

THE AI MUSIC PROBLEM

Why Machine Learning Conflicts With Musical Creativity

Christopher W. White

First published 2025

ISBN: 978-1-032-95976-4 (hbk)

ISBN: 978-1-032-95975-7 (pbk)

ISBN: 978-1-003-58741-5 (ebk)

Chapter 6

INTERPRETING MUSICAL ARTIFICIAL INTELLIGENCE

CC-BY-ND 4.0

DOI: 10.4324/9781003587415-6



Routledge
Taylor & Francis Group
NEW YORK AND LONDON

6

INTERPRETING MUSICAL ARTIFICIAL INTELLIGENCE

“The purpose of life, after all, is to live it, to taste experience to the utmost, to reach out eagerly and without fear for newer and richer experience.”

Eleanor Roosevelt

“Oh, I see,” said the Tin Woodman. “But, after all, brains are not the best things in the world.”

“Have you any?” inquired the Scarecrow.

“No, my head is quite empty,” answered the Woodman. “But once I had brains, and a heart also; so, having tried them both, I should much rather have a heart.”

The Wonderful Wizard of Oz, L. Frank Baum

In a brief scene in George Orwell’s *1984*, the protagonist nonchalantly reports that the tune currently topping the novel’s dystopian pop charts was generated by a computational instrument “without any human intervention”— it was created by AI. The character’s description of the song is not positive. In his estimation, the music is sugary, vapid, shallow, and devoid of authentic emotional content.

However, as he walks along the streets of London, Orwell’s protagonist hears a woman absentmindedly singing the tune while she hangs laundry. The experience immediately transforms the song in his mind. In her “powerful contralto... the woman sang so tunefully as to turn the dreadful rubbish into an almost pleasant sound.” The previously vacant tune becomes meaningful and expressive. Her humanity flips the song

from a background saccharine sound wash to something artistic and enjoyable.

There are two explanations for why the woman's rendition of the song was more enjoyable than the AI's version. For one, she could have manipulated the tune to make it sound better. Perhaps she added some lilting rhythmic freedom that was missing in the original. Perhaps she shaped the melody to highlight the contours of the phrase, making high points louder and low points softer. Perhaps her voice just sounded better than whatever mechanized voice came out of Orwell's imagined AI. Or perhaps she even added her own edits to the melody, making it a more elegant sequence of notes.

My previous chapters have been addressing this first explanation. My discussions have all focused on how and why it's hard for an AI—and specifically machine-learned LLMs—to make musical content that's as convincing as that created by human musicians. In *1984*, we could imagine Orwell's AI suffering from these same shortfalls, being simply worse at making music than the singing laundress.

But there's a second explanation, and it's the one I imagine at play in this dystopian scene. I imagine the AI producing flawless artistic content. I imagine it's impossible to distinguish between human-written and machine-made music in Orwell's world. In this reading, *1984*'s protagonist found AI-composed music dull, gray, and emotionless simply by virtue of the fact it was made by something that wasn't human—a computer. However, when he saw a flesh-and-blood human singing the song's words and melody, the tune was transmogrified into something expressive. The very fact that a human entered the artistic process allowed Orwell's character to find meaning and enjoyment in the performance.

If AI overcomes all the obstacles outlined in the previous pages—if computers compose music entirely indistinguishable from human-made music—I believe we still won't enjoy it as much as we enjoy human-made content. As listeners, we won't emotionally connect with music that we know was made by a computer, and we'll yearn for music that definitively connects to the human experience. This chapter explores the human component of music and its implications for AI's development.

I suggest that there are two main reasons we won't connect with AI-generated music: 1) AI programs are not truly creative, at least not in the same way that humans are, and 2) AI does not experience the human world. I'll argue that for music to be heard as meaningful and expressive—for it to be understood as *Art*—it needs both these components, and machine learning systems are not capable of incorporating these components. Both points have been argued for decades or even centuries, and I'll lay out some of the historical background behind each thread. I'll discuss two main retorts that scholars and engineers have often used

against these arguments. First, that deep learning models will soon be able to reliably create content so impressive that it will be indistinguishable from human creativity, and second, that future machines will be so sophisticated that they truly can experience the world like a human. In response, I'll argue that a particular type of creativity—what I'll dub *intentional artistic creativity*—is tethered to human-to-human connections and lived bodily experiences, and these relationships can't be replicated. In particular, I'll suggest that debates about generative AI often fail to take the connection between audience and creator into account, and that this connection is particularly evident in music. When we focus on the role that human-to-human contact plays in musical expression, and creativity in general, we start to see clear limits around the kind of content that generative AI will be able to produce convincingly.

