Maybe it will even inspire you to join!

6.1 Introduction

Humans are naturally curious beings, intelligent and (mostly) social – the precise core ingredients that make a good scientist. This means that everybody can contribute to science! Citizen science brings this idea to life by engaging volunteers, often non-scientists, in accessible yet meaningful research activities. From monitoring wildlife populations or measuring air quality to searching for extraterrestrial signs of life, enthusiastic citizen scientists collect and analyse a huge volume of data, contributing to scientific data-driven insights. Some of the projects even invite participants to act as actual scientists, articulating research questions, performing experiments, taking measurements (see Figure 6.1), analysing their findings and interpreting the results.

In fields like environmental research and nature conservation, this collaborative approach has proved invaluable in helping fill critical data gaps in a cost-effective manner. Besides the efficiency, it also fosters a sense of environmental stewardship, crowdsources creative out-of-the-box solutions and raises public awareness of pressing ecological challenges.

The fast development and increased availability of technology in recent years have opened new possibilities for citizen science, in terms of practical observation and analysis tools as well as community-forming. This has resulted in the exponential growth of citizen science initiatives, especially since 2000.

In this chapter, we reflect on how the citizen science movement has been transforming the relationship between science and society and explore its potential for the future. We start with a short history (Section 2), then discuss its impact on



FIGURE 6.1 Equipped with a torch and landing net, a group of nature enthusiasts visit pools in and around the city of Eindhoven. They record observations of amphibians and monitor aquatic plants in the pools. Every pool will be checked several times in the period from February to the end of June. This happens at the end of the day at late dusk since amphibians are nocturnal; in one evening the group walks (or cycles) a route and stops by five to eight pools. Aquatic plants are also monitored; the municipality wants to know about the spread of invasive plant species sold by garden centres and spread via garden ponds into our environment. The counts of amphibians are written down on paper in the field and at the end of the season the data of each pool for all the visits is collected on a form and sent to RAVON (the national agency that keeps record of trends in aquatic life in the Netherlands).

society at large and on individuals (Section 3). Next, we shift to the scientific perspective, highlighting how citizen participation enriches research and discovery (Section 4). Finally, we conclude with key takeaways and reflections on the future of citizen science (Section 5).

We highlight some citizen science projects to illustrate the main concepts and challenges, but this is definitely not a complete overview! This would be too difficult in the very large, fragmented and highly dynamic global and national landscape.

6.2 A short history

citizen science

Although the term ‘citizen science’ was only coined in the 1990s (Irwin, 1996), the involvement of volunteers in scientific activities like data collection goes back much further. The first recorded initiative in Europe that seems to fit the definition dates from 1736: *The Woodland Nature Trust* in the UK asked thousands of volunteers to record ‘the signs of the seasons’. This project has gathered a very large following and is still active today under the name *Nature’s calendar*.¹

Over the next 200 years, the number of European citizen science programmes slowly increased to about 20 in 1950 and about 50 in the 1990s. With the internet as a democratizing force and digital recording technologies from the 2000s, an explosion of citizen science projects followed. Now we have several hundred large initiatives running, and many spontaneous short-term initiatives too. Figure 6.2 gives an impression of this explosive evolution in the world and in the Netherlands. Notably, since around 2016 the number of active projects has appeared to stagnate and even decrease. This could signal a certain saturation of the topics and the public, and the start of a stabilization period, in which the existing citizen science initiatives could become more well known, established and impactful.

6.2.1 Citizen science in the Netherlands

The first serious Dutch citizen science project was PhenoWatch² (1868), mapping the phenology of plants, birds and insects to forecast hay fever and for other medicinal benefits. This initiative is still running today and has enlarged its scope to include climate change monitoring.

Collecting data about the natural world has, in fact, been the focus of most citizen science communities in the Netherlands. Most work is being done through several specialized organizations like *FLORON* (Floristisch Onderzoek Nederland, since 1988) – which is thoroughly monitoring flowering plants, *SOVON* (Stichting Ornithologisch Veldonderzoek Nederland) Sovon Vogelonderzoek Nederland, since 1973) – monitoring breeding, migratory and wintering birds, and *RAVON* (Reptielen Amfibieën Vissen Onderzoek Nederland, since 1987) – monitoring reptiles, amphibians, and fish. The more recent platform *Waarneming.nl* (since 2004) collects observations of the whole ‘web-of-live’. *FLORON*, *SOVON* and *RAVON* are established organizations, with a long history and strict observations protocols, while *Waarneming.nl* is representative of a younger organization, with free input formats, validated by both digital technology (mainly AI) and expert knowledge. The amount of data collected on *Waarneming.nl* is substantial. In 2023 more than 100,000 citizen scientists recorded almost 14 million observations (of which approximately 10 million were validated).

1 <https://naturescalendar.woodlandtrust.org.uk/>.

2 <http://natuurkalender.nl/>.

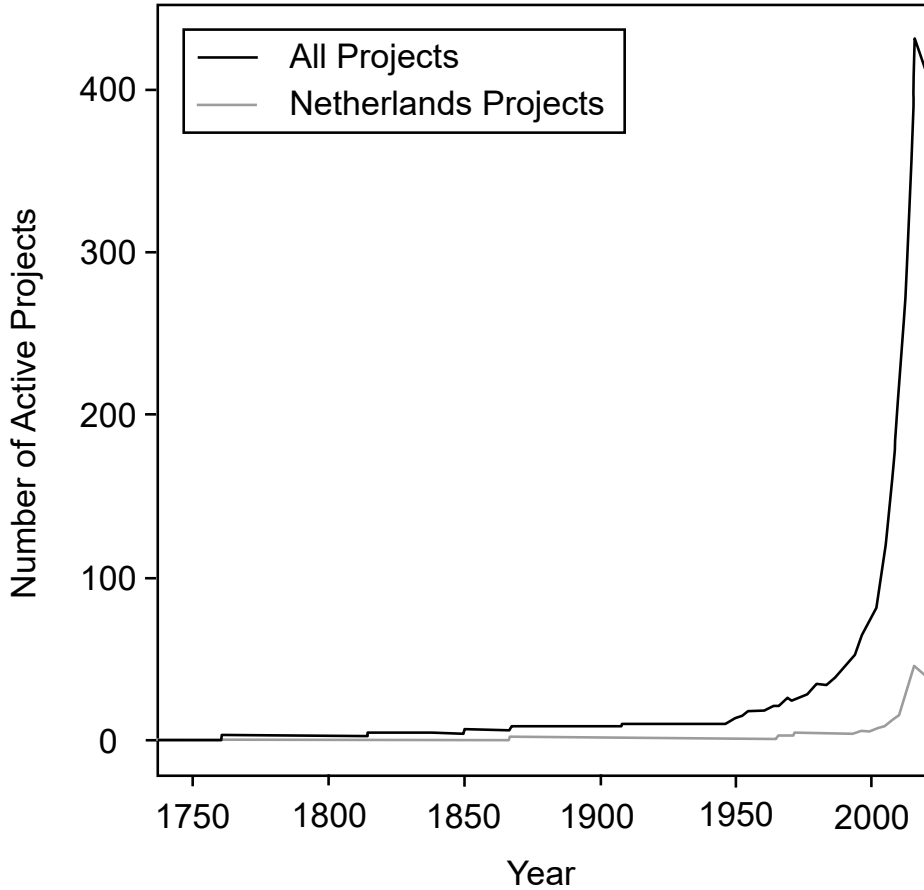


FIGURE 6.2 The number of active environmental citizen science projects through the years, according to the inventory maintained by the European Commission (JRC, 2018)

A complementary type of initiative, focusing on education just as much as gathering data (or more!), are citizen science programs organized by knowledge institutions like Naturalis Biodiversity Center, Wageningen University, etc. Programmes like Discover Urban Nature³ or Beach Shell Count⁴ are usually short term, targeting families, supported by clear instructions and attractive materials. The oldest and best-known project aimed at the general public is the twice-a-year recurrent National Bird Counting,⁵ where all Dutch people are invited to observe their garden (or balcony) for half an hour and report which birds they see. Over 100,000 people participate. The result is a ranking of most popular birds, on all aggregation levels from neighbourhood to whole country. In the last edition, the top three in the

³ <https://www.naturalis.nl/verborgenstadsnatuur>.

⁴ <https://www.naturalis.nl/en/museum/whats-on/beach-shell-census>.

⁵ <https://www.vogelbescherming.nl/tuinvogeltelling>.

Netherlands were: 1) House Sparrow (*Passer domesticus*); 2) Great Tit (*Parus major*); and 3) Blue Tit (*Parus caeruleus*). This first project (2003) was followed by many similar initiatives – like the National Bee Counting,⁶ and most recently the National Whale Counting,⁷ which had its first edition on 7 September 2024 and attracted much-needed attention for whale species at the Dutch coast.

On a smaller scale there are also individual researchers that reach out to the general public for a specific research project. A nice example is the PhD project of Auke-Florian Hiemstra, on ‘animals as architects in the Anthropocene’; Hiemstra studies structures made by animals, like nests and casings. Nowadays birds often use plastics to build their nests. Why do animals use these new materials? To answer this question, he appeals to existing citizen science initiatives to collect photographs of nests of Eurasian Coots (*Fulica atra*) made with plastics in the canals of Leiden (IVN, 2024).

6.2.2 *Now it's all digital*

Nowadays, all citizen science projects use digital support to build communities and aggregate the efforts, in the form of social media (e.g. Facebook) groups, dedicated websites, apps, digital magazines. Some prominent examples of very active global online projects are gbif.org, which aggregates biodiversity data worldwide, or sensor-community.com, which monitors air quality through a network of DIY sensors. In the Netherlands, *Waarneming.nl* has its own app *ObsIdentify*, that facilitates species identification with AI, location sharing, easy recording and validation of observations.

The strongest impact of the digital revolution is perhaps felt in the more specific situations and topics. Monitoring sea slugs species, for instance, is one such specialized niche where online tools facilitate the building of a small but enthusiastic community (Blauwtipje⁸). Digital platforms are also empowering remote and marginalized communities. Forest people, for example, historically isolated and exploited because of language and literacy barriers, are now getting a voice through apps customized for citizen science (Smith, 2022).

Platforms like Zooniverse,⁹ Iedereenwetenschapper,¹⁰ and SciStarter¹¹ function as marketplaces for new citizen science initiatives or projects. Organizations can submit their research projects requesting volunteers, and interested people can filter and select which ones to join. Besides the dedicated organizations, most universities and research institutes studying environmental and ecological issues

6 <https://www.nationalebijtelling.nl/>.

7 <https://rugvin.nl/nationale-walvistelling/>.

8 <https://blauwtipje.nl>.

9 <https://zooniverse.org>.

10 <https://www.iedereenwetenschapper.be/>.

11 <https://scistarter.org/>.

also actively employ, support, curate and supervise citizen science projects. In the Netherlands, such hubs can be found for instance at Wageningen University, Naturalis Biodiversity Center, and the Netherlands Institute of Ecology (NIOO).

6.3 Impact on society

It is widely assumed that citizen science practices have a significant positive impact on society at large, by directly targeting Sustainable Development Goals (SDGs), increasing people's engagement with environmental and ecological issues, and even local and global policy making.

6.3.1 Impact on education

Without a doubt, most people learn best by performing hands-on activities, and citizen science is a very powerful example of how serious education topics can be brought to life. By involving people in real-world research, these projects teach participants not just about the specific topic at hand but also about scientific methods and critical thinking, and spark curiosity about the enchanting natural world.

The impact of citizen science on education is, in fact, a new social research topic in itself. Studies have focussed, for instance, on Bioblitz events. This is a citizen science practice that began in 1996 in Washington, D.C., and has since spread globally, including to the Netherlands; in a BioBlitz, volunteers work with scientists to identify and record as many species as possible within a specific area and time. Studies have shown that BioBlitz events can really foster a sense of stewardship and a deeper connection to the environment, especially among young people, who already have an affinity for nature, particularly for the animal kingdom. Also, this effect is lasting, as BioBlitz participants typically remain active citizen scientists for months after the event (Meeus *et al.*, 2023).

Bioblitz

Indeed, children benefit most from being taught scientific concepts in fun, interactive ways. By counting birds, identifying insects, or determining the quality of river water, they learn not only biology, but also maths, physics, data analysis. It's a win-win situation – kids get to play and learn simultaneously, while science benefits from their fresh perspectives and eager participation.

A small point of concern, however, is that while these practices lead to knowledge accumulation and some increased awareness of environmental and ecological topics, they do not necessarily lead to long-term behavioural change (Jordan *et al.*, 2011). Changing is challenging anyway (see also Chapter 4), and there seems to be a missing link between knowing and doing. People should be convinced that their actions matter (Jordan *et al.*, 2011). We also know that factors like participation frequency and project design can positively impact attitudes towards environmental action, but more research is needed to fully understand the process and improve the behavioural change outcomes (Conrad and Hilchey, 2011).

6.3.2 Impact on Sustainable Development Goals

The idea of engaging the public in scientific research naturally aligns with global efforts to tackle challenges like climate change, biodiversity loss and water scarcity. As shown in Figure 6.3, some SDGs attract more citizen science projects than others. Nature conservation, for instance, is a top focus worldwide, reflecting a widespread public interest in preserving our planet’s ecosystems.

In the Netherlands, the distribution of citizen science projects is quite similar to the global trend, with a few notable differences. There is an emphasis on SDG 17 (Strengthening Global Partnerships) and SDG 6 (Ensuring Water Availability and Sustainable Management). The first points to the long-standing Dutch attitude of outreach and openness towards international collaboration. The latter could be explained by the country’s unique relationship with water management and its historical struggle against flooding; this expertise is now useful in other parts of the world. Another notable difference between the global and Dutch SDG focus is terrestrial and marine biodiversity (SDG 15 and SDG 14). In the Netherlands, although these two SDGs together still account for the largest impact overall, the national efforts lag behind the global ones and this is especially remarkable in the context of NL being ‘champion biodiversity loss’ (NOS, 2024). Finding ways to accelerate public awareness and action on biodiversity decline in the Netherlands is therefore an urgent challenge.

The greatest contributions of citizen science projects are to SDGs related to life on land, sustainable cities, health and water (Fraisl *et al.*, 2020). In high-income countries, the participation and impact of citizen science projects (especially on SDGs 3, 11 and 13) is definitely higher than in low- and middle-income countries, where immediate local needs take priority over national environment monitoring programmes. To fully realize the potential of citizen science for SDGs, there is a

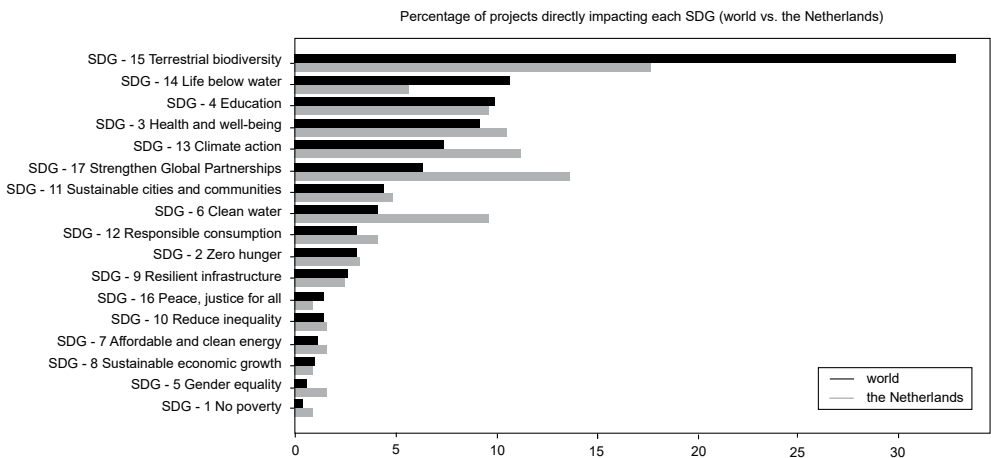


FIGURE 6.3 Impact of all citizen science projects on SDGs

need to demonstrate its value, build partnerships, and align citizen science with SDG indicators and targets at both country and more local levels (Fraisil *et al.*, 2020). In the spirit of ‘what you measure, you manage’, the question arises: How can we quantify SDG targets in order to make progress more concrete and increase the impact of various initiatives?

6.3.3 *Impact on policy and decision-making*

Citizen science has huge potential to strengthen the two-way avenue between citizens and (environmental) policymaking. On the one hand, it empowers citizens to express their voice and influence policy from the ground up; on the other hand, it gives policymakers the tools to support their decisions in a transparent and objective way, for better public acceptance.

Zooming in on the first perspective, data collected by worried citizens can be an objective basis on which to advocate for nature and environmental protection and to influence legislation at local and national levels. Civil society organizations (CSOs) often use the data to hold industries and governments accountable for illegal or harmful activities, such as illegal deforestation, poaching and mining. ‘Extreme citizen science’¹² (Smith, 2022) even takes the idea a step further and aims at environmental justice, i.e. giving voice to (geographically or otherwise) marginalized communities, which are often overlooked by decision-makers, or even nature itself. An inspiring example is how a low-literate hunter-gatherer population in Congo – supported by a well-designed icon-based app¹³ (Sapelli, 2020) – collected information about illegal logging and poaching. The goal was to protect their livelihood and the forest by facilitating a healthy shared forest management system with the local authorities. The app also collected their indigenous traditional ecological knowledge about protecting the environment, as a rich source of environmental management ideas for policymakers.

In developed countries, the initiative to use citizen science lies mostly with the authorities. The distance between citizens and local authorities is easier to bridge and the question is somewhat reversed: Can citizen science be used to make the policy decisions public-informed, or at least more transparent to the public, and increase their acceptance? For instance, if the city centre needs to be closed to cars, it is helpful if the citizens understand the benefits for the environment and air quality. Many Dutch municipalities actively encourage citizens to gather data on environmental topics – invasive species, wildlife, air quality, noise pollution etc. – with the primary goal of engaging the public in the local community and its governance.

Still, despite promising successes, a lot more research, experimentation and learning is needed for people and policymakers to truly be able to use citizen science for democratic engagement.

12 <https://www.ucl.ac.uk/geography/ECSAnVis/en/>.

13 <https://www.sapelli.org/>.

6.3.4 *The citizen*

Acting as a citizen scientist is guaranteed to bring the fulfilment of contributing to a higher purpose and the joy of connecting with nature and with other people (Meeus, 2023). Check The Citizen Science Song by Monty Harper¹⁴ for a moment of inspiration and the story of Jos in CASE 6.1.

In theory, everybody can be a citizen scientist, and this is one of the most naturally inclusive activities. Technology even helps break through the barriers of illiteracy and help the environment with the experience and deep insights of wise, caring locals, as shown by the extreme citizen science movement (ECSAnVis, 2020). However, currently, the citizen science participants are not representative of the general population, as they are disproportionately white, mostly male, highly educated, affluent, middle-aged and older (Pateman *et al.*, 2021).

CASE 6.1 INTERVIEW WITH JOS VAN DE KERKHOF, SPRING 2024

Against all odds

Jos (66) remembers that in his childhood lapwings (*Vanellus vanellus*) returned each spring in large numbers to breed in the meadows around his family farm. Now, only one pair is left. Jos doesn't like that, as he feels a deep connection to nature. 'I don't want to lose the sound of spring,' he says. But intensive farming and increased predator prevalence have thrown lapwings' habitat, and nature in general, seriously out of balance.

To help these birds survive against the unfavourable odds, Jos volunteers his time to talk to farmers for land access, finds and counts nests, and tries to save them by placing them in visible baskets so they can be safely and temporarily moved during farming activities. Together with other volunteers, they also set up electric fences and keep vegetation low to protect the baby lapwings from predators. Technology like drones with infrared cameras, apps etc. could help with monitoring, but these volunteers still stick to traditional methods.

This work connects to an old Dutch tradition of finding the first lapwing egg each year. This practice is now restricted by law, for the birds' protection. The unintended consequence is that this also limits people's bond with nature, eventually making it easier to harm exactly the habitats it wants to protect.

Jos and most of his fellow volunteers are in the 45–70 age range. 'It's hard to get young people interested,' he admits. Nature moves slowly, unlike the fast pace and fast rewards of modern life. A spectacular nature experience usually means that you have to spend many hours in the field. 'Young people, grown up with smartphones and used to instant access of information, often lack this patience,' he argues. But it's worth persevering because if we lose touch with nature, we will lose much more than the lapwings and the sound of spring.

14 <https://youtu.be/sAbLTDFv5Cw?si=oPRXZtpvDul3tHoR>.

Mobilizing young generations to become engaged in environmental protection seems hard. Research (Stapelton *et al.*, 2005) suggests that to increase its chances of success, a citizen science project aimed at young people should contain at least some of these ingredients: virtual or augmented reality apps; gamification; clear and fast rewards; social connections; an emphasis on skills development. Luckily, school education curricula are increasingly including such nature projects and therefore guarantee that all young people get some exposure to these important topics.

6.4 The science

While it is clear by now that citizen science is a powerful tool for raising public awareness and engagement with critical natural issues, questions remain about how seriously it intersects with actual professional research. Concerns about data quality, scientific rigour, and ethical considerations are very legitimate and will be discussed shortly in this section.

6.4.1 Citizen science is science ...

Considering that 25% of citizen science projects lead to scientific peer-reviewed publications (Davis *et al.*, 2023), and that currently more than 10,000 scientific publications make use of GBIF (Global Biodiversity Information Facility: a digital platform which holds more than 3 billion records worldwide and over 100,000 datasets), we can safely state that citizen science has proved its value. The greatest impacts are seen in biological studies of global climate change, such as landscape ecology and macro-ecology, as well as in research on rare and invasive species and ecosystems (Dickinson *et al.*, 2012).

Obvious benefits for research are cost-effectiveness and increased data collection as well as sharing, and search capabilities. Biodiversity monitoring on a large scale, and at high spatial and temporal resolutions, as performed by GBIF or *Waarneming.nl*, would not be possible without the contribution of thousands of volunteers. The same holds for air quality monitoring, where high resolution data is essential for drawing valid conclusions due to the mostly rather localized nature of pollution. The ‘*Snuffelfiets*’ project, where people measured air quality while biking to their destination, collected 25 million measurements and provided valuable insights into the air quality and pollution hotspots in Utrecht (Hendrix *et al.*, 2021).

While most individual contributions are fuel for the actual research, and not noteworthy on their own, once in a while there is a serious breakthrough – like the discovery of a new ‘*voorwerp*’! In 2007, Dutch citizen scientist Hanny van Arkel, while volunteering for Galaxy Zoo, found a mysterious astronomical object that puzzled the astronomers for about a decade afterwards. Since its nature was unclear at the time, it got the name ‘*Hanny’s Voorwerp*’. Later, similar objects acquired the nickname ‘*voorwerpjes*’ (little objects). Other impressive examples are the discovery of a new ant species by volunteers in the AntWeb-4 project, in 2019, and the

discovery of a new class of galaxies, the ‘green peas. Also, the first proof of insect decline (by 75%!), described in the landmark paper of Hallmann *et al.* (2017), relied on the systematic data collection effort by nature enthusiasts for 27 years. Finally, probably the most spectacular moment in this category was when players of the online game FoldIt¹⁵ managed to decipher the structure of a protein-cutting enzyme involved in the replication of the AIDS virus. This achievement had eluded scientists for over a decade. The citizen scientists’ solution offered insights that could lead to new drug development for HIV (Khatib *et al.*, 2011).

When such citizen scientists are involved in more complex aspects of research, such as data/image analysis, model development, and hypothesis testing, and if they contribute significantly to the research process, they are sometimes included as co-authors on publications. This trend is actually increasing (Gunnell *et al.*, 2018), indicating a growing recognition of what non-professionals can contribute to scientific results, and a sign of the ‘democratization of science’.

6.4.2 ... *but tread carefully!*

Non-expert involvement can also sometimes lead to loss of scientific rigour, loss of data quality, and potential biases. When sensors are employed in large numbers, such as for measuring air or water quality with citizen science projects, they are usually low cost and low precision, leading to imprecise results and therefore loss of scientific rigour. Ecological and other observation data often contain spatial biases, as observers tend to favour easily accessible areas, such as those near roads, or areas with a high biodiversity, such as protected areas (Boakes *et al.*, 2010; Johnston *et al.*, 2020).

To mitigate these risks, methods like sensor calibration, standardized protocols, clear guidelines and peer reviews must be in place. More recently, AI tools are being employed for tasks like identification of species from the submitted photos and automatic data cleaning (such as filtering out irrelevant records like ‘selfies’), and this has increased overall data quality (Loftian *et al.*, 2021).

Luckily, data can often also be cleaned and corrected after collection. Sensor calibration involves comparing measurements from the low-cost sensors with those from high-precision sensors, and subsequently applying a correction factor for all the data. With image data, AI can find and discard incorrectly labelled entries. Recent research (Cretois *et al.*, 2021; Jacobusse and Jongejans, 2024) is looking for more precise ways to improve data quality, such that the large citizen science datasets available in a variety of fields become more reliable and therefore more useful to science. Detecting and correcting biases in citizen science data remains a serious and interesting challenge.

15 <https://fold.it/>.

6.4.3 *Ethical considerations*

At the intersection of the social and scientific domains, both highly ethically sensitive, citizen science is bound to have its ethical challenges as well. The biggest one is probably *location privacy*: citizens are often asked to share their locations, which can be a random forest, but can also be their home address, if the sensor they are using is placed in private spaces such as their backyard. Despite data often being anonymized to some degree, possible personal information leaks remain a point of attention and require clear protocols and informed consent on how personal data is used.

Another concern is that publicizing the *location of endangered species* (Red List) may increase dangers like poaching, exploitation or disturbance, therefore unintentionally harming the species that these initiatives are aiming to protect. Mitigation includes limiting access to this kind of sensitive information, spatially blurring the locations or presenting only aggregated results to the public.

Finally, *volunteer rights* are also an important point, as participants have the right to know how their collected data is used and be given proper credit for their (mostly unpaid) work. Moreover, the activities need to be inclusive, ensuring that participants from diverse backgrounds have equal chances to benefit from the learning experience. So far, as pointed out in 6.3.4, this seems to be very much the case.

6.4.4 *How to design a successful citizen science project*

Citizen science projects can be trusted to produce sound results when they are carefully designed and executed. Even if they do not produce sound results, they can still be considered successful in terms of generating awareness, learning and community engagement.

The key factors in most successful citizen science projects turn out to be (Robinson *et al.*, 2020; Brown and Williams, 2018):

- A purposive project* – if possible, design a purposive project, where volunteers are assigned a specific part of the study, like an area and time to cover, therefore preventing many of the biases;
- Clarity* – in the objectives and all communication;
- Protocols* – including good training;
- Data quality control* – for instance with peer-, AI- or professional reviews;
- Ethical and encouraging engagement* – by aligning the project goals with the participants interest, using friendly material, acknowledging contributions and offering incentives.

6.4.5 *Participate*

If this overview was inspiring, look for a science project near you. Some good places to start are: SciStarter,¹⁶ where you can find thousands of initiatives worldwide, and where you can also publish your own project; sensor.community,¹⁷ if especially interested in air quality; GBIF¹⁸ worldwide, or Waarneming.nl¹⁹ in the Netherlands, if especially interested in observing nature data.

6.5 Conclusions and where we go next

Science is a human endeavour – a self-correcting, ongoing process that constructs an evolving narrative about the world. This quest for understanding is endless, as new discoveries often introduce previously unknown questions. For centuries, science was the domain of scholars, an exclusive group of ‘insiders’, while for most citizens, it remained distant and inaccessible. This does not reflect fundamental human traits like curiosity, sociality and intelligence. Citizen science breaks with this tradition by connecting anyone interested to contemporary research and questions, fulfilling a basic human need for exploration and participation. Technological advances, such as smartphones and the Internet, which bridge time and space, have made this inclusive engagement possible.

However, this is only part of the story. The world today faces a complex and multifaceted environmental crisis, including climate change, pollution, and the loss of biodiversity due to human activities. Science must serve society in developing and implementing sustainable solutions to these pressing challenges. Public engagement in research through citizen science is one of the most effective ways to create meaningful impact. Citizen science can raise awareness of our fragile environment and encourage people to contribute to its healing through innovative ideas and behavioural changes. The collective intelligence of the public is crucial to successfully propel such a much-needed large-scale sustainability transformation.

Also, citizen science itself has become an interesting new research topic with questions like: What are the main motivations to participate? What is the link between citizen science activities and long-lasting behavioural change?

For all the above reasons, although the rapid growth of new citizen science initiatives may be slowing, the overall interest and uptake should continue to grow – for the sake of humanity’s future.

16 <https://scistarter.org/finder>.

17 <https://sensor.community/en/>.

18 <https://www.gbif.org/citizen-science>.

19 <https://waarneming.nl/>.

References

- JRC, 2018. An inventory of citizen science activities for environmental policies. European Commission, Joint Research Centre (JRC) [Dataset] PID: <http://data.europa.eu/8gh/jrc-citsci-10004>.
- IVN – Florian, A. 'Wij biologen moeten onze stem laten horen'. (n.d.). <https://www.ivn.nl/leren-over-de-natuur/bioloog-auke-florian-wij-biologen-moeten-onze-stem-laten-horen/>.
- Brown, E.D. and Williams, B.K., 2019. The potential for citizen science to produce reliable and useful information in ecology. *Conservation biology*, 33:561–569. <https://doi.org/10.1111/cobi.13223>.
- Conrad, C.C. and Hilchey, K.G., 2011. A review of citizen science and community-based environmental monitoring: issues and opportunities. *Environmental monitoring and assessment*, 176:273–291. <https://doi.org/10.1007/s10661-010-1582-5>.
- Cretois, B., Simmonds, E.G., Linnell, J.D.C., Van Moorter, B., Rolandsen, C.M., Solberg, E.J., Strand, O., Gundersen, V., Roer, O. and Rød, J.K., 2021. Identifying and correcting spatial bias in opportunistic citizen science data for wild ungulates in Norway. *Ecology and evolution*, 11:15191–15204. <https://doi.org/10.1002/ece3.8200>.
- Davis, L., Zhu, L. and Finkler, W., 2023. Citizen science: is it good science? *Sustainability*, 15:4577. <https://doi.org/10.3390/su15054577>.
- Dickinson, J.L., Shirk, J., Bonter, D., Bonney, R., Crain, R.L., Martin, J., Phillips, T. and Purcell, K., 2012. The current state of citizen science as a tool for ecological research and public engagement. *Frontiers in ecology and the environment*, 10:291–297. <https://doi.org/10.1890/110236>.
- Fraisl, D., Campbell, J., See, L., Wehn, U., Wardlaw, J., Gold, M., Moorthy, I., Arias, R., Piera, J., Oliver, J., Masó, J., Penker, M. and Fritz, S., 2020. Mapping citizen science contributions to the UN sustainable development goals. *Sustainability science*, 15. <https://doi.org/10.1007/s11625-020-00833-7>.
- Gunnell, J., Golumbic, Y., Hayes, T. and Cooper, M., 2021. Co-created citizen science: challenging cultures and practice in scientific research. *Journal of science communication*, 20:Y01. <https://doi.org/10.22323/2.20050401>.
- Hallmann, C.A., Sorg, M., Jongejans, E., Siepel, H., Hofland, N., Schwan, H., Stenmans, W., Müller, A., Sumser, H., Hörrén, T., Goulson, D. and De Kroon, H., 2017. More than 75 percent decline over 27 years in total flying insect biomass in protected areas. *PLoS ONE* 12:e0185809. <https://doi.org/10.1371/journal.pone.0185809>.
- Hendricx, W., Wesseling, J., Huitema, M., and Janssen, G., 2021. Monitoring van de luchtkwaliteit ... op de fiets! <https://www.samenmeten.nl/documenten/monitoring-van-luchtkwaliteit-op-fiets>.
- Irwin, A., 1996. Citizen science: a study of people, expertise and sustainable development. *Choice reviews online*, 34:34–0267. <https://doi.org/10.5860/choice.34-0267>.

- Jacobusse, G.W., and Jongejans, E., 2024. Non-detection during excursions by citizen scientists modeled as a function of weather, season, list length, and individual preferences (p. 2024.09.30.615418). *bioRxiv*. <https://doi.org/10.1101/2024.09.30.615418>.
- Jordan, R.C., Gray, S.A., Howe, D.V., Brooks, W.R. and Ehrenfeld, J.G., 2011. Knowledge gain and behavioral change in citizen-science programs. *Conservation biology: the journal of the society for conservation biology*, 25:1148–1154. <https://doi.org/10.1111/j.1523-1739.2011.01745.x>.
- Khatib, F., DiMaio, F., Cooper, S., Kazmierczyk, M., Gilski, M., Krzywda, S., Zabranska, H., Pichova, I., Thompson, J., Popović, Z., Jaskolski, M. and Baker, D., 2011. Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature structural and molecular biology*, 18:1175–1177. <https://doi.org/10.1038/nsmb.2119>.
- Lotfian, M., Ingensand, J., and Brovelli, M.A., 2021. The partnership of citizen science and machine learning: benefits, risks, and future challenges for engagement, data collection, and data quality. *Sustainability*, 13:8087. <https://doi.org/10.3390/su13148087>.
- NOS, 2024, May 22. Nederland kampioen biodiversiteitsverlies, 'de natuur staat er dramatisch voor'. <https://nos.nl/artikel/2521490-nederland-kampioen-biodiversiteitsverlies-de-natuur-staat-er-dramatisch-voor>.
- Pateman, R., Dyke, A., and West, S., 2021. The diversity of participants in environmental citizen science. *Citizen science: theory and practice*, 6:9. <https://doi.org/10.5334/cstp.369>.
- Robinson, O.J., Ruiz-Gutierrez, Viviana, Reynolds, M.D., Golet, G.H., Strimas-Mackey, M., and Fink, D., 2020. Integrating citizen science data with expert surveys increases accuracy and spatial extent of species distribution models. *Diversity and distributions*, 26:976–986. <https://doi.org/10.1111/ddi.13068>.
- Stapleton, C., Smith, E., and Hughes, C.E., 2005. The art of nurturing citizen scientists through mixed reality. *Fourth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'05)*, 2–11. <https://doi.org/10.1109/ISMAR.2005.58>.

The role of technology in human–nature connectedness

Case studies on citizen participation

Derk Jan Stobbelaar and Jetske G. de Boer

Abstract

In this chapter we describe how technology has alienated us from nature. The current social-technical-ecological system (STES) tends to drive us towards techniques where connection with nature is more difficult. However, techniques can also play a role in restoring human–nature connectedness in (at least) two ways: (1) directly by restoring interest in nature through e.g. new materials, experiences, emotions and/or knowledge (case study 1 and 2); and (2) indirectly by organizing people around an environmental topic (case study 3). Citizen participation may help to counter this alienation and to accelerate the transition to sustainability. Finally, we reflect on the potential for designing technology that contributes to restoring human connectedness to nature or to the reinvention of our relationship with nature.

7.1 Introduction

Over half of the world population lives in cities and this proportion continues to increase. For many urban citizens, the physical separation from nature is large. It takes effort, time and money to get there. Even when we do not strictly define nature as areas that are not influenced by humans, distances to green spaces may be large for many citizens (e.g. Soliman *et al.*, 2017). Furthermore, technology and digitalization are booming and modern society relies on them. It is, perhaps, more difficult for people to imagine life without technology than to imagine life without nature. Yet, nature is essential for life on earth, for our well-being and for our future. We tend to forget that modern societies rely on nature, just as ancient societies did. We are part of it. We forget because most of us no longer feel connected to nature (Beery *et al.*, 2022). Interestingly, we rarely think of the many ways in which technology influences this connection and our behaviour with respect to nature. In this chapter, we reflect on the role of technology in our connection with nature. Let us consider some hypothetical situations to illustrate why we should think about this topic – situations that are likely to commonly arise (See EXAMPLE 7.1).

EXAMPLE 7.1

Rethinking human–nature connectedness

We zoom in on a train travelling between two cities in the Netherlands (Figure 7.1), one of the most densely populated countries in the world, where the landscape is dominated by agriculture. While much of this landscape is ‘green’, most passengers do not notice it. Their eyes are on their smartphones or their laptops. They are playing games, watching videos, and active on social media. Some passengers are studying or working. Only a few people look outside, only a few are engaged in conversations with others. In this situation, digital technology that captures our attention has reduced human connection with the physical and social environment. We zoom out and consider another, contrasting, situation: a group of people sitting quietly on the forest floor before going on a walk in the woods. They are forest-bathing. Being in the forest has positive effects on mental and physical health, exposing participants to a beneficial microbiome, reducing stress levels, calming their minds. The positive effects of forest-bathing – restoration of concentration and mental well-being – are highlighted in attention restoration theory (e.g. Hansen *et al.*, 2017). In addition, this ‘nature-intervention’, in which technology does not play a role, is intended to restore the ancient relationship between humans and nature.

Here, we consider two different, contrasting situations.

It is Saturday morning and Mr Smith is reading the electronic newspaper on his tablet, while the automatic lawnmower has been mowing the lawn since 6 o’clock that morning. The lawn looks picture-perfect, not a dandelion in sight. Mr Smith has no idea that this also means that his garden is a dead-zone for insects.

Two blocks down, a group of neighbours are working together to sow flowers and plant fruit trees in the garden in the centre of the square in which they live. This ‘garden’ used to be maintained by the municipality: the grass was mown every week, and the hedges were pruned once a year. Today, this central garden is more biodiverse, and has become a meeting place where neighbours drink coffee together. They are proud to have organized themselves and have taken over maintenance of this public green space. As a result, these people feel more connected to nature and to each other.

human–
nature
connect-
edness

The hypothetical situations in Example 7.1 suggest that technology (unintentionally) decreases our connection with nature. Yet, restoring this connection and/or reinventing our relationship with nature is considered crucial for sustainable development (Richardson *et al.*, 2023). This is not only because nature and natural processes are the basis of our being, which is threatened by the biodiversity crisis and



FIGURE 7.1 Technology tends to reduce our connection with nature

SOURCE: PHOTO 1: [HTTPS://FREERANGESTOCK.COM/PHOTOS/120873/PEOPLE-ON-THE-TRAIN-LOOKING-AT-THEIR-MOBILE-PHONES.HTML](https://freerangestock.com/photos/120873/people-on-the-train-looking-at-their-mobile-phones.html), PHOTO 120873; PHOTO 2: [HTTPS://WWW.FLICKR.COM/PHOTOS/KEEPITSURREAL/46981596084](https://www.flickr.com/photos/keepitsurreal/46981596084)

consequences of global climate change, but also because nature is essential for our (mental) health. In the following sections, we first introduce and explain important key concepts and their relationships in the context of this chapter. Next, we provide detailed descriptions of three case studies in which technology potentially mediates human–nature connectedness at different levels of citizen participation. Finally, we attempt to look ahead, envisioning a future in which human–nature connectedness is reinvented, supported by technology.

7.2 Key concepts and their relationships

In this section we introduce the key concepts of our approach and the coherence of these concepts: socio-technical-ecological systems, technology, citizen participation and human–nature connectedness. We also introduce relationships between the key concepts, whereby the intermediating role of technology is emphasized (7.2.2).

7.2.1 Key concepts

7.2.1.1 Society

Modern society can be described as a system combining social, technological and ecological realms that are interconnected (e.g. Ahlborg *et al.*, 2019). Changes in one realm of a Socio-Technical-Ecological system (STES) affect the other realms (Figure 7.2). As such, these systems have leverage points at which the complete system can change. While changes at the societal level are generally slow, they are triggered by changes at the ‘street level’, the level where we live. Feedback loops between the ‘street level’ and the ‘societal level’ can accelerate system change. We use these two levels to make our point that organization (participation) of people and human–nature connectedness take place in a larger system, influenced by

Socio-Technical-
Ecological system
(STES)

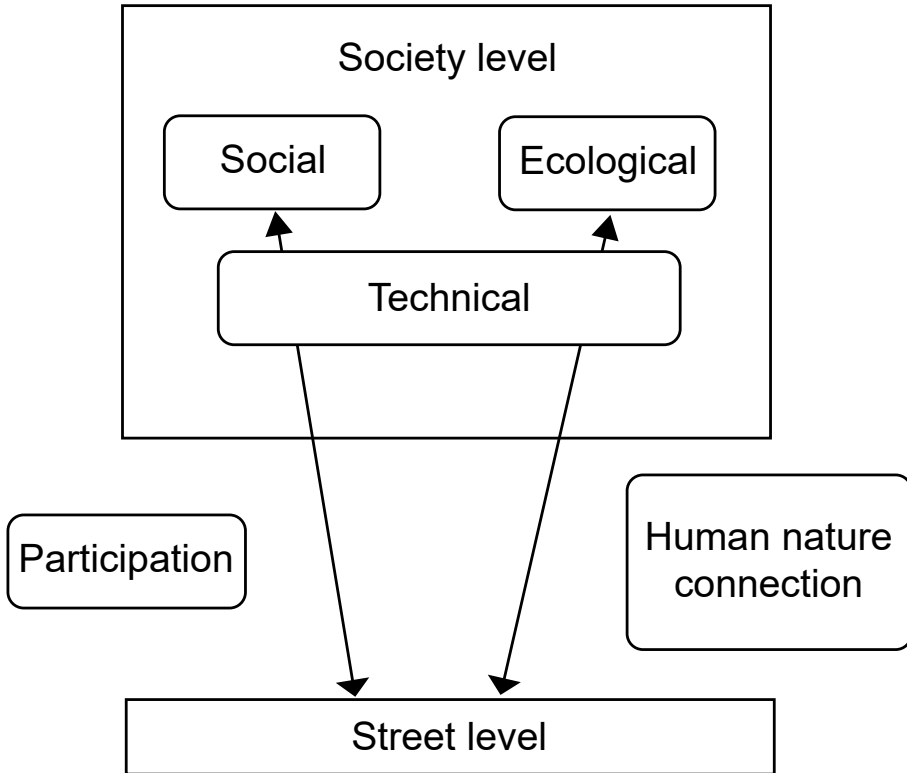


FIGURE 7.2 Theoretical model of the mutual influence between system level and daily life level

but also influencing daily life (Balászi *et al.*, 2019; Beery, 2023). Special attention should be given to the role of technology, because this often mediates the connection between the street level and the system level. For example, as also illustrated above, the introduction of smartphones (technology) changed our ways of communicating (human behaviour) and hence our social interactions (social realm). At the same time, the introduction of this technology contributed to reduced interaction with our physical environment and hence contributed to reduced nature connectedness (e.g. Richardson *et al.*, 2018; Richardson *et al.*, 2022, Balászi *et al.*, 2019). This in turn might change the ecological realm because we become less interested in taking care of the environment (see figure 7.2).

7.2.1.2 Technology

Technology is ‘the use of science in industry, engineering, etc., to invent useful things or to solve problems’, or ‘a machine, piece of equipment, method, etc., that is created by technology’ (Wernaart, 2022). In our case special attention should be paid to apps, podcasts, binoculars, (digital) maps and even serious games. Given the digital revolution (e.g. Wernaart 2022), we focus primarily on digital techniques in the case studies below. Technology is not usually developed with the specific

purpose of connecting humans and nature (or of encouraging citizen participation in nature conservation). However, technology designed for other purposes can be used in the context of enhancing citizen participation (e.g. media channels for communication, etc.).

Technology has moral aspects (Wernaart, 2022). New technologies change society, introducing new moral questions. Indeed, industrialization and urbanization are also effects of technological development. Wernaart (2022) describes the example of the industrial revolution bringing inequality into society. In our case, we can see that technology tends to widen the gap between human and nature (Soga and Gaston, 2016; Richardson *et al.* 2022). In our case, we interpret morality as the positive contribution to sustainability in general (or absence of) and to the effects on human–nature connectedness in particular.

moral aspects

7.2.1.3 Citizen participation

Citizen participation refers to citizen involvement in public decision-making. While interpretations differ, ‘group’, ‘organizations’ and ‘power’ are important keywords. ‘Citizens’ may be either individuals or organized communities and ‘participation’ may involve either observation, engagement or power. ‘The means of citizen participation include groups and formal organizations, meetings, inquiries, action, and technical assistance’ (Baum, 2001). In this definition, participation is a form of organization of individuals in a larger group. Arnstein (1969) puts more emphasis on the power aspect of participation, simply seeing citizen participation as a categorical term for citizen power. ‘It is the redistribution of power that enables the have-not citizens, presently excluded from the political and economic processes, to be deliberately included in the future.’ Citizen participation thus provides the means by which these individuals can induce significant social reform, enabling them to share in the benefits of the affluent society. Arnstein (1969) distinguishes three major levels of participation: information, consultation and engagement. These levels are also referred to as the ‘participation-ladder’, where self-organization is seen as the highest possible (sub)level.

The OECD (2022) says: ‘Citizen and stakeholder participation includes all the ways in which stakeholders (including citizens) can be involved in the policy cycle and in service design and delivery.’ This also requires efforts by public institutions to hear the views, perspectives, and inputs from citizens and stakeholders. Indeed, Ferreira *et al.* (2020) highlights the importance of citizen’s perceptions, preferences and perspectives for public participation and conclude that it is essential that managers and planners become more aware of this.

Interestingly, Ferreira *et al.* (2020) notes that, in most cases, participation processes in nature-based solutions were organized bottom-up and/or led by citizens (see section 7.3.3, the IJsselstein case study). Balancing this autonomy of social initiatives with expert knowledge in a context-specific way can lead to novel governance concepts (‘mosaic governance’), as proposed by Buijs *et al.* (2016).

7.2.1.4 Human–Nature Connectedness

Restoring Human–Nature Connectedness (HNC) is perhaps the biggest challenge of our time (Richardson, 2023). HNC has been emphasized as a key concept for leveraging sustainable changes in social-ecological systems (Balászi *et al.*, 2019, Ives *et al.*, 2018; Richardson *et al.*, 2022). When people really get attached to nature, then care for nature will be obvious. The question is, what can technologies do to support this attachment, thus extending a social ecological system (SES) to a social-technical-ecological system (STES)?

Several types of HNC are identified in the literature. One of the most used is HNC as described by Ives *et al.* (2018), who introduce material, experiential, emotional, cognitive and philosophical connections with nature (further explained in Table 7.2 in section 7.2.2). Interactions between these types of nature connectedness and feedback loops between types can offer stronger leverage potential in the sustainability transition. Structural change may often be necessary to enable successful implementation of interventions that reconnect people with nature and to realize the benefits. Ives *et al.* (2018) illustrate this with examples where a revised educational policy is needed to allow students greater interaction with nature as part of the curricula; reformed planning law is needed to increase biodiversity in cities and modified transport networks are needed to enable people to access natural areas easily. Reconnecting people with nature may therefore both affect and depend upon deep structural change.

The above-mentioned STES approach and the interaction between the societal and street levels make it clear that a certain level of participation (organization) should take place before human nature interactions can completely unfold. These participation processes can stimulate individual citizens to connect with nature (see also detailed case study descriptions).

7.2.2 *Relationships between key concepts*

7.2.2.1 The relationship between human behaviour and ecology (mediated by technology)

Richardson (2023 after Timothy Tayler, 2010) described men as the technological ape. We can no longer survive without technology: ‘We are born helpless and become bound to the technological culture we learn to survive in.’ But technology is progressively separating us from nature. In the early days of mankind, technology, such as bows and axes, was directly attached to survival and nature. Now, our technology is the pen and the computer. This separation is visible in all areas of life (Table 7.1).

Many daily activities now take place indoors as a result of technological advancements. That means that, at the very least, the physical connection with nature is diminishing. However, the effects are going deeper. Changes are taking place on all levels of HNC (Table 7.2). Although it is difficult to pinpoint exactly what has led to this decline in spontaneous outdoor activities, several possible triggers have

TABLE 7.1 Examples of changes through technology in various aspects of daily life

Areas of life	Used to be	Today
Work	On the land	Office
Living	Outside	City/apartment
Free time	Integrated in daily life	Movies, gaming, amusement park,....
Food	Directly from the land	Shop
School	Life	Building
Transport	Walking	Car

TABLE 7.2 Types of Human–Nature Connectedness (after Ives *et al.*, 2018) and examples of the role of technology

Types of HNC	Description	Role of technology
Material	Resource extraction and use (regional self-sufficiency)	Division of labour, increasing use of technology in farming
Experiential	Direct interaction with the natural environment, recreational activities in a green environment	Increasing urbanization, decreasing natural areas
Emotional	Feeling of attachments and affective responses to nature	Decreasing time in nature due to competition with technology
Cognitive	Knowledge or awareness of the environment, beliefs and attitudes	Increasing theoretical knowledge (indirect knowledge).
Philosophical	Perspectives or worldview of humanity's relationship to the natural world	Technology creates the feeling of mastery (of nature). It also creates the idea that technology will solve all problems (ecomodernism).

been identified. These include the rapid growth in the number and proportion of people living in urban areas (Turner *et al.*, 2004; Zhang *et al.*, 2014); technological advancements and the emergence of sedentary pastimes such as watching television, playing computer games, and browsing the internet (Pergams and Zaradic, 2006; Ballouard *et al.*, 2011); the overscheduling and micromanagement of children's lives (Clements, 2004; Hofferth, 2009). For many people today, outdoor experiences

of nature are being replaced by virtual alternatives (Pergams and Zaradic 2006; Hofferth, 2009; Ballouard *et al.*, 2011).

However, after this rather negative exposé of the role of technology in our separation from nature and our loss of nature connection, we can ask ourselves about positive possibilities. Can technology also help to create new connections to nature, possibly even support the reinvention of our relationship with nature?

Several studies have looked into this potential. Soliman *et al.* (2017) found that watching nature videos could increase nature relatedness (although this did not result in pro-environmental behaviour). Zabini *et al.* (2020) examined the relaxation effect of audio-visual stimulation using a forest video, in the absence of human – nature interactions resulting from the COVID-19 quarantine. The results showed a short-term decrease in anxiety levels after exposure to a forest video, but they did not indicate long-term changes. The virtual exposure to forest environments was effective in reducing human stress levels in people lacking direct nature contact. The authors concluded that virtual nature is unable to fully reproduce the effects of real nature, such as boosting of immune functions due to exposure to phytoncides. However, they also concluded that immersive virtual nature technologies could contribute positively to the physiological well-being of people who do not have direct access to nature, albeit with only short-term effects. As Matey (2017) described it: ‘Technological nature has its benefits; engaging with it makes us feel good by triggering our innate “biophilia,” a term for humanity’s inborn, primordial affiliation with the environment. For example, researchers have found that nature videos played in prisons drastically reduce violence amongst inmates, suggesting that nature’s relaxing influence translates through screens.’ Altrudi (2021) introduces a warning note on this point. Not everybody can be reached with technology. Most of those people reached by this means are already interested in nature (Altrudi, 2021).

So, the conclusion of this section is that technology has driven human beings away from nature, with a diminishing HNC as a consequence. Some technologies are already used to restore the connection between people and nature. However, the overall conclusion is that more change is needed in order to support a transition towards nature inclusive behaviour.

7.2.2.2 The relationship between participation and the social realm (mediated by technology)

Technology can be used to condition people. The critiques shared on social media are a good example of this role of technology. They steer people in certain directions in order to keep them occupied. People get ‘organized’ around subjects that keep them away from their direct surroundings.

Similarly, technology can also be used to organize people around nature, landscape and their home area, thus producing a force to counter technology that generally disconnects people from nature. There is considerable literature on citizen

participation in urban planning. The literature review of Ferreira *et al.* (2020) shows that technological methods and tools may be used to involve citizens and stakeholders as participants in Nature Based Solutions (a wider term than green infrastructure and other narrower terms). For example, GIS-based tools, digital twins and VR may enhance citizen participation in the strategic planning of green space in cities (see also the case study in IJsselstein). Questionnaires and surveys remain one of the most common tools in the participatory process. Nowadays, such methods often rely on technology, although they can be supplemented by approaches that are not necessarily technology-based, such as interviews, meetings or workshops.

This use of technology can also be seen in the work of organisations such as Ecó (previously known as AVAAZ), Greenpeace, Milieudefensie (Dutch Friends of the Earth) etc. – organisations that initiate signature campaigns around environmental problems. The Internet, with all its features, is very helpful in raising awareness and in starting a movement around (green) public issues. In addition, many organisations, such as E-NGOs, try to combat disconnection by offering educational programs and outdoor experiences (Beery *et al.*, 2023). Importantly, the effect of these programmes depends on the social and political context (STES) in which they operate.

7.2.2.3 The relationship between participation and human–nature connectedness

Often participation (in the sense of organization and social power) is necessary to enhance human–nature connectedness. Groups of people have to create and provide the means to make this relationship possible. It is important to understand how participation in such groups translates to the daily life of the participants, in particular, how their participation affects their connection with nature and the level of their nature-inclusive behaviour. Here, we will consider this relationship as ‘citizen science’ – research that is carried out in whole or in part by citizens. Examples include measurement of air or water quality or observation of animals or plants. This data can be collected and analysed by scientists. The largely voluntary contribution of citizens increases data availability and often increases temporal and geographic coverage. This is a relevant advantage in biodiversity research because the monitoring and identification of organisms takes a lot of time. Through citizen science, monitoring can be carried out at several different times of the year and at many locations. This is crucial in order to analyse the development of biodiversity and to link it to the functioning of ecosystems and (changes in) environmental factors. In the Netherlands, an estimated 25,000 people are active as volunteers in biodiversity monitoring. The vast majority of observations are recorded for birds. Lesser numbers of observations are received for vascular plants and butterflies (Wallis de Vries and Sparrius, 2022), a pattern that exists worldwide (Chandler *et al.*, 2017). Many studies based on citizen science have already been published and it can be concluded that this approach is very valuable for biodiversity research (e.g.

Chandler *et al.*, 2017; Koffler *et al.*, 2021; Peter *et al.*, 2019, 2021). This is especially the case when projects are well-structured and observations are made according to strict protocols (Galvan *et al.*, 2022; Johnston *et al.*, 2022).

In addition to generating data on biodiversity, citizen science can also be meaningful in other ways. It increases participants' knowledge of biodiversity and of species and their value (Peter *et al.*, 2019). In addition, participation in a citizen science project can contribute to greater involvement with nature, which may lead to a different attitude towards and connection to nature and, potentially, to changes in behaviour. Recent research shows that participants in citizen science projects in the field of biodiversity report that they can indeed communicate better about biodiversity and research, and that they also observe a change in their own behaviour in activities such as gardening (Peter *et al.*, 2021). Interestingly, Mattijssen and Terluin (2018) note that the relationship between citizen science and nature connectedness is reciprocal: increased HNC can result from participation in citizen science but is often also a driver to participate in such initiatives.

7.3 Case studies on citizen participation

Here we introduce three case studies with increasing emphasis on the role of citizen participation. All of the case studies make use of technology that helps people to become more strongly connected with nature (increased HNC). The first study is about walking apps. The next concerns the citizen science activities of the Dutch Butterfly Conservation foundation (De Vlinderstichting). The last study is about an environmental group in IJsselstein that built a GIS map that gives counterforce against the municipality and unites the inhabitants of IJsselstein in their activities around the special collection of fruit trees in the public green space.

7.3.1 *Walking apps*

There are different types of walking apps. Some direct attention to the environment, identifying items of interest and their locations. However, most help to measure achievement, such as the number of steps taken, or aim to support a certain state of mind (apps intended to improve physical or mental health).

In the Dutch National Prevention Agreement, the central government has made agreements with 70 social organizations aiming to achieve substantial health gains for the entire Dutch population by 2040 (Rijksoverheid, n.d.). One objective of these agreements is to reduce the percentage of overweight people from the current 50% to 38% of the population. The Dutch Brain Foundation (Hersenstichting) has therefore created the Ommetje app,¹ in collaboration with prominent brain scientist Erik Scherder. This app is designed with the intention of encouraging people

1 <https://www.hersenstichting.nl/ommetje/>.

to walk every day. The emphasis is on health benefits, not on experiencing nature. According to the Brain Foundation, the app has been downloaded 1.6 million times, 30% of those who have downloaded the app use it regularly (walked with the app for more than 66 days) and users have already walked a total of 272 million kilometres with the app. Each walk ends with a brain fact expressed by Erik Scherder. The research of De Bruin *et al.* (2021) stated that walking apps encourage people to go outside.

The success of the Ommetje app can be explained by the influencing techniques incorporated in it (Schooltink and Hippert, 2021; Kramer *et al.*, 2020). These include rewarding good behaviour (earning points), continuing pays off (more points for continued performance), social proof (everyone does it, including me) and rankings (you can measure your performance against others). These are all motivational techniques that can be built into apps.

This app is therefore a good example of how the government (via the Brain Foundation) can nudge people to participate in desired behaviour through smart information provision. This is participation at the level of information (OECD, 2017). People can share their experiences via Facebook or other social media. It is unclear to what extent this sharing has an effect on government campaigns, but this is probably more of an exchange between the participants. The level of consultation (Arnstein, 1969) is therefore not reached. As the app is mainly aimed at encouraging walking as an activity, there is less or no attention directed at the environment. In short, the app has a very limited design, but perhaps that is also part of its success.

The Klompenpad app (clog path) is an exception to the many apps such as the Ommetje app that focus solely on personal health and development.² This app focuses on walking several routes (clog paths) and provides information about the area, not only about the ecological environment but also about the cultural-historical aspects of the landscape. This app can be compared to a digital travel guide, with all kinds of useful extras such as the option to keep track of where you have walked and options to share the routes with others via WhatsApp and email. The literature seems to suggest that more knowledge of the socio-ecological environment increases appreciation (emotional bond). Runaar (2020) supports this, identifying that people came to appreciate the landscape more after seeing videos of farms.

In terms of HNC, the Ommetje app can provide some experiential connectedness, when the walker is able to see his/her surroundings (Table 7.3). The Klompenpad app adds the possibility of cognitive connectedness and some emotional binding through the stories about the landscape that are included in the app. Material HNC could be promoted by connecting with sites such as Wildplukwijzer³.

² See <https://klompenpaden.nl/>.

³ <http://www.wildplukwijzer.nl/>.



FIGURE 7.3 Case study 1, walking apps may be designed to stimulate people to connect more strongly to nature, e.g. by providing information and stories

7.3.2 *Citizen science organized by the Dutch Butterfly Conservation Foundation (De Vlinderstichting)*

The Dutch Butterfly Conservation foundation (De Vlinderstichting)⁴ is a public organization aiming to protect butterflies, moths and dragonflies in the Netherlands. They do this by collecting data and developing information and knowledge about where these insects can be found and by protecting and improving their habitat. Educating the public about butterflies is an important pillar of the organization because awareness is crucial for conservation. The foundation also coordinates the national monitoring of butterflies, moths and dragonflies, based almost entirely on citizen science. The Dutch Butterfly Conservation foundation started butterfly monitoring with Statistics Netherlands (CBS) in 1990, dragonfly monitoring in 1998 and moth monitoring in 2012. Statistics Netherlands checks and analyses the citizen science data, identifies national and provincial trends, provides information for the WWF Living Planet Report and reports to meet the requirements of the EU Bird and Habitat Directive (Article 17). Furthermore, butterfly data are combined with (citizen science) data on other species groups to calculate indicators

⁴ <https://www.vlinderstichting.nl/>.

then presented by Environmental Data Compendium Netherlands.⁵ Currently, citizen scientists monitor all Dutch butterfly species along 800 transects across the Netherlands and moths at 800 points using LED-buckets (see below).

Citizen scientists that work with the Dutch Butterfly Conservation foundation use a variety of methods and technologies that have or have not been specifically developed for this purpose. Day-active butterflies and dragonflies are monitored along transects with a specific protocol which specifies details such as the length of the transect, required weather conditions and the frequency of monitoring. Transect-monitoring requires a substantial time investment and a relatively high level of species-knowledge and expertise of the volunteer. The citizen scientist that walks a particular transect does not use technology specifically developed for this purpose (although binoculars may be used to identify butterflies from a distance). Data are collected on an online platform and participants have the option to submit their records in one of three ways: on paper forms, through the website or with an app on their mobile phones.

Monitoring of macro-moths is also done by registered volunteers, using LED-buckets (large buckets fitted with a specific light) to attract these insects during the night. The LED-buckets are set once every two weeks. Moths that enter the bucket are counted and identified the next morning. This approach requires a high level of expertise on the part of the citizen scientists because there are many more moth than butterfly species in the Netherlands and they are more difficult to identify. Moth monitoring with LED-buckets is also done by approximately 100 farmers on their own farmland. In this case, volunteers submit photographs to the Dutch Butterfly Conservation foundation, which then uses automatic image recognition for species identification. Automatic image recognition is relatively well developed for this group of insects and difficult cases are checked by experts of the Dutch Butterfly Conservation foundation. Participating farmers can also use the app Obsidentify⁶ to easily identify moths.

Flexible counting is a more user friendly and less time-consuming option than transect walks or systematic moth monitoring. It allows people to occasionally monitor butterflies along transects using a specific app. ButterflyCount⁷ is a European app that can be used for occasional counts of all butterfly species over 15-minute periods or for observations of a specific species, such as the wall brown (*Lasiommata megera* or argusvlinder in Dutch) or the dragonfly *Stylurus flavipes* (river clubtail or rivierrombout in Dutch).

These different approaches to citizen science require different levels of (individual) participation and levels of expertise from the volunteers. These levels relate to

5 <https://www.clo.nl/en>.

6 <https://waarneming.nl/apps/obsidentify/>.

7 <https://www.vlinderstichting.nl/>.

the planned primary outcomes of the citizen science, ranging from data collection to education or awareness. On the one hand, expert volunteers contribute important data on butterflies, moths and dragonflies that feeds directly into the *Netwerk Ecologische Monitoring* (NEM – Dutch Ecological Monitoring Network),⁸ which is the backbone of terrestrial biodiversity data informing policy in the Netherlands. Thus, these citizen scientists participate at the consultation level because the data that they collect may influence political decisions. On the other hand, anyone can report their incidental findings of these insects on the website of the Dutch Butterfly Conservation foundation. While these citizen science efforts are mostly individual, the Dutch Butterfly Conservation encourages or invites participants to monitor together with other people, either as an organised workgroup or with friends or colleagues. They also offer to help in connecting participants with other volunteers that are active in the same region.

In the case of the Dutch Butterfly Conservation foundation, there appears to be a trade-off between the level of expertise required for the different citizen science approaches and the technology required and also between the level of citizen participation and the technology required. Today, the citizen science data, the data that are most important in informing policy and in influencing decision-making, are those collected by the most experienced citizen scientists who use methods that require the least technology (transects of butterflies and point measurements of moths, where technology is only used in the data collection platform but not by the citizen scientist). Nevertheless, technology enables the involvement of a wider, less-experienced, group of people as citizen scientists. In this case, technology can be important in helping people to become more connected to nature and to start climbing the citizen participation ladder by becoming more educated and informed on how citizen science data may influence policy.

We see several possibilities for how participating in citizen science may influence HNC (Table 7.3). At the cognitive level, participants may vary as described above, but almost all citizen scientists will gain knowledge of species (recognition) and possibly also of measures for species conservation or nature restoration. This gain may possibly extend to philosophical HNC where participation influences participants' thinking about our role in protecting instead of exploiting nature. Experiential HNC is automatically influenced because citizen scientists spend time outside and careful observation is inherent to the activity. Finally, it is important to consider emotional HNC because participants may report pessimistic thoughts (see Peter *et al.*, 2021).

8 <https://www.netwerkecologischemonitoring.nl/>.



FIGURE 7.4 Case study 2, citizen science organized by the Dutch Butterfly Conservation foundation. Buckets with LED-light are used by citizen scientists to record moths in the Netherlands, while butterflies may be recorded along transects.

7.3.3 *IJsselstein case: fruit trees in the public green space*

There is an active citizens' initiative in IJsselstein – Climate Neutral IJsselstein (KNIJ – Klimaatneutraal IJsselstein).⁹ The initiative provides a range of activities related to climate, public greenery and the living environment. This case study focuses on fruit in public spaces. IJsselstein has a unique collection of fruit trees in public spaces. The former city landscape architect planted all the apple, pear and hazelnut varieties known to him in alphabetical order in the new residential areas of IJsselstein. The municipality of course knows this but was not aware of the uniqueness of the collection. The collection is fun, but is also complicated because every tree is unique and the falling fruit causes a nuisance (insects, dirty sidewalks, throwing by young people – common problems with fruit in public spaces).

KNIJ wants to achieve several goals with these fruit trees: (1) Preservation and strengthening of the cultural-historical and ecological value of public space in IJsselstein, and (2) Strengthening of the bond between residents and the fruit in the

⁹ <https://www.klimaatneutraalijsselstein.nl>.

TABLE 7.3 Ways in which the different types of HNC may be influenced in the three case studies

Types of HNC	Description	Walking apps	Citizen science	IJsselstein
Material	Resource extraction and use (regional self-sufficiency)	Indirectly through information on (agricultural) land-use (e.g. in Klompenpad app) or specific add-ons such as Wildplukwijzer		Use of fruits in the public space
Experiential	Direct interaction with the natural environment, recreational activities in the green environment	Through seeing the environment	Spending time outside, careful observation	Activities attached to the fruit (walking, pruning, harvesting, processing)
Emotional	Feeling of attachments and affective responses to nature	Possible when walking apps are combined with stories of the local socio-ecological environment	Emotions related to (not) finding certain species. Pessimism can occur with increased knowledge of biodiversity decline.	Sense of belonging: IJsselstein is a fruit city
Cognitive	Knowledge or awareness of the environment, beliefs and attitudes	Possible through apps or other means that provide information along walking routes	Increased species knowledge and knowledge of restoration measures	GIS map, information boards, tours, folders, research
Philosophical	Perspectives or worldview on humanity's relationship to the natural world		Human influence on biodiversity decline and our role in protecting nature or species (groups)	Nature as a partner



FIGURE 7.5 Case study 3, IJsselstein is a municipality in the Netherlands with a large collection of fruit trees in the public green space. Technology, including a website and GIS application, provides insight into the uniqueness of this collection and helps to coordinate activities associated with it.

public space. They want to achieve this by: (1) Having volunteers manage the fruit trees, (2) Making agreements about the use of the fruit, (3) Documenting the locations of the different varieties. In order to achieve the last goal, a website has been created with a GIS application. This is not a static map. It can be updated with new knowledge about the varieties. With these activities and goals, all types of HNC can be addressed or reached (Table 7.3).

Citizens' initiatives are often said to be unreliable (Hassink *et al.*, 2016), in the sense that they often change composition and that their timelines are flexible. In this case, however, the situation appears to be reversed. In just a few years, the citizens' initiative had to deal with as many as fifteen civil servants who were responsible for the fruit trees in public spaces. The citizens' initiative had to explain the value of the fruit again and again. The species website was a powerful tool (power through knowledge) indicating that the citizens' group knew what they were talking about. The technology (website) therefore ensured participation. The map also has an internal function, because it makes it easier to make agreements about management and use: it helps to organize activities.

In addition, the website was also a very useful way to inform residents of how special the fruit collection was. This, together with information boards, e-mails, lectures and conversations in the field, led to a better understanding by the residents of how the fruit can be used (material level, see Ives *et al.*, 2018) and what is needed to manage it (experiential, other level). Now that the fruit is being used, the nuisance is also less. In other words, technology helped to get from the participation level information (government sending information to citizens) to consultation (two-way relationship) and also, already, to some level of engagement.

Igalla *et al.* (2019) state that the three most important factors for the success of a citizens' initiative are: a diverse network, organizational capacity, and government support. The digital map helps with all three. The digital map is used in discussions with other parties and the municipality, it also helps internal coordination. Naturally, the reverse also applies: a citizens' initiative must already have a certain degree of organization in order to be able to create such a digital map.

7.3.4 *How is technology mediating citizen participation and human–nature connectedness?*

In general, technology reduces HNC because the STES that forms the current society, tends to favour technology that widens the gap between man and nature. When people do not feel connected to nature, they are also not much inclined to take care of nature, leading to more loss of biodiversity. This results in a negative feedback-loop because the distance to nature increases, and its quality deteriorates further as is evident from the current biodiversity crisis. Restoring or reinventing HNC is considered crucial in the transition to a sustainable STES (see, for example, Riechers *et al.*, 2021).

Interestingly, our case study descriptions show that digital technology can play a role in restoring HNC. The walking apps primarily stimulate experiential and cognitive connectedness. These aspects may be further enhanced by digital technology such as gamification. For example, Laato *et al.* (2022) show that location-based games can benefit people who go to the forest but also benefit forest owners and even the forestry industry. Treasure hunts that make use of GIS locations, also called geocaches,¹⁰ present another example of technology that can increase experiential HNC in a wide range of people. Citizen science mostly increases cognitive connectedness but may also contribute to experiential connectedness. Here, technology may play a role especially in addressing a group of people that have less expertise and who, through participation, will look at nature more intensively and precisely. By using competitive elements and challenges, nature-oriented apps such as ObsIdentify (by, for example, awarding badges when data are entered for specific groups of animals or plants) can draw an even wider audience. In the case study of fruit trees in IJsselstein (KNIJ) four types of HNC are addressed (Table 7.3).

¹⁰ <https://www.geocaching.nl/>.

Connecting with fruit trees and nature in their direct neighbourhood can also contribute to the sense-of-place of citizens, which is an important leverage point in reconnecting with nature (Riechers *et al.*, 2021).

Perhaps the case studies we discussed also influence the philosophical HNC, but we could not assess that. However, we note that feedback loops exist between the types of HNC. More cognitive connectedness often means more emotional connectedness and vice versa. A high score on material, experiential, emotional and cognitive connectedness often leads to a feeling of partnership with nature and hence philosophical connectedness. This means that it may not matter where you start when you want to increase HNC in a population. However, the effects will be largest when several interacting aspects of HNC are addressed together. For example, by providing experience of the environment, a certain emotional bond can develop. When material, experiential, emotional and cognitive HNC develop, there is a chance that the philosophical position will be taken that nature is a partner. Riechers *et al.* (2021) also recognize the importance of interconnection between various leverage points in restoring HNC.

In the three case studies that we discussed, we see that participation in the sense of organization and counterpower is an important element of the development of technology that increases HNC. The Dutch Butterfly Conservation foundation and KNIJ organized the expertise of many people to be able to develop this (moral) technology and to make it meaningful in terms of actual citizen participation. Here, digital technology can also help citizens climb the participation ladder. Gaming elements, such as challenges and sharing information can help people to take the first step by making them feel part of a group. Hence, we suggest that technology can be used to change the existing system (STES) in a more nature-inclusive and sustainable way.

7.4 Outlook to the symbiocene

The symbiocene is a future vision of how humans may interact with the planet in a non-destructive way. This term was introduced by Albrecht in 2011 and is based on the word symbiosis, which means organisms of different species living together. If humans are to live in harmony with all other beings, we need to be connected with nature. To restore HNC with the help of technology, we suggest two things need to happen. The first is that technology that supports HNC should be promoted, as the citizen science example of the Dutch Butterfly Conservation foundation shows. Digital technology can attract new groups of people because the technology helps to detect the insects and thus helps people look at their ecological environment in a new way (see also Table 7.4). The second is that we should strive for the moral design of technology. This means that HNC choices are integrated as a design principle as we described in the example of the walking apps. The success factors of the walking app, for instance the game element, could be expanded to explore the

TABLE 7.4 Future areas of life enhancing HNC

Areas of life	Used to be	Today	Future
Work	On the land	Office	
Living	Outside	City / apartment	Green cities
Free time	Integrated into daily life	Movies, amusement park,....	Nature parks
Food	Directly from the land	(Super)stores	Country shops
School	Life	Building	On site
Transport	Walking	Car	Bicycle

TABLE 7.5 The potential role of technology in the symbiocene

Types of HNC	Description	Role of technology in the symbiocene
Material	Resource extraction and use (regional self-sufficiency)	Technology helps to explain the source and origin of materials.
Experiential	Direct interaction with the natural environment, recreational activities in the green environment	Technical support for nature inclusive society, green cities (green walls, etc.).
Emotional	Feeling of attachments and affective responses to nature	Technology can help free up more spare time, which can be spent in nature. Behavioural knowledge (how to tempt people to do something) directed towards HNC built in technology (see walking apps)
Cognitive	Knowledge or awareness of the environment, beliefs and attitudes	Technical sources (newspapers, apps, tv) educate people about nature.
Philosophical	Perspectives on or worldview of humanity's relationship to the natural world	Moral technology, implying that technology is always designed to improve HNC.

surrounding, for example by adding a geocaching element as described above. Such game elements have at least two effects: (1) they urge people to study the surroundings, and (2) they create a greater urge to continue.

To reinvent our connection with nature and move towards a sustainable future, people and/or organizations are needed to create technology such as apps with which they can direct others (participate) towards a connection with nature (via technology). To make this happen, a counterforce is needed. Civil society groups can help to build this counterforce, either by participating in government programmes or through self-organization, as the example of IJsselstein shows. These civil society groups can also use technology to reach their goals, including a stronger HNC. They can also promote technology that supports HNC. They often do that already, but using other arguments (for example, reduction of pollution). We suggest that they can add the argument of HNC to their repertoire.

Changes on a larger scale must take place if we want to establish substantial changes in HNC (Table 7.4). These changes support the realization of the visions of a sustainable world. So, the optimistic message is that, if we walk the path of greening our world, HNC will increase in the future. (see table 7.5 for a future perspective). The pessimistic message is that greening of the world needs an a priori stronger HNC. Maybe the realistic message is that both elements need each other in order to grow. In the end, this can lead to a completely new STES with moral technology and (subsequently) a larger place for ecology. Importantly, it has not yet been stated that when technology leads to more human–nature connectedness, this also leads to nature inclusive behaviour. However, the literature claims that there is a positive relationship between these two (Richardson, 2020). New technology should, therefore, try to include incentives to act, as well as the incentives to connect on a sensory and emotional level.

References

- Altrudi, S., 2021. Connecting to nature through tech? The case of the iNaturalist app. *Convergence*, 27(1), 124–141. <https://doi.org/10.1177/1354856520933064>.
- Arnstein, S.R., 1969. A ladder of citizen participation. *Journal of the American Planning Association* 35(4), 216–224.
- Balázsi, Á., Riechers, M., Hartel, T., Julia Leventon, J. and Joern Fischer, J., 2019. The impacts of social-ecological system change on human-nature connectedness: a case study from Transylvania, Romania. *Land use policy*, 89, 104232. <https://doi.org/10.1016/j.landusepol.2019.104232>.
- Baum, H.S., 2001. Citizen participation. In: *International encyclopedia of the social & behavioral sciences*.
- Beery, T., Stahl Olafsson, A., Gentin, S., Maurer, M., Stålhammar, S., Albert, Ch., Bieling, C., Buijs, A., Fagerholm, N., Garcia-Martin, M., Plieninger, T. and Raymond, Ch.M., 2023.

- Disconnection from nature: expanding our understanding of human–nature relations. *People and nature*, 2023;470–488. <https://doi.org/10.1002/pan3.10451>.
- De Bruijn, G., Dallinga, J.M. and Deutekom M., 2021. Predictors of walking app users with comparison of current users, previous users, and informed nonusers in a sample of Dutch adults: questionnaire study. *JMIR Mhealth Uhealth* 2021;9:e13391 <https://doi.org/10.2196/13391>.
- Hansen, M.M., Jones, R. and Tocchini, K., 2017. Shinrin-Yoku (forest bathing) and nature therapy: a state-of-the-art review. *Int. J. Environ. Res. Public Health* 14:851 <https://doi.org/10.3390/ijerph14080851>.
- Hassink, J., Salverda, I., Vaandrager, L., Van Dam, R. and Wentink, C., 2016. Relationships between green urban citizens' initiatives and local governments. *Cogent social sciences*, 2:1, 1250336, <https://doi.org/10.1080/23311886.2016.1250336>.
- Igalla, M., Edelenbos, J. and Van Meerkerk, I., 2019. Citizens in action, what do they accomplish? A systematic literature review of citizen initiatives, their main characteristics, outcomes and factors. *ISTR*, 2019–10, <https://doi.org/10.1007/s11266-019-00129-0>.
- Ives, C.D., Abson, D.J., Wehrden, H., Dorninger, C., Klaniecki, K. and Fischer, J., 2018. Reconnecting with nature for sustainability. *Sustain. Sci.* 46, 30.
- Kahn, P.H. Jr., 2011. *Technological nature: adaptation and the future of human life*. Cambridge MA, MIT press.
- Kramer, J.N., Künzler, F., Mishra, V., Smit, S.N., Kotz, D., Scholtz, U., Fleisch, E., Kowatsch, T., 2020. Which components of a smartphone walking app help users to reach personalized step goals? Results from an optimization trial. *Annals of behavioral medicine*, Volume 54, July 2020, Pages 518–528, <https://doi.org/10.1093/abm/kaaa002>.
- Laato, S., Pynnönen, S. and Kangas, A., 2022. Gamification of forest-going: Opportunities for players, landowners and the forest industry CEUR Workshop Proceedings (CEUR-WS.org), 6th International GamiFIN Conference 2022 (GamiFIN 2022), April 26–29, 2022.
- Matej, A., 2017. <https://qz.com/1048433/technology-is-changing-our-relationship-with-nature-as-we-know-it>.
- Mattijssen, T.J.M. and I.J. Terluin 2018. Ecologische citizen science. Een weg naar grotere maatschappelijke betrokkenheid bij natuur? WOT-technical report 117, Wageningen University and Research, Wageningen.
- OECD, 2022. OECD guidelines for citizen participation processes, OECD public governance reviews, OECD Publishing, Paris, <https://doi.org/10.1787/f765caf6-en>.
- Peter, M., Diekötter, T., Höffner T. and Kremer, K., 2021. Biodiversity citizen science: outcomes for the participating citizens. *People & nature* 3: 294–311 <https://doi.org/10.1002/pan3.10193>.
- Richardson, M., Hussain, Z., and Griffiths, M.D., 2018. Problematic smartphone use, nature, connectedness and anxiety. *Journal of behavioural addictions*, 7, 109–116.
- Richardson, M., Dobson, J., Abson, D.J., Lumber, R., Hunt, A., Young, R. and Moorhouse, B., 2020. Applying the pathways to nature connectedness at a societal scale: a leverage points perspective, *Ecosystems and people*, 16:1, 387–401, <https://doi.org/10.1080/26395916.2020.1844296>.

- Richardson, M., 2023. Reconnection: fixing our broken relationship with nature. Pelagic Publishing, London.
- Riechers, M., Pătru-Dușe, I.A. and Balázsi A., 2021. Leverage points to foster human–nature connectedness in cultural landscapes. *Ambio* 50:1670–1680. <https://doi.org/10.1007/s13280-021-01504-2>.
- Rijksoverheid (Dutch Government), n.d. <https://www.rijksoverheid.nl/onderwerpen/gezondheid-en-preventie/nationaal-preventieakkoord>.
- Runhaar, H., 2020. Hoe waarden Nederlanders boeren natuur? *Landschap* 2020/1, 11–19.
- Schooltink, T. and Hippert, B., 2021. 13 beïnvloedingstechnieken die het succes van ommettes verklaren. <https://www.frankwatching.com/archive/2021/03/25/beinfluencestechnologies-ommetje-app-succes/>.
- Soga, M., and Gaston, K.J., 2016. Extinction of experience: the loss of human–nature interactions. *Frontiers in ecology and the environment*, 14(2), 94–101.
- Soliman, M, Peetz, J. and Davydenko, M., 2017. The impact of immersive technology on nature relatedness and pro-environmental behaviour. *Journal of media psychology* 29, 8–17.
- Taylor, T., 2010. The artificial ape. How technology changed the course of human evolution. St. Martin's press, New York.
- Wernaart, B.F.W., 2022. An introduction to moral design and technology. In: Wernaart, B.F.W. Moral design and technology. Wageningen Academic Publishers. https://doi.org/10.3920/978-90-8686-922-0_1.
- Zabini, F., Albanese, L., Becheri, F.R., Gavazzi, G., Giganti, F., Giovanelli, F., Gronchi, G., Guazzini, A., Laurino, M., Li, Q., *et al.*, 2020. Comparative study of the restorative effects of forest and urban videos during COVID-19 lockdown: intrinsic and benchmark values. *Int. J. Environ. Res. Public Health*, 17:8011. <https://doi.org/10.3390/ijerph17218011>.

Food ethics and technology

Towards food innovation with crowdsourced ethics

Bart Wernaart, Sonja Floto-Stammen, Marieke van Vliet, Anika Kok and Natalia Naranjo Guevara

Abstract

Food technology is at the core of a new moral revolution. It can alter the way we think about cultural acceptance of certain foods, our perception of what is healthy, and our perception of what it is fair to consume. A number of promising and even breathtaking food technologies will (can) drastically change the way we produce and consume food. Think of insect-based foods, 3D printed foods or cultivated meat. Such innovations have the potential to bring about sustainable change. However, they can disrupt entire production chains. They force us to rethink the way we produce, distribute and consume food. This of course has consequences for citizens in everyday life and is reason enough for an intense but polarized social debate: How do we want to eat sustainably and fairly in the future? For example, is it more animal-friendly to replace a conventional burger with an insect-based burger, or does that just make us mass murderers? And if we can produce meat without emissions and without animal husbandry, is there still a reason to be a vegetarian or vegan? Despite the consequences that new technologies bring, the end consumer is rarely involved in the design process. In this chapter, we will propose how to use crowdsourced ethics as input for reshaping our food system by building a food tech moral lab. We discuss tools that companies can use to continuously align values with citizens and co-design future food technologies. This should contribute to a better public debate about our food transition and give the food industry and policymakers' tools for including public moral perceptions in the design process of new food technology.

8.1 Introduction

The agri-food sector is one of the most critical economic sectors globally, contributing significantly to employment, trade, and GDP growth (Worldbank, 2022). However, it is also responsible for a considerable share of greenhouse gas emissions, causing a substantial environmental impact (Crippa *et al.*, 2021). With resources becoming increasingly scarce, there is an urgent need to transform the food system and create more sustainable ways of producing and consuming food. Furthermore, as the world's population continues to grow, the demand for food will only increase,

environmental
impact

exacerbating the need for change (United Nations Environment Programme, 2022). Creating major shifts in the food system that align with the 17 Sustainable Development Goals (SDGs; UN, 2023) and the Paris Agreement¹ is a highly complex challenge that requires thoughtful action. With that in mind, the concept of sustainability, by definition, encompasses the triple bottom line of ecological, societal, and economic dimensions, which can lead to significant trade-offs (Purvis, 2019). In the context of the food system, the complexity of the interconnected factors makes it difficult to predict the consequences of changes made in one domain on another domain. From a citizen perspective, the ethical considerations of these trade-offs are becoming increasingly important. As people become more aware of the impacts of food production and consumption on the environment, society, and the economy, they are beginning to demand more transparency and accountability from food producers and retailers. Ultimately, finding a balance between the three dimensions of sustainability and ensuring that ethical considerations are taken into account is crucial for creating a more sustainable and equitable food system.

Sustainable decision making and ethics are closely linked because both concepts are concerned about how individual actions and choices have impacts on the environment, society and future generations. On the one hand, ethical decision making involves thinking about how decisions impact others, how citizens should treat others fairly, and how they can act in a way that is consistent with their values and beliefs. Sustainable decision making, on the other hand, involves taking actions that are economically, socially, and environmentally responsible and that do not compromise the ability of future generations to meet their own needs. The acceptance of new food technology is often closely linked to ethical considerations. Food technologies can raise ethical concerns related to human health, animal welfare, and the environment, which can impact an individual buyer's attitudes towards these products (Frewer *et al.*, 2011). For example, individuals may be hesitant to adopt new plant-based protein products if they perceive them to be highly processed or unnatural, or if they are unsure about the safety or nutritional value of these products (Siegrist and Hartmann, 2020). Similarly, new technologies such as cultivated meat or genetic engineering may raise concerns about animal welfare, environmental impact, or the potential long-term effects on human health.

In order to overcome these concerns and increase acceptance, food technology designers need to consider the ethical implications of their products and develop transparent communication strategies that address them. To create a more inclusive and sustainable food system, it is crucial to involve citizens in the process of developing and implementing new technologies for food production. Crowdsourced ethics can play a meaningful role in designing food technology that supports the

ethical decision
making

sustainable
decision
making

crowdsourced
ethics

¹ The Paris Agreement is a legally binding international treaty on climate change. It was adopted by 196 Parties at the UN Climate Change Conference (COP21) in Paris, France, on 12 December 2015. It entered into force on 4 November 2016.

transition towards sustainable food production. By involving a diverse range of stakeholders, including citizens, scientists, and policymakers, in the design process, it is possible to ensure that emerging food technologies align with societal values and expectations. Crowdsourcing can help to identify and address ethical concerns, such as the impact of new technologies on the environment, animal welfare, and human health. It can also help to foster greater trust and transparency in the food system, which is essential for citizens' confidence and support. By leveraging the collective wisdom of a broad range of stakeholders, crowdsourced ethics can enable the food sector and policymakers to design more sustainable and socially responsible food technologies that are more likely to be accepted and adopted by society. Ultimately, this can accelerate the transition towards a more sustainable and equitable food system that meets the needs of present and future generations.

8.2 Food ethics and technology

8.2.1 *Applied food ethics*

The world is currently facing several challenges in the food system, including food insecurity, environmental degradation, and public health concerns. The conventional food production and consumption system has contributed significantly to these challenges by relying heavily on animal-based products, monoculture farming, excessive use of chemicals, and even destroying or at least deteriorating natural processes that are essential for our food (e.g. pollination and a healthy soil). As the global population grows, and climate change accelerates, there is an urgent need for a transition towards a sustainable food system. The transition requires a significant shift in our dietary patterns and the way we produce food. One of the key solutions to this problem is leveraging technology to produce sustainable and nutritious food. The global food tech market was worth \$220.32 billion in 2019, according to Emergen Research (2022), and is estimated to grow to \$342.52 billion by 2027. Food tech is increasing food production to help reduce the rate of hunger and feed the world. Agriculture is becoming more automated with the use of digital and advanced technology to produce food and raw materials with smart farming. Innovative technologies offer promising alternatives to traditional animal-based protein sources. All these technologies have the potential to reduce the environmental footprint of food production, enhance food security, and improve public health outcomes. Here are some examples:

- Genetically modified organisms. GM technology modifies a plant's genes to help it become disease resistant and grow in areas not favourable for production. GM technology is used in large crops such as rice, wheat and corn.
- Meat industry technology. AI is effective in poultry production where it helps detect health issues with birds based on the sounds they make. AI robots can work at poultry farms to collect eggs or assist with slaughter.

food
transition

innovative
technologies

- Crop monitoring. Along with the use of drones, AI can detect pests and diseases in crops. Digital apps – such as AgroPestAlert, Farm Scout Pro and IPM Toolkit – can help detect pest infestation and changing soil conditions to prevent large losses.
- 3D food printer. Food printers can create food – such as pizzas, snacks and candy – at a faster pace. AI helps design the layers and structure of the food by placing one ingredient at a time. This could eliminate waste, as leftover ingredients can be reused.
- Plant-based. Plant-based foods are made from ingredients such as peas and grains and are designed to mimic the taste and texture of meat, dairy, or other animal products.
- Insects: Insects are a source of protein, functioning as an alternative to traditional livestock farming, as insects require less land, water, and feed to produce the same amount of protein. In Europe crickets, mealworms and grasshoppers have been allowed in food since February 2023 (EFSA, 2023).
- Cultivated: Cultivated meat, also known as cultured meat, is produced by growing animal cells in a laboratory. It is designed to mimic the taste and texture of traditional meat. It belongs to the group of novel foods and has not yet received EU approval (EFSA, 2023).
- Fermented proteins: Fermentation is a process in which microorganisms break down sugars to create milk-like substances that are suitable for use as dairy alternatives in a variety of food products.

The design process of food technology development doesn't usually include the voice and participation of citizens, leading to a disconnect between food producers and society. This lack of citizen involvement results in products that do not fully meet their needs, preferences, and expectations. As a result, food products may be less appealing, which may result in a lack of acceptance of new developments in the food industry. The urgent need for sustainable food, whose success depends on the willingness of individuals to change their purchasing and eating behaviour, needs citizen participation. When it comes to the ethical and moral implications of food technologies, there are several important considerations to keep in mind. One of the key concerns is the potential impact of these technologies on human health and well-being. For example, some food technologies may be associated with negative health outcomes, such as increased risk of obesity, heart disease, or other chronic conditions. It is therefore important to carefully evaluate the safety and efficacy of any new food technologies before they are introduced to the market. Another ethical concern relates to the environmental impact of food technologies. Some food production methods may be associated with negative environmental externalities such as pollution, deforestation, or habitat destruction. It is important to consider the sustainability of any new food technologies and to ensure that they are aligned with broader efforts to mitigate the impacts of climate change and protect biodiversity. Finally, there may be ethical implications related to the social

and economic impacts of food technologies. For example, some technologies may contribute to the concentration of power and wealth in the hands of a few large food companies, while others may have negative implications for farmers or other stakeholders in the food supply chain. It is important to consider these potential impacts and to ensure that food technologies are developed in ways that are equitable and inclusive. Overall, the ethical and moral implications of food technologies are complex and multifaceted. It is important to carefully consider these implications when developing and promoting new food technologies, and to engage in ongoing dialogue and debate around these issues with all stakeholders involved in the food system.

8.2.2 *Theoretical approaches in food ethics and innovation*

descriptive
ethics

Food ethics can be considered from a descriptive and a normative ethics perspective. In descriptive ethics, human moral behaviour is analysed and explained and answers the question of how individuals or organizations ‘do’ ethics. In normative ethics, we answer the question of how individuals or organizations ‘should do’ ethics. Both approaches have their own methods and frameworks. In descriptive ethics, moral decision making by individuals is usually centred around various stages. The most commonly used model is proposed by James Rest (1986). He introduces four steps we typically go through before taking moral action: 1) recognizing the moral issue; 2) making a moral judgement; 3) establishing moral intent; and 4) moral acting. In other words: an individual first becomes aware of a moral problem; then evaluates the problem by reaching a moral conclusion; after that the actor becomes determined to act according to that conclusion; and finally behaves by implementing the moral decision itself (Crane *et al.*, 2019: 139–140). In descriptive ethics we usually assume that these four steps are of a relativist nature, meaning that each individual, depending on various factors, will go through these steps differently. The factors that influence the way individuals go through these stages are subgrouped in different ways. The most commonly used distinction is between individual, situational and issue-contingent factors (Jones, 1990; Craft, 2013; Schwartz, 2016). Individual factors are inherent to the individual that goes through the stages of moral decision making. These are typically demographics (gender, age, family situation, profession), but also relate to moral capacity (e.g. personality traits and life experience). Situational factors are typically organizational factors that influence the individual that operates with or within this organization. These are usually the formal and informal systems that altogether shape the moral infrastructure of an organization. These could be codes of conduct, ethical reporting guidelines and bonus rules (formal) but also work climate, shared traditions, or commonly shared company culture (informal). To conclude, there are issue-related factors. These relate to the characteristics of the moral issue that presents itself: the magnitude of the consequences, level of social consensus, probability of the effect, temporal immediacy, proximity, and the concentration of the effects.

When individuals are confronted with a food-related moral dilemma, the four stages of moral decision making, as well as the three influencing factors, offer an analytical structure to explore the moral viewpoints of citizens. When we consider normative food ethics, normativity is typically applied on two main themes: What to eat and what not to eat? And the environmental/societal consequences of food production. In any of these cases we see a complex interplay of values, which includes consumer autonomy, animal welfare, transparency, social justice (equality), health (food safety) and environmental prosperity. A normative framework that analyses the interplay of these values per moral dilemma is the food ethics matrix, proposed by the Food Ethics Council (www.foodethicscouncil.org), situated in the UK. Using this matrix results in an overview of core values per stakeholder in a given food-related moral issue. An ethical issue is analysed according to the values of food producers, food marketers, consumers and society members in the context of well-being, autonomy and fairness (Mepham, 2000, 2010, 2013). Well-being loosely represents a utilitarianism approach and is results-driven; autonomy refers to Kant's categorical imperative; and fairness is an interpretation of the principle of equality. This way, values of various stakeholders can be mapped in light of different approaches in normative ethics and can serve as input for ethical decision-making. In technology, ethical considerations are increasingly important. Especially the value sensitive (Friedman and Hendry, 2019) or moral design (Wernaart, 2022) of new technology is high on the agenda. The main challenge here is to align the values of designers, that of technology in itself, and that of society (Van de Poel, 2015). The three can have differences that lead to moral challenges. One element that is often overlooked here is that technology in itself is never ethically neutral and can have a moral charge that is different from its designers or users. Good examples are various innovations in agriculture that help optimize food production with the aim of ending hunger, while some of these innovations appear to have a very negative (and sometimes unforeseen) effect on the environment. Perhaps the most notable example is the famous Haber/Bosch process, where the intended values are very far away from the actual, realized values in society. When we focus on private undertakings, ethics plays a role in the managerial and organizational aspects of running a business. This is mostly referred to as 'business ethics' although different terms are used as well, such as Corporate Social Responsibility (CSR), or sustainable entrepreneurship. In this chapter, we will stick to the words business ethics. In this field, we have observed that since the '60s there has been a shift from a shareholder to a stakeholder approach, and later a network approach (Freeman, 2010; Brand and Blok, 2019). In short this means that not only the norms and values (and interest) of shareholders are considered in ethical decision making within a private organization, but also those who are affected by or can affect the interest of the business. The systematic mapping of ethical considerations of the network surrounding a business can help in taking better decisions that are accepted and desired by the network the business operates in. The above theories are well-known approaches

normative
ethicsfood ethics
matrix

business ethics

in both descriptive and normative ethics, with selective applications in food innovation and technology. What is so far lacking in academia and applied sciences is a comprehensive analysis that bridges the various approaches and methods in the field of food tech innovation, and that leads to a useful applied method to guide ethical decision-making in food technology design.

8.3 Citizens or consumers?

Numerous studies have investigated the ways in which food marketing strategies, such as labelling, advertising, and packaging, can influence individuals' perceptions of sustainable food and shape their purchasing behaviour (Grunert *et al.*, 2014; Julia and Hercberg, 2017; Vanclay *et al.*, 2011). Presenting health and environmental information on food products during the purchasing process could assist individuals in making purchasing decisions that are both healthier and more sustainable (Anastasiou, Miller, and Dickinson, 2019; Brown, Harris, Potter, and Knai, 2020; Macdiarmid, Cerroni, Kalentakis, and Reynolds, 2020). However, the drivers and motivation of food consumers behind their willingness to change their eating habits are also linked to normative and ethical considerations (Franzo and Mc Laren, 2020). The success of new food technology depends on the willingness to change purchasing habits and eating behaviour. Therefore, it is important to include citizens' perspectives in the design process; they need to be co-designers. In this context, the attitude-behavioural intention gap is often spoken about (Vermeir and Verbeke, 2006), whereby consumers want to do something about the environment, or animal welfare, but they do not act accordingly. This is also known as the citizen-consumer paradox. However, other research reasons that there is no distinct citizen or consumer, and that they are much more interlinked (Brom *et al.*, 2006; de Bakker and Dagevos, 2012). People buy their food for reasons of both price and enjoyment, as well as principles and ideals, which could be called citizen-inspired consumption behaviour (de Bakker and Dagevos, 2012). In this chapter we use the term 'citizen', bearing in mind that each citizen is (almost) per definition a food consumer, and taking into consideration the wide moral concerns of many individuals that play a prominent role in the public debate around food innovation. On top of that, it is worth noting that, to date, approaching the required food transition from the perspective of consumer research (or consumer behaviour), as discussed in this section, has not yielded the expected results. Therefore, we will shift the focus of attention from consumer perspectives to citizens' perspectives, at the same time recognizing there is a grey area between them, especially as concerns food consumption.

attitude-
behavioural
intention gap

8.4 A moral data city hunt on food technology: towards an action-based research agenda

The above leads to the following research question: How can we enable the food sector and policymakers with crowdsourced ethics in food technology design to support the transition to a more sustainable food system?