The Lovelace Objection, and the definition of *creativity*

Chapter 1 began with a vignette featuring Ada King, Countess of Lovelace, one of the greatest mathematical minds of the 19th century. She was the only child of the poet Lord Byron's marriage to Lady Anne Byron, and from an early age, she was fascinated by mathematics, scientific engineering, and—potentially because of her father's reputed “insanity”—the human brain. Although a consummate member of the aristocracy both through her parentage and through her marriage to William King, the Earl of Lovelace, she worked within these fields throughout her short life. She collaborated with Charles Babbage, the engineer who theorized the first digital computer, and in this capacity she was tasked with translating and annotating an 1842 article by Luigi Federico Menabrea, an Italian professor of mechanics who outlined Babbage's machine for the Royal Academy of Sciences at Brussels. While Menabrea's original explanation was slightly over 8,000 words, Lovelace's notes and annotations would top 20,000. Famously, she described how one might feed the machine certain presets for an iterative task, creating what is generally considered the first computer program.¹

It is writing these very notes where we joined Ada Lovelace in Chapter 1. In the first of these notes (“Note A”), she described the future potential of the computer. She theorized that any task that can be represented digitally, that is, in the “abstract science of operations,” can be translated into a computational function. Singling out music, she writes:

[If] the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations [i.e., encoded in computer-readable formats], the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent.

Later in her notes, however, she dismisses the potential of such machinery to ever exhibit true intelligence, stating that the output of such algorithms would simply reorder the components put into them. In other words, computers can only ever process the content given to them, and anything they generate would be a reshuffling of their source content rather than anything creative or new.² In Note A, she claims a computer “weaves algebraic patterns” just as the “loom weaves flowers and leaves” into fabric using presets and human-made designs, while in Note G, she writes:

The Analytical Engine has no pretensions whatever to originate any thing [sic]. It can do whatever we know how to order it to perform. It can follow analysis; but it has no power of anticipating any analytical relations or truths. Its province is to assist us in making available what we are already acquainted with.

In 2024, the Recording Industry Association of America filed a lawsuit against the music-generating AI companies Udio and Suno. This lawsuit follows many similar ones filed by authors, newspapers, visual artists, and graphic designers against a number of AI companies, and the list of such suits gets longer every day. These lawsuits allege precisely what Lovelace argued 181 years prior: Whenever a computer generates content, that content is not *new* and *original* but rather the reordering of the computer’s source material. Following this logic, when 21st-century machine learning models train using the work of authors, musicians, or graphic designers, the resulting content is a reconstitution of the original dataset of text, music, or images.

In Lovelace’s estimation—and in the allegations of these lawsuits—a computer cannot be *creative*. In Chapter 2, I relied on a basic definition of “creativity” with two essential criteria: that a creative act must 1) make new, novel content that 2) conforms to some series of norms rendering that content understandable and valuable to an audience. In that discussion, I worked under the assumption that a process that did not directly replicate its source material *verbatim* was generating something novel and was therefore “creative.”

Lovelace, however, suggested a higher bar. To her, processing and reorganizing source material is just reshuffling, even with the complex processes we see used in contemporary neural networks.³ To Lovelace, computers are not creative. Instead, she imagined some deeper transformative process at the root of true human creation, some spark that produces original content fundamentally distinct from any previous creation.

This argument makes a lot of sense, given how I’ve described AI’s learning processes. In Chapter 2, I outlined how these systems identify patterns in some datasets, and in Chapter 4, we saw how these patterns

are represented as mere strings of references within a computer's digital mind. Given that these patterns are regurgitated from the source content, they are not "creative" under Lovelace's logic. Instead, they are better understood as presets determined by the original dataset. This is the *Lovelace Objection* to artistic generative AI. Here, computational creativity will always be limited to reshuffling its source material, regardless of how complex its reorganizing process is. The Lovelace Objection also moves beyond the definition of creativity I adopted in Chapter 2. It assumes that true human creativity involves something more than reordering source data, some more holistic and synthetic method of creating content.⁴

Before complicating this objection and diving into various rebuttals, I want to outline the basics of a second historical objection to generative AI, that of the mid 20th-century neuroscientist, George Jefferson.