To achieve this, we propose to use a method of crowdsourced ethics: the moral data city hunt method. By gathering opinions from a diverse group of people, we aim to gain a deeper understanding of the values and principles that inform ethical decision making. Our approach involves exploring not only the ethical judgements made by individuals but also the underlying values that drive those decisions. We are interested in examining how different ethical considerations are weighted and prioritized by individuals, and how this varies across different groups and contexts. Through this research, we aim to identify areas where public concerns or values may conflict with industry or regulatory priorities, and to develop strategies for addressing these tensions. Ultimately, we believe that a better understanding of ethical decision making in new food technologies will be essential for ensuring that these technologies are developed and implemented in ways that are consistent with public values and ethical principles.

crowdsourced
ethics

In previous research, we have introduced the ‘moral data city hunt method’ in applied studies on responsible drones or a city mobility app (van Veen and Wernaart, 2022; Wernaart *et al.*, 2023). The moral data city hunt method is a way to solicit citizen input, as it not only utilizes participatory and visual research methods but also encompasses a linguistic approach. By combining participatory, visual and linguistic approaches, complex topics are illustrated and discussed in an accessible way with those whom it eventually concerns: the citizens (Cornwall and Jewkes, 1995; Mitchell *et al.*, 2017; Vaughn and Jacquez, 2020). One example of a visual research method to be incorporated is photo-elicitation, where images are used to guide a research experiment. One of the many advantages is that visual images are more likely to trigger reflection, association, feelings, emotions, memories and information than a stand-alone verbal conversation (Collier and Collier, 1986; Glaw *et al.*, 2017; Harper, 2002). The results of the moral data city hunt give a sense of what citizens consider important, but also of the values that underlie them, and how these values relate to one another. This can help in the design process by ‘front-loading ethics’ as it is called in Value Sensitive Design (van den Hoven 2007). Since the translation of values into design requirements is context-dependent (van der Poel, 2013), there is a need for a solution where companies and policymakers can themselves incorporate the citizens voice into their design process.

The moral data city hunt method centres on two research approaches:

1. We use augmented utilitarianism (descriptive ethics) (Aliman and Kester, 2022). Citizens interact with prototypes of future technology and record their preferences in moral programming, in a moral (food) lab.

moral data city
hunt method

2. We ask citizens what they think of this moral programming of the prototype and use a Personal Value Dictionary (based on Ponizovskiy *et al.*, 2020) to analyse their wording.

Moral
Food Lab

A 'Moral Food Lab' allows for a structured and participatory approach to exploring ethical issues in the development of food technologies and consists of five phases: interviews, design, interaction, analysis, and translation. In the Interviews phase, we engage with stakeholders and experts in the field to gather insights and perspectives on the ethical dimensions of food technologies. This helps to inform the subsequent phases of the process. In the design phase, the insights gathered from the Interviews can be used to design a set of ethical scenarios related to food technologies. These scenarios are incorporated in a moral lab, and are designed to stimulate ethical reflection and deliberation among citizens. In the Interaction phase, we stimulate large-scale citizen participation in these moral labs. This phase provides an opportunity for citizens to participate in the ethical debate surrounding food technologies and to contribute to the development of ethical guidelines and principles.

linguistic
analysis

In the second part of the moral data city hunt method, we will perform a linguistic analysis of the reaction of the participants based on the linguistic analysis of Ponizovskiy *et al.* (2020). The reaction is analysed based on a Dutch or English library of words that are connected to the 10 values of Schwartz based on his Theory of basic values (Schwartz, 2012). The approach assumes that the language people speak is a behavioural expression of their corresponding values (Ponizovskiy *et al.*, 2020). This bottom-up approach has shown that people's own words are a better way (Boyd, 2015) of capturing their day-to-day behaviours than traditional self-reporting questionnaires (Rokeach, 1973; Schwartz 1994 and Gouveia, 2014).

Ultimately, the analysis is translated to design principles that can serve as input for further technology development, in alignment with the values of those whom it concerns: citizens.

8.5 The Graz experiment

How better to show people how the moral food lab works than to let them experience it? This is what we did at the STS conference in Graz, May 2023. Our experiment was composed of two main parts. First, we asked attendees of the conference to respond to three moral food dilemmas:

- A. Would you rather kill 1 cow or 10,000 crickets?
- B. Would you rather eat something that tastes good or something that is healthy?
- C. Do you prefer cultural values and traditions or innovative adaptations in your food?




FIGURE 8.1 The value profiles of two people that participated in the moral food lab. Note that the size of the words is in no way correlated with the number of times the words were mentioned

For each of these three questions people were asked to put a sticker on a line from left to right that corresponded to one of the scenarios. Second, after they had placed three stickers, we asked them about how they would see the future food system based on the three questions. This interview was recorded so that we could perform a linguistic analysis on the spoken texts.

For each person this resulted in three answers, quantified to a value between -100 and 100. Next to that the interview was analysed according to the Personal Value Dictionary of Ponizovskiy *et al.* (2020). This analysis resulted in a value profile of each person based on the words they used in the interview (see for an example, figure 8.1).

8.5.1 Overall results

Overall, 18 people joined the moral food lab. This is a very small sample. The main goal was not to collect a large sample, but rather to test our research approach, and find out if our proposed methodological steps can lead to meaningful data output that can be used to guide food transition through technology in a way that is in alignment with citizen’s values.

We did not ask about demographics during the experiment; however, we suspect the group we interviewed was very heterogenic and does not represent our population as a whole. First, the value data was normalized between participants. Because

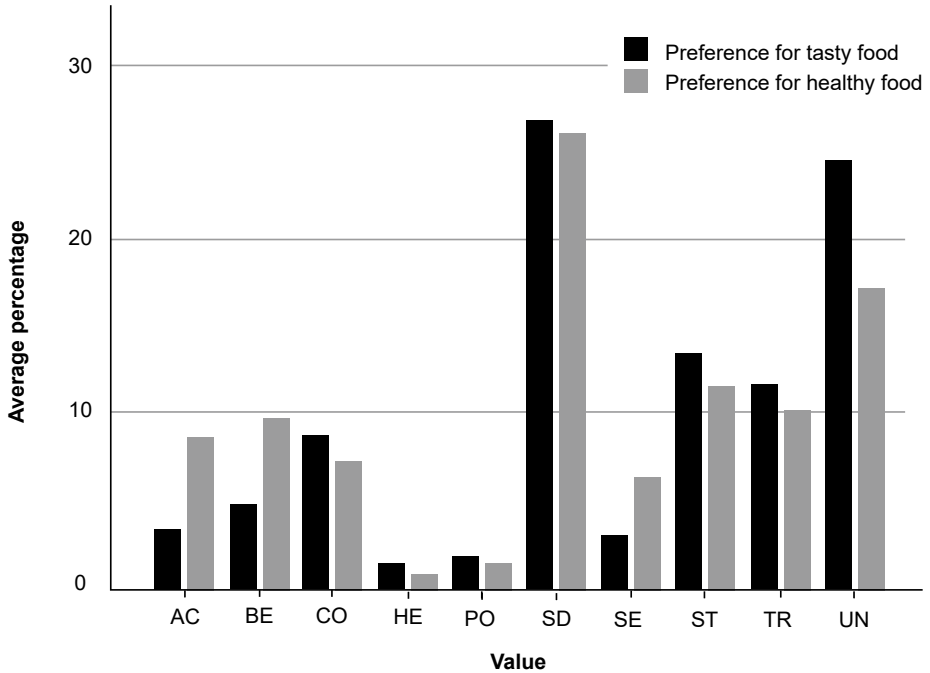


FIGURE 8.2 The average percentage of values people spoke in, grouped by how they responded to question B

of the nature of the questions and also the way we asked people to sketch their opinion of the future food system, some values were triggered more than others. To correct for that we subtracted the average value profile from each individual value profile, so that we could compare the relative value profiles to each other. We analysed the data based on the following two questions: 1. Did people who responded on one side of the moral dilemma have different value profiles from people who responded on the other side of the moral dilemma? 2. Do people with a value profile in which the dominant value has a personal focus respond differently to the three moral dilemmas than people with a dominant value that has a social focus?

Figure 8.2 shows the non-corrected percentages of values people spoke in. Notice that self-direction (SD) and universalism (UN) have the highest values, which could be explained by the nature of the subject or the way we addressed the question. Words that were typically used for self-direction (SD) were 'decision', 'choice', 'think' or 'know'. Words that were typically used for universalism were 'balance', 'culture' or 'environmental'. Question 1 only revealed minor differences between the different sides of the moral dilemmas. There was one minor difference in question B. People that were leaning more towards something that tastes good used more words related to stimulation (13.4%) than people who were leaning more towards something that is healthy (11.7%, $p = 0.089$).

For Question 2, people were placed in different groups based on the dominant value they spoke in. The groups were based on values with a personal focus on the one side: achievement, power, hedonism, stimulation, self-direction. And values with a social focus on the other side: security, conformity, tradition, universalism, benevolence. (Schwartz *et al.*, 2012, Figure 8.2). For each group the average response to the question was calculated.

The data was mostly too noisy to draw firm conclusions. The difference between the two groups for dilemma B were near significant ($p = 0.08$), where people with a dominant value with a personal focus leaned more towards healthy foods (10.11) and people with a dominant value with a social focus were rather neutral (1.22).

For future moral data city hunts we learned that it is important to be very neutral in asking the questions so that there is no value triggered in simply repeating the question or by doing what the interviewer asked (namely giving your opinion, triggering words like think, know and opinion). In general, we could say that the procedure we tested can lead to meaningful insights and sheds a unique light on how values are perceived to be at conflict amongst citizens. This can be translated to meaningful design-input in food innovation, and leads to a better value alignment between food businesses and society when taking the urgent necessary steps for food transition.

8.6 Future perspectives

We are witnessing a paradigm shift that goes beyond mere functionality and profit to redefine the nature of our food system. The moral design of food technology becomes a guiding principle based on ethical values and societal well-being. Sustainability becomes the cornerstone as innovative solutions prioritize the conservation of our planet and its resources. Social justice takes precedence as equal access to nutritious and culturally appropriate food becomes a fundamental right. Health and nutrition are prioritized through advances that promote well-being and address public health. Animal welfare is given due consideration to ensure humane treatment and minimize suffering throughout the production process. Transparency and citizen empowerment are paramount, promoting trust through clear information and comprehensive decision-making.

In this context of transformative change, we envision an idealistic scenario where proteins are extracted from the air we breathe, and meat is grown from nutrient solutions that nourish stem cells. This ground-breaking vision presents a tempting possibility – a world where global food production takes place within planetary boundaries. By integrating innovative technologies, we are able to print delicious, nutrient-rich food, reuse bio-waste streams, and reduce food waste to near zero. By using insect and seaweed-based foods in various innovative forms, we are witnessing an ambitious transformation that has the potential to reduce the incidence of

diabetes and coronary heart disease. This idealistic progress is made possible by the mindful development of new technologies that aspire to align with the ethical requirements of society. In this ambitious future, our food will ideally become a harmonious blend of culinary pleasure, environmental and animal welfare considerations, and improved public health, while enabling food security for the whole planet.

However, it is crucial to recognize that the direction of this transformation should ultimately be the choice of individuals and communities. People should have the agency to determine how this shift towards a more sustainable, just, and conscientious food system unfolds. While technology and ethical considerations are significant drivers of change, citizens should play a central role in shaping this future. This empowerment allows for a more democratic and inclusive approach to food transformation, ensuring that the choices made align with the values and preferences of the diverse global population. In this way, we can collectively forge a path forward that reflects the collective will and aspirations of our society.

References

- Aliman, N.-M., and Kester, L., 2022. 4. Moral programming. In *Moral design and technology* (pp. 63–80). Wageningen Academic Publishers, Wageningen, the Netherlands. https://doi.org/10.3920/978-90-8686-922-0_4.
- Anastasiou, K., Miller, M., and Dickinson, K., 2019. The relationship between food label use and dietary intake in adults: a systematic review. *Appetite*, 138:280–291.
- De Bakker, E., and Dagevos, H., 2012. Reducing meat consumption in today's consumer society: questioning the citizen-consumer gap. *Journal of agricultural and environmental ethics*, 25:877–894. <https://doi.org/10.1007/s10806-011-9345-z>.
- Boyd, R., Wilson, S., Pennebaker, J., Kosinski, M., Stillwell, D., and Mihalcea, R., 2021. Values in words: using language to evaluate and understand personal values. *Proceedings of the International AAAI Conference on Web and Social Media*, 9:31–40. <https://doi.org/10.1609/icwsm.v9i1.14589>.
- Brand, T. and Blok, V., 2019. Responsible innovation in business: a critical reflection on deliberative engagement as a central governance mechanism, *Journal of responsible innovation*, 6:4–24, DOI: 10.1080/23299460.2019.1575681.
- Brom, F.W.A., Visak, T., and Meijboom, F., 2006. Food, citizens, and market: the quest for responsible consuming. *Understanding consumers of food products*, 610–623.
- Brown, K.A., Harris, F., Potter, C., and Knai, C., 2020. The future of environmental sustainability labelling on food products. *The lancet planetary health*, 4(4), e137–e138.
- Buskermolen, D.O., and Terken, J., 2012. Co-constructing stories: a participatory design technique to elicit in-depth user feedback and suggestions about design concepts. *Proceedings of the 12th Participatory Design Conference: Exploratory Papers, Workshop Descriptions, Industry Cases*, 2:33–36. <https://doi.org/10.1145/2348144.2348156>.

- Collier, J., and Collier, M., 1986. *Visual anthropology: photography as a research method*. UNM Press, Albuquerque, NM, USA.
- Cornwall, A., and Jewkes, R., 1995. What is participatory research? *Social science & medicine*, 41:1667–1676. [https://doi.org/10.1016/0277-9536\(95\)00127-S](https://doi.org/10.1016/0277-9536(95)00127-S).
- Craft, J.L., 2013. A review of the empirical ethical decision-making literature: 2004–2011. *Journal of business ethics* 117: 221–259. <https://doi.org/10.1007/s10551-012-1518-9>.
- Crane, A., Matten, A., Glozer, and Spence, L., 2019. *Business ethics* (5th ed.). Oxford University Press, Oxford, UK.
- Crippa, M., Solazzo, E., Guizzardi, D. *et al.*, 2021. Food systems are responsible for a third of global anthropogenic GHG emissions. *Nature food* 2:198–209. <https://doi.org/10.1038/s43016-021-00225-9>.
- Emergen Research, 2022. Food tech market size, share, trends, by technology, by application, and by region – Forecast to 2027. Emergen Research. <https://www.emergenresearch.com/press-release/global-food-tech-market>.
- Freeman, E., 2010. *Stakeholder theory*. Cambridge University Press, Cambridge, UK.
- Frewer, L.J., Bergmann, K., Brennan, M., Lion, R., Meertens, R., Rowe, G., ... and Vereijken, C.M.J.L., 2011. Consumer response to novel agri-food technologies: implications for predicting consumer acceptance of emerging food technologies. *Trends in food science & technology*, 22:442–456.
- Friedman, B. and Hendry, D.G., 2019. *Value sensitive design, shaping technology with moral imagination*. MIT press, Cambridge, UK.
- Glaw, X., Inder, K., Kable, A., and Hazelton, M., 2017. Visual methodologies in qualitative research: autophotography and photo elicitation applied to mental health research. *International journal of qualitative methods*, 16:1. <https://doi.org/10.1177/1609406917748215>.
- Gouveia, V.V., Milfont, T.L., and Guerra, V.M., 2014. Functional theory of human values: testing its content and structure hypotheses. *Personality and individual differences*, 60:41–47. <https://doi.org/10.1016/j.paid.2013.12.012>.
- Harper, D., 2002. Talking about pictures: a case for photo elicitation. *Visual studies*, 17:1. <https://doi.org/10.1080/14725860220137345>.
- Jerome, K., and Vanclay, J.S., 2010. Customer response to carbon labelling of groceries [J]. *Journal of consumer policy*, 6:52–55.
- Jones, T.M., 1991. Ethical decision making by individuals in organizations: an issue-contingent model. *Academy of management review* 16: 366–395.
- Julia, C., and Hercberg, S., 2017. Nutri-Score: evidence of the effectiveness of the French front-of-pack nutrition label. *Ernahrungs Umschau*, 64:181–187.
- Macdiarmid, J.I., Cerroni, S., Kalentakakis, D., and Reynolds, C., 2021. How important is healthiness, carbon footprint and meat content when purchasing a ready meal? Evidence from a non-hypothetical discrete choice experiment. *Journal of cleaner production*, 282:124510.
- Mepham, B., 2000. A framework for the ethical analysis of novel foods: the ethical matrix. *Journal of agricultural and environmental ethics* 12: 165–176. <https://doi.org/10.1023/A:1009542714497>.

- Mepham, B., 2010. The ethical matrix as a tool in policy interventions: the obesity crisis. In: Gottwald, F.T., Ingensiep, H.W. and Meinhardt, M. (eds) *Food ethics*. Springer, New York, NY, USA, pp. 16–30.
- Mepham, B., 2013. Ethical principles and the ethical matrix. In: Clark, J.P. and Ritson, C. (eds) *Practical ethics for food professionals: ethics in research, education and the workplace*. John Wiley and Sons, Ltd., Hoboken, NJ, USA, pp. 39–56.
- Mitchell, C., Lange, N.D., and Moletsane, R., 2017. *Participatory visual methodologies: social change, community and policy*. SAGE.
- Ponizovskiy, V., Ardag, M., Grigoryan, L., Boyd, R., Dobewall, H., and Holtz, P., 2020. Development and validation of the Personal Values Dictionary: a theory-driven tool for investigating references to basic human values in text. *European journal of personality*, 34:885–902. <https://doi.org/10.1002/per.2294>.
- Purvis, B., Mao, Y. and Robinson, D., 2019. Three pillars of sustainability: in search of conceptual origins. *Sustain Sci* 14:681–695. <https://doi.org/10.1007/s11625-018-0627-5>.
- Rest, J.R., 1986. *Moral development: advances in research and theory*. Praeger, New York, NY, USA.
- Rokeach, M., 1973. *The nature of human values*. Free Press, New York, NY, USA.
- Schwartz, S.H., 1994. Are there universal aspects in the structure and contents of human values? *Journal of social issues*, 50: 19–45. <https://doi.org/10.1111/j.1540-4560.1994.tb01196.x>.
- Schwartz, S.H., 2012. An overview of the Schwartz theory of basic values. *Online readings in psychology and culture*, 2. <https://doi.org/10.9707/2307-0919.1116>.
- Schwartz, M.S., 2016. Ethical decision-making theory: an integrated approach. *Journal of business ethics* 139: 755–776. <https://doi.org/10.1007/s10551-015-2886-8>.
- Siegrist, M., and Hartmann, C., 2020. Consumer acceptance of novel food technologies. *Nature food*, 1:343–350.
- United Nations Environment Programme, 2022. *Emissions gap report 2022: the closing window – climate crisis calls for rapid transformation of societies*. Nairobi. <https://www.unep.org/emissionsgap-report-2022>.
- Van der Meulen, B. and Wernaart, B., 2022. Food law and regulatory affairs. In: Wernaart, B. and Van der Meulen, B. *Applied food science*. Wageningen Academic Publishers, Wageningen, the Netherlands. pp. 21–43. https://doi.org/10.3920/978-90-8686-933-6_20p.
- Vaughn, L.M., and Jacquez, F., 2020. Participatory research methods – choice points in the research process. *Journal of participatory research methods*, 1. <https://doi.org/10.35844/001c.13244>.
- Van de Poel, I., 2015. Values in engineering and technology. In: Gonzalez, W.J. (ed.) *New perspectives on technology, values, and ethics*. Boston studies in the philosophy and history of science, Volume 315, Springer, Cham, Switzerland, pp. 29–46.
- Van den Hoven, J., 2007. ICT and value sensitive design. In P. Goujon, S. Lavelle, P. Duquenoy, K. Kimppa, and V. Laurent (eds.), *The information society: innovation, legitimacy, ethics and democracy in honor of Professor Jacques Berleur S.J.* (pp. 67–72). Springer, New York, NY, USA. https://doi.org/10.1007/978-0-387-72381-5_8.

- van Veen, M., and Wernaart, B., 2022. Building a techno-moral city – reconciling public values, the ethical city committee and citizens’ moral gut feeling in techno-moral decision making by local governments. *OpenLivingLab Days Conference 2022*. p. 115.
- Vermeir, I., and Verbeke, W., 2006. Sustainable food consumption: exploring the consumer “attitude – behavioral intention” gap. *Journal of agricultural and environmental ethics*, 19:169–194. <https://doi.org/10.1007/s10806-005-5485-3>.
- Wernaart, B., 2021. Developing a roadmap for the moral programming of smart technology. *Technology in society*, 64:101466.
- Wernaart, B., 2022. Food ethics. In: Wernaart, B. and Van der Meulen, B. *Applied food science*. Wageningen Academic Publishers, Wageningen, the Netherlands. pp. 45–64.
- Wernaart, B., 2022. *Moral design and technology*. Wageningen Academic Publishers, Wageningen, the Netherlands.
- Wernaart, B., Vaznyte, R., van Velzen, M.B., and Dantuma, I.L., 2023. Responsible civilian drones. In *Applied human rights*. Wageningen Academic Publishers, Wageningen, the Netherlands. pp. 151–164. <https://doi.org/10.3920/978-90-8686-943-5>.
- World Bank, 2021. *Agriculture overview: development news, research, data*, World Bank.

How natural is our food?

How relevant is that word for the design of future food?

Niels Louwaars

Abstract

Ever since we were hunter-gatherers, we have learnt from nature to produce food increasingly efficiently. Hunter-gatherers collect their food directly from nature; we currently get our food from agriculture and domesticated crops and animals with significantly changed DNA, and more or less complex production and processing steps.

Even though food production and products are increasingly distant from nature, consumers have strong positive perceptions about 'natural' when it relates to food. This perception of naturalness is widely used in food marketing, but has also reached formal policies and regulation. Ethical debates are ongoing about soilless horticulture, genome editing and processed food, leading to various dilemmas. Where do we draw the line of 'naturalness' now, and how will that affect the future of our food? There is a complex relationship between the perception of 'natural' and the terms 'healthy' and 'sustainable'.

When looking at our food and how it is designed, it would be interesting to see how the definition of 'naturalness' evolves in the future. We conclude in our analysis that there is either no such thing as truly natural food, or that all our food is intrinsically natural since it is all derived from nature and natural processes. However, 'technology' often has a negative connotation when connected to food. Future food trends may need to take this perception of naturalness into account, but in the absence of clear definitions, the dilemmas will continue to create space for innovations.

9.1 Defining 'naturalness'

For most people 'naturalness' is a highly appreciated characteristic of many goods and food in particular, as exquisitely featured by the title of Allan Holland's paper: 'Nature – every drop of it – is Good' (Holland, 1996). It is deeply engrained in human psychology (Li and Chapman, 2012, Moscato and Machin, 2018, Román *et al.*, 2017). It particularly features in the context of food, such as in the discussions about genetic modification (Morris, 2001) and additives (Marley and Chambers, 2019). In the public perception there seems to be an intrinsic link between the terms 'natural' and 'healthy and sustainable', which makes it an important term in food marketing; however, without a proper definition it is bound to create confusion (Ferris, 2010).

The use of the term is therefore relevant, both in terms of ethics and in the future design of our food. The term ‘Natural foods’ is relevant to consumers, yet, like the term ‘green products’ (Durif *et al.*, 2010), it frustrates producers due to the lack of a formal definition of ‘natural’ or the even less delineated term ‘green’.

The terms ‘natural’ or ‘naturalness’ are not as straightforward as they may sound. Dictionaries show a wide range of meanings, from usual, spontaneous, untreated, to spontaneous and biological.

Dictionaries makes it clear that ‘naturalness’ is best explained when including an antonym: ‘the degree to which something appears to be natural rather than man-made’ (MacMillan¹), or Webster Dictionary: ‘occurring in conformity with the ordinary course of nature: not marvelous or supernatural’ (Webster²), or ‘as found in nature ...’, or ‘from nature; not artificial or involving anything made or caused by people’ (Cambridge³). The term is thus described through the antonym ‘contrived’, meaning artificial, manufactured, engineered and planned.

There are distinct differences in the perception of ‘naturalness’ when applied to food and medicines (Rozin *et al.*, 2004). When applied to food, the word ‘natural’ is commonly only connected to post-harvest operations: ‘food that has undergone minimal processing and contains no preservatives or artificial additives’,⁴ or ‘food that has had little processing (= preparation, change, or treatment of food) and nothing added to it’.⁵

perception

In a wide range of debates, however, the term naturalness is currently used in a much wider context, i.e. not just the state of the food itself (processing, additives), but also relating to the way it is produced. This opens up the use of the wide range of synonyms and antonyms of ‘naturalness’. In this chapter, we analyse the current ‘naturalness’ of our food and relate it to the trends and technical developments that are shaping our food and that will continue to do so in the future.

9.1.1 *The origins of our food*

Food supply is a basic prerequisite of all living species, including humans. Food sources shape the evolution of species to a significant extent. Humans developed into omnivores either by want or necessity and must have tested a wide diversity of natural products by trial and error for their palatability and possible toxicity creating a significant knowledge base that had to be transferred throughout generations. Such knowledge about food was largely communal, whereas medicinal knowledge was transferred from individual to individual successor in most cultures. Hunter-gatherers had to deal with the seasonality of many food stuffs. In nature,

1 <https://www.macmillandictionary.com/dictionary/british/naturalness> (accessed 1 July 2023).

2 <https://www.merriam-webster.com/dictionary/natural> (accessed 1 July 2023).

3 <https://dictionary.cambridge.org/dictionary/english/natural> (accessed 1 July 2023).

4 <https://www.merriam-webster.com/dictionary/natural%20food> (accessed 1 July 2023).

5 <https://dictionary.cambridge.org/dictionary/english/natural-food> (accessed 1 July 2023).

berries, nuts, roots, grains and game are not always available and despite romantic ideas about forest or nomadic life, food shortages must have been quite common at times.

domes-
tication

Farmers started to domesticate crops based on the products that they used to collect. They used their gradually increasing knowledge about soils, water, and plant diseases to adapt their natural environment as well as the plants and animals themselves, forcing humans to settle close to their crop production fields. This led Harari (2014) to suggest that early humans did not domesticate their crops, but rather that the crops domesticated humans, i.e. led them to adopt a sedentary life in permanent houses (domus) (Harari, 2014). This co-evolution of farming and social organization of humans has continued over the past millennia at increasing speed. Preserving and storing food creates more and more refined – in both meanings of the word – food for the increasing human population. Pickled cucumbers, salted herrings, cooked and sealed red beet and many more examples create food that can be consumed during the hunger season and that tastes different from the fresh product. The grinding of grains for baking bread or brewing beer was a food processing technique already applied before farming was ‘invented’ (Diamond, 1997; Harari, 2014).

Our food is thus, both in terms of the method of production and the content of the food itself, based on nature, but increasingly distant from it.

9.1.2 *Food production: learning from nature*

Food is of paramount importance for survival. Humanity has evolved because of its capacity to learn from nature and use that knowledge to transform the natural resources and environment to serve its needs. That transformation has been going on at least since the dawn of agriculture some 11,000 years ago. Concerns focused for a long time on household and community food security and innovations to adapt to new conditions when communities migrated. Gradually, and notably from the 1960s onwards, concerns increased with the complexity of food systems, including emerging risks associated with environmental degradation, geopolitical dependencies and climate change. Morality, which has always played an important role in food, in history through taboos and religion, is claiming an increasingly important role in debates about the future of our food. Perceptions of naturalness are an important issue in this respect.

Learning from nature has always been the driving factor behind innovation in our food production systems and our food itself. Learning that harvesting, selecting and replanting seeds gradually produced better crops, with plump grains that don't shed like the ancestors of the crops in natural conditions, that grow better when other plants (weeds) are removed, and that taste better than their often bitter or even toxic relatives. This ‘domestication’ has produced in various locations and moments in history different crops and uses. For example, humans in the Fertile Crescent gave us barley, wheat and lentils; in Latin America potato, beans, tomato,

maize, cassava and cocoa; and in China rice and soybean. Africa is the home of oil palm, coffee and sorghum. Such domestication continues until the present day. (Østerberg *et al.*, 2017).

Also, in terms of production we have increasingly used concepts of nature in order to move crop production systems away from actual natural ecologies. A major development in agronomy followed after Justus von Liebig who identified the chemical needs of plants in 1840, leading to the development of chemical fertilizers and eventually to soilless horticulture using rockwool or just water as substrate. Covering vegetable crops with glass to protect seedlings in spring has developed into computerized greenhouses and vertical farming under artificial lighting that fully recirculate water and where reared antagonists control plant pests.

agronomy

Chemical pest control has a very long history of protecting crops against pests and diseases. The use of sulphur dates back some 4500 years. Plant-based insecticides, such as Derris and Pyrethrum, were identified around 1750, followed by many others (Hikal *et al.*, 2017) and were largely replaced in the late 20th century by chemically produced pesticides. Biological pest control obtained by releasing enemies of the pest started in the late 19th century with the introduction of the Vedalia beetle *Rodolia cardinalis* from Australia to control cottony cushion scale in the USA.⁶ This is now a multimillion-dollar industry, rearing insects, parasitic nematodes and micro-organisms to ensure that several crops can now be grown in greenhouses without the use of any other crop protection measures but natural enemies of the major pests.

chemical
pest controlbiological pest
control

Important lessons from nature affecting the evolution of how our crops themselves were drawn by Camerarius (1694), who proved sexuality in plants which started an understanding of cross-breeding, by Mendel (1864) who formulated the laws on heredity which spurred scientific plant breeding, and Watson and Crick (1952) who described the structure of DNA. Many others identified novel uses of these and other basic concepts, which is now the basis of a thriving plant breeding sector using genetics, a variety of DNA techniques and big data/artificial intelligence to support plant breeding, i.e. in the creation of diversity and selecting from that diversity. Over the millennia, crops have changed in such ways that consumers would hardly be able to recognize the species from which they were originally derived; not just Teosinte developing into maize, and the grasses that created our cereals, but also the cabbage family with its many uses, and the tiny Mexican berries that we now call tomatoes.

crops

So, agriculture itself developed through subsequent steps in learning from nature and applying such knowledge effectively for the benefit of food production. Does this make our food unnatural? Is there a line that we crossed? Did we cross this when we started planting seeds and tending our gardens, or when we changed our

6 <https://ipm.ucanr.edu/agriculture/citrus/cottony-cushion-scale/>.

crops (and animals) in such a way that they would not survive without human care? Let's look at it in more detail.

9.1.3 *Dilemmas in agronomy with regard to naturalness*

The term 'natural' creates a wide range of dilemmas when applied to food production, especially when 'natural' is equated with 'good' and 'sustainable'.

copper sulphate

The use of copper sulphate, a naturally occurring chemical, in crop protection (Tamm *et al.*, 2022) is very effective against a range of plant pathogens, for example in grapes. But despite its 'naturalness', and on that basis approved in organic wine production, it is clearly not the most environmentally friendly compound as heavy metals like copper easily accumulate in the soil.

pyrethrum

The naturalness of 'pyrethrum' may also be a cause for debate, since the insecticide may be harvested from flowers in Kenya and is used way outside the context of its natural presence in intensive agriculture all over the world. The next question is whether the chemically produced pyrethrum compound can be considered as natural as the same chemical product extracted from plants, or whether a chemical derivative of the compound selected for improved effectiveness can still be considered 'natural'. Taking this dilemma to a more extreme form is the use of the insecticidal property of *Bacillus thuringiensis*, a bacterium that is quite common in many soils (de Maagd, 2015). It can be applied to crops directly from the soil or selected and purified to improve efficacy as a commercial spray. The same insecticidal property of the bacterium may also be introduced in the DNA of crops through transgenesis, i.e. genetic engineering.

crop production

That 'naturalness' creates some dilemmas in terms of sustainability and may also be illustrated by the substrates we use in crop production. The use of soil as substrate is prescribed by the rules of organic agriculture in Europe (EU, 2018). Other substrates like rockwool, hydroponics and aquaponics (Friscella *et al.*, 2021) can have clear sustainability advantages as they can prevent damage by soil-borne diseases and pests in greenhouse horticulture without the use of pesticides. They can also reduce the use of the soil-improver peat, a substrate that is harvested from nature at a rate that is unsustainable in several source-countries.

So unnatural substrates for crop production can have sustainability advantages, and some plant-based crop protection products can have clear sustainability challenges. Agriculture is either not natural, because it involves a wide range of human interventions, or it is all natural, because all interventions are based on learning from nature. And the relationship between naturalness and sustainability is dubious. How do we explain the phrase 'Nature – every drop of it – is good' from the introduction of this paper?

9.1.4 *Dilemmas in plant breeding with regard to naturalness*

breeders

Throughout history, farmers and breeders have adapted crops to their needs, initially through selection only. Breeders then found a variety of ways to find and create diversity to select from:

- Introduction of diversity from other regions to cross it with the local gene pool (from ca. 1800), something that would not happen in nature (Louwaars and Burgaud, 2016);
- The use of the natural concept of heterosis by creating hybrids using controlled crossing of selected parent lines in 1925 (Duvick, 2001);
- The use of mutagenesis using ionizing radiation as a way to create diversity in the 1920s (Stadler, 1928; Kharwal, 2012);
- Duplication of the number of chromosomes with natural colchicine in the late 1930s (Eigsti, 1938);
- Chemical mutagenesis on plants/tissues and in cell cultures in the 1940s (Auerback and Robson, 1944);
- Embryo rescue, preventing abortion in interspecific crosses in the 1950s (Nitsch, 1951);
- Tissue culture techniques for quick propagation which took off in the 1960s (Murashige and Skoog, 1962);
- Doubled haploids using cultured gametes to accelerate homozygosity since the early 1970s (Kasha, 1974);
- Cell fusion and protoplast fusion with the aim of combining genomes in the 1970s (Carlsen and Smith 1972);
- Transgenesis: the transfer of functional genes from one species to another in the 1980s (Herrera-Estrella *et al.*, 1983), which some claim is not unnatural (Wickel and Li, 2019);
- Molecular marker-assisted selection in the late 1980s (Paterson *et al.*, 1988);
- Cisgenesis: the transfer of functional genes within (or between crossable) species in the 2000s (Jacobsen and Schouten, 2007);
- Targeted mutagenesis: the targeted cutting and/or replacement of base pairs (ZFN, TALEN) (Shukla *et al.*, 2009) and techniques based on Clustered Regularly Interspaced Short Palindromic Repeats (CRISPR-Cas) (Jinek *et al.*, 2012);
- Genomic selection (Nakaya and Isobe, 2012) increasingly supported by artificial intelligence to analyse big data.

All these methods require human intervention, so in the strictest sense will not yield ‘natural’ products in the ‘Cambridge definition’ above. Crossing plants from different origins requires a person to bring them together so is not something that would easily happen in nature.

When we consider all inventions that are created by observing natural phenomena and using that knowledge in novel ways, then all the above could be considered ‘natural’. CRISPR-Cas was developed by observing that certain DNA sites in bacteria (van der Oost *et al.*, 2009), the Clustered Regularly Interspaced Repeats, are in fact a library of identifiers of viruses that the bacterium uses to quickly respond to a virus attack, followed by its capacity to cut and destroy them. That knowledge can now incite targeted mutations in higher organisms due to faults created by the natural repair mechanism, or guide such repairs thereby intentionally creating more specific alterations. So it is highly technological but is based on knowledge

of the natural processes. Still, in focus groups ‘perceived naturalness was found to be the main reason for obtaining different levels of acceptance, not only between gene-editing and genetic modification but across all breeding techniques examined’ (Nales, 2023). There are different perceptions with respect to levels of unnaturalness (Rozin, 2005) among such breeding processes. The fact that all these methods are derived from ‘learning from nature’ does not preclude value judgements.

9.1.5 *Dilemmas in food processing*

The term ‘natural food’ is most commonly used in terms of aspects of food processing (Chambers *et al.*, 2018, Evans *et al.*, 2010). The most extreme opponents of processing are the followers of the concept of raw foodism, which when fully implemented carries a range of health issues (Runn, 2007). Apart from raw vegetables, and in some cultures raw fish and meat, all our food is cooked, fried or otherwise processed and therefore may not be considered inherently unnatural.

factory-
processed food

The claim of unnaturalness is commonly related to factory-processed food – refined white sugar as opposed to honey or less refined brown sugar, and constituted food like cookies, soft drinks and tinned soup. In factories, individual constituents like proteins, oils or sugars may be added, as well as preservatives, and taste and colouring compounds, which could be ‘natural’, e.g. colouring from berries, or chemically produced copies of natural compounds such as ascorbic acid (vitamin C). Also here, clear lines are difficult to draw and dilemmas occur (Aschemann-Witzel *et al.*, 2019) when consumers want to choose on the basis of their ‘naturalness’ evaluation (Battacchi *et al.*, 2020). Marketeers know the importance of packaging in this consumer assessment (Binninger, 2017).

meat
replacement
products

An example is vegetarian meat replacement products whereby there is no animal suffering, but which are highly processed. Of course, consumers may also choose the less processed plant products which don’t have a ‘meat-like’ structure and which require some more work to prepare in the kitchen. The replacement is less easy with vegetarian ‘milk’, processed from almonds, oats or other products, especially when vitamins and minerals are to be added to make it more comparable to dairy. Concerns about such ‘processed food’ are thus weighed against animal welfare and sustainability. One other concern of vegetarian users of such ‘milk’ is that the press cake that is left at the end of the production process is used as feed in the animal production industry.

biotechnology

The use of preservatives reduces food waste, which is an important sustainability issue. Vinegar, sugar and salt may be considered natural preservatives though. Interestingly, there is little debate about the use of biotechnology in food processing, genetically changed microorganisms that are used in cheese making, beer brewing and all kinds of fermentation processes.

9.1.6 *More or less natural?*

So, all our food is natural, or all of it is totally unnatural. The crops we eat do not look like the natural plants that they were derived from through intentional and

often highly technical human interference, so our food is not natural. Our plants and animals were derived from nature and were changed through breeding processes that humans learnt from nature, so our food is clearly not unnatural (artificial). Our plants are grown in highly controlled conditions, with a lot of human interference in soil preparation, weeding, and pest control, using increasingly complex and computerized machines, but all agronomy is based on studying nature. Most of our food is processed at home or in increasingly complex factories using heat, fermentation and other technologies, which are to a very large extent based on learning from nature. Attempts to integrate the different 'naturalness' components in one system, the Food Naturalness Index' (Sanchez-Siles, *et al.* 2019) have not yet been rolled out on a large scale.

So, drawing a clear line under what which we could call 'quite natural' is difficult, especially for scientists. Ethics can support the drawing of such lines as shown by the biotechnology debates (Louwaars and Jochemsen, 2021). The organic movement, originating from an environmental perspective banning chemical fertilizers and crop protection products, currently puts Health, Ecology, Fairness, and Care' at the centre of its philosophy,⁷ based on the views of the Biodynamic (brand name Demeter) sector, 'a holistic, ecological, and ethical approach to farming, gardening, food, and nutrition, rooted in the work of philosopher and scientist Dr Rudolf Steiner, whose 1924 lectures to farmers opened a new way to integrate scientific understanding with a recognition of spirit in nature.'⁸

ethics

Based on the 'care' principle and the 'partner of nature' perspective of this movement, the intrinsic value of the organism (and the cell) is considered important. Extensive discussions in organic and bio-organic circles (Wilbois *et al.*, 2012) formally resulted (IFOAM, 2017) in the adoption of the concept that breeding methods where humans interfere in(!) the cell to change the DNA are not compatible with organic principles and values. These include cell fusion and mutation breeding, and any products of 'genetic engineering'. Marker-assisted selection and genomic selection are welcomed though as these techniques do not alter the plant (IFOAM, 2017). The biodynamic movement also opposes the use of hybrids, even though the word 'natural' or 'naturalness' is rarely used in their communication.

Consumers seem to understand that organic *and biodynamic* food is more natural (Hemmerling *et al.*, 2016), healthier (Siipi, 2013; Lusk, 2019) and more sustainable, but all these three components are so complex that such claims cannot be made across the board, as shown in previous sections in this paper (Amos *et al.*, 2014, 2019).

On the other hand, some religious groups approve of such breeding techniques as long as the species barriers are not crossed (Jochemsen, 2000, 2008), and environmental groups (Greenpeace, 2021) see benefits in using such laboratory techniques, but only for selection purposes. The European Group on Ethics in Science

⁷ <https://www.ifoam.bio/>.

⁸ <https://www.biodynamics.com/what-is-biodynamics> (accessed 1 July 2023).

and New Technologies (European Commission, 2021) sees few ethical concerns in the application of genome editing in farm animals or plants, but they do identify possible socio-economic concerns with regard to their potential to change the structure of the breeding sector and farming as such.

culture

The ‘culture’ perspective exists alongside the above ethical perspectives. This may be linked to the values of ‘local-for-local’ principles supporting both local food culture and reducing excessive transport of food products, as exemplified by for example the ‘slow food movement’ (Schneider, 2008). Also here, some dilemmas appear with respect to sustainability (when products can more efficiently be produced elsewhere) and ethics, as identified above. What if, for example, the Sangiovese grape can be ‘saved’ by making it resistant to diseases and tolerant to drought using genome editing? Local food culture is gaining importance in certain groups in the population, opposing globalized pasta, burgers and fries, and some welcome and others loathe the internationalized curries and rice dishes in the diversity of exotic restaurants on offer. Local-for-local does not necessarily link to local food culture though. When will quinoa, grown in the Dutch polders, become part of the Dutch food culture, in the same way as tomatoes which also originate from Latin America have become a key component of Italian cuisine?

Both the globalization of food culture, and the emergence of a range of concepts like natural, sustainable and organic, have multiplied the food choice for consumers, notably for those that have the financial means to choose.

9.2 The future of food and the choices and dilemmas for consumers

Technological developments, demographic and global dependencies will continue to change food habits and food choices. How the perception of naturalness will play out with cultured meat (Siegrist and Sütterlin, 2017), with food made of products produced by algae, or with any innovations in food processing, cultivation and reproductive material, is an important question for the designers of new foods.

There is proof that the arguments around the concept of ‘naturalness’ play a role in public debates and in policy and regulation. Such arguments play an important role in policy debates about genetic engineering (currently notably genome editing and ci-genesis), about pesticides (how will ‘green crop protection products’ be assessed?) and about alternative production and processing methods (the Italian Government suggests banning cultured meat on the basis of food culture arguments). The regulators may be guided by their specialists in ethics after the food safety assessments under the rules and procedures of Novel Food (in the EU⁹) have been cleared. Differences in perception among countries (Rozin *et al.*, 2012) may thus influence regulatory decisions.

⁹ https://food.ec.europa.eu/safety/novel-food_en.

Apart from regulation, consumer behaviour will also be important for the success of foods produced in novel ways (Lusk *et al.*, 2014). Will cultured meat be considered natural enough because it is created with ‘real’ animal cells rather than extrusion, recombination technologies including a wide range of additives to create vegetarian meat? Will the opposition to calling vegetarian meat-replacement products ‘meat’ also apply to such products?

consumer
behaviour

Akin *et al.* (2019) claim that there are spillover effects of ‘earlier’ debates on food and food technologies, such as GMO to future ones (in their case nano, but equally relevant for other novel foods), so all aspects of the debates need to be taken into account.

9.2.1 *Neophobia*

Neophobia, the strong psychological resistance to new things is coined in the debate, setting known (natural) products against novel ones (Siipi, 2013, Siegrist and Hartmann, 2020). Neophobia is analysed as either a personality trait (Pliner and Hobden, 1992) or a state, measured in food preference tests (Hobden and Pliner, 1995). In a review study, Rabadán and Barnabéu (2021) conclude that a downward trend over time in the level of food neophobia is in general connected to increased education, income and urbanization. They do consider food neophobia as a key variable in adoption of new foods (e.g. genetically modified foods) or novel food production techniques (e.g. nanotechnology).

9.2.2 *Justice and trust*

Neophobia is certainly not the only reason for opposition to new technologies and new types of food. Uncertainties about technologies themselves (‘opening Pandora’s box’), lack of participation (‘kept in the dark’) and questions about who will use and benefit from the technologies (the rich?, large corporations?) are arguments that have played a role, next to health concerns in the debate on radiation of food in the 1990s (Roberts, 2014; Degreef, 2018) and nanotechnology this century (Davies *et al.*, 2010), and indeed in genetic modification. So, naturalness is not the only aspect determining acceptance or rejection of a new food technology.

This objection to ‘going against nature’ reflects “wider unease about science, about technological modernity, and about hubris.” (Macnaghten, 2014). While Siegrist and Sütterlin (2017) point out the ‘trait’ aspect of neophobia such that ‘consumers rely on symbolic information when evaluating foods, which may lead to biased judgments and decisions’ (Siegrist and Sütterlin, 2017), natural science arguments are not the only reason why people adopt or reject our future foods.

9.2.3 *Balancing act for consumers*

Apart from deep psychological opposition to new foods, consumers are given the complex choice of balancing their conflicting interests. Roman *et al.* (2017) point to the interest consumers have in saving cooking time (convenience food) and at the

conflicting
interests

same time avoiding processed food. Siipi (2013) focuses on the dilemma of cultured meat, which can be considered environmentally and animal friendly. Discussions about whether 'organic' is a sustainability concept or basically an ethical concept that foregoes certain environmental sustainability options, contribute to the complexity of making 'the right' choice.

The marketing concept of 'green' products may facilitate such choices compared to 'natural', as it is understood to combine different components of the 'positive choice'. Choosing 'green' not only gives a positive feeling to the consumer (Maniatis, 2016), but also triggers positive social behaviour (Mazar and Zhong, 2010). For example, Cavaliere and Ventura (2018) found that the use of (any) technology to increase the shelf life of products was rejected by many respondents despite the clear sustainability benefit. Also here, the lack of definition and the resulting possibilities for using the term 'greenwashing' creates different responses among consumers (Durif *et al.*, 2010).

9.3 Conclusion

Past experiences do provide important information about the future foods that will enter the market. Food is never totally natural or unnatural. Consumer choices depend on their knowledge but also on socio-cultural values, and in the end different and often competing individual and collective values and goals compete for priority. A complicating factor in the analysis is that most consumer research is done on the perception of 'natural' and 'green' products, and not on actual buying behaviour (Bravo *et al.*, 2022; Barbu *et al.*, 2022).

Developers and designers of new products or products made with new processing technologies, growing or genetic technologies, start with a challenge. New technology in food is hardly ever a buying argument, in contrast to new technologies in cell phones, cars or medicines. This means that the extra benefits need to outweigh the natural reluctance of the average consumer. Sustainability may be an important argument, or animal welfare, or price.

Another important argument is that 'naturalness' is an important concept despite its unclear definition, but that other arguments can become dealbreakers. Who owns the technology, who benefits? Do we understand (the need for) the technology, and are suppliers transparent? Can we follow the ethics surrounding the technology? Food is special: we are comfortable with all kinds of biotechnology in medicine and allow doctors to inject them directly into our body; we are also comfortable with the biotechnology producing our beer and bread. But many people have a problem when the term DNA comes into play in discussions about our food and arguments that nature shuffles, breaks and moves DNA. The fact that we have been doing this for a hundred years or so is not convincing. It is not the product that counts but the way it has come about, as Otsuka (2021) found in genome

editing discussions. The term ‘natural’ can mean many things – it can both support and frustrate innovation. Marketeers have identified the value of the term – for scientists it is too ill-defined to be used. A general conclusion of this study was already formulated many years ago by Levy Strauss:

the selection of foods is made not only on physiological requirements, perceptual and cognitive mechanisms, but also on the basis of cultural and social representations, that result in additional constraints on what can and cannot be eaten, what is liked and what is disliked.’ As Strauss puts it: ‘things must not only be good to eat but also good to think. (Fischler, 1980)

References

- Akin, H., S.K. Yeo, C.D. Wirz, D.A. Scheufele, D. Brossard, M.A. Xenos, and Corley, E.A., 2019. Are attitudes toward labeling nano products linked to attitudes toward GMO? Exploring a potential ‘spillover’ effect for attitudes toward controversial technologies. *Journal of responsible innovation* 6:50–74. <https://doi.org/10.1080/23299460.2018.1495026>.
- Amos, C., Hansen, J., and King, S., 2019. All-natural versus organic: are the labels equivalent in consumers’ minds? *Journal of consumer marketing*, 36:516–526. <https://doi.org/10.1108/JCM-05-2018-2664>.
- Amos, C., Pentina, I., Hawkins, T., and Davis, N., 2014. “Natural” labeling and consumers’ sentimental pastoral notion. *Journal of product & brand management*, 23:268–281. <https://doi.org/10.1108/JPBM-03-2014-0516>.
- Aschemann-Witzel, J., Varela, P., and Peschel, A.O., 2019. Consumers’ categorization of food ingredients: do consumers perceive them as ‘clean label’ producers expect? An exploration with projective mapping. *Food quality and preference*, 71:117–128. <https://doi.org/10.1016/j.foodqual.2018.06.003>.
- Auerbach, C., Robson, J.M., 1944. Production of mutations by allyl-isothiocyanate. 154:81. <https://doi.org/10.1038/154081a0>.
- Barbu, A., Catană, Ș.-A., Deselnicu, D.C., Cioca, L.-I., Ioanid, A., 2022. Factors influencing consumer behavior toward green products: a systematic literature review. *International journal of environmental research and public health*. 19:16568. <https://doi.org/10.3390/ijerph192416568>.
- Battacchi, D., Verkerk, R., Pellegrini, N., Fogliano, V., and Steenbekkers, B., 2020. The state of the art of food ingredients’ naturalness evaluation: a review of proposed approaches and their relation with consumer trends. *Trends in food science & technology*, 106:434–444. <https://doi.org/10.1016/j.tifs.2020.10.013>.
- Binninger, A.-S., 2017. Perception of naturalness of food packaging and its role in consumer product evaluation. *Journal of food products marketing*, 23:251–266. <https://doi.org/10.1080/10454446.2014.885868>.

- Bravo, A., Vieira, D., Rebello, T.A., 2022. The origins, evolution, current state, and future of green products and consumer research: a bibliometric analysis. *Sustainability*, 14:11022. <https://doi.org/10.3390/su141711022>.
- Camerarius, R.J., 1694 *Ueber das geschlecht der pflanzen (De sexu plantarum epistola.)* 1694. Reprint 2023 by Gyan Books Pvt. Ltd., Delhi, India.
- Carlsen, P.S., Smith, H.H., Dearing, R.D., 1972. Parasexual interspecific plant hybridization. *Proceedings of the National Academy of Sciences of the USA*, 69:2292–2294. <https://doi.org/10.1073/pnas.69.8.2292>.
- Cavaliere A., Ventura, V., 2018. Mismatch between food sustainability and consumer acceptance toward innovation technologies among millennial students: the case of shelf life extension. *Journal of cleaner production* 175, 641–650, <https://doi.org/10.1016/j.jclepro.2017.12.087>.
- Chambers, E., Chambers, E. and Castro, M., 2018. What is “natural”? Consumer responses to selected ingredients. *Foods*, 7:65.
- Davies, S.R., Macnaghten, P., 2010. Narratives of mastery and resistance: lay ethics of nanotechnology. *Nanoethics*, 4:141–151. <https://doi.org/10.1007/s11569-010-0096-5>.
- Degreef, F., 2018. Understanding new food technologies and trust in food: framing analysis of food additives and food radiation (1960–1995). In: Namaste, N. and Nadales, M. Who decides? Competing narratives in constructing tastes, consumption and choice. Brill, Leiden, the Netherlands, pp. 223–246 https://doi.org/10.1163/9789004365247_014.
- De Maagd, R., 2015. *Bacillus thuringiensis*-based products for insect pest control. In: Lugtenberg, B. (Ed.) *Principles of plant-microbe interactions; microbes for sustainable agriculture*. Springer Cham Heidelberg New York Dordrecht London, pp. 185–192. <https://doi.org/10.1007/978-3-319-08575-3>.
- Diamond, J., 1997. *Guns, germs and steel. The fates of human societies*. WW Norton & Co, New York-London, USA-UK.
- Durif, F., Boivin, C., and Julien, C., 2010. In search of a green product definition. *Innovative marketing*, 6(1).
- Duvick, D.N., 2001. Biotechnology in the 1930s: the development of hybrid maize. *Nature reviews genetics*, 69–74. <https://www.nature.com/articles/35047587>.
- Eigsti, O.J., 1938. A cytological study of colchicine effects in the induction of polyploidy in plants. *Proceedings of the National Academy of Sciences of the USA*, 24:56–63. <https://doi.org/10.1073/pnas.24.2.56>.
- European Commission, Directorate-General for Research and Innovation, Executive summary of the opinion *Ethics of Genome Editing*, Publications Office, 2021, <https://data.europa.eu/doi/10.2777/949299>.
- European Union, 2018. Regulation (EU) 2018/848 of the European Parliament and of the Council of 30 May 2018 on organic production and labelling of organic products and repealing Council Regulation (EC) No 834/2007.
- Evans, G., de Challemaison, B., and Cox, D.N., 2010. Consumers' ratings of the natural and unnatural qualities of foods. *Appetite*, 54:557–563. <https://doi.org/10.1016/j.appet.2010.02.014>.

- Farris, A.L., 2010. The natural aversion: the FDA's reluctance to define a leading food-industry marketing claim, and the pressing need for a workable rule. *Food & drug law journal*, 65:403.
- Fischler, C., 1980. Food habits, social change and the nature/culture dilemma. *Social science information*, 19:937–953. <https://doi.org/10.1177/053901848001900603>.
- Fruscella, L., Kotzen, B., Milliken, S., 2021. Organic aquaponics in the European Union: towards sustainable farming practices in the framework of the new EU regulation. *Reviews in aquaculture*, 13:1661–1682; <https://doi.org/10.1111/raq.12539>.
- Greenpeace International, 2014. 'Smart breeding': biotechnology innovation that looks to the future but delivers today. <https://greenpeace.org/usa/news/smart-breeding-bio-technology-innovation/>.
- Harari, Y.N., 2014. *Sapiens. A brief history of humankind*. London, Vintage – Penguin Random House, New York, NY, USA, 498 p.
- Hemmerling, S., Canavari, M., and Spiller, A., 2016. Preference for naturalness of European organic consumers. *British food journal*, 118:2287–2307. <https://doi.org/10.1108/BFJ-11-2015-0457>.
- Herrera-Estrella, L., Depicker, A., Van Montagu, M. and Schell, J., 1983. Expression of chimaeric genes transferred into plant cells using a Ti-plasmid-derived vector. *Nature*, 303:209–213. <https://doi.org/10.1038/303209a0>.
- Hikal, W.M., Baeshen, R.S., and Said-Al Ahl, H.A., 2017. Botanical insecticide as simple extractives for pest control. *Cogent biology*, 3:1404274. <https://doi.org/10.1080/23312025.2017.1404274>.
- Hobden, K., and Pliner, P., 1995. Effects of a model on food neophobia in humans. *Appetite*, 25:101–113. <https://doi.org/10.1006/appe.1995.0046>.
- Holland, A., 2008: Nature – every last drop of it is good. In: Robin Attfield, *The ethics of the environment*. Routledge, London, UK. Chapter 6, 17 pp. <https://doi.org/10.4324/9781315239897>.
- IFOAM, 2017. Position paper. Compatibility of breeding techniques in organic systems. Bonn, IFOAM – Organics International. 22 pp. https://www.ifoam.bio/sites/default/files/2020-03/Breeding_position_paper_v01_web_0.pdf.
- Jacobsen, E., Schouten, H.J., 2007. Cisgenesis strongly improves introgression breeding and induced translocation breeding of plants. *Trends in biotechnology*, 25:219–223. <https://doi.org/10.1016/j.tibtech.2007.03.008>.
- Jinek, M., Chylinski, K., Fonfara, I., Hajer, M., Doudna, J.A., Charpentier, E.A., 2012. Programmable dual-RNA-guided DNA endonuclease in adaptive bacterial immunity. *Science*, 337, 816–821. <https://doi.org/10.1126/science.1225829>.
- Jochemsen, H., 2000. *Toetsen en begrenzen (Testing and limiting: an ethical and political assessment of modern biotechnology)*; Scientific Institute ChristenUnie: Amersfoort, The Netherlands.
- Jochemsen, H. An ethical assessment of cisgenesis in breeding late blight resistant potato. *Potato Res.* 2008, 51, 59–73. <https://doi.org/10.1007/s11540-008-9090-5>.
- Kasha, K.J., 1974. Haploids in higher plants: advances and potential. In *Proceedings of the First International Symposium*, Guelph, ON, Canada, 10–14.

- Kharwal, M.C., 2012. A brief history of plant mutagenesis. In: Shu, Q.Y., Forter, B.P., Nakagawa, H., (Eds) *Plant mutation breeding and biotechnology*, CABI: Wallingford, UK; FAO: Rome, Italy, 2012; pp. 21–46.
- Li, M., and Chapman, G.B., 2012. Why do people like natural? Instrumental and ideational bases for the naturalness preference. *Journal of applied social psychology*, 42:2859–2878. <https://doi.org/10.1111/j.1559-1816.2012.00964.x>.
- Louwaars, N., Burgaud, F., 2016. Variety registration: the evolution of registration systems with a special emphasis on agrobiodiversity conservation. In: Halewood, M., (Editor), *Farmers' crop varieties and farmers' rights – challenges in taxonomy and law*, Earthscan: London, UK, Chapter 6.
- Louwaars, N.P. and Jochemsen, H., 2021. An ethical and societal analysis for biotechnological methods in plant breeding. *Agronomy* 11:1183. DOI: 10.3390/agronomy11061183.
- Lusk, J.L., Roosen, J., and Bieberstein, A., 2014. Consumer acceptance of new food technologies: causes and roots of controversies. *Annual review of resource economics*, 6:381–405. <https://doi.org/10.1146/annurev-resource-100913-012735>.
- Lusk, J.L., 2019. Consumer perceptions of 'natural' foods. *Food technology*, 73:42–47.
- Macnaghten, P., 2004. Animals in their nature: a case study on public attitudes to animals, genetic modification and 'nature.' *Sociology*, 38:533–551. <https://doi.org/10.1177/0038038504043217>.
- Marley, T. and E. Chambers, 2019. The influence of colorants, flavorants and product identity on perceptions of naturalness. *Foods*, 8:317. <https://doi.org/10.3390/foods8080317>.
- Marris, C., 2001. Public views on GMO's: deconstructing the myths. *EMBO reports* 2:545–548. <https://doi.org/10.1093/embo-reports/kve142>.
- Mendel G., 1865. Versuche über pflanzenhybriden. Presented at the meetings of 8. February und 8. March 1865, printed in the Proceedings of the association for Nature Research in Brno. Press of the Association, Vol 4 S.3–47 1865.
- Moscato, E.M., and Machin, J.E., 2018. Mother natural: Motivations and associations for consuming natural foods. *Appetite*, 121:18–28. <https://doi.org/10.1016/j.appet.2017.10.031>.
- Murashige, T. and Skoog, F., 1962. A revised medium for rapid growth and bio assays with tobacco tissue cultures. *Physiol. Plant.* 15:473–497. <https://doi.org/10.1111/j.1399-3054.1962.tb08052.x>.
- Nakaya, A. and Isobe, S.N., 2012. Will genomic selection be a practical method for plant breeding? *Ann. Bot.* 110:1303–1316. <https://doi.org/10.1093/aob/mcs109>.
- Nales, P., and Fischer, A.R.H., 2023. Breeding by intervening: exploring the role of associations and deliberation in consumer acceptance of different breeding techniques. *Public understanding of science*, 0(0). <https://doi.org/10.1177/09636625231168087>.
- Nitsch, J.P., 1951. Growth and development in vitro of excised ovaries. *Am. J. Bot.* 38:556–576. <https://doi.org/10.1002/j.1537-2197.1951.tb14860.x>.

- Østerberg, J.T., Xiang, W., Olsen, L.I., Edenbrandt, A.K., Vedel, S.E., Christiansen, A., Landes, X., Andersen, M.M., Pagh, P., Sandøe, P., Nielsen, J., Christensen, S.B., Thorsen, B.J., Kappel, K., Gamborg, C. and Palmgren, M., 2017. Accelerating the domestication of new crops: feasibility and approaches. *Trends in plant science* 22:373–384. <https://doi.org/10.1016/j.tplants.2017.01.004>.
- Otsuka, Y., 2021. Consumer movements confronted by naturalness in gene editing in Japan. *East Asian science, technology and society: an international journal* 15:24–45. <https://doi.org/10.1080/18752160.2021.1877442>.
- Paterson, A.H., Lander, E.S., Hewitt, J.D., Peterson, S., Lincoln, S.E., and Tanksley, S.D., 1988. Resolution of quantitative traits into complete Linkage map of restriction fragment length polymorphisms. *Nature*, 335:721–726. <https://doi.org/10.1038/335721a0>.
- Pliner, P. and Hobden, K., 1992. Development of a scale to measure the trait of food neophobia in humans, *Appetite* 19:105–120. [https://doi.org/10.1016/0195-6663\(92\)90014-W](https://doi.org/10.1016/0195-6663(92)90014-W).
- Pliner, P., Pelchat, M. and Grabski, M., 1993. Reduction of neophobia in humans by exposure to novel foods. *Appetite*, 20:111.
- Rabadán, A. and Barnabéu, R., 2021. A systematic review of studies using the Food Neophobia Scale: conclusions from thirty years of studies, *Food quality and preference*, 93:104241. <https://doi.org/10.1016/j.foodqual.2021.104241>.
- Roberts, P.B., 2014. Food irradiation is safe: half a century of studies. *Radiation physics and chemistry* 105:78–82. <https://doi.org/10.1016/j.radphyschem.2014.05.016>.
- Román, S., Sánchez-Siles, L.M., and Siegrist, M., 2017. The importance of food naturalness for consumers: results of a systematic review. *Trends in food science & technology*, 67:44–57. <https://doi.org/10.1016/j.tifs.2017.06.010>.
- Rozin, P., Spranca, M., Krieger, Z., Neuhaus, R., Surillo, D., Swerdlin, A., and Wood, K., 2004. Preference for natural: instrumental and ideational/moral motivations, and the contrast between foods and medicines. *Appetite*, 43:147–154. <https://doi.org/10.1016/j.appet.2004.03.005>.
- Rozin, P., Fischler, C., and Shields-Argelès, C., 2012. European and American perspectives on the meaning of natural. *Appetite*, 59:448–455. <https://doi.org/10.1016/j.appet.2012.06.001>.
- Rynn, B., 2007. Raw foodism. In Andrew F. Smith. *The Oxford companion to American food and drink*. Oxford University Press, Oxford, UK. pp. 493–494.
- Sanchez-Siles, L.M., Michel, F., Román, S., Bernal, M.J., Philipsen, B., Haro, J.F., Bodenstab, S., and Siegrist, M. (2019). The Food Naturalness Index (FNI): an integrative tool to measure the degree of food naturalness. *Trends in food science & technology*, 91:681–690. <https://doi.org/10.1016/j.tifs.2019.07.015>.
- Schneider, S., 2008. Good, clean, fair: the rhetoric of the slow food movement. *College English*, 70:384–402. <http://www.jstor.org/stable/25472277>.
- Shukla, V.K., Doyon, Y., Miller, J.C., DeKolver, R.C., Moehle, E.A., Worden, S.E., Mitchell, J.C., Arnold, N.L., Gopalan, S., Meng, X., *et al.*, 2009. Precise genome modification in the crop species *Zea mays* using zinc-finger nucleases. *Nature*, 459:437–441. <https://doi.org/10.1038/nature07992>.

- Siegrist, M., and Hartmann, C., 2020. Perceived naturalness, disgust, trust and food neophobia as predictors of cultured meat acceptance in ten countries. *Appetite*, 155:104814. <https://doi.org/10.1016/j.appet.2020.104814>.
- Siegrist, M., and Sütterlin, B., 2017. Importance of perceived naturalness for acceptance of food additives and cultured meat. *Appetite*, 113:320–326. <https://doi.org/10.1016/j.appet.2017.03.019>.
- Stadler, L.J., 1928. Genetic effects of X-rays in barley. *Proceedings of the National Academy of Sciences of the United States of America*, 88:186–187.
- Tamm, L., Thuerig, B., Apostolov, S., Blogg, H., Borgo, E., Corneo, P.E., Fittje, S., de Palma, M., Donko, A., Experton, C., Alcázar Marín, E., Morell Pérez, A., Pertot, I., Rasmussen, A., Steinshamn, H., Vetemaa, A., Willer, H., Herforth-Rahmé, J., 2022. Use of copper-based fungicides in organic agriculture in twelve European countries. *Agronomy*, 12:673. <https://doi.org/10.3390/agronomy12030673>.
- Van der Oost, J., Jore, M.M., Westra, E.R., Lundgren, M. and Brouns, S.J.J., 2009. CRISPR-based adaptive and heritable immunity in prokaryotes. *Trends in biochemical sciences* 34:401–407. <https://doi.org/10.1016/j.tibs.2009.05.002>.
- Watson, J., Crick, F., 1953. Molecular structure of nucleic acids: a structure for deoxyribose nucleic acid. *Nature*, 171:737–738. <https://doi.org/10.1038/171737a0>.
- Wickell, D.A. and Li, F.-W., 2019. On the evolutionary significance of horizontal gene transfers in plants. *New Phytol.* 225:113–117. <https://doi.org/10.1111/nph.16022>.
- Wilbois, K.P., Baker, B., Raaijmakers, M. and Lammerts van Bueren, E.T., 2012. Values and principles in organic farming and consequences for breeding approaches and techniques. In: Lammerts van Bueren, E.T. and Myers, J.R. (editors). *Organic crop breeding*. Wiley, Hoboken, NJ, USA, pp. 125–138 <https://doi.org/10.1002/9781119945932>.

How to apply green AI in practice?

Moving from FLOPs to CO₂ footprint

Qin Zhao and Gerard Schouten

Abstract

Artificial intelligence (AI) is revolutionising industry and transforming society. However, the rapid growth of AI has also raised concerns about the environmental impact of this technology, as the energy consumption required to train and use AI models is increasing exponentially and can be significant. To address this issue, a new research field called *Green AI* has emerged, *which focuses on developing and deploying AI models that are both effective, efficient, and environmentally sustainable*. In this chapter, we specifically explore how to apply Green AI in practice, i.e. translate it into manageable principles that can be applied by professionals, researchers and students while building AI models. We propose an ontology of five domain-independent Green AI best practices: (i) monitor and report energy consumption as well as CO₂ footprint, (ii) select your model architecture with care, (iii) adapt a data-centric AI approach, (iv) apply a smart hyperparameter tuning strategy, (v) retrain your model with care. We have applied the above practices in two case studies: (i) a societal case focusing on measuring biodiversity, in particular wildflower abundance and richness ‘in the wild’ with an AI-based object detection approach, and (ii) a circular economy scenario focusing on automatic e-waste classification and outlier detection. This first exploration creates confidence that it is worthwhile having a ‘how-to’ guide for AI practitioners to take Green AI into account while building applications.

10.1 Introduction

The tremendous progress and fast development of AI in today’s information age enables human beings to benefit from the vast amount of data and realize breakthroughs for the economy, society, and the environment. Domains in which AI have already had a considerable impact include healthcare, finance, energy supply, and transportation. According to Joppa and Herweijer (2020) the usage of AI for environmental applications can contribute up to 4.4% of global GDP and save up to 4.0% of all greenhouse emissions worldwide by 2030.

However, as AI systems are becoming significantly more complex, their carbon footprint can no longer be ignored. Contemporary models typically contain

carbon
footprint

millions or even billions of parameters and run on clusters of GPUs, sometimes for thousands of hours. Strubell *et al.* (2020) have studied the carbon emissions of state-of-the-art Natural Language Processing (NLP) models and has found that training large NLP models emits an amount of carbon dioxide equivalent to a trans-American flight. Moreover, as a spoiler-alert, the impact of ChatGPT and similar Large Language Models (LLMs) on our climate has been discussed in many recent blogposts.¹ The concept of ‘Green AI’ by Schwartz *et al.* (2020) has become a top priority on the agenda of the AI research community. The goal of Green AI research is to provide concepts and tools to monitor, report, analyse and eventually reduce the carbon footprint of the whole AI lifecycle.

AI lifecycle
models

AI lifecycle models have many similar variants. An early example, known as named Cross Industry Standard Process for Data Mining (CRISP-DM)² is presented by Wirth and Hipp (2000). In such an AI lifecycle, depicted in Figure 10.1, we usually start with a problem definition phase together with business or associated partners. Once the objectives and requirements are defined, data acquisition will take place, including discovering available data sets, improving data quality, and deriving initial insights from the data and perspectives. Model development, training, evaluation, and refinement will be performed in the next phase. While data discovery is usually the most time-consuming process, iterative model development and refinement is the most resource exhaustive phase per actor (typically a data scientist or AI engineer). Often different model architectures need to be trained, compared, and fine-tuned to achieve the best performance. Next, the inference and deployment phase brings the optimized model into production and exposes it as an AI service to the outside world. In this phase the scale of usage is the main determinant of energy consumption. Finally, the stage of performance monitoring measures the accuracy of the trained model over time while being used by many end users or customers in real-world situations. Data drift might for instance lead to less effective models that need to be retrained to maintain a high-performant service. In today’s cloud platforms, tooling solutions, a.k.a. Machine Learning Operations (MLOps), are available to automate the AI life cycle, from streamlining data collection to carefully designing retrain intervals (Heck *et al.*, 2021).

Training and inference of AI are computationally intensive, albeit in different ways. The computational effort of AI models is usually expressed in the hardware-independent metric Floating Point Operations (FLOPs) and typically refers to one forward pass in a neural network (see Annex). Training effort is on a scale with the number of AI experiments (hyperparameter tuning strategy) that are carried out to ‘arrive at’ an accurate model. Inference effort is on a scale with the number of requests or ‘model calls’ by the community of users. Along with the cost

1 <https://www.tudelft.nl/en/stories/articles/sustainable-artificial-intelligence-from-chatgpt-to-green-ai>.

2 Data mining is a process to extract useful data from a larger set of raw data. AI takes this one step further and uses the data to solve cognitive problems associated with human intelligence.

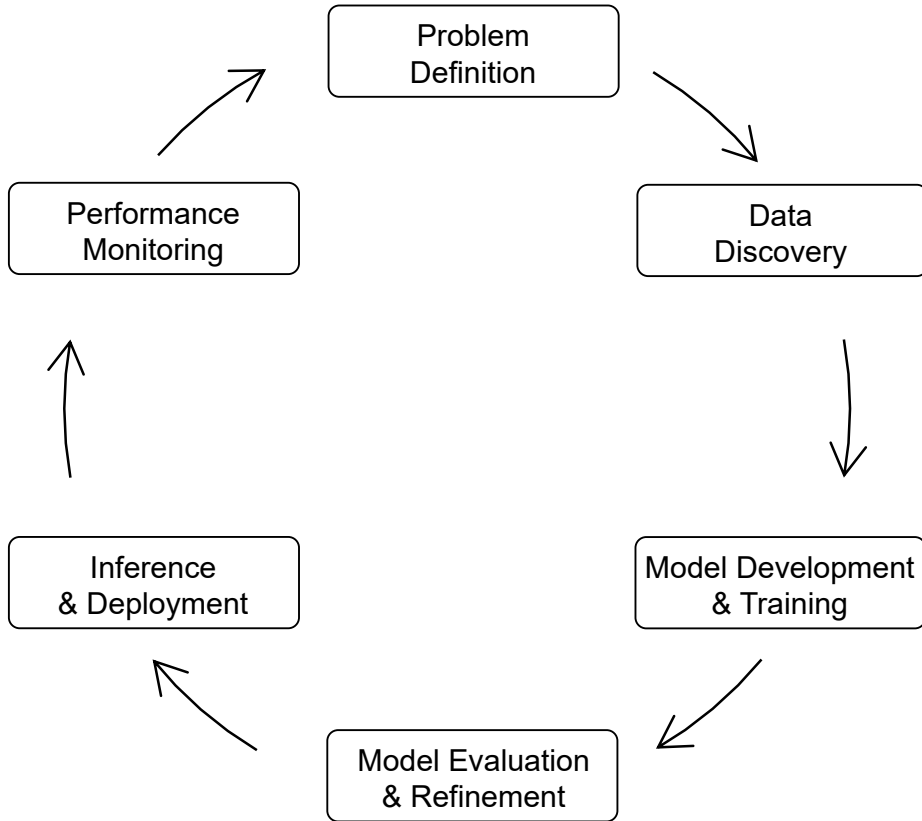


FIGURE 10.1 Typical AI lifecycle model

of hardware and consumed electricity, the carbon emission already expands from tons to several hundred tons (and even more if we consider the fast evolution of the latest LLM models). Developing and utilizing AI technologies in a sustainable way has become a major concern in a world where action is much needed to mitigate climate change. Therefore, carbon emissions of AI throughout its entire lifecycle need to be monitored, reported, and communicated transparently, and methods should be developed to benchmark and assess the AI lifecycle’s environmental impact.

In section 2 we give an overview of the background and related work and briefly sketch the existing energy-related metrics used to assess the cost of the models. Section 3 presents a recipe for how to move from FLOPs to carbon emissions of models. Section 4 discusses an ontology of Green AI best practices and provides guidelines on how to apply them. Section 5 presents two study cases and shows the achievements of applying Green AI best practices to these projects. Finally, section 6 concludes that AI practitioners should be aware of the carbon emissions associated with AI design and experimentation activities and use Green AI best practices to create sustainable solutions.

10.2 Background and related work

Metrics for Green AI aim at different aspects. An overview is presented in Table 10.1. FLOPs are used as a hardware-independent measure, with the aim to benchmark foundation models. Energy efficiency relates directly to the energy (in Joules) consumed by a model while running on specific hardware. CO₂ emission has a direct link with the impact on our environment. Note that CO₂ emission takes into account whether traditional fossil-based fuels or alternative green energy sources (e.g. solar, or wind energy) are exploited, and hence depends on the place on Earth where the model is running, or trained for that matter.

The most influential literature on Green AI for the above metrics is briefly discussed below.

TABLE 10.1 Metrics of Green AI

FLOPs	Floating point operations, a number that expresses the complexity (in terms of size or scale) of an algorithmic process
FLOP/s	Floating point operations per second, a unit that relates to the performance (or computational power) of a computer
Energy efficiency	A measurement of how much energy in Joules is consumed per unit (such as FLOP/s)
CO ₂ emission	A measurement of total greenhouse gas emitted, expressed in terms of the equivalent measurement of carbon dioxide

10.2.1 Evolution of models in terms of FLOP/s

With the fast development of AI in recent years models are getting more complex; consequently, the energy consumption of these models increases exponentially, causing a notable environmental impact. OpenAI did extensive research on the growth trend for computing power in terms of FLOP/s needed for different AI models developed in recent years. According to OpenAI (2018) AI models have doubled the FLOP/s used every 3.4 months since 2012. This trend is depicted in Figure 10.2 on a log scale of petaflops/s-day³ and has been coined ‘Moore’s Law of the AI era’.

Moore’s Law
of the AI era

This trend won’t stop with today’s fierce competition and rapid progress in LLMs since ChatGPT (and other generative AI tools) launched so successfully in November 2022. The need to manage the risks of generative AI in a broader sense has been highlighted by researchers, scientists, and industry leaders (see e.g. Bengio, *et al.* 2023).

³ FLOP/s of performing 10¹⁵ neural net operations per second-day.




FIGURE 10.2 The computational effort of AI models increases according to Moore's Law of the AI era

10.2.2 Energy efficiency of models

A detailed analysis of the training cost and algorithmic efficiency of AI models including the latest development of vision models was carried out by Hernandez and Brown (2020). They focused on the amount of computation used to train models with the huge ImageNet dataset. Although both algorithm efficiency and hardware efficiency have been improved a lot recently, the growing demand for better performance is still driving the AI research community to build larger and more complex models. Desislavov *et al.* (2021) analysed energy efficiency with a greater number of deep neural networks with various hardware components in a longer time frame. Thompson *et al.* (2022) reported the computational demands of several deep learning applications, showing that progress in them is strongly reliant on increases in computing power.

Compared to training cost, there are a few studies reflecting on the inference cost. Canziani *et al.* (2016) compared accuracy, memory footprint, parameters, operations count, inference time and power consumption of 14 models trained on ImageNet. A similar study by Li *et al.* (2016) measured energy efficiency, Joules per image, for a single forward and backward propagation iteration. This study benchmarked four Convolutional Neural Networks (CNNs) on different hardware architectures such as different CPU and GPU configurations. Both publications analyse model efficiency, but they do this for very concrete cases.

10.2.3 Carbon emission of models

Some recent work has started the discussion on carbon emissions and the sustainable usage of AI models. Strubell *et al.* (2020) demonstrated the issue of carbon and energy impacts of training large NLP models by evaluating estimated power usage and carbon emissions for a set of case studies. Their results led to the conclusion that we need to reduce the carbon footprint of developing and running AI models. Schwartz *et al.* (2020) defined the term ‘Green AI’ as ‘AI research that yields novel results while taking into account the computational cost’. Henderson *et al.* (2020) proposed a framework for tracking real-time energy consumption and carbon emissions and create a leaderboard to incentivize energy-efficient research. Platforms such as Huggingface provide tools utilizing CodeCarbon made by Budennyaya *et al.* (2022) to calculate CO₂ emissions when performing training or pre-training.

CodeCarbon

10.3 Conversion sequence: from FLOPs to CO₂ footprint

Green AI

The term ‘Green AI’ refers to research that promotes measurement of energy efficiency for algorithms or models as a widely accepted evaluation metric alongside model accuracy. Scientific literature shows that a gamut of metrics related to energy efficiency is available. In this section we list the most used metrics and advocate that we should move from FLOPs to CO₂ footprint.

- i. FLOPs: The number of floating-point operations provides a direct estimation of the amount of work by the computational process (OpenAI, 2018). It is agnostic to the hardware on which the model is run. Most authors typically report the computation required to go through one forward pass of a neural network.
- ii. Number of parameters: This measure is closely correlated with the amount of memory consumed by the model. As a result, different models with a similar number of parameters often perform different amounts of work (Canziani, 2016).
- iii. Elapsed time: Authors report the time to train the whole neural network, or the time needed for inference, i.e. the time it takes to execute one forward pass. This measure is highly influenced by factors such as the underlying hardware, other jobs running on the same machine, and the number of cores used (Jeon and Kim, 2018).
- iv. Energy: The energy consumption of training or inferencing of an AI model can be obtained either by a calculation based on FLOPs, number of parameters, the elapsed real time, or by measurement with power meters on CPUs and GPUs (Desislavov *et al.*, 2021).
- v. Carbon emission: To calculate carbon emission, the carbon intensity, i.e. how many grams of carbon dioxide (CO₂) are released to produce a kilowatt hour

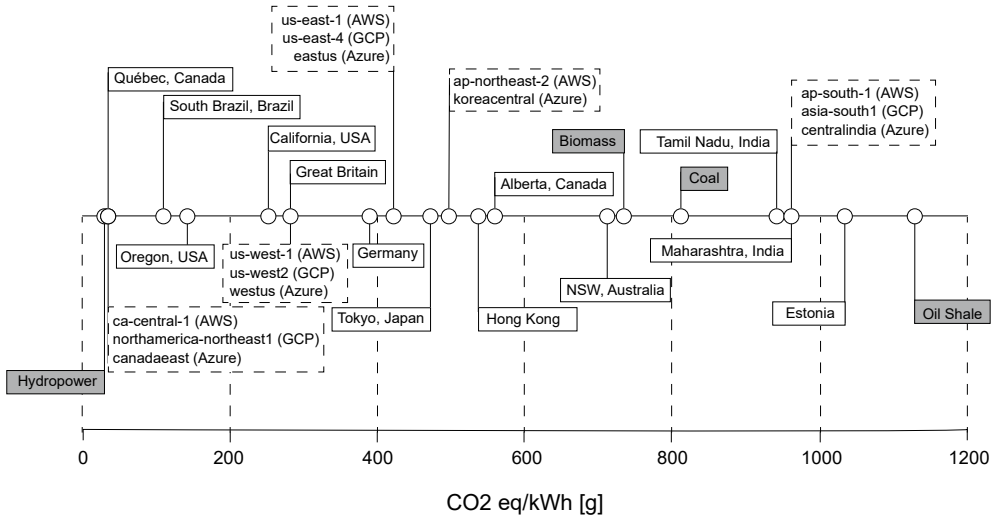


FIGURE 10.3 Carbon intensity for an assortment of locations

(kWh) of electricity, is collected from the local grid, and used to multiply power estimation of training or inference process. It is essential to know that the local energy grid makes a huge difference when calculating carbon emissions. Figure 10.3 shows the CO₂ intensity (in 2014/2015) for an assortment of cloud-provider regions and energy production methods (Henderson *et al.*, 2020). It is clear that running an AI job in Quebec is much cleaner than the other regions.

10.4 Green AI best practices

10.4.1 Monitor and report CO₂

As AI practitioners we have a responsibility to develop sustainable AI. The first responsibility is to monitor and report the energy consumption as well as the carbon footprint of the AI model during training and inference phases. Several tools or packages have been developed for this purpose, such as CodeCarbon (Budennyya at al., 2022), CarbonTracker (Antony *et al.*, 2020) and ExperimentImpactTracker (Henderson *et al.*, 2020). Those packages have utilized a publicly available framework and can easily be integrated with the development code for tracking both energy consumption and carbon emission.

CodeCarbon
CarbonTracker
Experiment
ImpactTracker

10.4.2 Select your model wisely

It has often been a belief in AI community that to achieve greater accuracy larger and deeper models are needed with huge amounts of data to train them. However,




FIGURE 10.4 Benchmark of CNN models for image classification

larger models come with a higher financial and environmental cost. What's more, it is not always true that larger and deeper models achieve greater accuracy. Some researchers have done extensive benchmarking for the AI tasks, such as imaging classification, with various deep learning neural network models. Image classification is a fundamental task in vision recognition that aims to understand and categorize an image as a whole under a specific label. As shown in Figure 10.4 by the study of Canziani (2016), fourteen image classification models have been compared in a computation vs accuracy graph. In this figure, the x-axis represents the total number of operations in Giga FLOPs. The y-axis is the top-1 accuracy, i.e. the highest probability of correct prediction of the labelled images. Bubble size refers to the number of parameters of each model.

Figure 10.4 shows some remarkable and useful patterns, such as:

- ResNet-101 has doubled the number of FLOPs compared to ResNet-50. Note that this leads only to a minor improvement in accuracy (2%).
- For a top-1 accuracy of approximate 70% you can use several models. Note that ResNet-34 (with even a slightly higher score) has one-sixth of the FLOPs compared to VGG-16, thereby delivering more value with less energy consumption and a lower CO2 footprint.

In summary, it is always good practice to check the model benchmark and pick up the one that is most sufficient in terms of accuracy with cost efficiency for the application. In today's education on AI, the focus is still very much on accuracy and not yet on energy efficiency and carbon footprint. Our message reflects that there should be a careful per-case consideration of the trade-off between accuracy and efficiency/emission.

10.4.3 Adopt a data-centric AI approach

The quality of data has a huge impact on how well the whole system works, and how we fuel AI models is crucial to their success. Data-centric AI has been an emerging discipline that systematically deals with data quality to build AI systems. In Figure 10.5, we show as an educational example trend lines for noisy data and clean data. Note that, in order to reach the same level of performance, the amount of high-quality data is only a small fraction of the amount of low-quality noisy data. This practice is also strongly supported by AI pioneers, see for instance




FIGURE 10.5 Clean data improves prediction accuracy

‘Andrew Ng, AI minimalist: the machine-learning pioneer says small is the new big’ (Strickland, 2022).

It is a critical and essential process to carefully check the dataset, remove data noise and balance the data as much as possible. We have also practised data-centric AI in our projects, and have obtained significant accuracy improvement and timing reduction, and consequently obtained significant energy and emission savings.

10.4.4 *Tune your hyperparameters in a smart way*

Model training is a process through which a model learns its parameters (often billions). Besides this, every model also has hyperparameters (usually a few) that it cannot *learn*, but can be *tuned* for. The process of tuning hyperparameter values is called hyperparameter tuning. It is a vital aspect of increasing model performance.

There are several hyperparameter tuning strategies, such as grid search, random search and Bayesian search or optimization.

- Grid Search – A grid or set of hyperparameters are defined and every possible combination is used for training a model. This is exhaustive and computationally expensive and is used when the hyperparameter search space is restricted.
- Random Search – Instead of a grid, statistical distribution of each hyperparameter is provided and the number of iterations can be controlled. This is a suitable strategy for larger search spaces.
- Bayesian Optimization – A sequential model-based optimization that uses the results from previous iterations to decide the next hyperparameter value candidates.

Bayesian optimization methods are more efficient because they select hyperparameters in an informed manner. By prioritizing hyperparameters that appear more promising from past results, Bayesian methods can find the best hyperparameters in less time (in fewer iterations!) than both grid search and random search. Therefore, it is preferable to use a Bayesian optimization strategy for hyperparameter tuning.

10.4.5 *Retrain with care*

Once the model is trained and put into production for a while, the model might need to be retrained with the following observations.

- The model’s performance metrics have deteriorated.
- The distribution of the prediction is different from those observed during training.
- The training data and the live data diverge, that is, the training data is no longer a good representation of the real world.

However, retrain is a very costly process, and has a financial, operational and environmental impact. Retrain intervals are closely associated with business cases and should be carefully analysed and determined to reduce the total cost including the emission and the environmental impact.

10.5 Case studies

We embrace 'AI for good' cases in various applied research projects. The first example we present is about biodiversity loss in the Netherlands. This is a major environmental issue, caused by e.g. habitat loss, intensive farming or pollution. Monitoring and classifying wild flowering plants with AI will help us better understand the changes in the biodiversity and react to it. The second example is about building a more circular economy with the use of AI. In recent decades tons of E-waste (electronic waste) have been produced worldwide. The de-assembly and recycling process is widely enforced but is still largely a manual process and could be greatly improved with AI technology.

10.5.1 *Automatic wildflower monitoring*

To better understand biodiversity, it is necessary to initiate and automate large-scale monitoring programs supported by AI technology. Monitoring means *identifying* and *counting* objects of interest. In this case study we applied object detection algorithms for monitoring flowering plants 'in the wild' (Heck and Schouten, 2023; Schouten *et al.*, 2024). Wildflowers are an essential component of biodiversity. They provide many eco-system services, such as medicine, building materials and food; they keep our soil healthy, purify water and mitigate climate change to a large extent by absorbing greenhouse gases and significantly lowering temperatures in cities. To this end, a unique expert-annotated reference dataset with over 2000 high resolution images, each covering approximately 1m² of soil, has been collected around the city of Eindhoven. This Eindhoven Wildflower Dataset (EWD) holds 160 flowering plant species and contains images of roadsides, rich-weed grasslands, marshland, and urban green areas (Schouten *et al.*, 2024). As with many biological datasets that are collected 'in the wild', EWD has a long-tailed distribution. Common species are overrepresented and rare or inconspicuous species are underrepresented.

We selected a state-of-the-art R-CNN object detection algorithm (Ren *et al.*, 2015) with a Resnet50 backbone algorithm from the PyTorch library and trained it with EWD images. During the AI experimentation phase, we found out that the mean average precision (mAP) varied considerably over the species and is affected to a large degree by the skewed distribution. By carefully pre-processing the image data and creating a balanced subset from the original long-tailed EWD dataset, the training time was reduced from 10 hours to 35 minutes, thereby greatly reducing the energy consumption (and CO₂ emission) of the training process, while the accuracy (mean average precision) even increased from 0.68 to 0.82.

This case demonstrates the importance of a data-centric AI approach. Making sure to train models with high-quality data is a major factor in achieving Green AI.




FIGURE 10.6 Clustering the E-waste dataset into device groups

10.5.2 *Sorting E-waste for disassembly*

This circular economy project focuses on using AI, and in particular a cluster algorithm, to identify E-waste devices with removable batteries.

Batteries, especially lithium batteries, need to be carefully recycled because they can catch fire with the slightest damage. While new-generation phones and tablets are designed with non-removable batteries to make the product tighter, slimmer and waterproof, a small proportion of old phones and gadgets still contain removable batteries. Currently, during the recycling process all E-waste devices are mixed, and then manually sorted on the conveyor belt. To distinguish devices that contain removable batteries from those that don't is a time-consuming and highly unsafe job for local workers. AI technology has the potential to improve this process. To achieve this, a dataset is collected, containing E-waste devices spanning from smart-watches to phones and tablets, with their physical dimensions as well as images.

The brute-force approach is to train an advanced image classification algorithm, such as a CNN (Convolutional Neural Network), and sort the devices automatically. A smarter approach is to use non-visual clues, in particular the dimensions (width and height) of the devices. By checking the data carefully, it was found that the 'outliers' correspond to rarely used old phones with removable batteries. An additional insight is that even if the device type is not known beforehand, it is common knowledge that smart watches, phones and tablets are intrinsically of different sizes. By

applying the DBSCAN clustering algorithm to this dataset, the data can be organized into separate groups corresponding to device types, as shown in Figure 10.6. This data pre-processing and cluster approach provides strong support for automatic decision-making for E-waste sorting.

This case study clearly illustrates that a lot can be gained by picking the right AI approach. An obvious and straightforward solution for sorting or inspecting products is to apply computationally intensive vision AI. Here we demonstrate that a lightweight cluster model also works well. Hence, rethinking the obvious solution can make an excellent contribution to Green AI.

10.6 Conclusions

The digital transformation of society means that more and more data are being collected and fed into AI models. Model architectures are becoming larger and more complex, and at the same time more iterations for optimization are being executed to achieve the best performance. As a consequence, the calculations executed on computers incur a heavy energy and financial cost, and the CO₂ emissions of AI models are no longer negligible. As AI practitioners, educators, and researchers, we need to be aware of the CO₂ footprint of our AI design and experimentation activities. Specifically, we should inform ourselves and others about both the positive effects and negative environmental consequences of using AI. Therefore, we should be able to measure and explicitly state the emissions of models, and finally we should take steps to reduce the carbon emissions during the development and deployment of AI solutions.

Annex: Explaining training and inference with an artificial neural network

To train a model, researchers usually pre-process the data, define a model architecture, optimizing strategy and loss function, and then take an iterative process to loop over the model network with the number of epochs. A neural network is made of multiple neurons and these neurons are stacked into layers. The connections between the layers occurred through the parameters of the network. Calculations are carried out in forward propagation and backward propagation through each layer in the network until the loss function is minimized.

Unlike training, inference doesn't re-evaluate the layers and parameters of the neural network. Inference applies the trained neural network model or a forward pass to a new unknown dataset, and outputs the prediction based on the accuracy of the model. Minimizing latency issues during the inference process can pose a challenge for getting the system to make decisions in real time.

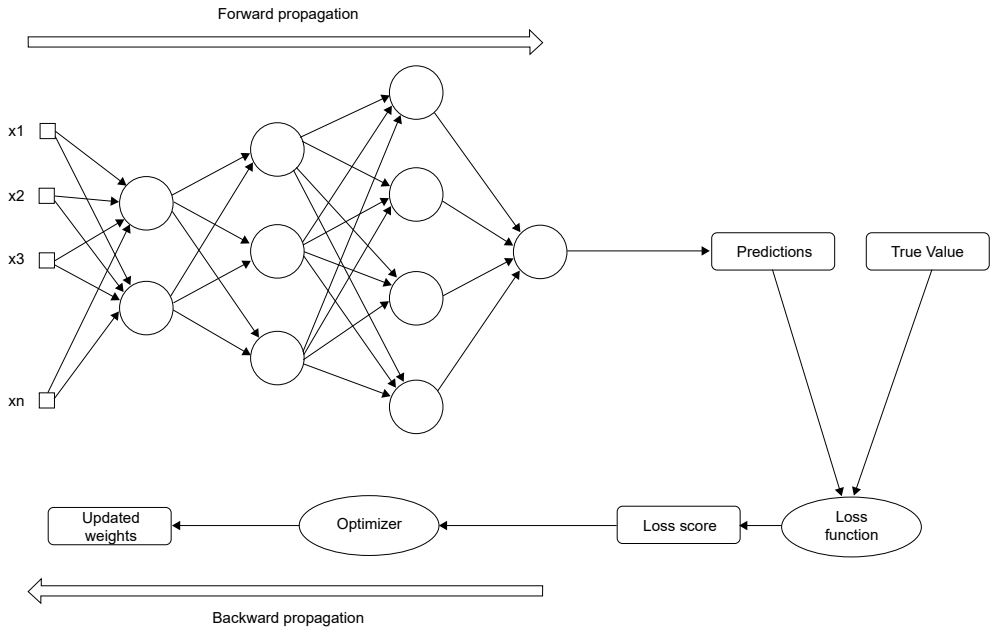


FIGURE 10.7 Training process of a neural network

References

- Anthony, L.F.W., Kanding, B. and Selvan, R., 2020. Carbontracker: tracking and predicting the carbon footprint of training deep learning models. arXiv preprint arXiv:2007.03051. <https://github.com/lfwa/carbontracker>.
- Bengio, Y., Hinton, G., Yao, A., Song, D., Abbeel, P., Harari, Y.N., ... and Mindermann, S., 2023. Managing AI risks in an era of rapid progress. arXiv preprint arXiv:2310.17688.
- Budenny, S.A., Lazarev, V.D., Zakharenko, N.N., Korovin, A.N., Plosskaya, O.A., Dimitrov, D.V.E., ... and Zhukov, L.E.E., 2022. Eco2AI: carbon emissions tracking of machine learning models as the first step towards sustainable AI. In: *Doklady Mathematics* (Vol. 106, No. Suppl 1, pp. S118–S128). Pleiades Publishing, Moscow, Russia. <https://github.com/mlco2/codecarbon>.
- Canziani, A., Paszke, A. and Culurciello, E., 2016. An analysis of deep neural network models for practical applications. arXiv preprint arXiv:1605.07678.
- Desislavov, R., Martínez-Plumed, F. and Hernández-Orallo, J., 2021. Compute and energy consumption trends in deep learning inference. arXiv preprint arXiv:2109.05472.
- Heck, P., Schouten, G. and Cruz, L., 2021. A software engineering perspective on building production-ready machine learning systems. In: *Handbook of research on applied data science and artificial intelligence in business and industry*. IGI Global, Hershey, PA, USA, pp. 23–54.

- Heck, P. and Schouten, G., 2023. Defining quality requirements for a trustworthy AI wildflower monitoring platform. In: IEEE/ACM 2nd International Conference on AI Engineering–Software Engineering for AI (CAIN), pp. 119–126.
- Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D. and Pineau, J., 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of machine learning research*, 21:1–43.
- Hernandez, D. and Brown, T.B., 2020. Measuring the algorithmic efficiency of neural networks. arXiv preprint arXiv:2005.04305.
- Jeon, Y. and Kim, J., 2018. Constructing fast network through deconstruction of convolution. *Advances in neural information processing systems*, 31.
- Joppa, L. and Herweijer, C., 2020. How AI can enable a sustainable future. Microsoft Corporation. <https://news.microsoft.com/wp-content/uploads/prod/sites/53/2019/04/PwC-Executive-Summary.pdf>.
- Li, D., Chen, X., Becchi, M. and Zong, Z., 2016. Evaluating the energy efficiency of deep convolutional neural networks on CPUs and GPUs. In: IEEE international conferences on big data and cloud computing (BDCloud), social computing and networking (SocialCom), sustainable computing and communications (SustainCom), pp. 477–484.
- OpenAI, 2018. AI and compute. <https://openai.com/research/ai-and-compute>.
- Ren, S., He, K., Girshick, R. and Sun, J., 2016. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39:1137–1149.
- Schouten, G., Michielsen, B.S.H.T. and Gravendeel, B., 2024. Data-centric AI approach for automated wildflower monitoring. *PloS One*, 19:e0302958.
- Schwartz, R., Dodge, J., Smith, N.A. and Etzioni, O., 2020. Green AI. *Communications of the ACM*, 63:54–63.
- Strickland, E., 2022. Andrew Ng, AI minimalist: the machine-learning pioneer says small is the new big. *IEEE spectrum*, 59:22–50.
- Strubell, E., Ganesh, A. and McCallum, A., 2020. Energy and policy considerations for modern deep learning research. In: *Proceedings of the AAAI conference on artificial intelligence*, 34:13693–13696.
- Thompson, N.C., Greenewald, K., Lee, K. and Manso, G.F., 2020. The computational limits of deep learning. arXiv preprint arXiv:2007.05558.
- Wirth, R. and Hipp, J., 2000. CRISP-DM: towards a standard process model for data mining. In: *Proceedings of the 4th international conference on the practical applications of knowledge discovery and data mining*, 1:29–39.

Lessons learned from developing green software

Luís Cruz and Petra Heck

Abstract

Technology brings exciting opportunities to improve our interactions with the natural surroundings. However, that same technological development might also negatively impact the environment. Every new technology has a carbon footprint, whether from its construction or operation. And most technological developments require software systems, and more recently AI-based software systems. For these software systems to positively impact our environment, they need to be developed and operated with sustainability in mind, also called 'green' in the discipline of software engineering.