The Jefferson Objection, and the humanity of art

In 1949, George Jefferson, a professor of neurosurgery at the University of Manchester, published "The Mind of Mechanical Man" in the *British Medical Journal*. The article is a response to the new "calculating machines" ("computers" to us) that burst onto the scene in the handful of years since the end of World War II. Jefferson noted that computer engineering has many similarities to the human nervous system in the way that it sends, stores, and analyzes information through a bevy of interconnected electric pulses. These similarities made him wonder whether mechanical minds could have the capacity to mimic human cognitive processes. Jefferson summarized his conclusion:

Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain—that is, not only write it but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants... I conclude, therefore, that although [an] electronic apparatus can probably parallel some of the simpler activities of [the] nerve and spinal cord... it still does not take us over the blank wall that confronts us when we come to explore thinking, the ultimate in mind. Nor do I believe that it will do so.

To Jefferson, any computational system will be found lacking because it doesn't experience, interact, and feel the world around it. Importantly,

he doesn't seem to care whether it *can* “write a sonnet or compose a concerto,” but that it does so “because of thoughts and emotions.” If Lovelace focuses on the mechanical processes underpinning the creation of new content, Jefferson worries about the *cause* behind the process. He believes any art or music that AI generates will be found hollow because it lacks the experience and motivations that define a flesh-and-blood human mind.

Again, this argument makes sense, given what we've seen in previous chapters. In particular, Chapter 4 showed the level of abstraction and removal that a tokenization process engages in. Not only are deep learning machines focused on the patterns within a dataset rather than the underlying meaning of those patterns, but the items of data in an AI's “mind” are mere computational placeholders for whatever original content may have existed in its dataset. Machine-learned models, then, are several layers removed from the actual feeling—the pleasure, grief, and misery Jefferson describes—that might have inspired the dataset.

The *Jefferson Objection*, then, posits that human experience forms the essential foundation for creative and artistic content. By definition, computational systems can never replicate this experience. Because expression relies on human experience, AI-made visual art, music, dance, and literature will be hollow. Regardless of how complex, detailed, and human-like a computational system might become, because it does not share our lived experiences, there will always be a chasm between human- and AI-produced content.

As with the Lovelace Objection, there are plenty of pros and cons surrounding the Jefferson Objection. To begin engaging with these issues, I'll outline two famous rebuttals to these arguments, beginning with one from a founder of modern computing, Alan Turing.

Turing's retort

In Alan Turing's 1950 paper, “Computing Machinery and Intelligence,” he argued that computers will likely exhibit some form of intelligence at some point in the future. The paper is famous for proposing the *Imitation Game*, what later researchers would call a *Turing Test*. In its original form, the Turing Test involves a human having a conversation with another entity via text, typing on a computer, and seeing responses on the screen. The interlocutor is either a human or a text-generating chatbot, but its identity is unknown to the conversing human. If a chatbot fools the human into believing it's another human, it has won the Imitation Game. This paradigm has been extended and abstracted into several creative domains and is often used to judge whether some computer program produces content that's passable as a human creation. If a human audience

believes that some prose, poetry, image, or music generated by an AI was actually created by a human, that AI has “passed the Turing Test.”

Turing did not originally view his test as a method of AI quality control. Rather, he imagined his test as a way to determine whether a machine can exhibit some form of intelligence. In Turing’s estimation, we only base our belief that other humans are “intelligent” according to the content of the various sorts of communication we exchange with one another. We do not peer into the neurons within our fellow humans’ skulls to deduce whether they have a consciousness and intelligence similar to our own. No, we make this deduction by exchanging conversation, text, art, music, and so on. If someone produces intelligent content, they must be intelligent. They’ve passed a simple version of the human-to-human Turing Test.

Turing extended this logic to computer programs. If computers create content that *seems* to come from human-like intelligence, we can reasonably assume that the computer possesses such intelligence. By locating the definition of “thinking” in this way, Turing argued that a machine that reliably produces human-like content can “think.” In this logic, the underlying programming or engineering is beside the point. It doesn’t matter whether a program’s inner workings mimic neurology or whether its consciousness is structured exactly like that of a human. After all, given that we never glance into other humans’ consciousnesses or brains, those facts don’t really matter to our day-to-day definition of consciousness anyhow.⁵

Turing therefore has no time for the Lovelace and Jefferson Objections.⁶ It doesn’t matter what sorts of experiences or processes lie behind an AI, only that it *acts like* it’s intelligent. From this standpoint, an AI that reorders its source data so convincingly that its content seems “new” will counter the Lovelace Objection. An AI that creates content that seems like it’s experienced deep human emotion similarly undercuts the Jefferson Objection. And from where Turing sat in the middle of the 20th century, he saw no reason that sophisticated future computers would not be able to accomplish both tasks with flying colors.