This chapter explores various dimensions of sustainability in software system development, drawing on existing software quality frameworks. We highlight green software best practices for development and knowledge transfer. We examine AI-based software systems, emphasising the importance of energy efficiency and carbon impact in the next generation of intelligent systems. This entails considering decisions at different stages of the AI lifecycle, ranging from underlying design choices in training pipelines to selecting optimal hardware for training and serving models.

This chapter presents the intersection of green software, sustainable software engineering, and green AI as of major importance for future innovation. By prioritising sustainability in software development and AI, we can foster a more sustainable and eco-friendly future, with the potential to reduce energy consumption and mitigate the environmental impact of technology.

11.1 Innovation and environmental responsibility

In an era defined by technological marvels, our world stands at a crossroads where innovation and environmental responsibility converge. Software technology promises to transform how we interact with the natural world, offering the prospect of a sustainable future. For example, it is used to track climate patterns, biodiversity, water and air quality; it is also used in waste management systems to optimise waste collection routes, manage landfill operations, and so on. Yet, software takes its own toll in terms of impact in the environment. For example, running software worldwide requires massive amounts of electricity. It is estimated that, by 2030, it will be responsible for 13% of electricity consumption globally (EU, 2022). Hence,

we are faced with a formidable challenge: how can we harness the power of software without leaving a devastating ecological footprint in our wake?

To bend the curve of biodiversity loss, our attitude towards nature has to change fundamentally. This behavioural change can for instance be triggered by engaging smartphone apps that can recognise flora and fauna at species level. This helps raise awareness and create knowledge of our natural environment.

CASE 11.1

AI-enabled mobile application

Consider, a mobile application that is designed to store and analyse pictures of flora and fauna. It leverages an artificial intelligence (AI) model to identify species automatically and catalogue your pictures according to contemporary taxonomic tables. It includes both local storage, in the phone, and online storage, in a cloud server, so that you do not lose your records even if something happens to your phone.

One key feature of this app is uploading the pictures to a server in the cloud. There are several ways of implementing this feature. For example, one way could be that, as soon as a picture is taken, the app takes care of immediately uploading it. This way, the pictures are immediately backed up as soon as they are taken. Most users would appreciate having their pictures safely secured as soon as they are taken.

Imagine, however, that many of the app users enjoy going for a hike in the woods, where there is poor internet connectivity. While enjoying the hike, users also enjoy taking pictures of the interesting things they see along the way: fauna, flora, landscape, etc. However, that implies that the phone makes a new data connection with the server every time a picture is taken. With poor connectivity, it would probably mean that the phone would take such a long time to upload each picture that it would not be fast enough for all the pictures that the user is taking. To make things worse, the phone's battery level is getting low, but the app is eagerly trying to get the pictures uploaded and will not stop until the phone dies.

The scenario in CASE 11.1 motivates developers to find strategies to make sure their software is designed in the most energy efficient way. It challenges them to include strategies to test and monitor the energy efficiency of their code. On top of that, smartphones are now running AI features locally, without an internet connection. This means that we save energy by not using an internet connection, but we drain our battery to run these powerful AI models. The proliferation of large language models in 2021 takes this challenge to another level, yielding ever-growing models with billions of parameters. As we delve deeper into this chapter, we show that a

single training iteration of these AI models leads to a massive carbon footprint. This motivates a need for new standards and practices that make energy efficiency a *de facto* requirement for modern software.

In this chapter, we cover the various dimensions of sustainable software and AI engineering, providing insights into how these pressing issues are being tackled. We start by defining the concepts around software sustainability. Then we analyse the well-established software quality standard ISO 25000 and look at how it relates to the practical realms of developing green software and green AI. Along the way, we unveil best practices that not only guarantee the energy efficiency but also the collection of reliable energy measurements from software.

11.2 Sustainable software engineering

sustainable
software
engineering

Although sustainability is a widely used term, the definition of sustainable software engineering is not always clear. Within this chapter, we define it as ‘the discipline that studies the process of creating software systems that create value in the long term without hindering its surroundings’. Despite being short, this definition gives us a starting point from which to identify software systems that are not sustainable. For example, if a software system is carbon efficient and eco-friendly but does not create value, it is not sustainable. Moreover, if a software system is eco-friendly but was designed in a way that makes it difficult to maintain (e.g. fixing security vulnerabilities), it is also not sustainable. The examples are numerous.

Software sustainability can be divided into five major dimensions (Becker *et al.*, 2015), as illustrated in Figure 11.1: environmental, social, individual, technical, and economical. We explain these five perspectives of software sustainability below.

environmental
sustainability

Environmental sustainability relates to the long-term effects of software on natural systems. This dimension includes ecosystems, raw resources, climate change, water, pollution, waste, etc. This can be attributed for example to the power consumption of the infrastructure used to run the software (e.g., electricity used by data centres), the water consumption used to cool down supercomputers, the disposal of IT hardware to acquire new state-of-the-art replacements, and so on. The branch of software engineering that studies this dimension is known as green software.

social
sustainability

Social sustainability is concerned with societal communities (groups of people, organisations) and the factors that erode trust in society. This dimension includes social equity, justice, employment, democracy, public health, public well-being, and so on. There is no doubt that software systems have a tremendous impact in our society and on individual people nowadays. This impact is often positive, but it can also have some negative consequences. A recurrent example of lack of social sustainability lies in many modern social media platforms, where we have witnessed the exploitation of their feed algorithms to disseminate fake news and mislead public opinion.




FIGURE 11.1 The five dimensions of sustainable software engineering

Individual sustainability refers to the well-being of the people involved in the development of the software. This includes mental and physical well-being, education, self-respect, skills, mobility, and so on. Several factors such as work-life balance, social environment, job autonomy, and physical health affect the well-being of tech professionals. Most of the work in this dimension aims at making sure that all the collaborators in a software team have a motivating, safe, and healthy working environment that fosters their contributions to the software project. Examples include promoting diversity within teams, sponsoring gym membership, defining clear career paths, organising team building activities, and so on.

individual sustainability

Technical sustainability refers to the longevity of information, systems, and infrastructure and their adequate evolution with changing surrounding conditions. It includes concerns such as maintenance, innovation, data integrity, etc. This is a dimension that has been widely addressed within software engineering communities. The lack of technical sustainability means, for example, that a system is not scalable and thus does not follow the growth of customers. It can also mean that whenever a new feature needs to be implemented, it is very complicated to do it in a way that will not break other features.

technical sustainability

Economic sustainability focuses on assets, capital and added value. It includes wealth creation, prosperity, profitability, capital investment, income, and so on.

economic sustainability

This is an essential dimension for any organisation. The operation of a software system must not lead its investors into bankruptcy. Even when we talk about toy software systems, if their contributors cannot support the costs of having a system available to its users, it will soon have to be decommissioned.

11.2.1 *Green software*

All these dimensions are important and interact with each other. For example, environmental sustainability is important for social sustainability and should not harm the economic sustainability of a project. However, some of these five dimensions have become more important throughout the history of software engineering. It is not difficult to convince a software organisation that they need to worry about the technical and economic sustainability of their software. For example, tech organisations take great pains to ensure that their business model is economically sustainable. Start-ups go through different rounds of funding where economical sustainability is the main concern. For example, if a fintech company decides to become more environmentally sustainable and moves all its software systems to data centres that run on clean energy, these data centres will probably be more expensive. Hence, the company will have to come up with a business model that is able to accommodate these extra costs, perhaps by charging their customers slightly higher fees. If customers do not deem these extra costs to be reasonable, this company might face serious economic issues – or simply shut down the idea at the first opportunity. Technical sustainability is also quite popular and is second only to economic sustainability. In the early versions of a software product, prototypes and proof-of-concepts are the main artefact being used to assess whether the product might be (economically) sustainable. Right after this stage, stakeholders start thinking about technical sustainability. In other words, software sustainability is a systemic concept and one dimension cannot be studied without considering the others.

With that in mind, a new discipline is emerging in software engineering: green software. Green software covers activities across the whole software development lifecycle: planning, analysis, design, implementation, and maintenance. This includes, for example, making sure that the software is energy efficient, i.e. that it uses the least power – which is easier said than done (Cruz, 2021). With the emergence of this field, any software organisation or developer that wants to deliver high-quality software cannot succeed without considering Green Software practices.

11.2.2 *A quality perspective on green software*

There have long been standards (since 1991!) that define what high quality means. ISO 9126 and its successor ISO 25000 (Systems and software Quality Requirements and Evaluation: SQuaRE) define quality characteristics for software systems,

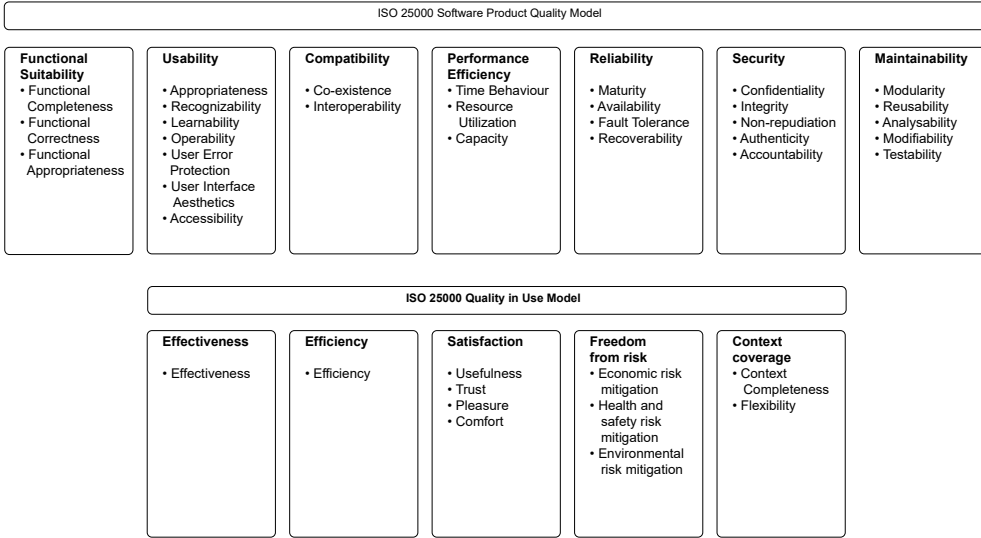


FIGURE 11.2 ISO 25000 quality model

including a measurement model with metrics. But it is only much more recently that the debate about green software has begun. So, one might think that green is a new quality characteristic that needs to be added to the ISO 25000 standards. However, as seen in Figure 11.2, which shows the quality characteristics for the inherent product quality and the emerging product quality (quality in use), ISO 25000 already includes the quality characteristics that define green software under the category 'Performance Efficiency'. In fact, ISO 25000 goes one step further and explicitly states under the category 'Freedom from risk' that the software system shall also mitigate any other environmental risk than just consuming unnecessary energy or storage.

So, in the light of ISO 25000, green software can be translated to software that scores high on the following two quality characteristics:

1. Performance efficiency: 'performance relative to the amount of resources used under stated conditions'.
 - a. time behaviour: 'degree to which the response and processing times and throughput rates of a product or system, when performing its functions, meet requirements'.
 - b. resource utilisation: 'degree to which the amounts and types of resources used by a product or system, when performing its functions, meet requirements'.
 - c. capacity: 'degree to which the maximum limits of a product or system parameter meet requirements (note: parameters can include the number of items that can be stored, the number of concurrent users, the

performance efficiency

environmental
risk mitigation

communication bandwidth, throughput of transactions, and size of database):

2. Environmental risk mitigation: ‘degree to which a product or system mitigates the potential risk to property or the environment in the intended contexts of use’.

One could argue that when developing performance-efficient software, environmental risk is also mitigated, but of course more needs to be done to fully mitigate any harm to the environment. In the remainder of this chapter, we will focus on the first characteristic, related to the energy consumption of software, because this is what green software has been mostly focusing on. Energy consumption is seen as the biggest environmental risk for software development, especially with the advent of AI-based software systems that require huge amounts of data storage and computing power.

11.3 Measuring energy consumption

There is no perfect way to measure energy consumption. It mostly boils down to a trade-off between simplicity and accuracy. In other words, the easiest approach tends to be less accurate, and the most accurate approach tends to be difficult to set up.

Let us start with the easiest approach. To estimate energy consumption, one can simply measure the time a software takes to execute a given task. If task A takes more time to run than task B, one can assume that A takes more energy than B. This assumption works well under well-defined conditions where we have observed this to be the case. For example, if we are comparing single-threaded Central Processing Unit (CPU) intensive algorithms, the algorithm that takes more time will be asking the CPU to execute more work. The more work from the CPU, the more energy is spent.

This assumption is challenged when the software uses a large set of resources with dynamic configurations that we cannot entirely predict beforehand: for example, a mobile application can use different sensors (GPS, accelerometer, camera, touch, etc.), different CPU modes (low-power or high-power units), different actuators (screen with different brightness colours, speakers, haptics, etc.). Modelling these different interactions is far from straightforward. Moreover, it is very unlikely that these interactions are identical across two different executions of the same task or use case.

As an alternative, the most accurate solution is to collect exact power measurements during the execution of the software – i.e. using a hardware-based approach. Using power monitoring tools, such as the Monsoon Power Monitor presented in Figure 11.3, it is possible to collect power data from the different hardware components and log power measurements per instant of time.




FIGURE 11.3 Energy monitoring setup for mobile app development

Unfortunately, this is far from easy. The figure above shows the power monitor connected to a smartphone. This setup required disassembling a smartphone and extracting its battery. That is suboptimal: mobile developers cannot be expected to recreate such a setup for the sake of improving energy efficiency.

An alternative is the use of software-based estimators of energy consumption. These estimators rely on indicators of the usage of different hardware resources (memory, CPU, and so on) to compute a value of the power consumption of a given software. A popular example is Intel® RAPL, which has been widely used in this area by researchers and tool developers. Depending on the scenario, these different alternatives need to be considered.

There are of course more challenges that need to be addressed before being able to measure software energy consumption. Software seldom runs in isolation: there is an operative system, other software applications, background processes, etc. that are running in parallel and contribute to the measured energy consumption.

To add to that, there are a multitude of other factors that also contribute to energy consumption: the power state of the different components (e.g., is the GPS already running or does it still need to be powered up?), their temperature, the temperature of the surroundings, and so on.

This means that if we measure the same software task twice, it is very unlikely that we will measure the same energy consumption.

There are several steps that need to be followed to collect reliable energy data. Below we have an extensive but not complete list of actions that should be taken before collecting energy data:

- zen mode – The first thing we need to ensure is that the only thing running in our system is the software we want to measure. Unfortunately, this is impossible in practice – our system will always have other tasks and things that it will run at the same time. Still, we must at least minimise all these competing tasks: all applications should be closed, and notifications should be turned off; only the required hardware should be connected (avoid USB drives, external disks, external displays, etc.); remove any unnecessary services running in the background (e.g., web server, file sharing, etc.); if you do not need an internet or intranet connection, switch off your network; cable is preferable to wireless – the energy consumption from a cable connection is more stable than from a wireless connection.
- freeze your settings – It is not possible to shut off the unnecessary things that run in our system. Still, we need to at least make sure that they will behave the same across all sets of measurements. Thus, we must fix and report some configuration settings. One good example is the brightness and resolution of your screen – report the exact value and make sure it stays the same throughout the measurement. A common mistake is to keep the automatic brightness adjustment on – this leads to major errors when measuring energy efficiency in mobile apps.
- warm up – Run dummy tasks a few times before keeping track of the measurements. Energy consumption is greatly affected by the temperature of your hardware. The higher the temperature, the higher the resistance of electrical conductors, leading to higher dissipation and consequently more energy consumption. If we start measurements right after the workstation comes back from sleep mode, it will probably be relatively cool. As soon as we start executing our energy tests, the temperature of the computer will rise until it reaches a plateau. This means that the first measurements will show less energy consumption precisely because they were the first. Typically, it is recommended to run a dummy task before starting to measure energy consumption. It can take as long as five minutes, depending on the scenario, and can be done by executing a CPU-intensive task, such as the Fibonacci sequence.
- repeat – Despite all efforts, subsequent energy measurements will lead to different results. It is therefore common practice to repeat executions and average the results. Depending on the scenario, the number of repetitions could go up to 30. This is the typical magic number in scientific studies because it enables statistical analyses. There is no golden rule: the best method is to adjust this number along the way based on the variability of the results.
- rest – Give the device a one-minute sleep between measurements. If we repeat the same experiment 30 times with no rest in between, our CPU will probably be

warmer in the last experiment than in the first. It is important to make sure that all measurements are executed under the same conditions. Hence, it is common practice to ensure a pause/break of a few seconds or minutes between repetitions. It can be more or less dependent on your hardware or the duration of your energy test.

- In some cases, we collect energy data to compare two versions of the software when performing the same use case. If we execute version A 30 times and version B another 30 times, there is a chance that between the first and the second half of measurements, the factors affecting the energy consumption have changed. Hence, the energy consumption of the two versions will also be different because the context has changed. To mitigate this issue, it is important to mix the executions of the different versions. shuffle
- It is well established that temperature affects energy consumption. We should not let differences in temperature affect our measurements. For example, it can happen that we measure the software over different periods of the day. This means that some measurements could happen during the night, with lower temperatures, or during the day, with higher temperatures. This bias has to be avoided at all costs to avoid misleading observations. This is done by controlling room temperature and collecting temperature data alongside energy data. With temperature data, one can discard energy measurements that might be misleading. keep it cool
- Although this sounds obvious, it tends to be challenging. If we manually run the software and test its features, we are introducing a new variable that may affect energy consumption. One cannot be sure that the interactions are exactly the same. automate executions

All these practices help improve the reliability of energy data collected. However, for practical reasons, they need to be taken with a grain of salt. There might be scenarios where overlooking some of these practices will not lead to significantly different results. As always, perfect is the enemy of good. Measuring something, even with flaws, is better than not measuring anything at all.

11.4 Best practices for green software development

Above, we discussed the difficulties inherent in measuring the energy consumption of a software application. It requires a great deal of effort and time to assess whether a given use case is energy efficient. In practice, developers may not be able to afford to follow such a meticulous process to collect energy. Hence, other alternatives ought to be considered.

This is where best practices or guidelines can play a major role. In other words, you do not need to measure every single line of code to have a certain degree of confidence that the code is energy efficient. As long as it follows energy efficiency

guidelines, energy measurements can be reserved for cases where it is deemed necessary.

Guidelines can provide advice at several levels of the software development process: at the code level – with code patterns, testing; at the server level – with regard to the deployment strategy, servers used, hardware; at the project management level – with practices ensuring skill diversity, knowledge transfer, refactoring iterations; and so on. In this section, we delve into existing design patterns for energy efficiency, we look at how the choice of programming language can affect the energy efficiency of a software system, and we showcase how the energy consumption of software can be addressed in the cloud.

11.4.1 *Green software design patterns*

It is not always necessary to measure the energy consumption of an application to understand that it will probably have issues in this area. If software developers are aware of design patterns for energy efficiency, they can already anticipate these scenarios. The good thing about energy issues is that they tend to be recurrent and will probably apply to different applications.

energy patterns

Design patterns for energy efficiency are typically known as energy patterns. They provide developers with advice on the best way to design their software code so that it runs with the minimum energy consumption. They are typically defined by (at least) three components: *context*, explaining where the pattern can be used; *solution*, explaining what can be done to improve energy efficiency in that context; and *examples*, with a few instances of software code where the pattern was applied. This proved to be useful when designing mobile applications (Cruz and Abreu, 2019).

Despite energy patterns being a useful resource for developers that want to build energy efficient tools, they can also be a wonderful tool for knowledge transfer. Imagine, for example, that in a software organisation with several teams, there is only one developer that is an expert in green software. This developer cannot be involved in too many projects at the same time, so most projects will not have anyone with expertise to improve the energy efficiency of their code. However, the developer is keeping a log tracking the main energy efficiency improvements that they have been making in their apps. At the end of the week, the developer documents these improvements and adds them to a catalogue of energy patterns that is accessible by other developers from other teams. These developers will then be able to learn from the expertise of the green software developer. Assuming that there is enough freedom for developers to explore this catalogue and apply some of the patterns in their own code, eventually the software organisation will have more developers with knowledge on energy-efficient coding practices at a very reasonable cost. In other words, education is an important aspect of green software and there are many easy ways to facilitate this.

11.4.2 *Green programming languages*

Green software is all about choices and trade-offs. When choosing the programming language for their next project, software developers have to consider a number of factors: experience of the team, code readability, the language community and ecosystem, existing libraries, and so on. To add to that, a group of researchers have recently studied the energy consumption of different programming languages (Pereira, 2021). As it turns out, the choice of language can have a massive impact on the energy consumption of the software.

The group of researchers took the existing benchmark for programming languages, called ‘The Computer Language Benchmarks Game’, and collected energy consumption metrics. The benchmark is used by experts from all over the world to rank programming languages according to their time and memory efficiency. The same tasks are solved across different programming languages and then results are compared.

When analysing the energy consumption data, the results were astonishing: popular programming languages such as Python were observed to be 76 times less energy efficient than top performing languages such as C. This is particularly interesting as Python is becoming more and more popular given its ease of use and available libraries for AI-based software development.

It is important to note that, in the case of using Python for AI development, most of it relies on third-party libraries that are often developed in C. This means that although we are coding in Python, our final software is probably using other programming languages under the hood and the overall energy efficiency is much better than one would expect. Nevertheless, it is important to be aware of these details – energy efficiency is not a given and our choices can make a big difference.

11.4.3 *Carbon-aware data centres*

Data centres continuously run a large number of complex everyday software tasks – for example, the so-called ‘cloud’ is nothing more than several large-scale software systems running in data centres. Naturally, it cannot be done without a massive ecological footprint from powering and cooling down all the servers running in the data centre. To reduce carbon emissions, many data centres are opting for renewable energy. However, renewable energy is not always available. Depending on demand and the availability of clean energy in the grid at a particular time, data centres may have to resort to less clean power sources. This is more severe during peak hours, when there is high demand and servers have to run at full capacity, requiring more energy to power the data centre.

Some of the tasks running during peak hours are urgent and need to be immediately executed. For example, when a bank system is processing a payment from a grocery store. Any small delay leads to poor customer experience. However, there are other tasks that do not really need to be executed at a precise second or minute.

For example, many AI training tasks take several hours to execute. In those cases, it might be okay to wait a couple hours before performing the task.

Software tasks that have flexibility in their execution time offer an excellent opportunity to reduce carbon emissions. For example, if the carbon intensity from the power grid is high – i.e. there is not enough clean energy in the grid – we can simply wait until there is enough clean energy before we execute that task.

This simple idea makes it easier to make sure a data centre runs mostly on clean energy. It has been tested in a few contexts. However, it still poses a few challenges: estimating the carbon intensity of the grid for the next few hours, predicting how complex a given task is, predicting how much time it takes to finish, and defining how much time it can be delayed. Making incorrect estimates of these numbers can lead to suboptimal results. For example, if we predict that the sun will be bright later in the day – meaning that solar panels will produce clean energy – and then it turns out to be cloudy, delaying tasks could be emitting more carbon than we anticipated. Nevertheless, preliminary results show great potential in this strategy.

With simple ideas like this, the carbon footprint of software systems running in data centres can be massively improved. However, the carbon footprint of software is always being challenged with new cutting-edge technologies that always require more energy and resources. One example is the new hype for AI technologies, which require a massive amount of data and computational resources.

11.5 Green AI

The advent of artificial intelligence (AI) is making the new generation of software intelligent. Software systems feature components that perform tasks that until now have only been accomplished by humans. It takes software use cases to another level, e.g. by enabling them to generate text, images, and recognise objects, faces, etc. It can even perform certain tasks better than humans – for example, AI models are widely used to detect fraudulent online transactions.

machine
learning

Machine learning (ML), the most popular form of AI, consists of feeding large amounts of data into a computational pipeline that will process it and extract patterns in order to generate a model – a task known as training. This model can then be used to take inputs from users and return meaningful answers – a task typically known as inference. Of course, these models are mostly useful when integrated with software. For example, the model itself does not provide a graphical user interface that allows users to interact with it. So, in general, AI models are being used without users even being aware of it. Behind every bank transaction, social media post, etc. there are multiple models running in the background.

Despite the numerous benefits of adding AI models to software systems, there is an inherent cost that needs to be addressed: energy consumption. These models require a great deal of energy to be developed and executed. For example, the

popular large language model ChatGPT consumed 550 tonnes of CO₂-eq to be trained (Patterson *et al.*, 2021). This is equivalent to driving 1000 cars for 1000 kms. To make matters worse, this is only related to a single training of the model – the cost of experimenting and tuning the training strategy took several iterations before converging to the final one. If we factor in the cost of running the model – i.e. interacting with the chat – this number will probably be astronomical. According to the latest numbers, the cost of running ChatGPT for a year is equivalent to 25 times the cost of training it (Chien *et al.*, 2023). When we think that AI is being considered a competitive technology for every business, it is not too big a leap to imagine these numbers quickly escalating.

This means that modern AI, despite being a powerful technology to combat climate change, has a carbon footprint of its own that cannot be ignored. Hence, when assessing the quality of a model, the answer can no longer be exclusively about its ability to come up with correct answers. Other metrics need to be looked at. AI is therefore becoming an important part of building sustainable software systems. It calls for a redesign of the processes we have for developing and monitoring AI systems, to ensure not only green software, but also green AI.

11.5.1 *A quality perspective on green AI*

Many have analysed the differences between rule-based software systems and (self-) learning software systems (Heck *et al.*, 2021).

With respect to the energy consumption of AI-enabled systems, four differences stand out:

1. AI-enabled systems consist of code, data and models (see Figure 11.4). This means we need to incorporate performance efficiency of data and models too.
2. Data is considered the new gold due to the hype surrounding data analytics in general and AI in particular. This encourages organisations to store huge amounts of data, with a considerable environmental cost, even if they are not sure they will use this data in the future.
3. AI-enabled systems are very computer-intensive because of the complicated mathematical calculations that are behind the machine-learning algorithms and the large datasets that they need to be applied to.
4. There is a focus on optimal models (see Figure 11.4), where stronger results are ‘bought’ by massively increasing computational power. Schwartz *et al.* (2020) call this ‘Red AI’.

So green AI is essentially green software on steroids. That is why it emerged as a separate research field and is being addressed in many industry standards and guidelines, including the upcoming EU AI Act. For example, the EU guidelines for trustworthy AI (EU, 2019) list seven requirements, including ‘Societal and environmental well-being’. According to this requirement ‘AI systems promise to help tackle some of the most pressing societal concerns, yet it must be ensured that this occurs in the most environmentally friendly way possible. The system’s




FIGURE 11.4 Development process for AI-enabled systems, including a data and model (ML) loop

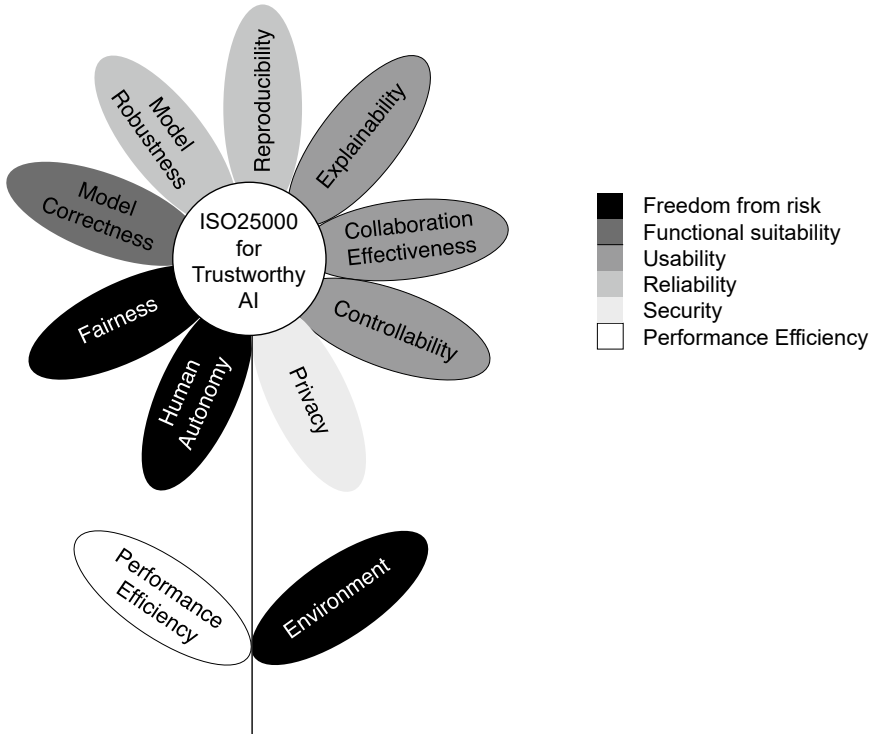


FIGURE 11.5 Green AI at the root of Trustworthy AI

development, deployment and use process, as well as its entire supply chain, should be assessed in this regard, e.g. via a critical examination of the resource usage and energy consumption during training, opting for less harmful choices. Measures securing the environmental friendliness of AI systems’ entire supply chain should be encouraged.’

Heck (2022) analysed the quality characteristics of AI systems and concluded that nine properties need to be added to ISO 25000: the leaves of the flower in Figure 11.5. However, from the description in the EU guidelines we can see that in fact green AI quality characteristics are the same as green software quality characteristics:

performance efficiency and environmental risk mitigation. According to those same guidelines, environmental well-being is a necessary condition for trustworthy AI. This vision is depicted in Figure 11.5, highlighting the importance of green AI in the development of future AI-enabled software systems.

As described in Figure 11.4, the AI development process covers a wide range of things that need to be addressed before converging to a fully functional AI system. Such a complex pipeline implies that to make green AI systems we need to address energy efficiency at all frontiers.

11.5.2 *Best practices for green AI*

Green AI as a field is still in its infancy but there are already some interesting approaches that deserve a mention. One is data simplification. Not all the data we collect brings something worthwhile to be learned. The way we have been dealing with issues in data quality is by increasing the size of our data. This way, any issues in the data are attenuated by other data points. However, if we increase the size of the data, we also increase the computation required to train the model. As shown by Yarally *et al.* (2021), simple techniques for feature selection or data size reduction can have a massive impact on energy consumption.

data
simplification

Other ingenious techniques rely on simplifying the model. Popular techniques of model simplification are quantization, pruning, and distillation (Van Steenweghen, 2023). Quantization consists of changing the data types used at model training – e.g. by reducing the number of bits used. Pruning consists of removing parts of a neural network that are not relevant for the performance of the model. Distillation consists of learning a new smaller model – also known as the student model – based on the knowledge of the original large model – also known as the teacher model. The downside is that some of the model simplification strategies can make model training more energy intensive, as they require extra computation for training. The benefits are mostly achieved in terms of the energy consumption of using the model. Another downside is the fact that applying some of these techniques is often not straightforward and requires domain expertise. These techniques come in different flavours and require tuning before they can produce improvements in energy efficiency. A good example of this is the project LLaMA.cpp: a port of Meta's large language model LLaMA that aims at running in low-powered devices such as a smartphone but running several model simplification techniques.

simplifying
the model

Another interesting angle is the pareto between the training pipeline and the training hardware. When they want to train their models, AI practitioners often opt for the best server available. This is a natural choice, because you want to use the best state-of-the-art tools available. However, as usual, we need the best tool for the right job. When we select the most powerful hardware, this does not necessarily give us better results. Research has shown that, depending on the model, a hardware setup with lower specs can lead to better energy efficiency without hindering accuracy metrics (del Rey *et al.*, 2023).

lower specs

Such an observation is in line with reports from Meta revealing that a vast portion of their AI systems are only using GPUs at 30% of their full capacity (Wu *et al.*, 2022). This is a major issue because it means that we need three times more GPUs to compensate for this lack of resource efficiency. Moreover, hardware has an embodied carbon footprint – i.e. the carbon emissions incurred by producing and shipping it – that is being wasted on stalled GPUs.

obsolete
hardware

Another issue arising from reliance on the latest hardware is that existing hardware soon becomes obsolete (del Rey *et al.*, 2023). The latest AI libraries are quickly dropping their support for older hardware. This does not necessarily mean that the hardware is unable to perform the required computational tasks. In some cases, it simply means that the cost of providing support to old hardware has become economically unsustainable. The consequences of this problem are twofold: 1) we are reducing the lifecycle of hardware without reducing its embodied carbon footprint, and 2) only practitioners that have the capacity to buy new state-of-the-art hardware are able to run the latest advancements in AI technology. This last consequence raises an issue in terms of AI democracy. We want AI technologies to be accessible to any group of tech enthusiasts and not just to a few big tech companies that have enough negotiating power and financial resources to get the hardware that runs the latest AI software.

11.6 Conclusion

This chapter lays the foundation for sustainable software engineering and all its different dimensions. The means to creating sustainable software requires a multi-factorial approach that considers trade-offs between different sustainability aspects: environmental, social, individual, technical, and economic. This is easier said than done: despite being covered by industrial software quality models, designing green software systems requires a number of non-trivial techniques to test and monitor the energy efficiency of software.

In this chapter, we have seen how energy patterns and other best practices can help build energy-efficient software by design. However, the advent of AI is challenging the software industry by requiring even more computational power to run state-of-the-art intelligence. The emerging guidelines for trustworthy AI make performance efficiency and environmental risk mitigation essential requirements.

However, more work needs to be done to enable a fully green landscape in the realm of intelligent software. This is a concern that involves developers, designers, policymakers, users, and society in general. We currently rely on software for everything – as software users, we have the right to ask for transparent and trustworthy software. When we download a new app or sign up for a new service, we want to know what is behind the curtain. All stakeholders have a role to play, and little by little, step by step, we will make all software green, together.

References

- Becker, C., Chitchyan, R., Duboc, L., Easterbrook, S., Penzenstadler, B., Seyff, N. and Venters, C.C., 2015. Sustainability design and software: the Karlskrona manifesto. In: 2015 IEEE/ACM 37th IEEE International Conference on Software Engineering (Vol. 2, pp. 467–476).
- Chien, A.A., Lin, L., Nguyen, H., Rao, V., Sharma, T., and Wijayawardana, R., 2023. Reducing the carbon impact of generative AI inference (today and in 2035). In: Proceedings of the 2nd workshop on sustainable computer systems (pp. 1–7).
- Cruz, L. and Abreu, R., 2019. Catalog of energy patterns for mobile applications. *Empirical software engineering*, 24:2209–2235.
- Cruz, L., 2021. Green software engineering done right: a scientific guide to set up energy efficiency experiments. <http://luisacruz.github.io/2021/10/10/scientific-guide.html>.
- del Rey, S., Martínez-Fernández, S., Cruz, L., and Franch, X., 2023. Do DL models and training environments have an impact on energy consumption? In: 49th Euromicro Conference on Software Engineering and Advanced Applications (SEAA), pp. 150–158.
- EU, 2019. Ethics guidelines for trustworthy AI. Brussels: European Commission. <https://data.europa.eu/doi/10.2759/346720>.
- EU, 2022. Questions and answers: EU action plan on digitalising the energy system. https://ec.europa.eu/commission/presscorner/detail/en/qanda_22_6229.
- Heck, P., Schouten, G., and Cruz, L., 2021. A software engineering perspective on building production-ready machine learning systems. In: Handbook of research on applied data science and artificial intelligence in business and industry (pp. 23–54). IGI Global.
- Heck, P., 2022. A quality model for trustworthy AI systems. <https://fontysblogt.nl/a-quality-model-for-trustworthy-ai-systems/>.
- Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.M., Rothchild, D., ... and Dean, J., 2021. Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350.
- Pereira, R., Couto, M., Ribeiro, F., Rua, R., Cunha, J., Fernandes, J.P., and Saraiva, J., 2021. Ranking programming languages by energy efficiency. *Science of computer programming*, 205, 102609.
- Schwartz, R., Dodge, J., Smith, N.A., and Etzioni, O., 2020. Green AI. *Communications of the ACM*, 63:54–63.
- Van Steenweghen, A., 2023. EasyCompress: automated compression for deep learning models. Master Thesis TU Delft, the Netherlands.
- Wu, C.J., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., Chang, G., Aga, F., Huang, J., Bai, C. and Gschwind, M., 2022. Sustainable AI: environmental implications, challenges and opportunities. In: Proceedings of machine learning and systems, 4:795–813.
- Yarally, T., Cruz, L., Feitosa, D., Sallou J., and Van Deursen, A., 2023. Uncovering energy-efficient practices in deep learning training: preliminary steps towards green AI. In: IEEE/ACM 2nd International Conference on AI Engineering–Software Engineering for AI (CAIN) (pp. 25–36).

A daily data workout!

Being in correspondence for a green data revolution

Danielle Arets and Jessie Harms

Abstract

With the expansion of AI-driven applications and services, it is becoming increasingly crucial to scrutinize the CO₂ emissions associated with AI. After all, both the usage and training of these applications consume a significant amount of energy. This is a topic that, as of yet, remains understudied and inadequately addressed in the political and public arenas. Consequently, there is a lack of policies, and consumers are largely unaware of their digital footprint. In this essay, we argue that, for a sustainable digitized future, it is crucial to develop frameworks that consider the energy usage of AI-driven tools and services. Additionally, engaging with artistic research is essential to help and inspire the public at large to stay ‘in correspondence’ with green AI developments.

12.1 Introduction

CASE 12.1

The Media Gym

What if you could feel the weight and energy consumption of your daily media interactions? How many kilocalories would you need to burn to offset your e-mails, surfing, and Netflix habits? With The Media Gym,¹ the Dutch design studio Cream on Chrome, designed a pop-up installation to foster a conversation about the ecological footprint of digital media, using a gym metaphor. The workout enables participants to confront their data usage and associated CO₂ emissions. Whether engaging in aerobics or lifting weights, the designers instigate a discussion about the digital work culture. At times, this involves additional weightlifting because colleagues are in online conference meetings, resulting in even more data consumption during the workout. By initiating the conversation in a humorous but, more importantly, tangible manner – literally feeling the pressure – Studio Cream on Chrome aims to encourage collective reflection on the digital footprint.

¹ <https://cream-on-chrome.com/mediagym>.




FIGURE 12.1 Leaflet Media Gym, Studio Cream on Chrome

Art installations like the Media Gym highlight a frequently overlooked issue: the connection between AI applications and energy consumption. As our dependence on AI-generated services increases, so will CO₂ emissions. The environmental impact of our digital behaviour has only recently become a subject of academic research (Van Wynsberghe, 2021; Verdecchia *et al.*, 2023). There is growing interest in using AI for sustainability, to create more efficient and green services, systems and tools in order to achieve Sustainable Development Goals (SDGs). However, the sustainability of AI itself is only now receiving attention. The primary focus of Green AI is on AI's carbon footprint, as the training, use and maintenance of AI require significant energy and result in CO₂ emissions.

Gaining accurate insights into the AI footprint requires intricate calculations to account for all energy-consuming elements. These insights are crucial for AI developers, companies and governments in order to understand AI's broader impact on sustainability. However, the current diversity among systems of measurement leads to confusion and to a lack of consistent regulations and policies (Verdecchia *et al.*, 2023). A more unified system and comprehensive understanding are essential to ensure a sustainable AI future.

To further advance Green AI initiatives, it is crucial not only to excel in complex calculations but also to bridge the gap between this specialized field and the general public. AI often remains opaque and distant, making it challenging for citizens to engage meaningfully.

We propose that involving artistic researchers can play a pivotal role in fostering broader citizen engagement, which is vital for the progress of Green AI. By incorporating artistic approaches, such as interactive installations or performances, Green

AI concepts can be communicated in ways that resonate with diverse communities and allow the public to be in correspondence with these developments.

12.2 Doing the maths

Among the first researchers looking into Green AI, Strubell *et al.* (2019) estimated that training a single large language transformer would cost about 300 tons of CO₂. The popular Chat GPT application is considered the most significant energy guzzler, with MIT Technology Review estimating that OpenAI's GPT-3 emits more than 575 metric tons of carbon dioxide during training² (Lacoste *et al.*, 2019). These figures were a revelation to the AI community, resulting in an increased focus on Green AI in research papers paying particular attention to the calculation of CO₂ emissions resulting from AI in the past three years (Henderson *et al.*, 2020; Verdecchia *et al.*, 2023)

However, doing the maths here is complicated. There are various ways to estimate the energy consumption of AI systems and services. Some studies log AI training time, while others focus on the use of trained models by end users or on data centre energy use. [A more technical approach for Green AI best practices can be found in chapter 10 and 11 of this book]. Furthermore, most studies are not bound to a specific context or algorithm, making a comparison of the numbers difficult.

laboratory
studies

Moreover, studies in this field revolve mainly around laboratory studies, and insights are primarily published in academic papers. A minority address readers in the industry and very few are aimed at the general public (Anthony *et al.*, 2020). Therefore, the impact of these calculations is largely limited or restricted to the academic field.

Carbontracker

According to Van Wynsberghe (2021), the two most common methods of calculation are the Machine Learning Emissions Calculator and the Experiment-Impact-Tracker framework. Both tools track training and usage time. Van Wynsberghe emphasizes that 'greater incentives are needed to encourage researchers and industry developers to measure and report such findings' (Van Wynsberghe, 2021: 216). She, therefore, prefers a tool like Carbontracker (Anthony *et al.*, 2020). This not only calculates but also looks forward, predicting energy use over time and allowing for the model training to be 'stopped, at the user's discretion, if the predicted environmental cost is exceeded' (Van Wynsberghe, 2021: 216).

directive models

Kaack *et al.* (2022) also stress the importance of steering the developments utilizing more directive models. They introduce a systematic framework for describing the effects of machine learning (ML) on greenhouse gas (GHG) emissions. The framework encompasses three categories: computing-related impacts, the

2 <https://medium.com/mllearning-ai/an-ai-model-that-is-energy-efficient-is-just-as-important-as-its-purpose-71d17822a183>.

immediate impacts of applying machine learning, and system-level impacts. It identifies priorities for impact assessment and scenario analysis and suggests policy arrangements for better understanding and shaping the effects of machine learning on climate change mitigation.

However, the researchers warn of shortcomings in their model, as both technical and societal developments continue to challenge their framework. Hence, alongside more forward-looking models, an ongoing social and political debate is necessary regarding sustainable AI in order to establish appropriate regulations and foster innovation in sustainable practices.

social and
political
debate

Van Wynsberghe (2021) proposes the establishment of *sustainable AI task forces* in governments. The study further suggests that the European Commission should develop a 'proportionality framework' to assess whether the training of an AI model for a specific task is proportional to the carbon footprint and overall environmental impact incurred by that training and tuning. These calls appear to be gaining traction. In the European Green Deal (European Parliament: Directorate-General for Internal Policies of the Union *et al.*, 2021) there is a strong emphasis on applying AI to achieve sustainability, although a focus on the sustainability of AI itself is also present. The former takes a more prominent position on the agenda.

European Green
Deal

Despite the lack of research into Green AI policies and related actions, there is also a serious lack of understanding about how the public at large engages with this topic. Currently, citizens remain at a considerable distance from these developments. This is problematic, especially given the potential societal impacts of AI driven tools, services and systems. There is a widespread call to engage citizens more meaningfully in dialogue and decision-making processes regarding the use of AI (Dollbo, 2023).

12.3 Correspondence with AI

Over the past three decades, deliberative democracy has gained traction as a model where citizens engage meaningfully in decision-making. The deliberative process takes various forms, including citizen assemblies, summits, Socratic dialogues, voting experiments, and more (Elstub and Escobar, 2019). The goal of deliberation is to stimulate an informed and coherent public understanding of complex issues, such as AI.

Bringing citizens together around the topic of AI is challenging due to its abstract and complex nature. Technological topics are often difficult to comprehend. Furthermore, the dominant, rational focus of deliberation processes tends to overlook the fears and anxieties associated with new technologies such as AI (De Droog-Arets, 2024). A growing number of scholars advocate the broadening of deliberation processes (Roeser and Pesch, 2024) to better address emotions and ideals, as these are important normative guides when dealing with AI.

The incorporation of artistic and design practices is a fruitful way to diversify deliberation processes. These can transform discussions into material and experiential forms (De Droog-Arets, 2024; DiSalvo, 2015). Making things tangible is a crucial way to enhance understanding, especially when dealing with complex, abstract and invisible topics such as AI. Tangible translations of the technology enable citizens to relate to the technology, to imagine potential, beneficial and harmful effects and to question the desirability of these developments (De Droog-Arets, 2024). This can enable all participants to join and contribute to the discussion of complex issues or possible futures.

Furthermore, artistic and design interventions could help the general public to be in ‘correspondence’ (Ingold, 2017) with AI. With this term, Ingold refers to our relationship with things as they present themselves to us, our experiences, and relationships, emphasizing the importance of being attuned to our surroundings and the entities within them. To be in correspondence, requires focusing on new forms of knowledge creation. To Ingold, our current focus is too much on knowledge creation in the Major Key (Ingold, 2018), which means a rational fixation on solving problems by bringing together, combining, and unravelling the particles and ‘doing the maths accordingly’. In the Major Key, there is a strong emphasis on intentionality, planning, and achieving predetermined goals. This perspective manages uncertainty through structured processes and systems.

In order to make ‘wise’ decisions on how to use and integrate AI, we also need to understand complex problems and societal developments, and allow ourselves to engage with creative ways of generating knowledge. This knowledge in the Minor Key involves engaging and being in correspondence, maintaining a responsive relationship with the world, like an ongoing conversation. Being in correspondence is especially important in relation to Green AI. To Ingold, the reciprocal relationship between humans and their surroundings, their tools and devices, is crucial to the creation of a more sustainable way of living. The Minor Key recognizes that not all uncertainties can be eliminated and that living with and adapting to the unknown is a valuable skill.

12.4 Art and design as a means to be in correspondence

How then can art help us to do this? Returning to Studio Cream’s Data Gym project on Chrome, this initiative encourages users to engage with their data habits. Beyond presenting facts and figures, users interact with their data usage through physical movement. This approach not only appeals to our rational understanding but also taps into our emotional and visceral connections with data, instigating reflections on the impact of our digital environments and behaviours.

The project demonstrates how the subject of Green AI can be made tangible and experiential by involving audiences in active participation. This activity shifts



FIGURE 12.2 Screenshot of taste workshop by E-missions

the focus from simply ‘doing the maths’ to linking the intensity of physical exercise with the amount of data consumed. Following Ingold, it bridges Major and Minor key thinking by emphasizing our profound relationship with data consumption and energy use, aiming to provoke meaningful questions about sustainability, about technology’s ecological footprint and about our personal activities.

CASE 12.2

E-missions project

The E-missions project is another noteworthy art project, developed by Dutch artists and scientists.³ According to the development team, the sustainability of AI is overlooked due to the lack of a tangible relationship between digital behaviour and the related consequences. The E-mission project developed a Sensing CO₂ tasting experience, where they convert CO₂, which is odourless and invisible to the human eye, into a visible sparkling soda. Depending on your data behaviour, this fizzy drink triggers the tongue – pleasantly or unpleasantly.

Furthermore, the E-missions group developed a ‘digital etiquette’, a tool to enable discussion, definition and deployment of a code of conduct to greenify the digital workstyle of an organization, contributing to its sustainability strategy by making it actionable. The etiquette that E-missions proposes is not a fixed cost-benefit analysis, but rather a tool for being in correspondence, to look at the world in flux and to continually question our behaviour and relationship with it.

The Sensing CO₂ workshop offers participants a taste of potential futures and encourages them to explore their desired relationship with sustainable data practices. This type of activity aligns with Ingold’s endorsement of using imagination to experiment and engage directly with current realities. It illustrates how both major

³ <https://www.e-missions.nl/en/about-e-missions/>.

and minor key thinking can coexist, enabling a broader audience to discuss developments in Green AI in a meaningful and tangible manner.

The act of tasting a possible future goes beyond mere speculation, it embodies Ingold's notion of 'seeking into' the future. The Sensing CO₂ workshop is designed to be enjoyable and has a low barrier to entry, making it accessible for all participants to join, contribute and discuss tasty futures.

CASE 12.3

The Climate Futures project

In the art-science project, Climate Futures, the researchers (Querubín and Niederer, 2022) engaged with machine learning from technical, critical and practical perspectives in order to explore new modes of engagement with the issue of climate change. The Climate Futures team applied generative AI to create climate fiction (Clifi) novels. As various scholars have pointed out, reading climate fiction can make people more open to climate issues, and this fiction helps to elicit emotions and living experiences surrounding climate change (Trexler, 2015; Schneider-Mayerson, 2018). In this project, AI is used to help us to be in correspondence with its developments.

This project exemplifies the concept of 'correspondence' where user engagement occurs on an emotional level. It uses AI as a tool to critically examine the environmental implications of AI. The collaboration between researchers and artists in this endeavour vividly illustrates Ingold's advocacy for interdisciplinary knowledge creation.

12.5 Allowing for fluidity

How are we to apply these Minor Key, art and design-driven practices in environments that are predominantly structured, goal-oriented, and focused on standardized methods (Major Key thinking)? In our research group, Designing Journalism, based at a university of Applied Sciences, we work with a diverse team of researchers, including AI experts, journalism researchers, anthropologists, designers, and artists. This allows for interdisciplinary collaborations, a mix of methods and innovative points of view. Being part of a large university requires us to emphasize accountability, deliverables, recording hours worked and academic output. However, this doesn't prevent us from cultivating a culture of experimentation that embraces learning from experiences and failures. Our dominant research approach, research through design, places significant emphasis on designing physical outcomes,

iteration and learning from (reflecting on) the outcomes and the processes. The experiences and knowledge (also being tacit or experiential knowledge) generated through the design, allow us to integrate both Major and Minor key insights.

Designers can play a crucial role in shaping how societies navigate through this balance between control and open-endedness. Faced with an uncertain future, designers can contend with both the desire for certainty and the need for flexibility. Designs that incorporate elements of flexibility and resilience can better support individuals in navigating an uncertain world. This requires a shift from simply focusing on reduction of risk towards enabling people to cope with and learn from uncertainty. Doubt also plays a critical role here. Allowing for doubt encourages a more reflective and adaptive approach to decision-making and action. Designers can help to foster a culture of questioning and continuous learning, essential for thriving in an unpredictable world.

Space is also an important aspect. It is important to allow for the creation of 'spaces' where new forms of knowledge can emerge. Lab environments that provide both physical and mental space for creation, experimentation, and play are vital (Asenbaum and Hanusch, 2021). Creative lab environments allow participants to move beyond discursive expression and to engage in embodied experiences and actions.

As an example, we currently collaborate with the Dutch National Climate Platform (a national advisory council) in a research initiative focused on analysing the imagery and metaphors used in journalistic discussion of climate change. Designers, journalism researchers, environmental experts and policymakers collaborate in the project. As part of our research, we've developed a mobile lab that will feature at various music festivals. In the lab, young adults will engage with journalistic images and metaphors on climate change in order to discover how to design and to explore more impactful climate narratives.

A centrepiece of our lab is a photo suit – a unique jacket adorned with thousands of photos, making it extremely heavy. Festival attendees will be encouraged to wear the suit and participate in discussions about the significant digital footprint associated with each digital image that they capture. This interactive experience aims to turn a complex concept into a tangible focal point for dialogue.

12.6 To conclude

In order to enhance Green AI developments, we need to look beyond the data to make wise decisions. Or, as Chat GPT itself aptly explains when asked about its energy use: 'developing sustainable AI is an ongoing process that requires a holistic approach and a commitment to minimizing negative impacts, while maximizing the positive contributions of AI technology to society. Collaboration with experts

and stakeholders, transparency and responsible decision-making are crucial elements of this endeavour⁴[28].

The art-driven projects mentioned above focus on Green AI developments by actively engaging with and questioning their purpose, materiality, and sensory qualities. They enable a diverse audience to generate knowledge through hands-on making, experimentation and engaging with AI. The tangible nature of these projects helps to demystify the elusive aspects of AI, making it more accessible and understandable.

Moreover, art projects foster a way of being that encourages audiences to remain attentive and responsive to our relationships with the world around us. This connection is crucial, especially with a technology like AI that is invisible and challenging to comprehend. In essence, these initiatives not only explore AI's complexities but also nurture a mindset that values engagement, reflection, and dialogue, contributing to a more informed and inclusive discourse on its role in society. Or, as Ingold would express it: helping us to seek into the future.

References

- Anthony, L.F.W., Kanding, B. and Selvan R., 2020. Carbontracker: tracking and predicting the carbon footprint of training deep learning models. *ArXiv:2007.03051*.
- Asenbaum, H. and Hanusch, F., 2021. (De) futuring democracy: labs, playgrounds, and ateliers as democratic innovations. *Futures*, 134, 102836.
- De Droog-Arets, D.J.A.M., 2024. Save the debate: through adversarial design, PhD dissertation, University of Technology Eindhoven.
- DiSalvo, C., 2015. *Adversarial design*. MIT Press, Cambridge, MA, USA.
- Dollbo, A., 2023. *Deliberating AI: why AI in the public sector requires citizen participation* (Master's thesis).
- Elstub, S. and Escobar, O. (Eds.), 2019. *Handbook of democratic innovation and governance*. Edward Elgar Publishing, Cheltenham, UK.
- European Parliament: Directorate-General for Internal Policies of the Union, Herold, A., Gailhofer, P., Urrutia, C., Braungardt, S., Köhler, A., Scherf, C. and Schemmel, J., 2021. *The role of artificial intelligence in the European Green Deal*, European Parliament. <https://data.europa.eu/doi/10.2861/882830>.
- Henderson P., Hu, J., Romoff J., Brunskill, E., Jurafsky, D. and Pineau, J., 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *ArXiv:2002.05651*.
- Ingold, T., 2019. Art and anthropology for a sustainable world. *Journal of the Royal Anthropological Institute*, 25(4), 659–675.

4 OpenAI; prompt 'developing sustainable AI'.

- Ingold, T., 2018. Evolution in the minor key. *Evolution of wisdom: major and minor keys*, 115–23.
- Kaack, L.H., Donti, P.L., Strubell, E., Kamiya, G., Creutzig, F. and Rolnick, D., 2022. Aligning artificial intelligence with climate change mitigation. *Nature climate change*, 12(6), 518–527.
- Lacoste, A., Luccioni, A., Schmidt, V., and Dandres T., 2019. Quantifying the carbon emissions of machine learning. *arXiv:1910.09700*.
- Querubín, N.S. and Niederer, S., 2022. Climate futures: machine learning from cli-fi. *Convergence*, 0(0). <https://doi.org/10.1177/13548565221135715>.
- Roeser, S. and Pesch, U., 2016. An emotional deliberation approach to risk. *Science, technology, & human values*, 41(2), 274–297.
- Schneider-Mayerson, M., 2018. The influence of climate fiction: an empirical survey of readers. *Environmental humanities* 10(2): 473–500.
- Strubell, E., Ganesh, A. and McCallum, A., 2019. Energy and policy considerations for deep learning in NLP. *ArXiv:1906.02243*.
- Trexler, A., 2015. *Anthropocene fictions: the novel in a time of climate change*. University of Virginia Press.
- Van Wynsberghe, A., 2021. Sustainable AI: AI for sustainability and the sustainability of AI. *AI ethics* 1:213–218. <https://doi.org/10.1007/s43681-021-00043-6>.
- Verdecchia, R., Sallou, J., and Cruz, L., 2023. A systematic review of green AI. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 13(4), e1507.

Added value of AI for studying urban plants

Barbara Gravendeel and Yannick Woudstra

Abstract

Wild plant species establish themselves spontaneously in cities all over the world in different microhabitats. Thriving in a city is by no means easy for plants as urban landscapes are hot, highly fragmented, severely disturbed and polluted. This has triggered urban plants to adapt to their environment in different ways. We will provide examples of adaptations that have evolved in different urban species and populations. Initially characterized as ‘weeds’, urban flora is increasingly appreciated for improving the quality of our living environment by providing shade during heat waves, reducing traffic noise, and retaining water during floods. Sadly, many urban designers and citizens suffer from ‘plant blindness’: they do not notice the diversity of their green surroundings and miss out on how urban flora enriches their lives. Artificial intelligence has huge potential to cure plant blindness, validate plant-based ecosystem services, and change perceptions of the added value of urban green spaces with multiple stakeholders. We describe how computer vision algorithms are currently being developed for (micro)species identification, quantification, and phenology assessments of plants. AI-based tools are also increasingly employed to study adaptations of urban plants. We will end this chapter by providing some suggestions for the practical applications of such tools.

13.1 Introduction

Over the past centuries, most human inhabitants of our planet have lived in small rural communities. In recent decades, this picture has changed as many people moved from rural to urban areas. As a result, more than half of us now live in cities. This process is called urbanization, and it especially occurs in high- and middle-income regions. The benefits of a city life for human inhabitants include a higher density of economic activity, shorter trade links, and a shared infrastructure (Ritchie and Roser, 2018).

wild plant
species

To wild plant species establishing themselves in cities there are benefits as well. And that is probably why cities nowadays generally host a higher number of plant species than surrounding rural areas of comparable size. The plant species composition among such urban and nearby rural areas is very different. Some native plants cannot survive in urban landscapes that are hot, highly fragmented, severely disturbed, and polluted (Johnson and Munshi-South, 2017). These species have disappeared. But new species, especially adapted to city life, have been arriving and

establishing themselves. The percentage of newly arriving species increases with city size (Johnson *et al.*, 2015), possibly because a larger city can harbour more micro-habitats hosting particular species.

Examples of typical micro-habitats in European cities include stone walls and the areas around or on ornamental trees. These micro-habitats all have their own challenges that select for very specific plant species assemblages. Below, we describe some of these in more detail.

micro-habitats

13.2 Urban flora

Urban environments are dense in impervious surfaces, which limits the available soil for plants to grow in and also generates hot and dry conditions (the urban heat island effect). Plants may grow in cracks in pavements and streets, but these surfaces are vigorously swept clean as these areas are usually situated in historical urban centres with lots of tourist attractions. Such conditions are highly unfavourable for seed germination. Plants that survive here are therefore usually heat resistant and drought tolerant (see Figure 13.1). On top of that, a relatively high percentage has ant-dispersed seeds. Typical examples include Mediterranean Spurge (*Euphorbia characias*) and Rock Fumewort (*Pseudofumaria lutea*).

cracks in pavement

Stone walls are a very hard substrate with only narrow penetrable areas in between the bricks that dry out very quickly. The stones warm up quickly in the sun and retain heat for a long time. Species such as Ivy (*Hedera* spp.) and Ivy-leaved Toadflax (*Cymbalaria muralis*), however, thrive on stone walls.

stone walls

Species in areas around ornamental trees can cope well with disturbance by trampling and eutrophication from urinating dogs. Typical species are dandelions (*Taraxacum*) (Wittig and Becker 2010; Panitsa *et al.*, 2020) and small nettle (*Urtica urens*) (Denters, 2020). Epiphytic plants found on ornamental trees often have bird- or wind-dispersed seeds. An example of the first category is Privet Honeysuckle (*Lonicera pileata*).

ornamental trees

13.3 Adaptations within urban plant populations

Shifts also occur within species growing in such environments through adaptive evolutionary changes in their traits. Compared to normal evolutionary time scales, where environmental changes happen over longer periods ranging from thousands to millions of years, the changes observed in urban environments happened over a short period of several hundred years. This demands an extremely rapid evolutionary response from species to survive in cities.

Many of such Human-Induced Rapid Evolutionary Changes (HIREC) have been reported for wild native species, particularly in animals (Lambert *et al.*, 2021). Evidence for urban plant adaptation is still limited but nonetheless strong,

Human-Induced Rapid Evolutionary Changes (HIREC)




FIGURE 13.1 Examples of plants in an urban environment. a) Bird's-foot Trefoil (*Lotus corniculatus*), b) Ivy-leaved Toadflax (*Cymbalaria muralis*), c) Kidney Vetch (*Anthyllis vulneraria*), d) White Clover (*Trifolium repens*), e) Dandelion (*Taraxacum officinale*), f) Yarrow (*Achillea millefolium*), g) Common Poppy (*Papaver rhoeas*), h) Ground Elder (*Aegopodium podagraria*), i) Wallflower (*Erysimum cheiri*)

providing hope that native species might adapt and become resilient to urban pressures. At the same time this research highlights that current plant species emerging as 'winners' in urban environments are mostly newly arrived ones that evolved elsewhere, such as introduced ornamentals and unintentional imports through global traffic. Here, we discuss evidence gathered on urban plant adaptations within populations.

The Urban Heat Island Effect (UHIE) is caused by the absorption and retention of heat by impervious stone surfaces combined with reduced albedo and the production of heat by anthropogenic sources (Rizwan *et al.*, 2008). Consequential higher temperatures have led to changed flowering times in city-dwelling plants

such as common ragweed (*Ambrosia artemisiifolia*) (Gorton *et al.*, 2018), dandelion (*Taraxacum*) (Mazumder and Kelly, 2021) and thale cress (*Arabidopsis thaliana*) (Schmitz *et al.*, 2022). A generally observed pattern is that flowering time is advanced in urban populations (Zipper *et al.*, 2016). This may be a physiological response to overall higher temperatures and increased exposure to artificial light. However, when such patterns are conserved in the same plants growing in a controlled environment, as has been tested for common ragweed and thale cress, this indicates natural selection. Earlier flowering may enable a longer growing season in which to fertilize and develop seeds and fruits, allowing the plant to invest more resources in offspring and thereby increasing its fitness. With higher temperatures, pollinating insects may emerge earlier too, providing a need for earlier flowering both from the plant and the pollinator point of view. However, this is not without its risks, as early flowering also exposes the vulnerable reproductive organs of plants to potential freezing.

An additional effect of urban heat in relation to flowering is the reduction of vernalization exposure. Plants generally need to be exposed to a continuous period of low temperatures (e.g. $\leq 4\text{ }^{\circ}\text{C}$) to activate flowering once temperatures increase again. This process is called vernalization and ensures the plant flowers in the right season when pollinating insects and sexual mates are abundant. The UHIE increases overall temperatures year round, including the winter period, and therefore causes vernalization periods to be shortened. This effect caused natural selection on reduced vernalization requirements in dandelions in the Dutch city of Amsterdam (Woudstra *et al.*, 2024), producing urban genotypes that can flower without or with reduced vernalization exposure.

Higher temperatures also present opportunities for plants as they increase the potential for growth. This can be particularly beneficial in late spring when newly germinated seedlings need to reach a decent size and build up a reserve to survive a hot summer. In dandelions, it was found that urban plant seedlings accumulate more biomass at higher temperatures ($\geq 26\text{ }^{\circ}\text{C}$) compared to rural seedlings (Woudstra *et al.*, 2024). However, extreme high temperatures (e.g. $38\text{ }^{\circ}\text{C}$, which are easily reached in urban environments during summer heat waves) are too much for most plants. Mechanisms to cope with heat stress have been revealed in creeping woodsorrel (*Oxalis corniculata*, Fukano *et al.*, 2023). Urban populations of this herb have red instead of green leaves due to a higher concentration of red pigments. Such red pigments alleviate damage induced by excessive sunlight under hot and dry conditions.

Habitat fragmentation is high in urban environments, where roads, buildings and other impervious structures seriously reduce the number of suitable places for plants to grow in. Plants therefore must make use of the limited green that is available to them on stone walls and areas around and on ornamental trees. Effective seed dispersal becomes an even more critical part of the life cycle. One simple solution is for the seed to land right next to the mother plant to ensure suitable soil to grow in.

reduction of
vernalization
exposure

higher
temperatures

habitat
fragmentation

Urban individuals of holy hawksbeard (*Crepis sancta*) in France make heavier seeds with fewer pappus hairs to enable nearby dispersal (Cheptou *et al.*, 2008).

limited space

The limited space for plants also makes the occurring plant species less abundant. The resulting lower availability of sexual mates, coupled with a lower abundance in pollinating insects, may favour asexual modes of reproduction (Van Drunen *et al.*, 2022). Although this has yet to be investigated in detail, it is already apparent that plants reproducing clonally through apomixis, such as dandelions (*Taraxacum*) and hawkweeds (*Hieracium*), or vegetative reproduction, such as bindweeds (*Calystegia* spp.) and yellow toadflax (*Linaria vulgaris*), are abundant in urban environments.

pollution

Pollution is exacerbated in urban habitats due to anthropogenic activity and the use of heavy metals in urban infrastructure. This leads to physiological responses to cope with oxidative stress. Upregulation of genes that reduce oxidative stress was found in urban populations of sweet alyssum (*Lobularia maritima*, Ben Saad *et al.*, 2023). Dandelions (*Taraxacum*) in the Italian town of Pisa were found to increase their phenolic compound production to cope with soil pollution (Vanni *et al.*, 2015). Salt spray is another form of urban pollution, particularly along roads where it often results in bare patches in the vegetation. This is expected to favour halophytic salt-tolerant plants such as particular clades of dandelions (Gilbert, 2012).

anti-herbivore
defence
strategies

Parallel to urbanization gradients, there are changes in the types and abundances of herbivorous animals. Plants adapt their anti-herbivore defence strategies to this. For instance, urban white clovers (*Trifolium repens*) reduce the production of toxic cyanogenic compounds when there are fewer herbivores around (Santangelo *et al.*, 2022). This allows the plant to invest more resources into growing and reproducing. In dandelions (*Taraxacum*), on the other hand, urban plants showed a stronger response to locust attack than rural plants by reducing early seed production (Pisman *et al.*, 2020). These particular dandelions delay their reproduction until the peak in herbivore attack in urban areas has passed.

13.4 Ecosystem services provided by urban flora

Urban plants are not just part of the urban ecosystem, as described above, they are also key to mitigating the negative effects of human presence (see also Figure 13.2). They provide shade, retain water, reduce traffic noise, prevent soil erosion, and filter air (Tan *et al.*, 2021). They also feed pollinators of crops and natural enemies of pests. In addition, urban flora increases property value and has positive psychological effects on citizens (Fineschi and Loretto, 2020; Stroud *et al.*, 2022) by for instance providing suitable outdoor areas for recreational and other social activities.



FIGURE 13.2 Impact of green design of private urban gardens on quality of living environment. Left: tiled backyard with overheated owner with irritated respiratory tract due to allergenic pollen released by ornamental olive shrubs. Right: backyard filled with non-allergenic trees, shrubs and herbs, providing shade, water retention, food to wild animals and a general feeling of well-being to owners

ILLUSTRATION BY ERIK-JAN BOSCH

13.5 Plant blindness

Unfortunately, there are still too few people aware of the huge added value of urban flora. The inability to notice plants and just take them for granted as irrelevant background is called plant blindness (Achurra, 2020). It is a dangerous affliction as it results in fully tiled front and backyards of citizens and urban designers with negative effects on the general quality of the living environment. Sufferers of plant blindness consider spontaneously spreading herbs as ‘weeds’ that need to be cleaned up, rather than as allies in our combat against heat stress and flooding. Newly arriving species are erroneously categorized as harmful to native species, whereas accumulating evidence shows that they are incorporated into local food webs (Schilthuizen *et al.*, 2016). Sufferers of plant blindness blame urban trees for taking up parking space, obstructing views, covering streets and cars with



FIGURE 13.3 Example of urban trees encouraging bird safaris. Migrating Bohemian Waxwings (*Bombycilla garrulus*) foraging for berries in a tree planted along the canal of a typical Dutch historical urban center with bird watchers enjoying the scene, while keeping a respectful distance

undesirable litter, or presenting a hazard because they can fall down during storms. Accumulating evidence shows, however, that urban trees, if properly maintained, can become very old and provide shelter and food to thousands of natural enemies of pests. They furthermore improve social cohesion in municipalities with group activities (Figure 13.3). Artificial intelligence (AI) could help in several ways to fight plant blindness and increase awareness of the added value of urban plants. Below, we will provide some examples.

13.6 How AI can help curb urban plant blindness and validate ecosystem services

plant
monitoring

Plant monitoring is mostly done by volunteers delivering their data to national species organizations such as FLORON. Data collection has traditionally focussed on nature reserves and rural areas. Collecting observations of wild plants in cities is

not a popular activity. Rather than strolling through quiet and healthy landscapes, urban naturalists need to collect their data amidst heavy traffic and exhaust fumes. Looking at plants while also keeping an eye on cars, cyclists and other pedestrians is dangerous. Furthermore, many urban plants are exotic as they are often garden escapees, requiring knowledge of horticulture on top of the native flora. For studying plant-animal interactions, urban naturalists need to go out at night as many urban animals are nocturnal, posing entirely different problems such as poor visibility. AI-based tools offer several solutions to these challenges.

First, such tools can greatly speed up identification of plants. Apps such as ObsIdentify, LeafSnap, iNaturalist, Pl@ntNet and FlowerPower can provide the name of a plant in just a few seconds. This is done by automatically running photographs against pictures in a reference database filled with images that have previously been validated by specialists. AI-based identification already works well for many plant families and species, particularly on flowering specimens. Less accurate identification of unpopular groups like grasses, vegetative or fruiting specimens, and horticultural varieties could be overcome by volunteers contributing extra pictures of these categories to expand reference datasets and further train computer algorithms. If submitting data through the VERA app, developed by FLORON, you can specify whether a species is wild, a rewilding garden escapee, or planted. Cities have the highest human population density per km² and therefore offer a huge opportunity for adding additional data needed to further improve AI-based identification of urban plants.

AI-based
identification

By organizing ObsIdentify challenges for school kids, and visitors of local natural history museums and botanical gardens, omissions in reference databases are currently being dealt with. With these challenges, awareness of urban plants is also slowly but steadily rising.

Big Data analyses are invaluable for increasing knowledge on naturalizing processes of newly arriving species in urban plant communities. Changes in species composition, flowering moments, insect visitors, and colouration can be monitored consistently over time in this way. Inventories by drones and photographs made from cars increase data collection along busy roads and other locations that are not very attractive or too dangerous for traditional plant monitoring. Diopsis insect cameras, and wildlife cameras capable of making photographs at night with infrared light help reveal new urban plant-animal interactions. Software such as Megadetector and platforms like Agouti and ARISE can scan thousands of images in just a few minutes for photographs in which animals are pollinating flowers or eating or dispersing fruits. Reference databases of these platforms and associated software are well stocked with nocturnal images of native fauna but could still benefit from more black and white images of native flora taken at night.

Big Data

Smart monitoring as described above not only helps to cure plant blindness. It also provides data in a systematic way to validate different ecosystem services and disservices of urban plants. Currently, such data collection has been focused

ecosystem
services and
disservices

on trees and a few ecosystem services and disservices only. The first mostly consisted of canopy cover, aesthetic contribution, and local nectar indices, whereas the latter mostly included allergy induction and production of undesirable organic debris blocking drainage systems (Stroud *et al.*, 2022). Collecting additional data for important but often neglected urban plant groups such as small herbs and weeds, which make up most of our street and roof flora, could help improve urban surface water quality by absorbing run-off water polluted by fine dust. Similarly, smart monitoring of aquatic urban plants can help improve urban water management by alleviating flooding. Smart monitoring also allows for combining services operating at very local scales, such as habitat provision for predators of pests, with services operating at larger scales, such as the impact of childhood green space exposure on adult psychiatric disorders, or the relationship between green space land cover and life expectancy.

13.7 Practical applications of AI-based tools for studying urban plants

By counting individuals of species upon their arrival, spread and establishment can be assessed and predicted through time. Doing so using a Big Data approach on AI-processed images could help detect potential invasive behaviour of newly arriving urban plant species and provide earlier warnings of an explosive increase in local plant distributions before it is too late for action. Reference databases of software like FlowerPower still need more metadata though, especially of the average number of flowers or inflorescences per individual, to reliably estimate the number of different plants present in individual aerial images.

To beat urban heat, green walls and roofs are increasingly being incorporated in new buildings. To optimally take care of the plants on such walls and roofs, AI-based monitoring of plant health can help save costs in watering and other maintenance, and detect new interactions among plant species, pollinators, seed predators and dispersers.

Effectiveness of planting and mowing regimes of urban vegetation can be assessed by comparing the diversity of local pollinator communities using Diopsis insect cameras (Van Klink *et al.*, 2022). Automatized identification and quantification of images of insects captured by these cameras can help identify local biodiversity hotspots. This information can help raise awareness of the added value of urban plants and assist designers and city planners in deciding where and how to stimulate urban greening and maximize ecosystem services.

13.8 Concluding remarks

In this chapter, we showed that wild plants are a vital part of urban ecosystems and are rapidly adapting to the challenges that we pose to them. They are essential allies in making urban environments healthier for both humans and urban wildlife. A wide range of ecosystem services is being provided by urban plants. So why is it still not commonplace to include wild plants in modern urban designs?

The answer lies in the lack of data proving the effectiveness of plant-produced ecosystem services and their economic justification. Smart monitoring with AI tools will provide such data. But evidence-based proof of the benefits and economic value of urban plants alone is not enough. Awareness of the added value of urban green spaces among different actors needs to be raised as well. Current perceptions of important stakeholders, ranging from homeowners to council members, vary from not seeing wild plants at all to considering them as harmful weeds. AI-based tools can help cure plant blindness, increase engagement with urban flora, and ultimately change perception of urban green spaces. Engagement with multiple stakeholders will be crucial to obtain a wider appreciation of urban plant-based solutions over building with asphalt and concrete only. Only then will we be able to keep our cities up to speed with ongoing climate change and population growth.

References

- Achurra, A., 2020. Plant blindness: a focus on its biological basis. *Frontiers in education*, 7. <https://doi.org/10.3389/educ.2022.963448>.
- Ben Saad, R., Ben Romdhane, W., Baazaoui, N., Bouteraa, M.T., Ben Hsouna, A., Mishra, A. and Čavar Zeljković, S., 2023. Assessment of the cadmium and copper phytoremediation potential of the *Lobularia maritima* thioredoxin 2 gene using genetically engineered tobacco. *Agronomy* 13:399. <https://doi.org/10.3390/agronomy13020399>.
- Cheptou, P.-O., Carrue, O., Rouifed, S. and Cantarel, A., 2008. Rapid evolution of seed dispersal in an urban environment in the weed *Crepis sancta*. In: *Proceedings of the National Academy of Sciences* 105:3796–3799. <https://doi.org/10.1073/pnas.0708446105>.
- Denters, T., 2020. *Stadsflora van de Lage Landen*. Fontaine Uitgevers, Amsterdam, The Netherlands.
- Diamond, S.E. and Martin, R.A., 2021. Evolution in cities. *Annual review of ecology, evolution and systematics* 52:519–540. <https://doi.org/10.1146/annurev-ecolsys-012021-021402>.
- Fukano, Y., Yamori, W., Misu, H., Sato, M.P., Shirasawa, K., Tachiki, Y. and Uchida, K., 2023. From green to red: urban heat stress drives leaf color evolution. *Science advances*, 9: eabq3542. <https://doi.org/10.1126/sciadv.abq3542>.

- Gilbert, O. 2012. The ecology of urban habitats. Springer Science & Business Media, Dordrecht, The Netherlands.
- Gorton, A.J., Moeller, D.A. and Tiffin, P., 2018. Little plant, big city: a test of adaptation to urban environments in common ragweed (*Ambrosia artemisiifolia*). In: Proceedings of the Royal Society B: Biological Sciences 285: 20180968. <https://doi.org/10.1098/rspb.2018.0968>.
- Johnson, M.T.J. and Munshi-South, J., 2017. Evolution of life in urban environments. Science 358:eaam8327. <https://doi.org/10.1126/science.aam8327>.
- Johnson, M.T.J., Thompson, K.A. and Saini, H.S., 2015. Plant evolution in the urban jungle. American journal of botany, 102:1951–1953. <https://doi.org/10.3732/ajb.1500386>.
- Lambert, M.R., Brans, K.L., Des Roches, S., Donihue, C.M. and Diamond, S.E., 2021. Adaptive evolution in cities: progress and misconceptions. Trends in ecology & evolution, 36:239–257. <https://doi.org/10.1016/j.tree.2020.10.011>.
- Mazumder, L. and Kesseli, R., 2021. Population structure, seasonal genotypic differentiation, and clonal diversity of weedy dandelions in three Boston area populations (*Taraxacum* sp.). Ecology & evolution, 11:10926–10935. <https://doi.org/10.1002/ece3.7870>.
- Panitsa, M., Iliadou, E., Kokkoris, I., Kallimanis, A., Patelodimou, C., Strid, A., Raus, T., Bergmeier, E. and Dimopoulos, P., 2020. Distribution patterns of ruderal plant diversity in Greece. Biodiversity and conservation, 29:869–891. <https://doi.org/10.1007/s10531-019-01915-4>.
- Pisman, M., Bonte, D. and De la Peña, E., 2020. Urbanization alters plastic responses in the common dandelion *Taraxacum officinale*. Ecology & evolution, 10:4082–4090. <https://doi.org/10.1002/ece3.6176>.
- Ritchie, H. and Roser, M., 2018. Urbanization. Published online at OurWorldInData.org. Retrieved from: <https://ourworldindata.org/urbanization>.
- Rizwan, A.M., Dennis, L.Y. and Liu, C., 2008. A review on the generation, determination and mitigation of urban heat island. Journal of environmental science, 20:120–128. [https://doi.org/10.1016/S1001-0742\(08\)60019-4](https://doi.org/10.1016/S1001-0742(08)60019-4).
- Santangelo, J.S., Ness, R.W., Cohan, B., Fitzpatrick, C.R., Innes, S.G., Koch, S., ... and Lampej, C., 2022. Global urban environmental change drives adaptation in white clover. Science, 375:1275–1281. <https://doi.org/10.1126/science.abk0989>.
- Schilthuizen, M., Santos Pimenta, L.P., Lammers, Y., Steenbergen, P.J., Flohil, M., Beveridge, N.G.P., Van Duijn, P.T., Meulblok, M.M., Sosef, N., Van de Ven, R., Werring, R., Beentjes, K.K., Meijer, K., Vos, R.A., Vrieling, K., Gravendeel, B., Choi, Y., Verpoorte, R., Smit, C. and Beukeboom, L.W., 2016. Incorporation of an invasive plant into a native insect herbivore food web. PeerJ, 4:e1954. <https://doi.org/10.7717/peerj.1954>.
- Schmitz, G., Linstädter, A., Frank, A.S.K., Dittberner, H., Thome, J., Schrader, A., Linne von Berg, K.-H., Fulgione, A., Coupland, G. and De Meaux, J., 2022. Local adaptation of life-history traits within urban populations of *Arabidopsis thaliana*. BioRxiv preprint. <https://doi.org/10.1101/2022.10.03.510679>.

- Stroud, S., Peacock, J. and Hassall, C. 2022. Vegetation-based ecosystem service delivery in urban landscapes: a systematic review. *Basic and Applied Ecology*, 61:82–101. <https://doi.org/10.1016/j.baae.2022.02.007>.
- Tan, J.K.N., Belcher, R.N., Tan, H.T.W., Menz, S. and Schroepfer, T., 2021. The urban heat island mitigation potential of vegetation depends on local surface type and shade. *Urban forestry & urban greening*, 62:127128. <https://doi.org/10.1016/j.ufug.2021.127128>.
- Van Drunen, W.E., and Johnson, M.T.J., 2022. Polyploidy in urban environments. *Trends in ecology & evolution*, 37:507–516. <https://doi.org/10.1016/j.tree.2022.02.005>.
- Van Klink, R., August, T., Bas, Y., Bodesheim, P., Bonn, A., Fossøy, F., Høye, T.T., Jongejans, E., Menz, M.H.M., Miraldo, A., Roslin, T., Roy, H.E., Ruczyński, I., Schigel, D., Schäffler, L., Sheard, J.K., Svenningsen, C., Tschan, G.F., Wäldchen, J., Zizka, V.M.A. and Bowler, D.E., 2022. Emerging technologies revolutionise insect ecology and monitoring. *Trends in ecology and evolution*, 37:872–885. <https://doi.org/10.1016/j.tree.2022.06.001>.
- Vanni, G., Cardelli, R., Marchini, F., Saviozzi, A. and Guidi, L., 2015. Are the physiological and biochemical characteristics in dandelion plants growing in an urban area (Pisa, Italy) indicative of soil pollution? *Water, air & soil pollution*, 226:124. <https://doi.org/10.1007/s11270-014-2256-7>.
- Wittig, R. and Bekker, U., 2010. The spontaneous flora around street trees in cities – a striking example for the worldwide homogenization of the flora of urban habitats. *Flora*, 205:704–709. <https://doi.org/10.1016/j.flora.2009.09.001>.
- Woudstra, Y., Kraaiveld, R., Jorritsma, A., Vijverberg, K., Ivanovic, S., Erkens, R., Huber, H., Gravendeel, B. and Verhoeven, K.J.F., 2024. Some like it hot: adaptation to the urban heat island in common dandelion. *Evolution Letters*, 8(6): 881–892. <https://doi.org/10.1093/evlett/qrae040>.
- Zipper, S.C., Schatz, J., Singh, A., Kucharik, C.J., Townsend, P.A. and Loheide, S.P., 2016. Urban heat island impacts on plant phenology: intra-urban variability and response to land cover. *Environmental research letters*, 11:054023. <https://doi.org/10.1088/1748-9326/11/5/054023>.

Adding contextual information to object detection models: a wildflower monitoring case

Georgiana Manolache and Gerard Schouten

Abstract

Understanding biodiversity trends is essential for preservation policy planning, and advanced computer vision solutions now enable large-scale automated monitoring for many biodiversity use cases. Wildflower monitoring presents unique challenges, such as inter-class similarity and intra-class variation of species. Indeed, close visual similarities in shape and colour may exist between different species, while wildflowers within a species may have significant visual differences. Moreover, flowers follow a growth cycle – from bud to fruit with a blooming stage in between – and look distinctly different over the year, while different species flower at different times of the year. Having access to flowering phenology means that more accurate predictions can be made. We propose a novel multimodal wildflower model, leveraging both high-quality, expert-annotated wildflower images and flowering phenology estimates. Moreover, we benchmark several data fusion models using two groups of common wildflowers that have high inter-class similarity and show that this multimodal approach significantly outperforms image-only baselines. With this approach, we aim to encourage the development of standards for automated wildflower monitoring as a step towards bending the curve of biodiversity loss.

14.1 Introduction

biodiversity loss

Habitat loss, pollution and climate change are the primary drivers of biodiversity loss, causing ecosystems to degrade, which in turn impacts human well-being (IPBES, 2019). We depend directly on other life forms: plants are a source of food and shelter for many animals, as well as unique sources of medicine, while trees provide building materials and act as carbon sinks. Flowering plants play a vital part in supporting ecosystems, attracting pollinators that in turn enable plants to develop seeds to produce offspring (Ollerton *et al.*, 2011). Nearly half of the world's known wildflowers are threatened with extinction according to a new study (Bachman *et al.*, 2024). The UK alone is reported to have lost 97% of its weed-rich meadows since the 1930s (Fuller, 1987). Wildflowers are under threat and therefore




FIGURE 14.1 Challenges of flower identification ‘in the wild’: 1) viewpoint variations (*Papaver rhoeas*); 2) occlusion (*Caltha palustris*); 3) clutter (*Achillea millefolium*); 4) light variation (*Leucanthemum vulgare*); 5) deformations (*Bellis perennis*); 6) intra-class variation (*Ficaria verna*); 7) inter-class similarity (*Bellis perennis*, *Leucanthemum vulgare*, *Matricaria chamomilla*)
PHOTOS ARE CROPS FROM EINDHOVEN WILDFLOWER DATASET (SCHOUTEN ET AL., 2024)

so are the pollinators they feed (Goulson *et al.*, 2015), thus impacting directly the food that we eat (Van der Sluijs and Vaage, 2016). It is crucial to understand biodiversity trends for preservation policy planning. However, due to the amount of effort and expertise required for conventional field monitoring, there are still large gaps in our knowledge. Furthermore, ad-hoc data collection in open citizen science platforms often results in biased data because of species overrepresentation and underrepresentation (Schermer and Hogeweg, 2018).

In the last few years, deep learning – with a focus on computer vision – has been playing an essential role in large-scale automated monitoring. Recent papers (Hicks *et al.*, 2021; Mann *et al.*, 2022; Schouten *et al.*, 2024) explore how object detection may be purposed for automating in-situ wildflower monitoring by identifying and counting species in images. As it turns out, wildflower monitoring combines a unique set of challenges for computer vision (Nguyen, 2016). As illustrated in Figure 14.1 these include: 1) viewpoint variations, 2) occlusion by other plants, 3) clutter, 4) light variation, 5) object deformations, as flowers are easily damaged, 6) intra-class variation, and 7) inter-class similarity. The latter two challenges are particularly hard for current models: some images of flowers in the same class may have significant visual differences, while at the same time visual similarities in shape and colour may exist between some species belonging to different classes (note that the wildflowers depicted in the last column of Figure 14.1 belong to three different species). This confusion makes it difficult even for humans to distinguish the species without deeper taxonomic expertise, and subsequently limits

deep learning

the performance of image-only wildflower monitoring models. Indeed, while the visual modality is rich and detailed, it lacks certain information available to human experts while doing fieldwork.

multimodal
approach

Eindhoven
Wildflower
Dataset
(EWD)

We therefore propose a multimodal approach that includes information on *flowering phenology* to help overcome the flower inter-class similarity and intra-class variation challenges for the Eindhoven Wildflower Dataset (EWD), an expert-annotated bird's-eye view high-resolution image dataset (Schouten *et al.*, 2024). Flowering phenology refers to the study of the timing of seasonal events in flowering plants over their growth cycle. Our approach leverages the fact that the visual characteristics as well as the presence of wildflowers at a given point in time highly depends on their estimated flowering time. In this chapter, we present a novel and large multimodal wildflower dataset and benchmark several multimodal models to encourage machine-learning research that may help us monitor, understand and ultimately preserve biodiversity.

More concretely, our contributions include:

- A benchmark of several fusion models demonstrating that multimodal approaches significantly outperform image-only methods.
- An open multimodal dataset for wildflower monitoring, uniquely combining high-quality annotated wildflower images from the Netherlands (Schouten *et al.*, 2024) with flowering phenology measurements from a public database (NDFF, n.d.).

The remainder of this paper is structured as follows. Section 14.2 discusses prior work in automated flower monitoring. Section 14.3 discusses some object detection and multimodality preliminaries, while section 14.4 defines, in formal language, the task of identifying and counting wildflowers using multimodal data. Section 14.5 details how the dataset was created. Section 14.6 describes the methodology and section 14.7 presents the benchmarks and the results. Section 14.8 discusses current limitations and section 14.9 concludes.

14.2 Flower recognition with computer vision

coding of
morphological
features

Wildflowers are intensely studied in biodiversity research since their clear visual characteristics, such as colour, shape and texture, enable large-scale monitoring (Tran *et al.*, 2018). Earlier methods to automatically identify wildflowers involve explicit coding of morphological features (colour, texture, shape) using hand-crafted image processing filters (Wäldchen and Mäder, 2015; Hong and Choi, 2012), but since the introduction of Convolutional Neural Network (CNN) architectures (He *et al.*, 2016; Krizhevsky *et al.*, 2012; Simonyan and Zisserman, 2014; Szegedy *et al.*, 2016), wildflower identification has mostly transitioned to deep learning. However, training a CNN with sufficient generalization capabilities to solve this

task requires a large, labelled dataset with significant intra-class diversity. The PlantCLEF challenge (Krishna *et al.*, 2020) targets the identification of plants (including flowers) with computer vision methods, holding over two million images of about 80,000 species. Another large-scale citizen science database for flower identification is iNaturalist (iNaturalist, n.d.), comprising over 65 million images of about 140,000 species. In these settings, however, usually a close-up of a single flower (or inflorescence) located at the centre of the image has been captured, rather than a wider image of a natural scene containing many wildflowers at the same time.

Recently, wildflower monitoring use cases have been introduced that deploy object detection to identify multiple flowers in an image, using bounding boxes (BBs) to annotate the images. Remote sensing data has also been used for a variety of downstream tasks such as detecting visiting pollinators on flowers (Tran *et al.*, 2018), detecting the blooming stages of a species (Årje *et al.*, 2019), and measuring flower phenology (Mann *et al.*, 2022). However, in these studies, there are no visually similar species within the data.

using bounding
boxes (BBs)

Other studies use datasets with a larger variety of species which do have inter-class similarity: the Oxford102 (Nilsback and Zisserman, 2008) and Jenna30 (Seeland *et al.*, 2017) datasets comprise both single and multiple flower instances in an image, annotated to investigate flower detection and classification for 102 and 30 species, respectively (Patel and Patel, 2020). Wildflower monitoring to estimate nectar sugar mass was done using images of 25 species in weed-rich grasslands in the UK (Hicks *et al.*, 2021), and flower detection from different image perspectives (entire flower, frontal, and lateral view) have been done for 10 species (Abbas *et al.*, 2022). Although not publicly available, drone-based image acquisition datasets have been collected in mountainous areas for 25 species in Switzerland (Gallmann *et al.*, 2022). In these studies, however, data and annotation quality varies wildly. For real-world images with many visually similar species, precisely labelled and annotated high-quality images are essential (Elphick, 2008; Farnsworth *et al.*, 2013). EWD offers a significantly larger collection of wildflower images comprising 160 species from various habitats in the Netherlands, with guaranteed in-situ high-quality expert-annotated images based on well-established annotation guidelines (Schouten *et al.*, 2024).

14.3 Object detection and multimodality preliminaries

Over the course of just a few years, deep learning has gained immense popularity. We briefly review the field of object detection within deep learning and its state-of-the-art model Faster R-CNN (Faster Region-based Convolutional Network) before discussing multimodal learning related work.



FIGURE 14.2 Image recognition and object detection comparison. Image recognition (first and third) labels the entire image, while object detection (second and fourth) localizes objects in an image by drawing bounding boxes around them and then labels them accordingly. Photos are crops from EWD images (Schouten *et al.*, 2024)

14.3.1 Object detection

Object detection is often confused with image recognition, the fundamental computer vision task. Figure 14.2 shows an example of the distinction between image recognition and object detection. Image recognition assigns one label to an entire image. A picture of a flower receives the label ‘flower’; a picture of three flowers also receives the label ‘flower’. Object detection, on the other hand, draws a BB around each object and labels each individual box. The model predicts where each object is and what label should be applied. In that way, object detection provides more information about an image than whole-image labelling, i.e. through its unique ability to locate objects within an image. Identifying the right object is still a classification task, localizing an object can be expressed as a regression task. A common approach to predicting a BB around an object is to predict the coordinates of the object’s centre (i.e. horizontal and vertical coordinates within the image) and height and width of the BB box; thus, four continuous numbers.

The major drawback of object detection, however, is the cost of annotating. To annotate objects’ BBs and labels for each image in a dataset is still the hardest and most time-consuming part of object detection. However, if the labels and BBs have been obtained for every object in an image, then an object detection model can be trained. Notable object detection models include Faster R-CNN (Ren, 2016), YOLO (You Only Look Once) (Redmon, 2016), or SSD (Single Shot Detector) (Liu *et al.*, 2016). Note that object detection is essential for many real-world tasks such as diagnosis (recognizing deficiencies), monitoring (counting objects in an image) and navigation (recognizing various objects in a scene and acting upon that information).

14.3.2 Faster R-CNN

Most of the current state-of-the-art object detection models are built on the groundwork laid by the Faster R-CNN model. A Faster R-CNN model has several stages, as shown in Figure 14.3. An entire image and a set of object BBs with corresponding

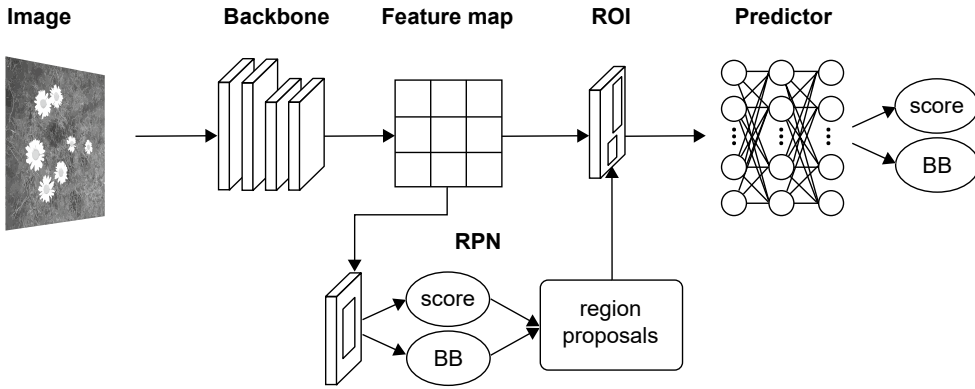


FIGURE 14.3 Diagram of the Faster R-CNN architecture

labels is passed as input. First, the image passes through the backbone dense convolutional neural network (CNN). Digital images consist of a series of numbers, arranged in a 3-layered 2D array (one layer for ‘red’, one layer for ‘green’, and one layer for ‘blue’). This is how a computer perceives an image, and CNNs build on this idea. The idea is that, given an array of numbers, the network is able to use low-level features such as edges and curves to generate abstract high-level features in the later convolutional layers. After passing through several convolutional and pooling layers (that basically perform filter operations and downsample), a feature map is generated. The feature map is a spatially dense tensor that represents the learned features of the image.

convolutional neural network (CNN)

Using the feature map and the image’s annotated bounding boxes, the Region Proposal Network (RPN) predicts whether there is an object or not and the bounding box of those objects. Initially, RPN generates anchor points at regular intervals over the image, which will act as a prior to predict object proposals. Predefined anchor boxes of different sizes and aspect ratios are placed at the centre of each anchor point. The network refines these anchor boxes throughout training to better match actual object positions and sizes. A basic geometrical concept for evaluating how well an object detection model predicts BBs is the Intersection over Union (IoU). This concept measures the area overlap between the predicted BB and the annotated or target BB, divided by the area of their union. An IoU close to the maximum value of 1 is considered very good, whereas values close to 0 mean both bounding boxes nearly do not intercept each other. For each anchor, the RPN predicts two set of parameters, namely the probability of the anchor containing an object (also known as an objectness score) and adjustments to the anchor boxes’ coordinates to match the actual object’s shape (four scores, one for each box corner).

Region Proposal Network (RPN)

Intersection over Union (IoU)

The region proposals are of different sizes, so a technique called Region of Interest (ROI) pooling is used to resize them before passing them through the Faster R-CNN network. The network flattens the proposed region feature map

Region of Interest (ROI)

into a fully connected layer and learns to label categories using the cross-entropy loss and ground truth boxes using the least square errors (L2) regression loss. The cross-entropy loss (or log loss) measures the performance of a classification model whose output is a probability value between 0 and 1. The L2 regression loss minimizes the error, which is the sum of all the squared differences between the true value and the predicted value.

14.3.3 Evaluation metrics explained

Object detection models are therefore more complex since they combine both location prediction and object classification. So, what happens if the model detects the correct class but at the wrong location? One approach is to define an IoU threshold. The object detection evaluation metrics need to know what a ‘correct detection’ and an ‘incorrect detection’ are. By setting an IoU threshold, the user can classify a detection as correct or incorrect. For example, we may consider that a prediction is correct only if the IoU is greater than a threshold of 0.5 and the predicted class is correct. The evaluation metrics can then be more or less restrictive as thresholds closer to 1 require almost perfect detections, whereas a threshold near 0 accepts even poor detections.

mean average
precision
(mAP)

mean average
recall (mAR)

The standard metrics used in the object detection tasks are the mean average precision (mAP) and mean average recall (mAR). These metrics are derived from the confusion matrix by counting correct and wrong classifications, according to the chosen IoU threshold (see Figure 14.9 for an example). Each detected BB belongs to one of the four components of the confusion matrix and is classified per class observation:

- True Positive (TP): a correct detection of the observed ground truth class.
- False Positive (FP): an incorrect detection of another class or background.
- False Negative (FN): an undetected bounding box of the observed ground truth class.
- True Negative (TN): a correct detection of no class/background (no bounding box placed).

Precision measures the accuracy of positive predictions, i.e. $TP/(TP+FP)$, while recall measures the ratio of positive predictions that are correctly detected by the classifier, i.e. $TP/(TP+FN)$. Precision and recall values range from 0 (lowest value) to 1 (best value). Ideally, both precision and recall should be high. The trade-off between precision and recall can be visualized in a so-called Precision-Recall (PR) curve.

Average precision (AP) shows the model’s performance by taking the area under the PR curve and computing the average precision at different recall values (i.e. 0, 0.1, ..., 1) for a given IoU threshold. In multi-class object detection (i.e. there are more than two classes), the AP is computed for each class and then the mAP is computed. The mAR is computed similarly, except that for average recall (AR) we average recall values over several IoU thresholds (i.e. [0.5, 0.95] with a 0.05 step), then compute mAR in multi-class object detection.

14.3.4 *Data fusion and multimodality*

Current computer vision research that is beyond state of the art involves multi-modal input data. Modalities can be divided into four main groups (Pawłowski *et al.*, 2023):

- *Tabular data*: observations are stored as rows and their features as columns.
- *Graphs*: observations are vertices, and their features are in the form of edges between individual vertices.
- *Signals*: observations are files of appropriate extension (images:.jpeg, audio:.wav, etc.) and their features are the numerical data provided within files.
- *Sequences*: observations are in the form of characters/words/documents, where the type of character/word corresponds to features.

Data fusion strategies – combining several of these modalities – can be categorised as *early fusion* and *late fusion* (Huang *et al.*, 2020).¹ The schemas of early and late fusion are depicted in Figure 14.4. Early fusion, also denoted as data-level or feature-level fusion, simply joins data or features into common feature space, usually by concatenation. Late fusion, or decision-level fusion, joins prediction results from unimodal models (such as probability logits or categorical outcomes in classification tasks) to achieve multimodal prediction. In this study, we focus on early feature-level fusion techniques as they usually surpass the unimodal counterparts of late fusion and offer more flexibility in the AI-pipeline (Cui *et al.*, 2023; Stahlschmidt *et al.*, 2022).

data fusion strategies




FIGURE 14.4 Data fusion techniques: (left) early fusion, and (right) late fusion

¹ Huang *et al.* (2020) also discern intermediate fusion. This is in essence similar to early fusion, but instead of just joining original data features, it learns high-dimensional features (typically through a neural network) before joining them into the common feature space.

In this study, we only leverage from one non-image modality. The common practice is to perform simple operations of concatenation, element-wise summation or element-wise multiplication, with the last two operations requiring feature vectors of different modalities to be converted into the same shape. It is difficult to compare the performance of different operation-based feature-level fusions, since different studies were done on different data with different settings. However, in general multimodal solutions significantly outperform solutions using a single modality, but it might yield inferior performance when learning complex interactions of different modalities.

Most multimodality studies are done in the domain of medical image analysis. To give some examples: Holste *et al.* (2021) compare these operation-based methods joining medical imaging with non-image data medical records for breast cancer classification. They extend the ResNet50 (He *et al.*, 2016) architecture with early fusion strategies. Their multimodal models outperform the unimodal model, with the early fusion strategy achieving superior performance. Chen *et al.* (2019) propose a multimodal architecture which is based on Faster R-CNN for cervical cancer classification. While they also fuse features from two types of medical images, they fuse non-image features extracted from clinical tests with learned image features by using one fully connected layer.

14.4 Task definition

Figure 14.5 provides an overview of the learning task we aim to solve. First, the flowering phenology is a graphical representation of the timing and duration of seasonal flowering for a species, with time (e.g. days, weeks, months) on the horizontal axis and a phenology estimate on the vertical axis (Haggerty and Mazer, 2008). Given a set of images, each with a creation date d , and a set of possible species s_1, \dots, s_n that we aim to detect, we extract the flowering phenology estimates for that point in time: $p_{s_1}^d, \dots, p_{s_n}^d$. We then aim to train a model m that, for each detected flower in the image, receives the (pre-trained) image features z_1, \dots, z_m and flowering




FIGURE 14.5 An overview of our multimodal object detection solution

phenology estimates $p_{s_1}^d, \dots, p_{s_n}^d$ as input, and predicts the vector $y = (y_{s_1}, \dots, y_{s_n})$ of probabilities for each possible species.

We aim to use annotated images and flowering phenology estimates in conjunction to classify species; thus, our task can be considered a supervised multimodal object detection problem. We use a pre-trained object detection architecture for image feature extraction and localization and attempt early feature fusion by concatenating non-image data (i.e. tabular data $p_{s_1}^d, \dots, p_{s_n}^d$) within the Faster R-CNN architecture. Our goal during training is to reduce misclassifications of visually similar species.

14.5 Dataset

Our dataset is constructed using data from two open data resources: wildflower images from the Eindhoven Wildflower Dataset (EWD) (Schouten *et al.*, 2024), an expert-annotated high-resolution image dataset, and flowering phenology estimates from the Dutch National Database Flora and Fauna (NDFD, n.d.). Both datasets are from the Netherlands. To the best of our knowledge, it is the first multimodal dataset allowing fine-grained differentiation of a set of visually similar wildflowers using phenology. Similar approaches consider spatio-temporal location data (Mac Aodha *et al.*, 2019; De Lutio *et al.*, 2021). However, their temporal information is limited by the presence-only data, which inherently lacks information about species growth cycle.

14.5.1 Image dataset

We fully leverage the image and annotation quality of EWD, which contains top-view flowering plant images with expert-annotated BBs around individual flowers. The images were collected in-situ from flower beds of approximately 1 m² taken (near-)vertically downwards at a height ranging from 1.5 to 1.9 m, in the region of Eindhoven, the Netherlands, over the years 2021 and 2022. Annotations are formatted in Pascal VOC (Visual Object Classes), containing BBs with their corresponding species labels stored in a human readable XML format. There are over 65,000 annotations for 160 species with a long-tailed species distribution.

Since the EWD images are significantly large (6720 × 4480 pixels) and today's most advanced Faster R-CNN architectures have a maximum input size that is significantly lower (1,333 pixels on either axis) (Li *et al.*, 2021), we slice the original images into 15 (5 × 3) tiles. Tiles without object annotations are disregarded. We employ a cutting approach that is aware of the presence of annotations in the image and attempts to find the least damaging way to cut it. A helper function is employed to divide the input images into slices that approximate the target size. Figure 14.6 shows an example of image sliced into 15 tiles with this heuristic approach. We train, validate and test models with image tiles in our experiments.



FIGURE 14.6 An example of an EWD image sliced into tiles, taken from Schouten *et al.* (2024). The dashed lines show equally sized tiles. Note the difference in the number of cut wildflowers (*Calyta palustris* in this case) between the two tiling schemes

14.5.2 Flower phenology data

We collect flowering phenology data from NDFE (NDFE, n.d.). Flowering phenology is a graphical representation of the timing and duration of blooming for a particular flowering plant over a specific period (Haggerty and Mazer, 2008). The data is modelled based on observations from the period 2000–2021. Observations are collected by volunteers and professionals and validated by experts. Only approved observations (over 2 million in total!) are included. Observation counts from each year are converted into a circular format and then averaged. The peak and median are then calculated from the averaged circular density data, as well as the day numbers on which 10%, 20%, 80% and 90% of the observations are made. A box plot is then generated from this data for each species. The flowering phenology is therefore determined on the basis of a percentile value (Van der Hak, 2022).

14.5.3 Data alignment

The phenology data is aligned with the image data based on the exact time that the image was taken, as described in section 14.4. To analyse whether the phenology data matches the wildflowers in the images, we selected two groups of common wildflowers; each group contains three species with high inter-class similarity. Figure 14.7 shows a sample of the selected species, while Table 14.1 describes the details of the two groups.




FIGURE 14.7 Selected flower species grouped by visual similarity. Group 1: Buttercup (aggregate), *Caltha palustris*, *Ficaria verna*; Group 2: *Bellis perennis*, Chamomile (aggregate), *Leucanthemum vulgare*. Photos are crops randomly sampled from EWD

TABLE 14.1 Selected flower species dataset description. Subspecies visually indistinguishable in the field are merged in the EWD dataset: *Ranunculus acris* and *Ranunculus repens* are labelled as Buttercup (aggregate), while *Matricaria chamomilla* and *Matricaria maritima* are labelled as Chamomile (aggregate)

	Species	Objects	Images	Image tiles	Dates
Group 1	<i>Buttercup (aggregate)</i>	4192	190	1188	Apr 23,26,27; May 9,12,13,14,15,16,17,18,19,20,21,24,28,29,30; Jun 2,4,5,6,7,9,10,17,18,26; Jul 9,26, Aug 18
	<i>Caltha palustris</i>	1201	50	432	Mar 23,24; Apr 3,4,6,9,10,14,15,23,26; May 9,12,21
	<i>Ficaria verna</i>	817	31	208	Mar 16,23,24; Apr 2,14
Group 2	<i>Bellis perennis</i>	1429	22	213	Mar 16; Apr 2,8; May 1,9,15,19,21,28; Jun 4; Aug 4,26
	<i>Chamomile (aggregate)</i>	2202	45	369	Apr 21; Jun 1,4,5,6; Jul 2,3,26,28; Aug 6,18
	<i>Leucanthemum vulgare</i>	1424	48	327	May 10,11,14,16,28; Jun 1,5,7, 9,19; Aug 18; Sep 22

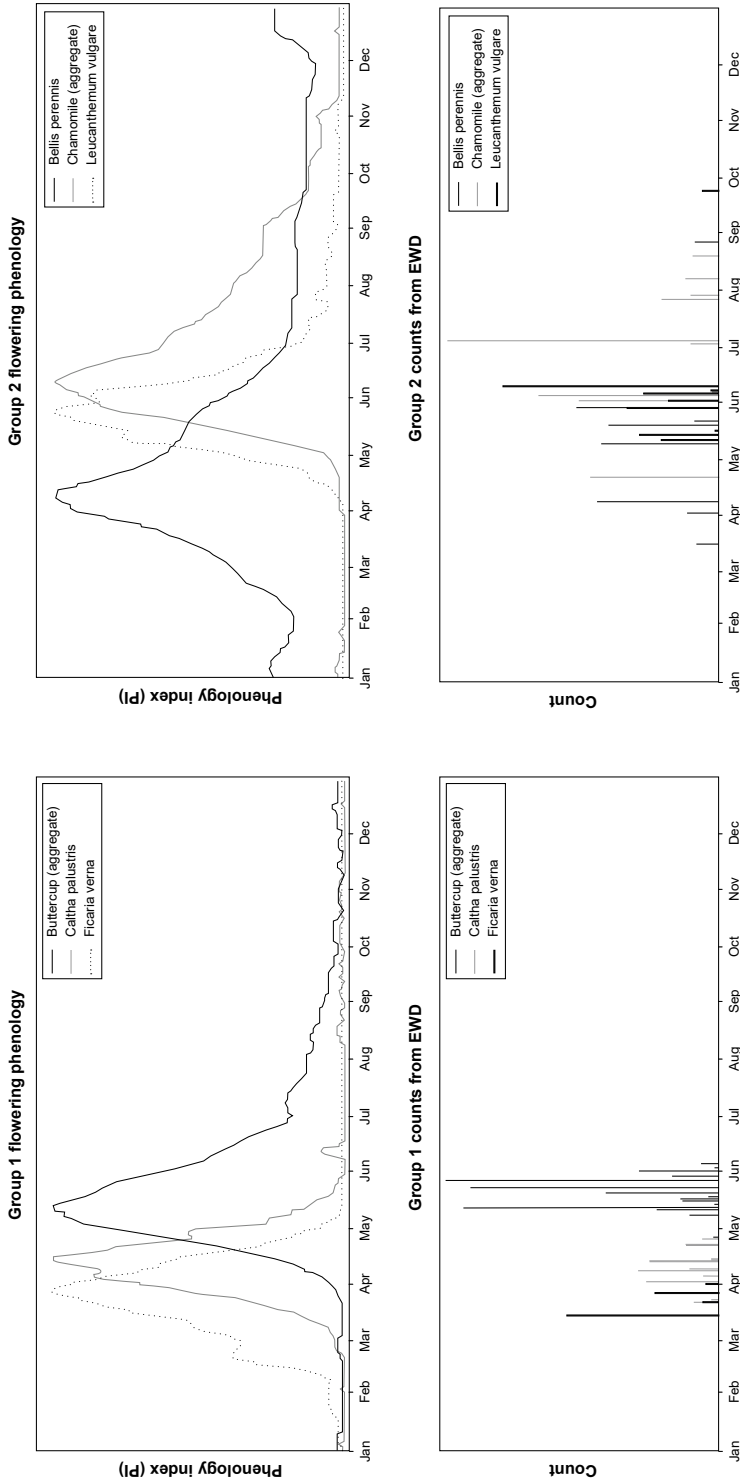


FIGURE 14.8 Data alignment overview for both groups: (top) flowering phenology estimates from NDFE for the selected species and (bottom) histogram of objects counts from EWD for the selected species. The horizontal axis is the day of year ranging from 1 to 365, while the vertical axis is (top) the phenological index, normalized from 0 to 1, and (bottom) an object count

Figure 14.8 shows an overview of the flowering phenology (top) and histograms of the distributions by day of the year (bottom) for each species in each group. All species are denoted by their botanical name. Species in Group 1 belong to the *Ranunculaceae* family with yellow flowers in blooming stage, while species in Group 2 belong to the *Asteraceae* family with flowerheads (white radial flowers and yellow disc flowers in the centre) in blooming stage. We can indeed observe different peaks for each species. The merged species labelled Buttercup (aggregate) have statistically similar flowering phenology estimates. *Matricaria chamomilla* and *Matricaria maritima* have different estimates, with the latter recurring all year round. We select the *Matricaria chamomilla* flowering phenology estimates for the merged species Chamomile (aggregate). Species objects are counted from 2021 and 2022 images per day. More images of Buttercup (aggregate) and Chamomile (aggregate) have been collected in a single day. The observations and flowering phenology peaks correlate in Group 1. In Group 2 there is more overlap for observations and flowering phenology peaks.

14.5.4 Dataset splits

To avoid bias, we randomly selected images such that the total object amount matches that of the training-validation-test sets over all species. In object detection, a single image can contain observations of multiple objects as well as different types of objects, making it challenging to create a balanced dataset. If an image is selected because it contains a specific wildflower, it may also include other wildflowers incidentally. By providing the sample sizes for the training, validation and test sets, theoretically all permutations that achieve the exact desired numbers can be computed. However, this approach becomes exponentially time-consuming as the dataset grows larger. To address this issue, we use a trial-and-error approach, making numerous attempts and stopping early when a solution is found. With this approach, we randomly select images amounting to 550 objects for the training sets, 100 objects for validation set, and 50 objects in test set for each selected species.

14.6 Methodology

14.6.1 Models

14.6.1.1 Baselines

In line with the existing literature on flower monitoring (Gallmann *et al.*, 2022; Hicks *et al.*, 2021; Schouten *et al.*, 2024), we propose using the Faster R-CNN object detection model as our main image-only baseline. We compare this with other state-of-the-art object detection models such as SSD and YOLOv8. Furthermore, we test the vision capabilities of the state-of-the-art Multimodal Large Language Model (MLLM) with vision (GPT-4v) (OpenAI, n.d.) and LLaVA-1.5 (Liu *et al.*, 2023), which we do not fine-tune. We also investigate the classification performance from

Faster R-CNN
object detection
model

Multimodal Large
Language Model
(MLLM)

just the phenology features with XGBoost classifier (Chen and Guestrin, 2016), as well as from just image classification, using ResNet-50 (He *et al.*, 2016).

14.6.1.2 Feature fusion

We proposed feature-level fusion combining features from images and flowering phenology into a single high-dimensional feature vector using concatenation (Cui, *et al.*, 2023). To this end, we extended the Faster R-CNN architecture. Faster R-CNN works by first extracting feature vectors from a backbone CNN, after which the RPN (Region Proposal Network) generated region proposals in the image and a fixed-length feature vector for each region, which was in turn fed into successive fully connected layers, followed by two outputs: the species classification head and the BB location regression head. The classifier produced probability values of each proposed object belonging to n categories and one catch-all background category. The regressor head output four offsets (x , y , h , w) from the RPN more precisely, where (x, y) specified the values of the position of the left corner, and (h, w) the height and width of the window. The flowering phenology was passed into a 1-dimensional vector of length n in the classifier. The values of the flowering phenology feature vector were extracted from the flowering phenology graphs at a given image creation date and passed as input alongside the image. We concatenated feature vectors from each Faster R-CNN proposed region with the flowering phenology vector before feeding it to the classifier. During training, we computed the classification loss using cross-entropy. We used the image features from generated proposals for the BB regressor.

14.6.1.3 Learned feature fusion model

We also proposed to simultaneously learn features from flowering phenology, then combined them with the image features. Although simple and effective, feature concatenation might not exploit the complex correlation between the heterogeneous modalities (Cui *et al.*, 2023). We used the flower phenology 1-dimensional vector of length n categories to learn a vector of the same length as the image features, which we then combined. Thus, we processed the low-dimensional flowering phenology features by fully connected layers to the same dimension of image features before fusion. To understand how different fusion operations impacted predictive performance, we experimented with three fusion operations: 1) concatenation, 2) element-wise multiplication and 3) element-wise addition (Holste *et al.*, 2021).

14.6.2 *Experimental setup*

We trained separate models for each group for 30 epochs. We used five seeds for training-validation-test sets. Each experiment was run on training-validation-test sets from each seed and results reported the average. We only inferred the test set on the MLLMs. Training, validation and test sets followed the splitting described in Section 14.5. For the Faster R-CNN models the image features are learned with

ResNet-50 (He *et al.*, 2016) with pre-trained weights, which performed best in wild-flower detection studies (Patel and Patel, 2020; Schouten *et al.*, 2024). We used the open-source PyTorch framework and the Faster-RCNN (Ren, 2016) and SSD (Liu *et al.*, 2016) model, as provided in the TorchVision library. We used the Ultralytics framework for the YOLOv8 model. We leveraged the OpenAI API to utilize GPT-4v, the vision-enabled variant of the GPT-4 model and OllamaChat integrated with LLaVA v1.5, for the advanced image processing and analysis task. The Faster R-CNN classifier ResNet50 was fine-tuned for the image-only classification. Finally, for the phenology-only classification, we leveraged from the XGBoost library.

14.6.3 Metrics

We measured object detection performance using COCO metrics (Lin *et al.*, 2015), i.e. we reported the AP per class and the mAP over all classes at 0.50 and [0.50,0.95] interval with 0.05 step IoU threshold (see also Section 14.3.3). We also reported the coefficient matrix with confidence threshold over 0.75 and IoU threshold over 0.50. We reported the precision scores for the classification models after training. All scores were averaged over the five seeds. In addition, we used the non-parametric paired sample t-test to determine whether the different values across performance scores from each seed were statistically significant (at a significance level of 0.05) or occurred by chance.

14.7 Results

Table 14.2 offers a holistic view of each model performance. We observe that fusion models outperform image-only models for both groups. This highlights the value of flowering phenology for our task. Precision increases significantly with additional features for the learned feature fusion with elementwise concatenation and addition variants ($p < 0.1$). The learned feature fusion method likely captures more nuanced relationships between learned features from both modalities and optimizes the combination process, leading to improved performance compared to a simple concatenation approach. The variant using element-wise addition fusion had the best predictive performance, which could imply that the learned image and flowering phenology features were relevant and complementary.

Table 14.3 shows the AP scores per class for statistically different fusion models. Buttercup (aggregate) and Chamomile (aggregate) species have significantly higher AP scores in the learned feature fusion variants, which may indicate that there are significantly fewer missed detections and/or misclassifications. However, these scores describe the performance of object detection models based on both the accuracy of object localization and the ability to detect all instances of objects, and do not directly capture misclassifications. To better evaluate misclassification, Figure 14.9 shows the performance scores for each species in a confusion

TABLE 14.2 Test results on Group 1 and Group 2 for all models averaged over five seeds. Best results are shown in bold

		mAP	
	Model	@.5IoU	@[.50:.95]IoU
Group 1	SSD	0.20 ± 0.04	0.09 ± 0.02
	YOLOv8	0.70 ± 0.05	0.51 ± 0.03
	Faster R-CNN	0.77 ± 0.07	0.59 ± 0.06
	Feature fusion w/Faster R-CNN	0.88 ± 0.04	0.67 ± 0.04
	Learned feature fusion concatenation w/Faster R-CNN	0.89 ± 0.04	0.68 ± 0.04
	Learned feature fusion multiplication w/Faster R-CNN	0.89 ± 0.04	0.68 ± 0.04
	Learned feature fusion addition w/Faster R-CNN	0.90 ± 0.03	0.68 ± 0.04
Group 2	SSD	0.24 ± 0.05	0.13 ± 0.04
	YOLOv8	0.66 ± 0.11	0.48 ± 0.07
	Faster R-CNN	0.67 ± 0.11	0.48 ± 0.07
	Feature fusion w/Faster R-CNN	0.69 ± 0.13	0.51 ± 0.10
	Learned feature fusion concatenation w/Faster R-CNN	0.80 ± 0.07	0.58 ± 0.05
	Learned feature fusion multiplication w/Faster R-CNN	0.74 ± 0.14	0.54 ± 0.08
	Learned feature fusion addition w/Faster R-CNN	0.81 ± 0.08	0.68 ± 0.04

matrix averaged over five runs. We report baseline and elementwise addition learned feature fusion variant at confidence score over 0.75 and IoU threshold over 0.50. Interestingly, there are fewer missed detections in the fusion model at a 0.75 confidence. This may indicate that more robust and generalizable features are learned with the addition of flowering phenology. There are also significantly fewer misclassifications in Group 1. There are slightly more misclassifications between *Bellis perennis* and *Leucanthemum vulgare* in the multimodal variant. This is primarily caused by the training-validation-test splits which did not take into account image creation date distribution. There are more *Bellis perennis* image samples at the flowering phenology peak of *Leucanthemum vulgare*. In Group 1, data distribution for each species matches flowering phenology peaks, hence the significant class differentiation. Thus, we suggest balancing training, validation and test set splits also on dates.

TABLE 14.3 Test results per species class for the image-only baseline and the best performing feature fusion models averaged over five seeds

		mAP		
	Model	@.50IoU	@.75IoU	@[.50:.95]IoU
Buttercup (aggregate)	Image-only	0.68 ± 0.09	0.60 ± 0.09	0.50 ± 0.06
	Learned feature fusion concatenation	0.84 ± 0.03	0.78 ± 0.04	0.62 ± 0.02
	Learned feature fusion addition	0.83 ± 0.02	0.78 ± 0.03	0.62 ± 0.01
<i>Caltha palustris</i>	Image-only	0.75 ± 0.05	0.68 ± 0.04	0.57 ± 0.03
	Learned feature fusion concatenation	0.92 ± 0.05	0.83 ± 0.09	0.68 ± 0.05
	Learned feature fusion addition	0.93 ± 0.03	0.82 ± 0.07	0.69 ± 0.05
<i>Finca verna</i>	Image-only	0.88 ± 0.06	0.83 ± 0.10	0.70 ± 0.08
	Learned feature fusion concatenation	0.93 ± 0.03	0.86 ± 0.05	0.73 ± 0.05
	Learned feature fusion addition	0.93 ± 0.03	0.89 ± 0.06	0.73 ± 0.06
<i>Bellis perennis</i>	Image-only	0.75 ± 0.08	0.62 ± 0.04	0.52 ± 0.03
	Learned feature fusion concatenation	0.84 ± 0.08	0.68 ± 0.06	0.58 ± 0.03
	Learned feature fusion addition	0.85 ± 0.06	0.72 ± 0.07	0.59 ± 0.03
Chamomile (aggregate)	Image-only	0.57 ± 0.13	0.46 ± 0.10	0.40 ± 0.09
	Learned feature fusion concatenation	0.77 ± 0.08	0.64 ± 0.08	0.55 ± 0.06
	Learned feature fusion addition	0.78 ± 0.11	0.65 ± 0.08	0.56 ± 0.06
<i>Leucanthemum vulgare</i>	Image-only	0.68 ± 0.13	0.61 ± 0.11	0.52 ± 0.10
	Learned feature fusion concatenation	0.80 ± 0.05	0.70 ± 0.08	0.61 ± 0.05
	Learned feature fusion addition	0.79 ± 0.06	0.69 ± 0.09	0.61 ± 0.06

We also experimented with MLLMs, but we found lower performance for the task of object detection, as shown in Table 14.4. Recent inclusion of vision in MLLMs shows great potential indeed. However, the wildflower detection use case is still challenging for these models. While the largest MLLM, GPT-4v, can now identify and provide information about objects within images, it is still limited in object detection functionality. According to the OpenAI guide on GPT-4v, the model may give approximate counts for objects in images. After conducting some tests on GPT-4v, we found that the GPT-4v API and the GPT-4v web app version returned coordinates when given a direct prompt, but the coordinates were not correct. In both versions, just requesting counts worked better. However, when feeding the full resolution (6720×4480 pixels) EWD images, the model hallucinated. When feeding image tiles the counts were better approximated, but not for all images. We tested with the test datasets from each group. The current behaviour therefore shows that the model is capable of object detection but does not perform well. Recent approaches use the ‘classic’ computer vision model in tandem with GPT-4v as an object detection model. Thus, GPT-4v alone in its current state is not near object detection state of the art.

Open-source LLaVA is less efficient than GPT-4v as an object detection model. Similar to GPT-4v, when asked to generate the bounding box coordinates, the model seemed to hallucinate. Furthermore, the model is reluctant to give counts when flowers are cluttered and smaller. The model LLaVA also fails to correctly identify the visually similar species, which was expected as LLaVA is smaller than GPT-4v. However, unlike GPT-4v, the model can be fine-tuned. Ultimately, LLaVA is limited by hallucinations and weak in-depth reasoning common to many LLMs.

TABLE 14.4 Test results on Group 1 and Group 2 for all classification models averaged over five seeds, hence average precision (AP)

	Data	Model	AP
Group 1	Image	GPT-4v*	0.68 ± 0.11
	Image	LLaVA-1.5v*	0.27 ± 0.04
	Image	ResNet50	0.84 ± 0.02
	Phenology	XGBoost	1.00 ± 1.00
Group 2	Image	GPT-4v*	0.70 ± 0.19
	Image	LLaVA-1.5v* ResNet50	0.01 ± 0.01
	Image	XGBoost	0.86 ± 0.03
	Phenology		1.00 ± 1.00

* Models are not fine-tuned, only inferred on the test set.

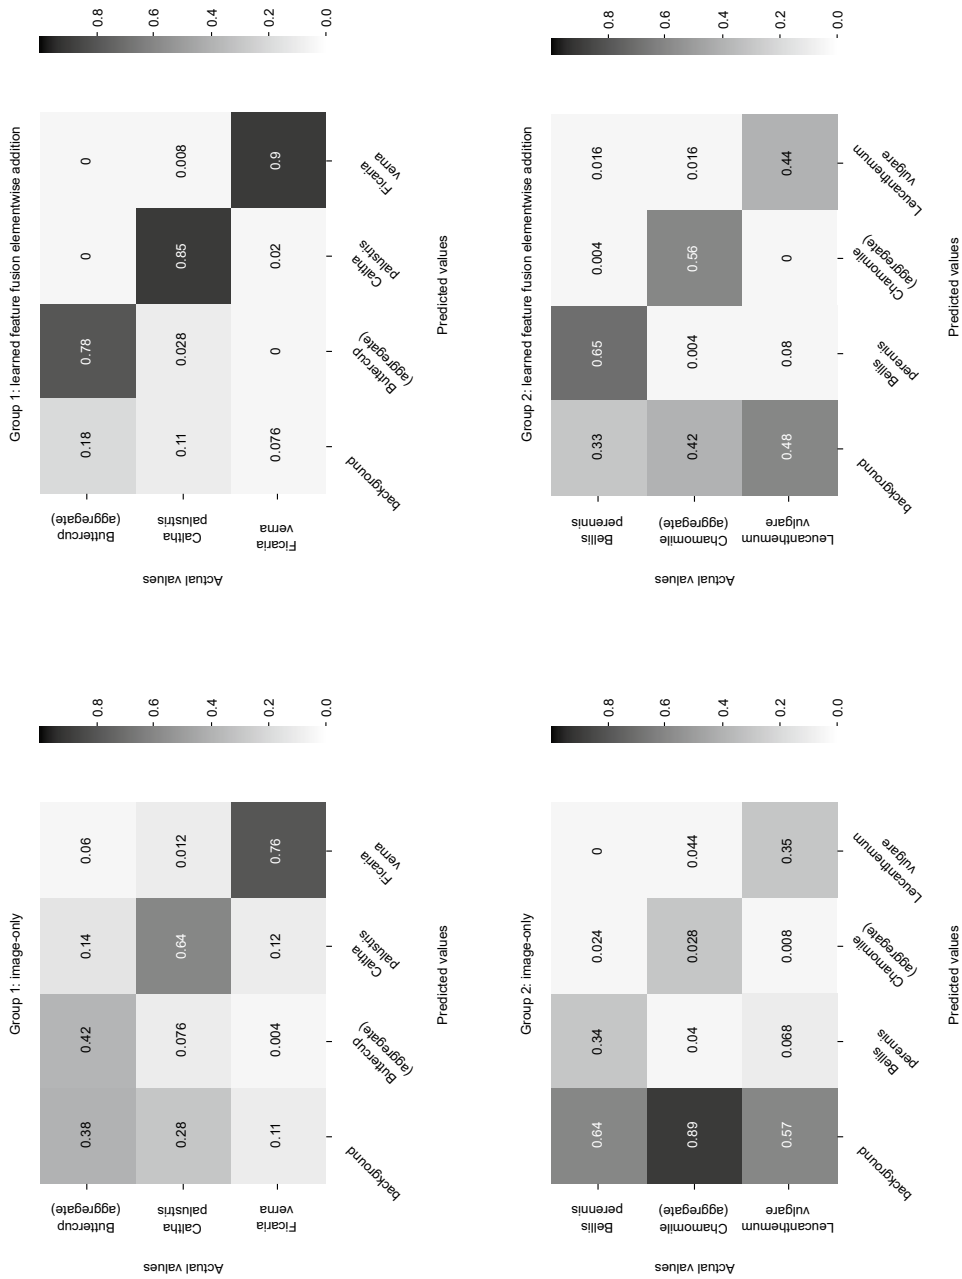


FIGURE 14.9 Confusion matrix with confidence threshold over 0.75 and IoU threshold over 0.50 for image-only and learned feature-level fusion elementwise addition models

We also report the classification performance on the selected species. We used the same train-validation-test splits and trained two classifiers, using images only and phenology only. ResNet50 was trained on 30 epochs. We also inferred state-of-the-art MLLMs on the test set. We feed the box images, not the entire image, to the classifiers and MLLMs. The classifier using phenology always predicts the correct class. Training an image classifier is also more efficient with higher performance scores. Object detection models are evaluated using metrics like mean mAP, which consider both classification accuracy and localization. These metrics are more stringent than those used for classification alone (such as accuracy or precision-recall), and as a result, they tend to reflect lower precision scores.

14.8 Limitations

Firstly, flowering phenology estimates do not take into account recent rapid climate change trends, modelling an average of all observations over the last 21 years. In our study we use image data from 2021 and 2022. Generally, climate change will advance the timing of seasonal events for the majority of flowering plants, which is already well documented (Geissler *et al.*, 2023). The flowering phenology estimates could perhaps be improved based on annual weather measurements for a better representation of the species encounter. Furthermore, reliable flowering phenology estimates may not be publicly available for species worldwide. Nonetheless, monthly flower counts from citizen science platforms – like iNaturalist or the Dutch platform waarneming.nl – may be used to compute flowering phenology estimates similarly (Van der Hak, 2022). Subsequently, the image dataset has imbalanced sample dates. This is a limitation shared with other datasets, and a common issue in in-situ data collection.

Furthermore, the dataset may not include sufficient images with different backgrounds, lighting conditions, and object poses. Data augmenting techniques such as random flipping, addition of noise, blur, contrast and brightness shifts can be leveraged to improve the robustness of the models (Rebuffi *et al.*, 2021). Another limitation of our study might be the small training setup. However, training with more image data may not always improve model performance (Zhu *et al.*, 2011). Nevertheless, the results of the multimodal models are still remarkable. Finally, the limited amount of annotated data currently available for object detection purposes related to wildflowers still proves to be a significant obstacle in current computer vision research.

14.9 Conclusion

We propose a multimodal dataset and benchmark for wildflower monitoring using their flowering phenology estimates. The multimodal dataset includes high-quality

annotated wildflower images and flowering phenology estimates from the Netherlands. We detail the process of collecting and using the flowering phenology estimates. We also present results for the dataset on a range of data fusion variants. Extensive experiments have corroborated the effectiveness of our multimodal approach in reducing misclassification. We are planning a next version of the dataset that will include flowering phenology that accounts for changes in climate over time. As this work is intended to directly impact wildflower monitoring, we hope such input will be valuable to researchers seeking to understand biodiversity and climate change, as well as policymakers interested in evaluating preservation priorities across different areas of land.

References

- Abbas, T., Razzaq, A., Zia, M.A., Mumtaz, I., Saleem, M.A., Akbar, W., ... and Shivachi, C.S., 2022. Deep neural networks for automatic flower species localization and recognition. *Computational intelligence and neuroscience*, 2022:9359353.
- Årje, J., Milioris, D., Tran, D.T., Jepsen, J.U., Raitoharju, J., Gabbouj, M., ... and Høye, T.T., 2019. Automatic flower detection and classification system using a light-weight convolutional neural network. In: *EUSIPCO Workshop on Signal Processing, Computer Vision and Deep Learning for Autonomous Systems*.
- Bachman, S.P., Brown, M.J., Leão, T.C., Nic Lughadha, E. and Walker, B.E., 2024. Extinction risk predictions for the world's flowering plants to support their conservation. *New phytologist*, 242:797–808.
- Chen, T. and Guestrin, C., 2016. Xgboost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 785–794.
- Chen, T., Ma, X., Ying, X., Wang, W., Yuan, C., Lu, W., ... and Wu, J., 2019. Multi-modal fusion learning for cervical dysplasia diagnosis. In: *IEEE 16th international symposium on biomedical imaging (ISBI 2019)*, pp. 1505–1509.
- Cui, C., Yang, H., Wang, Y., Zhao, S., Asad, Z., Coburn, L.A., ... and Huo, Y., 2023. Deep multi-modal fusion of image and non-image data in disease diagnosis and prognosis: a review. *Progress in biomedical engineering*, 5:022001.
- De Lutio, R., She, Y., D'Aronco, S., Russo, S., Brun, P., Wegner, J.D. and Schindler, K., 2021. Digital taxonomist: identifying plant species in community scientists' photographs. *ISPRS journal of photogrammetry and remote sensing*, 182:112–121.
- Elphick, C.S., 2008. How you count counts: the importance of methods research in applied ecology. *Journal of applied ecology*, 45:1313–1320.
- Farnsworth, E.J., Chu, M., Kress, W.J., Neill, A.K., Best, J.H., Pickering, J., ... and Ellison, A.M., 2013. Next-generation field guides. *BioScience*, 63:891–899.
- Fuller, R. 1987. The changing extent and conservation interest of lowland grasslands in England and Wales: a review of grassland surveys 1930–1984. *Biological conservation*, 40:281–300.

- Gallmann, J., Schüpbach, B., Jacot, K., Albrecht, M., Winizki, J., Kirchgessner, N. and Aasen, H., 2022. Flower mapping in grasslands with drones and deep learning. *Frontiers in plant science*, 12:774965.
- Geissler, C., Davidson, A. and Niesenbaum, R.A., 2023. The influence of climate warming on flowering phenology in relation to historical annual and seasonal temperatures and plant functional traits. *PeerJ*, 11, e15188.
- Goulson, D. Nicholls, E., Botías, C., and Rotheray, E.L., 2015. Bee declines driven by combined stress from parasites, pesticides, and lack of flowers. *Science*, 347:1255957.
- Haggerty, B.P. and Mazer, S.J., 2008. *The phenology handbook*. University of California, Santa Barbara, P41.
- He, K., Zhang, X., Ren, S. and Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Hicks, D., Baude, M., Kratz, C., Ouvrard, P. and Stone, G., 2021. Deep learning object detection to estimate the nectar sugar mass of flowering vegetation. *Ecological solutions and evidence*, 2:e12099.
- Holste, G., Partridge, S.C., Rahbar, H., Biswas, D., Lee, C.I. and Alessio, A.M., 2021. End-to-end learning of fused image and non-image features for improved breast cancer classification from MRI. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3294–3303.
- Hong, S.W. and Choi, L., 2012. Automatic recognition of flowers through color and edge-based contour detection. In: *3rd International conference on image processing theory, tools and applications (IPTA)*, pp. 141–146.
- Huang, S.C., Pareek, A., Zamanian, R., Banerjee, I., and Lungren, M.P., 2020. Multimodal fusion with deep neural networks for leveraging CT imaging and electronic health record: a case-study in pulmonary embolism detection. *Scientific reports*, 10:22147.
- iNaturalist (n.d.). iNaturalist. Retrieved April, 2024 from <https://www.inaturalist.org/taxa/47125-Angiospermae>.
- IPBES (2019). *Global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services*.
- Krishna, N.H., Rakesh, M. and Ram Kaushik R., 2020. Plant species identification using transfer learning-PlantCLEF 2020. In: *CLEF (Working Notes)*.
- Krizhevsky, A., Sutskever, I. and Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., ... and Zitnick, C.L., 2014. Microsoft COCO: common objects in context. In: *Computer vision – ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. Springer International Publishing, pp. 740–755.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y. and Berg, A.C., 2016. SSD: single shot multibox detector. In: *Computer vision – ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*. Springer International Publishing, pp. 21–37.

- Liu, H., Li, C., Wu, Q. and Lee, Y.J., 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Mac Aodha, O., Cole, E. and Perona, P., 2019. Presence-only geographical priors for fine-grained image classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9596–9606.
- Mann, H.M., Iosifidis, A., Jepsen, J.U., Welker, J.M., Loonen, M.J. and Høye, T.T., 2022. Automatic flower detection and phenology monitoring using time-lapse cameras and deep learning. *Remote sensing in ecology and conservation*, 8:765–777.
- NDFD (n.d.). NDFD distribution atlas. Retrieved April, 2024 from <https://www.ver spreidingsatlas.nl>.
- Nguyen, T.T.N., Le, V.T., Le, T.L., Hai, V., Pantuwong, N. and Yagi, Y., 2016. Flower species identification using deep convolutional neural networks. In: *AUN/SEED-Net Regional Conference for Computer and Information Engineering*.
- Nilsback, M.E. and Zisserman, A., 2008. Automated flower classification over a large number of classes. In: *6th Indian conference on computer vision, graphics & image processing*, pp. 722–729.
- Ollerton, J., Winfree, R. and Tarrant, S., 2011. How many flowering plants are pollinated by animals? *Oikos*, 120:321–326.
- OpenAI, n.d. GPT-4 with vision capabilities. Retrieved June, 2024 from <https://openai.com>.
- Patel, I. and Patel, S., 2020. An optimized deep learning model for flower classification using NAS-FPN and faster R-CNN. *International journal of scientific & technology research*, 9:5308–5318.
- Pawłowski, M., Wróblewska, A. and Sysko-Romańczuk, S., 2023. Effective techniques for multimodal data fusion: a comparative analysis. *Sensors*, 23:2381.
- Rebuffi, S.A., Goyal, S., Calian, D.A., Stimberg, F., Wiles, O. and Mann, T.A., 2021. Data augmentation can improve robustness. *Advances in neural information processing systems*, 34:29935–29948.
- Redmon, J., 2016. You only look once: unified, real-time object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Ren, S., 2015. Faster R-CNN: towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*.
- Schermer, M. and Hogeweg, L., 2018. Supporting citizen scientists with automatic species identification using deep learning image recognition models. *Biodiversity information science and standards*.
- Schouten, G., Michielsen, B.S.H.T. and Gravendeel, B., 2024. Data-centric AI approach for automated wildflower monitoring. *PLoS One*, 19:e0302958.
- Seeland, M., Rzanny, M., Alaqraa, N., Wäldchen, J. and Mäder, P., 2017. Plant species classification using flower images – a comparative study of local feature representations. *PLoS One*, 12:e0170629.
- Simonyan, K. and Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

- Stahlschmidt, S.R., Ulfenborg, B. and Synnergren, J., 2022. Multimodal deep learning for biomedical data fusion: a review. *Briefings in Bioinformatics*, 23:bbab569.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. and Wojna, Z., 2016. Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826.
- Tran, D.T., Høye, T.T., Gabbouj, M. and Iosifidis, A., 2018. Automatic flower and visitor detection system. In: *26th European Signal Processing Conference (Eusipco)* (pp. 405–409).
- Van der Hak, D.D., 2022. Phenology diagram generator. Retrieved May, 2024 from https://github.com/Desharin/phenology_diagram_generator.
- Van der Sluijs, J.P. and Vaage, N.S., 2016. Pollinators and global food security: the need for holistic global stewardship. *Food ethics*, 1:75–91.
- Wäldchen, J. and Mäder, P., 2018. Plant species identification using computer vision techniques: a systematic literature review. *Archives of computational methods in engineering*, 25:507–543.
- Zhu, X., Vondrick, C., Ramanan, D. and Fowlkes, C.C., 2012. Do we need more training data or better models for object detection? In: *BMVC*, 3.

ARISE: a Dutch dataspace connecting nature and people

Elaine van Ommen Kloeke, W. Daniel Kissling, Julian Evans, Chantal Huijbers, Jacob Kamminga and Gerard Schouten

Abstract

Biodiversity is declining worldwide at an unprecedented rate. The densely populated country of the Netherlands is even one of the forerunners exhibiting this dramatic decline. In recent years, however, concern about the environment has moved decisively from niche to mainstream. In this chapter we introduce ARISE (Authoritative and Rapid Identification System for Essential biodiversity information), a government-funded research infrastructure that connects nature and people. The ambition of ARISE is to enable recognition of all natural species in order to monitor on a large scale end-to-end and near real-time biodiversity, thereby helping to 'bend the curve' of biodiversity decline. To do so, ARISE (i) provides an open data platform to collect field-captured samples and digital observations of all living organisms for species level recognition, and (ii) offers tools and services to challenge and engage researchers, policymakers and citizens to create the data to enable insights that help us to better understand biodiversity in relation to our environment and human activities. ARISE is linked to the Dutch national SURF infrastructure for data management and HPC capacity. Tools comprise: (i) an AI repository for species recognition models; (ii) smart annotation services for images and sound; (iii) dashboards, leaderboards and maps; and (iv) an array of sensors to capture multicellular species in their natural environment. As such, ARISE will be the 'one-stop-shop' or marketplace for species recognition services and non-invasive biodiversity monitoring.

15.1 Biodiversity at risk – what can we do to 'bend the curve'?

Biodiversity encompasses all kinds of life – the variety of animals, plants, fungi, and even microorganisms like bacteria that make up our natural world – as well as the diverse ecosystems they create through complex interactions between such living species. Humans are an integral part of this 'fabric of life'. Since the landmark paper of Hallmann *et al.* (2017) – which convincingly shows a 75% decline in insect biomass over the last three decades in nature protection areas in Germany – the topic of biodiversity has gradually shifted into the spotlights of policymakers and researchers. It has been estimated that over a million species are headed

drivers of
biodiversity
loss

for extinction (Tollefson, 2019). As stated in the alarming Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services (2019) report, the five major drivers of biodiversity loss are: (i) changes in land and sea use; (ii) climate change; (iii) pollution; (iv) direct exploitation of organisms; and (v) invasive species. These drivers are largely a result of anthropogenic activities (e.g. Kehou *et al.*, 2017).

biological
processes

Why is biodiversity important? To answer this question, we paraphrase from the booklet *'Biodiversity at risk: today's choices matter'* of the National Academies (2022): 'The strong and ancient connections between humans and other living species mean that we cannot really separate ourselves from the ecosystems within which we evolved. Although humans have come to dominate many of Earth's ecosystems, we still rely on these connections, making conserving and restoring biodiversity a matter of survival.' Biological processes make our planet liveable, mitigate climate change (Palmer, 2021, Hawkins *et al.* 2023), provide humans with food, clothing, medicines and building materials (e.g. Desborough & Keeling, 2017; IPBES, 2019); and finally, biodiversity protects our health and to a large extent supports our economies. One in five people rely on wild species for food and income. It has been estimated that the overall value of ecosystem services to human well-being equates to more than double the global gross domestic product (Costanza, 2014).

values and
behaviour

Our attitude towards nature is deeply rooted in our underlying values and behaviour. Western culture and philosophy have a long tradition of positioning humans above nature; people are encouraged and even obliged to control and exploit nature (Harrison, 2019). This attitude is counterproductive to halting biodiversity decline. Behavioural change is needed. In other words, a transition towards a nature-positive society that involves policymakers, conservationists, researchers, citizens, and organizations (Leclère *et al.*, 2020). The EU has acknowledged the urgency of this much needed transformation by articulating a biodiversity strategy¹ (as part of the Green Deal, a package of measures to make Europe climate neutral by 2050) aiming to (i) restore nature; (ii) improve the relationship between humans and nature; and (iii) reverse ecosystem degradation. A key success factor in implementing this strategy is reliable monitoring programs, i.e. biodiversity observation networks (Pereira *et al.* 2022; Gonzalez *et al.* 2023), and the development of so-called Essential Biodiversity Variables (EBVs). EBVs are essential for assessing, comparing and predicting the effects of human actions and spatial interventions (Navarro *et al.*, 2017).

ARISE

ARISE is the technical answer to large-scale monitoring; it provides an infrastructure and toolset to enable a cost-effective way of achieving the above vision. The ambition of ARISE is to become the central hub for all species recognition data and end-to-end services in the Netherlands. ARISE adopts an open science policy

1 https://environment.ec.europa.eu/strategy/biodiversity-strategy-2030_en.

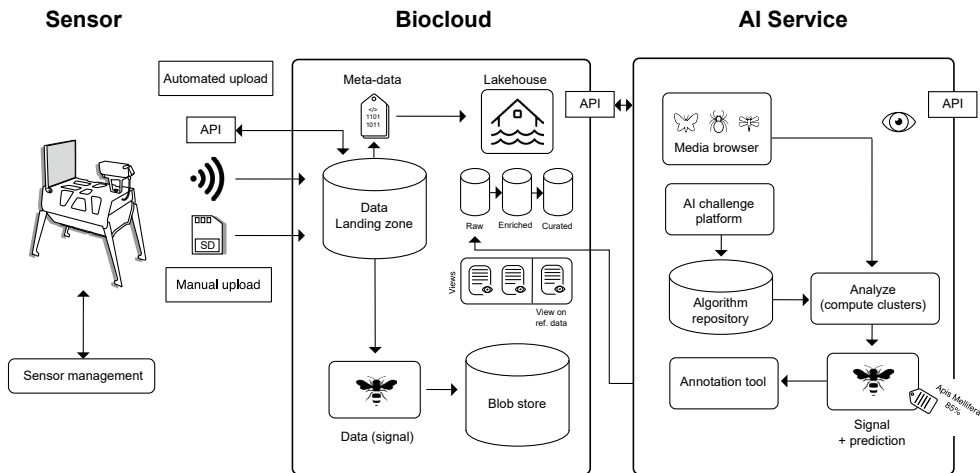


FIGURE 15.1 Overview of the end-to-end architecture in ARISE

(for data, code and publications), adheres to the FAIR data principles,² and respects data ownership by design. By building an *open* dataspace to store field observations (images, sound, radar) of all living organisms ARISE promotes and supports biodiversity research for everyone, including citizens. Biodiversity services derived from the data are developed and offered to researchers, policymakers, and companies, thereby connecting society with the challenge of biodiversity decline. Measuring the occurrence of species through time enables, for instance, estimation of soil and water quality through biological indicators. It also serves to answer ecological questions that are relevant for agricultural systems (such as the interplay between crop pollinators, pest species and their natural enemies). Furthermore, recognition of invasive species through sensor networks provides early warning signals for policymakers.

In this chapter we further detail what ARISE is and explain with tangible examples what automated species recognition for large-scale biodiversity monitoring entails. We start by sketching the landscape of measuring biodiversity in the field (Section 2). Next, as visualized in Figure 15.1, we focus on the end-to-end pipeline of (i) digital biodiversity sensors (Section 3); (ii) managing diverse and large amounts of data in a state-of-the-art data lakehouse system (Section 4); and (iii) digital species identification using AI, in particular deep learning (Section 5). We end this chapter with conclusions and lessons learned from the ARISE initiative so far and present an outlook for automated species recognition and biodiversity monitoring.

2 <https://www.go-fair.org/fair-principles>. FAIR stands for Findability, Accessibility, Interoperability, and Reuse of digital assets.

15.2 Measuring biodiversity in the field – a brief history and future outlook

observing
species

It all starts in the field. Which creatures are out there, and how can we capture them with sensors? Understanding our natural living world and, in a way, measuring the state of biodiversity has been done for centuries. The first observations on species and their behaviour date back to the ancient Greek period, over time becoming more and more prominent in the age of enlightenment during the European Renaissance period through the naturalist movement ('History of biology', n.d.). Observing species in the field can tell us if a population is doing well or not, help us predict how an ecosystem is functioning or if we should expect trouble in the future for other species or even sense the effect of bigger underlying phenomena such as climate change. In this light let's have a look at one of the largest species groups that comprises biodiversity: insects. Insects comprise between 40–80% of biodiversity and perform all sorts of useful functions. For instance, they serve as food sources for many other species such as birds, fish and small mammals. They are also important for our food production system, as many crops depend on insect pollination, and they play an important role in nutrient recycling to keep the soil healthy. A drastic loss of insects can therefore have a serious effect on other species and even entire ecosystems. Recent studies have shown declines in insect populations in the Netherlands with 75% over the past three decades (Hallmann *et al.* 2017; Van Klink *et al.* 2020, Wagner *et al.* 2020). Keeping track of how well these species are doing is therefore more important than ever.

manual and
labour-
intensive
methods

Collecting observations with the goal of accessing the state of biodiversity still relies on using manual and labour-intensive methods, ranging from on-the-spot observations by experts in the field, to setting traps, such as malaise and pan traps (Figure 15.2, top), in order to catch species in a particular area and time frame (Montgomery *et al.* 2021). These conventional invasive methods are used by long-running programmes, often through dedicated species foundations and ecological monitoring networks to answer trend-related scientific research questions. In the Netherlands one of the most well-known initiatives is NEM-net,³ focussing on the International Union for Conservation of Nature (IUCN) red list of threatened species. This initiative has been monitoring since 1999, and data is stored in the national database for flora and fauna (NDFP) and by Statistics Netherlands (CBS in Dutch) for policy purposes. Additional data is also captured through more recently established species foundations or volunteering networks such as Observation International.⁴ On an international level, observational data and other biodiversity-related data are captured through information portals, or so-called data repositories, with the Global Biodiversity Information Facility⁵ (GBIF.org)

3 <https://edepot.wur.nl/532548>. NEM stands for the Dutch name 'Netwerk Ecologische Monitoring'. This means – roughly translated – Network for Ecological Monitoring'.

4 <https://www.waarneming.nl> or <https://www.observation.org>.

5 <https://www.gbif.org>.





FIGURE 15.2 Manual data collection and identification (top) versus a fully automatic AI-powered solution for monitoring biodiversity (bottom)

being one of the largest with over 2.5 billion occurrence records, more than 90,000 datasets, 2130 publishing organizations, and an ever-growing number of peer-reviewed papers using the data (Saran *et al.*, 2022).

Even with all these long-standing and thorough initiatives it is currently impossible to know the full extent of our biodiversity. With the traditional methods we cannot answer the question ‘what lives here?’, as many species are inconspicuous, too small to be observed or in places not reachable for humans; so-called hidden biodiversity. One such place is the natural world underground, where almost all terrestrial plant life starts, including the crops we eat. It can also be the case that the human observer simply isn’t in the right place at the right time or trapping methods attract some species more than others. Next to this, the traditional methods are usually extremely labour-intensive and hence can only be done in small areas, for a limited number of species and a few times per year (or less). This is where new technology can make a difference and complement the long-standing traditional methods through high-frequency observations, and at times that otherwise might not be feasible. Both DNA and digital sensor technology, capturing e.g. images, sound and radar, have the potential to capture species presence and answer all sorts of research and policy-related questions on a large scale and in semi-automated ways (Figure 15.2, bottom), and especially in combination with AI technology to make sense of the vast amount of data that these methods generate. ARISE can help us answer questions such as: Which prey are nesting birds eating this season? How is this species group doing compared to last year? Can we increase biodiversity if we change our agricultural practices or restore our ecosystems? Or even more

limitations

new technology

generally: What types of land use are adversarial or beneficial for biodiversity? Is an invasive species gaining in territory? More examples will be given in the following sections.

15.3 Digital biodiversity sensors – from smart nest boxes to insect soundscapes

Automated biodiversity monitoring with digital sensors is becoming increasingly feasible (Besson *et al.*, 2022). Compared to traditional biodiversity monitoring methods, stationary digital sensors such as digital cameras, microphones and radars allow high-frequency observations of species without observer disturbance, in remote areas or extreme environments, with relatively little labour and comparably low costs (Kissling *et al.*, 2018). Moreover, rapid advances in technology mean that sensors have become more affordable, can now be smarter and more autonomous, and remain in the field for longer periods of time. Within ARISE, the setting-up, running and maintenance of several digital sensors are being trialled across several monitoring demonstration sites within the Netherlands (current focus on three sites, see Figure 15.3). ARISE is thereby testing the use of digital sensors in a range of different habitats and conditions and demonstrating their innovation potential for biodiversity research and monitoring. The following sensors are currently deployed (as of October 2023):

15.3.1 *Sensors in the field*

Wildlife cameras: Two types of wildlife cameras are deployed. The Browning 2021 Spec Ops Elite HP4 triggered by passive infrared (PIR) represents a more traditional wildlife camera because it operates with batteries and requires data to be downloaded manually via an SD card. In contrast, the Snyper Commander 4G Wireless works more autonomously with a solar panel and 4G data transmission. It is either PIR triggered or operates in time-lapse mode and is deployed with either a wide (100°) or regular (52°) lens. The wildlife cameras are typically mounted near the ground (20–50 cm) to detect ground-dwelling birds and mammals (with PIR trigger), or on tall poles (e.g. 2 m) to count larger flocks or groups of animals (with time lapse).

Audio loggers: AudioMoth, a small, low-cost and full-spectrum audio logger is deployed for passive acoustic monitoring. It is capable of recording uncompressed audio of both audible and ultrasonic frequencies to microSD cards. The device is operated with batteries and can be configured with a certain sample rate and recording schedule. In ARISE, the audible frequency is currently used to record bird song and the ultrasonic frequency to monitor urban rats.




FIGURE 15.3 Diversity of digital biodiversity sensors tested in ARISE. (a) Location of the three ARISE monitoring demonstration sites in the Netherlands and the deployed sensors to monitor biodiversity non-invasively and remotely. (b) Different sensors and their data volumes

Insect cameras: Several insect cameras are deployed. This includes the solar-powered DIOPSIS camera⁶ which attracts various groups of insects to a yellow screen, takes pictures when motion is detected and at regular intervals (every 10 seconds), day and night. Images and camera information can be sent to a server through the 4G network. Another insect camera in use is the Automated Moth Trap (AMT) which uses multiple light sources to attract moths at night (Bjerge *et al.*, 2021). The AMT version runs on mains power and within ARISE extra connectivity through a wired network was added to allow automatic data transmission. A third insect camera is a solar-powered insect detect camera trap⁷ which uses an artificial flower platform for automated recording of flower-visiting insects (e.g. hoverflies). These insect cameras are assembled with low-cost off-the-shelf hardware components and can be combined with open-source software (Sittinger, 2022). For ARISE, the insect detection camera traps were modified to be solar powered and to transmit data automatically via 4G.

Nest box cameras: To monitor birds and the diet of chicks during the breeding season, intelligent nest box cameras were developed by a start-up company (Ecomoni⁸) in connection with ARISE. These cameras are motion-triggered through active infrared and automatically make a video recording when a bird enters the nest, for instance allowing the detection of food items that parents bring into the nest box for their offspring. The cameras can be operated with different lenses (63 °, 75 ° or 120 °) and currently work with batteries or mains power and use 4G or wifi for data transmission.

Radar and aerial camera: The Swiss Birdscan MR1 Radar⁹ to monitor bird, bat and insect movements is also being tested through ARISE. The sensor uses a pulsed radar that emits beams vertically across a conically-shaped field to estimate the number of flying targets at an altitude of 30–1000 m (or up to 2000 m for large targets). Migration traffic rates for specific altitude layers are then computed. Wing beat characteristics of targets are used for target identification, which are typically groups of species that have similar wing beat frequencies (e.g. waders, songbirds, swifts, insects). The radar operates with mains power, and data transmission is through a wired network. To improve information on the species composition of the aerial fauna, a new and innovative aerial facing bird camera with a zoom lens (AeroecologyCam AC1) is used in conjunction with the Birdscan MR1 radar to gather image series of detected targets (e.g. small birds up to 400 m, large birds up to 1000 m). The AeroecologyCam AC1 also operates with mains power and data are transmitted through a wired network.

6 <https://diopsis.eu/en/>.

7 <https://maxsitt.github.io/insect-detect-docs/>.

8 <https://www.ecomoni.nl/>.

9 <https://swiss-birdradar.com/systems/radar-birdscan-mr1/>.

15.3.2 *Sensor autonomy, data pipelines and performance monitoring*

ARISE puts the emphasis on deploying sensors as autonomously as possible. Moreover, the development of automated data pipelines and the remote monitoring of sensor performance is a key focus. The currently deployed sensors broadly fall into three groups with increasing levels of autonomy: (i) traditional sensors, (ii) sensors with automatic transmission, and (iii) smart sensors. Traditional sensors such as the Browning wildlife camera and the AudioMoth require regular field visits to change the batteries and to manually collect data from the SD cards. To transmit the data from these sensors, which are currently some of the most commonly used sensors in ecological monitoring, ARISE has developed software to bulk download data from SD cards to a local storage device (e.g. laptop) and then upload the data to the sensor portal using an API. The sensors and their metadata must be registered manually. Moreover, sensor performance needs to be checked manually in the field to see if data are being correctly recorded.

Sensors with automatic transmission are more autonomous. For instance, the Snyder 4G wildlife camera is able to transmit data over 4G and can be powered with solar panels. This allows such sensors to remain in the field for extended periods of time with less human effort for repeated control and support. The data are transmitted over 4G to an external storage platform (e.g. FTP server or cloud service), from which they are automatically downloaded. Such cameras can also transmit a daily report of performance metrics, which allows users to remotely monitor battery status and on-board data storage over time. Problems can be remotely diagnosed, which reduces the amount of time in the field. These sensors and their metadata must still be registered manually in the sensor portal.

The most autonomous sensors are devices with on-board computational power capable of performing a variety of tasks. Most sensors deployed by ARISE fall into this category (e.g. insect cameras, nestbox cameras, radar and the aerial camera). These sensors can transmit data directly over 4G or wifi, rather than first transmitting them to an external repository. This allows them to attach a wide range of metadata directly to the media files, which can be automatically assigned to the correct sensor. The sensors' on-board computing typically results in high energy consumption which requires large solar panels (e.g. DIOPSIS system), large batteries (e.g. nest box cameras) or even mains power (e.g. Birdscan MR1 Radar, AeroecologyCam, AMT). However, they are more flexible in the kind of performance information they can transmit. This includes not only metrics such as battery status, but also information on CPU temperature or detailed logs produced by code running on-board. These smart sensors are capable of automatically registering themselves and their metadata to the sensor portal. They can also receive information from the infrastructure, allowing them to change settings remotely while the sensor is in the field.

15.4 All that data – managing a data lakehouse before it becomes a swamp

Innovative technologies such as digital sensors can greatly enhance our ability to conduct large-scale biodiversity monitoring and thus improve our understanding of which species live where. Yet, these technologies also produce vast amounts of data that need to be managed adequately for users to be able to access and analyse it. Deploying one sensor of each type described in the previous section would generate ~70 GB of data per day. As most studies would implement an array of sensors over an extended period of time, this quickly adds up to several terabytes per day (e.g. 10 sensors of each type = 4.2 TB/day). A comprehensive network of biodiversity sensors throughout the country could easily generate > 200 TB of data per month consisting of > 14M records.¹⁰ Besides the actual media files, a data management system also needs to capture metadata such as information about the location, the sensors, deployments of a sensor, media type, projects, users, environmental information and other associated data (Figure 15.4). The combination of unstructured and structured data and the volume, variety and velocity with which data is generated means that the use of such technology, and therefore ARISE, needs a big data system to handle all this data.

data
management
technologies

The big data era has prompted a revolution in data management technologies (Harby and Zulkernine, 2022). Traditional data warehouses are sufficient for storing large amounts of structured data but are incapable of dealing with a large variety of unstructured data arriving at different velocities, leading to the implementation of data lakes over the last two decades (Khine and Wang, 2018). These data lakes, however, showed critical issues with processing data fast enough to prevent inconsistencies and errors resulting in unprocessed data in so-called data swamps. This initiated the development of a new solution in which desired features of data warehouses and data lakes are combined while addressing their weaknesses, resulting in the data lakehouse architecture which is rapidly becoming the industry standard (Armbrust *et al.*, 2021). This architecture provides low-cost storage in an open data format in combination with performance features such as indexing, caching, query optimization, and data versioning. The data lakehouse architecture has been implemented in other domains (e.g. biomedical research: Begoli *et al.*, 2022; maritime monitoring: Park *et al.*, 2023) and is also used by ARISE as a data management system.

¹⁰ Numbers based on 100 insect cameras, 100 wildlife cameras, 100 sound recorders and 5 bird radars deployed for 6 months.

15.4.1 Biocloud: the brain that holds all ARISE information

The Biocloud is the brain of the ARISE data management system (see Figure 15.1). It includes multiple layers, each serving a unique purpose in the data processing pipeline (Figure 15.4). The landing zone is our primary data intake layer. The Biocloud supports various data formats and utilizes a greedy ingestion process, meaning that all received data from a source is stored. The landing zone essentially acts as a staging area where the raw data is first collected before further processing. Once data is collected in the landing zone, the metadata is processed to the raw layer. Here, the data is stored in Delta tables in a Parquet file format, while its original structure is maintained. Delta tables provide the ability to perform update, delete, and merge operations on datasets, which can significantly improve the performance of the data processing pipeline (Armbrust *et al.*, 2020). After the raw layer, the data is moved to the enriched layer.

This layer acts as a transitional stage where the data is validated, cleaned and harmonized. The data is again stored in Delta tables, but it is now in a more refined, harmonized form that is ready for detailed analysis or further processing. This enrichment can involve cleaning up data anomalies, harmonizing disparate data formats, and enriching data with additional context or metadata. This results in one table per data entity (e.g. location, sensor, algorithm) that includes all data of that entity from different data sources. The last layer in our data pipeline is the curated layer. At this stage, the data is tailored into the format and shape that the end-users need. This can involve aggregating data, applying business rules, or transforming data to meet specific reporting or analytical requirements. The curated layer delivers data in a readily usable format, making it easier for business users, data analysts, and data scientists to extract insights without worrying about data clean-up or transformation tasks. Blob data such as images, sounds and videos, are also ingested




FIGURE 15.4 Overview of the Biocloud architecture with the different layers of processing the original data sources from raw to enriched and curated data for future use and access

into the landing zone, and then moved to S3 storage. The segregation between blob data and metadata allows for efficient querying and management of unstructured data while keeping the overall architecture optimized for performance.

Data stored in the Biocloud will be accessible through the ARISE platform as well as through machine-to-machine connections with other research infrastructures. The ARISE platform will provide a range of end user services such as search and browse through all publicly available ARISE data, view data from user-registered sensors, annotate data, and use AI algorithms to identify species from digital media (see also Section 5). While ARISE is mostly focused on services enabling species identification, other research infrastructures provide a wealth of information such as environmental data, human observations of species and global DNA barcodes, which can complement ARISE data for a broader perspective on the state of our ecosystems. Our goal is to connect with infrastructures such as the Global Biodiversity Information Facility (GBIF), Catalogue of Life, Barcode of Life Data System (BOLD), and Distributed Systems of Scientific Collections (DiSSCo) to contribute to the global assessment of biodiversity.

15.4.2 *FAIR by design*

The ARISE infrastructure is a complex system connecting and serving multiple systems, stakeholders, and user communities. Our design principles are aimed at building a system that is scalable and flexible, but also secure and reliable. Three of our guiding principles are:

Adhere to the FAIR principles: Data generated, analysed, stored and managed through ARISE can form the basis for many different avenues of biodiversity research. It is important to ensure provenance and full traceability of data objects so that analyses can be replicated, and data can be reused. Therefore, the architecture adheres to the FAIR principles as much as possible to make the data Findable, Accessible, Interoperable and Reusable (Wilkinson *et al.*, 2016). We implement Persistent Identifiers (PIDs), which are unique and unchangeable labels assigned to each digital object ensuring their accessibility and traceability over time (De Smedt *et al.*, 2020).

Follow and use international standards and practices: ARISE is dedicated to fostering open development by creating an extensible infrastructure that adheres to recognized standards. Data standards are key for consistency, efficiency and maintenance of high data quality by reducing the risk of errors and inaccuracies. Some well-known standards for biodiversity data include Darwin Core (Wieczorek *et al.*, 2012) and the Ecological Metadata Language (EML) (Fegraus *et al.*, 2005). More recently, GBIF presented a unified common data model supporting more complex types of biodiversity data,¹¹ and also more detailed standards for camera traps have

11 <https://www.gbif.org/composition/7AZMWZtLvrnfYbFpUdF83I/diversifying-the-gbif-data-model-webinar-resources>.

been published (Bubnicki *et al.*, 2024). ARISE will align with these standards as much as possible to promote interoperability and reusability of data.

Design for evolution: The ARISE infrastructure aims for a long life span; a time in which biodiversity research and technologies will evolve, so ARISE needs to be able to grow and evolve as well. We are building a system that will facilitate new types of sensors, data types, reference datasets or AI algorithms. The architecture is designed with scalability in mind allowing growth in volumes of data, number of data sources, users, service requests and connections with other infrastructures.

15.5 Powering AI for biodiversity

Digital biodiversity sensors generate vast amounts of data (see Figure 15.3), necessitating AI for analysis. ARISE enables non-AI experts to use AI for data interpretation and bridges the divide between computer science and ecology. Accurate algorithms are essential; errors in analysing biodiversity data can hinder crisis response. ARISE aims to identify every species in digital media, like images and audio, constantly improving accuracy. Furthermore, ARISE empowers researchers to evaluate, compare, and innovate AI for biodiversity. Its open, standardized approach ensures that breakthroughs are shared, adopted, and built upon, fostering a community of collaborative growth.

15.5.1 *The daunting challenge of AI in biodiversity*

Automatically identifying species using sensors (e.g. recording images or sound) and AI is a tremendous challenge. Biodiversity datasets have a long-tailed distribution and contain many examples of common species and few examples of rarer species. This constitutes a bias in the training of identification algorithms (van Horn & Perona, 2017) and is a crucial issue because rarer species are often the most important ones to monitor. Furthermore, there are millions of potential species to recognize, while the number of experts that can identify them is small and declining (Engel *et al.*, 2021; Greeff *et al.*, 2022).

The road to achieving the ARISE vision is not without its hurdles:

1. Quality data access: Ensuring high-quality training data to train high-performance algorithms
2. Dynamic digital framework: Crafting an adaptive digital species identification infrastructure that can evolve, allowing for updates and development of identification algorithms while preserving a clear provenance trail for automatic identifications
3. Employing modern AI: Developing the expertise and tools to leverage cutting-edge AI developments for species identification
4. Societal integration: Ensuring the infrastructure's relevance and upkeep in society for the long term

15.5.2 *The promise of ARISE*

The diversity of species contrasts sharply with the few experts capable of identifying them. Taxonomists and ecologists are therefore vital for integrating their knowledge into AI models. Ecologists frequently annotate excessive, non-essential data for AI learning, which could be more effectively focused on essential annotation tasks, like under-represented species. ARISE offers enhanced data annotation technologies and interactive AI training. ARISE aims to intertwine human expertise with evolving AI algorithms. While species observations were once manual, automation now aids detection and classification. Routine data can be accurately auto-classified, but complex data demands human validation to refine training data. In this collaborative model, as experts annotate, the AI simultaneously refines its capabilities as shown in Figure 15.5. With every iteration, AI handles broader datasets and highlights intricate data for human review. The objective is twofold: maximize expert efficiency and boost AI's proficiency in detailed species identification.

15.5.3 *Harnessing ARISE: the architecture behind advanced species identification*

Designed for researchers aiming to analyse biodiversity through digital sensors, the species identification architecture transforms heterogeneous 'raw' data into structured labelled data. When analysing data, a researcher typically faces one of three scenarios:

three scenarios

1. An available algorithm that performs adequately for fully automated reporting.
2. An available algorithm that requires enhancement or additional features (e.g. species granularity).
3. A lack of an algorithm tailored to the specific research goal.

While ideally only Scenario 1 would be needed for full automation, Scenarios 2 and 3 are commonly encountered, necessitating additional steps. As ARISE evolves, an increasing number of research objectives will fit into Scenario 1.

Using ARISE involves:

1. Connecting sensors or uploading media and organizing it.
2. Automatically analysing data for species detection and species identification, using:
 - a. Existing AI algorithms, or
 - b. New algorithms if current ones are unsuitable or outdated.
3. Finetuning, evaluating, and annotating new data or correct misclassifications.
4. Adjusting the algorithm for the specific domain and re-evaluating.
5. Iterating the above until satisfactory results are achieved.
6. Downloading final analysis results.

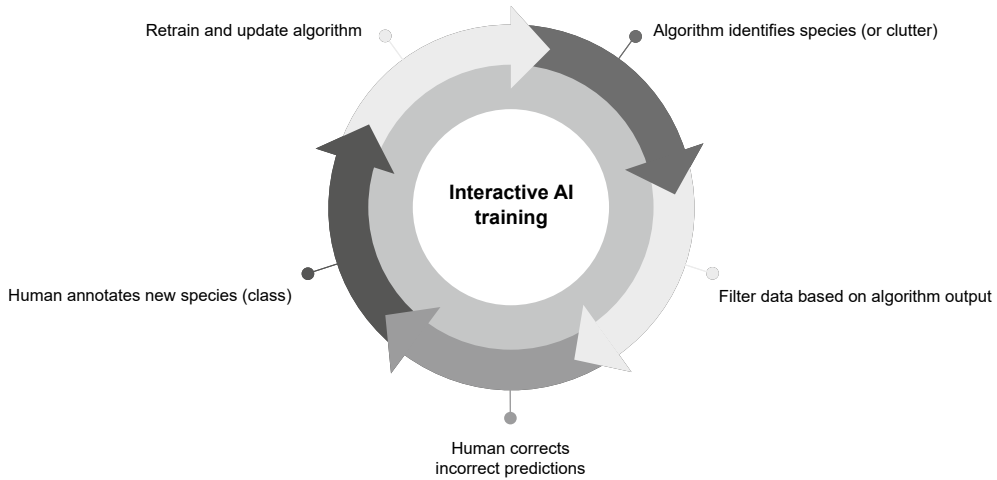


FIGURE 15.5 The active learning cycle of advanced species identification in ARISE

TABLE 15.1 Mapping of processes and scenarios to ARISE components

Process	Scenario 1	Scenario 2	Scenario 3	ARISE components
Data collection and storage	x	x	x	Sensor portal + Biocloud
Data visualization and organization	x	x	x	Media browser
Organization of algorithms and automatic deployment for data analysis	x	x	x	Algorithm repository
Data annotation for cleaning, training, and evaluation		x	x	Media browser + annotation tool
Algorithm evaluation and comparison		x	x	AI challenge platform
Algorithm design and training			x	AI challenge platform

Automated species identification processes, outlined in Table 15.1, remain consistent across media types, e.g. images and sounds. However, certain processes are specific to Scenarios 2 and 3 since they are automated in Scenario 1.

The key components of advanced species identification in ARISE are:

- *Media browser*: A central tool for experts to interact, visualize, label, and filter their data based on metadata like sensor ID, deployments, or analysis output. This component aids in creating datasets for analyses and annotations.
- *Annotation tool*: Beyond the media browser, ARISE integrates with external tools like Label Studio and Intel Geti, facilitating transitions to specialized annotation sessions and active learning. All annotations are synced with ARISE for data organization, training, and algorithm evaluation.
- *Algorithm repository*: ARISE hosts a range of open-source algorithms, from source code to automated model serving. Users can quickly deploy algorithms on platforms like Amazon cloud or SURF's Snellius. Transfer learning and fine-tuning principles enable existing algorithm reuse for related research objectives. If no suitable algorithm is found, ARISE's annotation tools assist in creating training data, and the AI challenge platform can be used to develop an algorithm.
- *AI challenge platform*: To accommodate evolving AI and diverse research needs, ARISE offers a challenge platform focused on biodiversity. Researchers can initiate and submit challenges, supplying necessary training and evaluation data acquired through ARISE. A dedicated evaluation container reviews submitted algorithms, with results displayed on a leaderboard. Successful algorithms can then be integrated into the algorithm repository for subsequent analyses.

15.6 Path forward and new challenges

As we embark on this journey towards autonomous biodiversity monitoring, the collaboration between technology, taxonomy, ecology, and biodiversity research has never been more crucial. In the Netherlands, ARISE offers a promising path forward, championing the cause of open science, expert involvement, multi-datatype monitoring, and the unparalleled capabilities of AI. The future beckons, and with tools like ARISE, we will deepen our understanding of the complexity of nature and have better tools to track progress towards achieving conservation management targets and biodiversity policy goals.

Although ARISE solves many of the issues related to large-scale non-invasive data collection as well as data interpretation by using AI, there are also new challenges ahead:

- How do we ensure that people start using ARISE? Building an end-to-end data-space is not enough. How do we engage researchers, citizens and companies to use this vast amount of data to answer research, management and policy questions related to biodiversity?

- Is our architecture open enough? Can we easily add other types of sensors? For instance, drones that can automatically fly transects and capture photos at waypoints of a field with wildflowers. How many sensors do we need to estimate population sizes of species or the biodiversity that drives ecosystem functions and services?
- The advancement of novel technologies for understanding and measuring biodiversity should be driven by the scientific inquiries that researchers are tackling as well as by the requirements of society and policymakers. Do we have a process in place to capture (and implement in an agile way) these new requirements?

References

- Armbrust, M., Das, T., Sun, L., Yavuz, B., Zhu, S., Murthy, M., ... and Zaharia, M., 2020. Delta lake: high-performance ACID table storage over cloud object stores. In: Proceedings of the VLDB endowment, 13:3411–3424.
- Armbrust, M., Ghodsi, A., Xin, R. and Zaharia, M., 2021. Lakehouse: a new generation of open platforms that unify data warehousing and advanced analytics. In: Proceedings of CIDR, 8:28.
- Begoli, E., Goethert, I. and Knight, K., 2021. A lakehouse architecture for the management and analysis of heterogeneous data for biomedical research and mega-biobanks. In: IEEE International Conference on Big Data, pp. 4643–4651.
- Besson, M., Alison, J., Bjerge, K., Gorochowski, T.E., Høye, T.T., Jucker, T., ... and Clements, C.F., 2022. Towards the fully automated monitoring of ecological communities. Ecology letters, 25:2753–2775. <https://doi.org/10.1111/ele.14123>.
- Bjerge, K., Nielsen, J.B., Sepstrup, M.V., Helsing-Nielsen, F., and Høye, T.T., 2021. An automated light trap to monitor moths (*Lepidoptera*) using computer vision-based tracking and deep learning. Sensors, 21:343.
- Bubnicki, J.W., Norton, B., Baskauf, S.J., Bruce, T., Cagnacci, F., Casaer, J., ... and Desmet, P., 2024. Camtrap DP: an open standard for the FAIR exchange and archiving of camera trap data. Remote sensing in ecology and conservation, 10:283–295.
- Costanza, R., De Groot, R., Sutton, P., Van der Ploeg, S., Anderson, S.J., Kubiszewski, I., ... and Turner, R.K., 2014. Changes in the global value of ecosystem services. Global environmental change, 26, 152–158.
- Desborough, M.J., and Keeling, D.M., 2017. The aspirin story – from willow to wonder drug. British journal of haematology, 177:674–683.
- De Smedt, K., Koureas, D. and Wittenburg, P., 2020. FAIR digital objects for science: from data pieces to actionable knowledge units. Publications, 8:21.
- Engel, M.S. Ceriaco, L.M., Daniel, G.M., Dellapé, P.M., Löbl, I., Marinov, M., ... and Zacharie, C.K., 2021. The taxonomic impediment: a shortage of taxonomists, not the lack of technical approaches. Zoological journal of the Linnean Society, 193:381–387.
- Fegraus, E.H., Andelman, S., Jones, M.B. and Schildhauer, M., 2005. Maximizing the value of ecological data with structured metadata: an introduction to ecological metadata

- language (EML) and principles for metadata creation. *Bulletin of the Ecological Society of America*, 86:158–168.
- Gonzalez, A., Vihervaara, P., Balvanera, P., Bates, A.E., Bayraktarov, E., Bellingham, P.J., ... and Torrelio, C.Z., 2023. A global biodiversity observing system to unite monitoring and guide action. *Nature ecology & evolution*, 1–5.
- Greeff, M., Caspers, M., Kalkman, V., Willemse, L., Sunderland, B.D., Bánki, O. and Hogeweg, L., 2022. Sharing taxonomic expertise between natural history collections using image recognition. *Research ideas and outcomes*, 8:e79187.
- Hallmann, C.A., Sorg, M., Jongejans, E., Siepel, H., Hofland, N., Schwan, ... and De Kroon, H., 2017. More than 75 percent decline over 27 years in total flying insect biomass in protected areas. *PLoS One* 12;10:e0185809. <https://doi.org/10.1371/journal.pone.0185809>.
- Harby, A.A. and Zulkernine, F., 2022. From data warehouse to lakehouse: a comparative review. In: *IEEE International Conference on Big Data*, pp. 389–395.
- Harrison, P., 2019. Reformation of science. *Aeon Magazine*. <https://aeon.co/essays/how-protestantism-influenced-the-making-of-modern-science>.
- Hawkins, H.J., Cargill, R.I., Van Nuland, M.E., Hagen, S.C., Field, K.J., Sheldrake, M., ... and Kiers, E.T., 2023. Mycorrhizal mycelium as a global carbon pool. *Current biology*, 33:R560–R573.
- History of biology. https://en.wikipedia.org/wiki/History_of_biology.
- IPBES, 2019. Global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. Brondizio, E.S., Settele, J., Diaz, S. and Ngo, H.T. (editors). IPBES secretariat, Bonn, Germany.
- Kissling, W.D., Ahumada, J.A., Bowser, A., Fernandez, M., Fernández, N., García, E.A., ... and Hardisty, A.R., 2018. Building essential biodiversity variables (EBVs) of species distribution and abundance at a global scale. *Biological reviews*, 93:600–625. <https://doi.org/10.1111/brv.12359>.
- Kehoe, L., Romero-Muñoz, A., Polaina, E., Estes, L., Kreft, H. and Kueimmerle, T., 2017. Biodiversity at risk under future cropland expansion and intensification. *Nature ecology & evolution*, 1:1129–1135.
- Khine, P.P. and Wang, Z.S., 2018. Data lake: a new ideology in big data era. In: *ITM web of conferences*, 17:03025. EDP Sciences.
- Leclère, D., Obersteiner, M., Barrett, M., Butchart, S.H., Chaudhary, A., De Palma, A., ... and Young, L., 2020. Bending the curve of terrestrial biodiversity needs an integrated strategy. *Nature*, 585:551–556.
- Montgomery, G.A., Belitz, M.W., Guralnick, R.P. and Tingley, M.W., 2021. Standards and best practices for monitoring and benchmarking insects. *Frontiers in ecology and evolution*, 8:579193.
- National Academies of Sciences, Engineering, and Medicine, 2022. Biodiversity at risk: today's choices matter. <https://doi.org/10.17226/26384>.
- Navarro, L.M., Fernandez, N., Guerra, C., Guralnick, R., Kissling, W.D., Londono, M.C., ... and Pereira, H.M., 2017. Monitoring biodiversity change through effective global coordination. *Current opinion in environmental sustainability*, 29:158–169.

- Palmer, L. 2021. How trees and forests reduce risks from climate change. *Nature climate change* 11:374–377.
- Park, S., Yang, C.S. and Kim, J., 2023. Design of vessel data lakehouse with big data and AI analysis technology for vessel monitoring system. *Electronics*, 12:1943.
- Pereira, H.M., Junker, J., Fernández, N., Maes, J., Beja, P., Bonn, A., ... and Zuleger, A.M., 2022. Europa Biodiversity Observation Network: integrating data streams to support policy. ARPHA preprints. <https://doi.org/10.3897/arphapreprints.e81207>.
- Saran, S., Chaudhary, S.K., Singh, P., Tiwari, A., and Kumar, V., 2022. A comprehensive review on biodiversity information portals. *Biodiversity and conservation*, 31:1445–1468.
- Sittinger, M., 2023. Insect Detect – software for automated insect monitoring with a DIY camera trap system. Zenodo. <https://doi.org/10.5281/zenodo.7793296>.
- Tollefson, J., 2019. One million species face extinction. *Nature*, 569:171.
- Van Horn, G. and Perona, P., 2017. The devil is in the tails: fine-grained classification in the wild. <https://doi.org/10.48550/arXiv.1709.01450>.
- Van Klink, R., Bowler, D.E., Gongalsky, K.B., Swengel, A.B., Gentile, A. and Chase, J.M., 2020. Meta-analysis reveals declines in terrestrial but increases in freshwater insect abundances. *Science*, 368:417–420.
- Wagner, D.L., 2020. Insect declines in the Anthropocene. *Annual review of entomology*, 65:457–480.
- Wieczorek, J., Bloom, D., Guralnick, R., Blum, S., Döring, M., Giovanni, R. ... and Vieglais, D., 2012. Darwin Core: an evolving community-developed biodiversity data standard. *PLoS One*, 7:p.e29715.
- Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A.,.... and Bouwman, J., 2016. The FAIR guiding principles for scientific data management and stewardship. *Scientific data*, 3:1–9.

Index

- 2030 Agenda for Sustainable Development 14
- 3D food printer 123
- 8R model, Latouche 15

- absolutist, normative theories 70
- Agenda 21 13
- Agenda for Sustainable Development, 2030 14
- Agent-Deed-Consequence Model 73
- Agreement, Paris 121
- agriculture, automated 122
- agronomy 139
- AI act, EU 69, 72
- AI system, objectives 72
- AI technology, species observation 237
- AI, General Purpose 69
- AI, Generative 69
- AI, Green 181
- AI, probability of harm 69, 72
- AI, Red 181
- AI, risk management procedure 72
- AI, trustworthy 181
- AI-based identification, plant monitoring 203
- AIXR system 76
- AIXR system risks 70
- AIXR systems, high risk 77
- ancient Greek, species observation 236
- anthropogenic activities 234
- anti-herbivory defence strategies 200
- approach, shareholder 11, 16
- approach, stakeholder 11
- ARISE 203, 234
- Arise, Key Components 248
- ARISE, processes 248
- ARISE, scenarios 246
- Arnstein, Sherry 101
- artificial intelligence (AI) 180
- attitude-behavioural intention gap 126
- audio loggers, species observation 238
- Augmented Utilitarianism 73, 127
- augmented virtuality 69
- authority, moral 71
- automated biodiversity monitoring 238
- automated vehicles 75
- Average Precision (AP) 214
- avoidance, tax 21

- Big Data, plant monitoring 203
- Bioblitz 87

- biodiversity 233
- biodiversity crisis 98
- biodiversity loss 208
- biodiversity loss, drivers 234
- biodiversity monitoring, automated 238
- biodiversity observation networks 234
- biodiversity research 210
- biodiversity strategy, EU 234
- biological pest control 139
- biological processes 234
- biomass, insect 233
- biophilia 104
- biotechnology 142
- Bounding Boxes (BBs) 211–213, 217
- breeders 140
- business ethics 125

- cameras, Diopsis insect 203
- cameras, wildlife 203
- capacity, software 173
- carbon-aware datacenters 179
- carbontracker 188
- care, principle 143
- Carson, Rachel 8
- categorical imperative, Kant 125
- cell fusion 141
- Central Processing Unit (CPU) 174
- chain responsibility 21
- chatGPT 188
- chemical pest control 139
- chromosomes, duplication 141
- cisgenesis 141
- citizen participation 101
- citizen participation, definition 101
- citizen science 82, 84, 105
- citizen science database 211
- citizen science projects 84
- citizen-consumer paradox 126
- climate change 138
- Club of Rome 9
- coding of morphological features 210
- colchicine, natural 141
- Commission, European 189
- comparative approach, meta ethics 73
- competition, regulatory 21
- consumer behaviour 126
- consumer behaviour, food 145
- controlled crossing of parent lines 141

- convolutional Neural Network (CNN) 210
- copper sulphate 140
- Corporate Social Responsibility (CSR) 125
- Covid 19 pandemic 13, 104
- cradle to cradle 18, 20
- crisis, biodiversity 98
- CRISPR-Cas 141
- crop monitoring, drones 123
- crop production 140
- crops, evolving 139
- crowd-sourced ethics 121, 127
- cultured gametes 141
- culture, food 144

- Darwin Core 244
- data fusion strategies 215
- data lakes 242
- data management technologies 242
- data warehouses 242
- data, flowering phenology 218
- datacenters, carbon-aware 179
- dataspace, open 235
- decision making, ethical 121
- decision making, sustainable 121
- Declaration, Stockholm 9
- deep learning 209
- defence strategies, anti-herbivory 200
- degradation, environmental 122, 138
- degrowth 15
- descriptive ethics 124
- design process, food technology 123
- design, moral 125
- design, value sensitive 125
- development 98
- development, sustainable 98
- Dickens, Charles 5
- digital recording technologies 84
- digital sensor technology, species observation 237
- digital sensors, species observation 238
- Diopsis insect cameras 203
- disposal, waste 20
- DNA technology, species observation 237
- donut economy 15
- doubled haploids 141
- downcycling 20
- drones, crop monitoring 123
- dyadic morality 74

- Ecological Metadata Language (EML) 244
- economy, donut 15
- ecosystem 234
- Eindhoven Wildflower Dataset (EWD) 210–211, 217
- Embryo rescue 141
- energy consumption measurement 174
- energy patterns, software design 178
- engineering development process 182
- engineering, software 172
- engineering, sustainable software 170
- engineers, moral dashboard 72
- Enlightenment 7
- environmental degradation 122, 138
- environmental footprint, food production 122
- environmental justice 89
- environmental risk mitigation, software 174
- environmental sustainability 170
- Essential Biodiversity Variables (EBVs) 234
- ethical decision making 121
- ethics, crowd-sourced 127
- ethics, descriptive 124
- ethics, food 143
- ethics, normative 125
- EU, biodiversity strategy 234
- European Commission 189
- European Green Deal 189
- European Group on Ethics in Science and New Technologies 144
- European renaissance 236
- Expected Moral Value Theory 73
- Experiment-Impact-Tracker framework 188
- exposure, vernalisation 199
- extended reality 68–69
- extinction of species 234

- factory-processed food 142
- FAIR principles 244
- farming, smart 122
- Faster R-CNN (Faster Region-based Convolutional Network) 211–212
- Faster R-CNN model 212
- Faster R-CNN object detection 221
- fermented proteins 123
- financial crises, 2008 12
- Floristisch Onderzoek Nederland (FLORON) 84
- flowering phenology 210
- flowering phenology data 218
- flowering plants 208
- food consumption 122
- food culture 144
- food ethics 143

- Food Ethics Council 125
- food ethics matrix 125
- food insecurity 122
- food lab, moral 128
- food printer, 3D 123
- food processing 142
- food processing, biotechnology 142
- food production 122
- food production, environmental footprint 122
- food systems, complexity 138
- food tech market 122
- Food technologies 121
- food technologies, moral implications 123
- food technology, design process 123
- food transition 122
- food, factory-processed 142
- food, insects 123
- food, nutritious 122
- food, sustainable 122
- food, toxicity 137
- foods, future 146
- foods, natural 137
- foods, Plant-based 123
- fragmentation, habitat 199
- Freedman, Milton 16
- Freeman, R. Edward 11
- future foods 146

- gametes, cultured 141
- General Purpose AI 69
- generative AI models 69
- genetic modification 136
- Genetically Modified Organisms (GMOs) 122
- genomic selection 141
- geopolitical dependencies 138
- Global Biodiversity Information Facility (GBIF) 236
- global warming 13
- governance, mosaic 101
- green AI 181
- green programming languages 179
- green software 172, 179
- greenhouse gas emissions 188
- greenwashing 16, 22

- Haber/Bosch process 125
- habitat fragmentation 199
- hard cases, moral 71
- health, public 122
- helix, quadruple 18
- heterosis 141

- high stake moral problem 72
- Human–Nature Connectedness (HNC) 98, 102
- Human-Induced Rapid Evolutionary Changes (HIREC) 197
- hunter-gatherers 137

- incarnation, waste 20
- individual sustainability 171
- insect biomass 233
- insect cameras, Diopsis 203
- insect cameras “species observation, insect cameras”, species observation 240
- insects, food 123
- insecurity, food 122
- Intergovernmental Panel on Climate Change (IPCC) 13
- International Union for Conservation of Nature (IUCN) 236
- Intersection over Union (IoU) 213–214
- invasive methods, species observation 236
- ionizing radiation 141
- IPBES report 2019 234
- islands, urban heat 197
- ISO 9126 172
- ISO 25000 172–173, 182

- justice, environmental 89

- Kant, categorical imperative 125
- Kyoto Protocol 13

- lab-grown meat 123
- ladder of Lansink 20
- Lansink, Ad 20
- Latouche, Serge 15
- Liebig, Justus von 139
- low stake moral problem 71

- machine learning (ML) 180, 188
- Machine Learning Emissions Calculator 188
- market, food tech 122
- matrix, food ethics 125
- mean average precision (mAP) 214
- measurement, energy consumption 174
- meat industry technology 122
- meat replacement products 142
- meat, Lab-grown 123
- meta ethics, comparative approach 73
- meta-ethical models 73
- microhabitats, European Cities 197

- millennium goals 13
- mixed reality 69
- modification, genetic 136
- molecular marker-assisted selection 141
- monitoring, plant 202
- Monsoon Power Monitor 174
- moral aspects, technology 101
- moral authority 71
- moral dashboard for engineers 72
- moral data city hunt method 127
- moral decoration 22
- moral design 125
- moral food lab 128
- moral hard cases 71
- moral implications, food technologies 123
- moral problem, high stake 72
- moral problem, low stake 71
- moral problems 70
- moral problems, nature 70
- moral programming 127
- moral solutions 70
- moral truth 70
- morality, dyadic 74
- morphological features, coding 210
- mosaic governance 101
- Multimodal Large Language Model (MLLM) 221
- mutagenesis 141
- mutagenesis, targeted 141

- National Database for Flora and Fauna (NDDF) 236
- natural colchicine 141
- natural foods 137
- Naturalis Biodiversity Center 87
- naturalness 136
- naturalness, perception 137
- NEM-net 236
- neophobia 145
- Nest box cameras, species observation 240
- Netherlands Institute of Ecology (NIOO) 87
- normative ethics 125
- normative theories, absolutist 70
- normative theories, relativist 70
- nutritious food 122

- Observation International 236
- omnivores 137
- open dataspace 235
- OpenAI 188
- ornamental trees 197

- pandemic, Covid 19 13
- Paris Agreement 121
- Paris Climate Agreement 13
- participation, citizen 101
- performance efficiency, software 173
- Personal Value Dictionary 128
- pessimism, technological 7
- pest control, biological 139
- pest control, chemical 139
- PhenoWatch 84
- Pikkety, Thomas 14
- plant blindness 201
- Plant monitoring 202
- plant monitoring, AI-based identification 203
- plant monitoring, Big Data 203
- plant species, wild 196
- Plant-based foods 123
- PlantCLEF challenge 211
- plants, flowering 208
- platonic conflict 70
- pollution 200
- Power Monitor, Monsoon 174
- principle, care 143
- probability of harm, AI 69, 72
- processing, food 142
- product lifetime 20
- proteins, fermented 123
- protoplast fusion 141
- public health 122
- public-interest entities 17
- pyrethrum 140

- quadruple helix 18

- R10-model 20
- Radar and aerial camera, species observation 240
- radiation, ionizing 141
- Raworth, Kate 15
- recycling 20
- recycling, downcycling 20
- recycling, upcycling 20
- Red AI 181
- Region of Interest (ROI) 213
- Region Proposal Network (RPN) 213
- regulation, XR 69
- regulatory competition 21
- relativist normative theories 70
- renaissance, European 236
- Reptielen Amfibieën Vissen Onderzoek Nederland (RAVON) 84

- resource utilisation, software 173
responsibility, chain 21
R-hierarchy 20
risk management procedure, AI 72
R-ladder 20
- Scherder, Erik 106
science fiction 6
science, citizen 82, 84
shareholder approach 11, 16
Silent Spring 8
Sinclair, Upton 6
smart farming 122
social ecological system (SES) 102
Social sustainability 170
social washing 22
Socio-Technical-Ecological system (STES) 99, 102
Socio-Technological Feedback (SOTEF) loop 74–75
software design, Energy Patterns 178
software engineering 172
software sustainability 170
software, capacity 173
software, environmental risk mitigation 174
software, green 179
software, Performance Efficiency 173
software, resource utilisation 173
software, sustainable engineering 170
software, time behaviour 173
Sovon Vogelonderzoek Nederland 84
species observation, AI technology 237
species observation, ancient Greek 236
species observation, Audio loggers 238
species observation, digital sensor technology 237
species observation, digital sensors 238
species observation, DNA technology 237
species observation, insect cameras 240
species observation, Nest box cameras 240
species observation, Radar and aerial camera 240
species observation, species observation 236
species observation, wildlife cameras 238
species, wild plant 196
SSD (Single Shot Detector) 212
stakeholder approach 11
Statistics Netherlands (CBS) 236
Stichting Ornithologisch Veldonderzoek Nederland (SOVON) 84
Stockholm Declaration 9
Stockholm Declaration, principles 9
sulphur 139
sustainability, Environmental 170
sustainability, Individual 171
sustainability, Social 170
sustainability, Technical 171
sustainable AI task force 189
sustainable decision making 121
sustainable development 98
Sustainable Development Goals (SDGs) 13, 87–88, 121
sustainable entrepreneurship 125
sustainable food 122
sustainable software engineering 170
systems, food consumption 122
systems, food production 122
- targeted mutagenesis 141
tax avoidance 21
technical sustainability 171
technological ape 102
technological, technological 7
technologies, digital recording 84
technologies, food 121
technology, definition 100
technology, meat industry 122
technology, moral aspects 101
technology, value alignment 125
theory of basic values 128
time behaviour, software 173
toxicity, food 137
trade deflection 21
transition, food 122
Trebeck, Katherine 15
trees, ornamental 197
triple bottom line 121
trustworthy AI 181
truth, moral 70
- UN, World Commission on Environment and Development 9
upcycling 20
urban flora, ecosystem services 200
Urban Heat Island Effect (UHIE) 197–198
utilitarianism, augmented 127
- value alignment, technology 125
value sensitive design 125
values, theory of basic 128

- vernalisation exposure 199
- Verne, Jules 6
- virtual reality 69

- Waarneming.nl 84
- Wageningen University 87
- walking apps 106
- waste disposal 20
- waste pyramid 20
- waste, incarnation 20
- Wells, H.G. 6
- wild plant species 196

- wildflower monitoring 209
- wildflowers 208
- wildlife cameras 203
- wildlife cameras, species observation 238
- Williams, Jeremy 15
- Woodland Nature Trust, UK 84
- World Commission on Environment and
Development, UN 9

- XR regulation 69

- YOLO (You Only Look Once) 212

Welcome to the Anthropocene, the era in which humans have put a tangible mark on our planet. But also the era in which humans and technology have the potential to shape the necessary transitions towards a sustainable world. Technology and nature are often considered as two opposing phenomena. However, they are increasingly intertwined, for better or worse. In this book, we explore how technology and nature relate to one another in the moral design of new, green technology.

This book is relevant for IT and engineering professionals, business leaders and policy makers with (green) innovation in their portfolios and students of (applied) science who are interested in either sustainable and green design of technology or in the application of technology – with an emphasis on AI and IT – to create a greener, more sustainable world. The chapters have been written by experts and leading researchers in an attractive, accessible, and practical writing style. Each chapter offers colourful examples and challenges the reader to critically think through moral decision-making and the design of innovations considering our planet's perspective. This is a conceptual change in values. Nature should not be considered as a resource: it is the fabric of life that makes our own existence possible.

Bart Wernaart is Professor of Moral Design Strategy and Chair of the Centre of Expertise Sustainable and Circular Transitions at Fontys University of Applied Sciences, The Netherlands. He earned his Ph.D. in the field of international human rights law and has extensive experience in teaching legal subjects in business schools. He is the head of Research at Fontys Business and Communication.

Gerard Schouten is Professor of AI & data at Fontys University of Applied Sciences and Chair of the Centre of Expertise AI for Society. He earned his Ph.D. in the field of cognition and perception, and has extensive experience in applied research and teaching data- and AI-related topics.



BRILL.COM
ISBN: 978-90-04-71167-9