The Haugeland-Chalmers retort

Some theorists go even further. They suggest that AI doesn’t just act like human consciousness, but that it can feature the same inner workings as a human’s consciousness. To these thinkers, computers are theoretically capable of learning, remembering, thinking, feeling, and producing in much the same way as humans, and technologies like neural networks and LLMs go some distance in this direction.

In 1985, the cognitive philosopher John Haugeland coined the term *Good Old Fashioned Artificial Intelligence* (GOFAI) to describe the majority of 20th-century computer modeling. At the heart of GOFAI is

the idea that an analyst can pinpoint specific ways objects or concepts are represented within a computer program. For instance, if you were to examine the computational framework of a GOFAI designed to generate images, you could precisely identify the individual parts of the program responsible for creating images of teddy bears and those responsible for skateboards. Similarly, in a music-generating GOFAI, it would be possible to specify exactly which parts of the program are responsible for creating specific musical elements such as scale degree 1, scale degree 5, half notes, quarter notes, downbeats, and phrase lengths.

Haugeland contends that the human brain and mind do not operate like GOFAI. It's notoriously challenging to identify the particular neurons associated with particular cognitive concepts or actions.⁷ Unlike a library where each book has a specific place on a specific shelf, the human brain is not compartmentalized into discrete slots for each thought or idea. Instead, human cognition resembles ever-changing cloud formations, congealing into shapes only to dissolve and rearrange into new configurations. Brain regions interact dynamically, resonating with various networks of neurons to generate thoughts, memories, or emotions.

Haugeland argues that more recent and complex computational architectures like neural networks mirror the abstract and interconnected nature of the human brain. He refers to these systems as New Fangled AI (NFAI). Unlike traditional AI systems, the complex neural networks that form the basis of deep learning systems do not maintain straightforward one-to-one correspondences. For example, within the computational architecture of Google's chatbot Gemini, there is no single location where the concept of a "sonnet" resides. Nor can one pinpoint a specific node in the neural network of the image generator Dall-E where the shape of a teddy bear is stored. The same goes for trying to locate a distinct place in the Udio musical AI that represents scale degree 1. Instead, these concepts exist as part of a complex web of interlaced connections that extend deep within the network's layers.

To Haugeland, the amorphous and complex networks of connections within modern AI systems closely resemble the dynamics of human cognition. This logic suggests that contemporary deep learning models like LLMs, with their enormous computational architectures, are structured in ways that parallel human thought processes. This design allows them to approximate how the human mind operates, making them increasingly capable of mimicking human reasoning and learning.

The philosopher David Chalmers adds a further layer to this argument. Chalmers emphasizes a key aspect of human cognition: our minds integrate various senses and modes of expression in composing our thoughts. For instance, red and orange remind us of warm temperatures, and that association is an important part of how and where those colors are used in

images. A picture of a cookie might remind us of a particular taste associated with a specific childhood memory, while a musical gesture might be informed by the types of dance moves it often accompanies. Such associations will contribute to how a painter or composer might use these ingredients. Colors, cookies, and musical gestures gain meaning by reaching across different senses, types of experience, and modes of expression.

Chalmers contrasts this with the domain specificity of current machine-learning models. Chatbots, for instance, operate primarily with linguistic data, image generating AIs handle visual inputs and their descriptions, and music-making LLMs focus on auditory elements. These systems are constrained within one or two mediums and therefore cannot make complex connections across different senses and modes of experiencing the world. An image-generating AI might know how to place a picture of a cookie in a larger tableau to evoke a sense of nostalgia, but it wouldn't know anything about the cookie's taste or its reference to times past.⁸

However, Chalmers argues that recent advancements in AI technology are gradually overcoming these limitations by integrating multiple types of knowledge. He introduces the concept of an *LLM+*, a more advanced model that learns connections between several domains. Such a system would enhance the AI's ability to generate more sophisticated and nuanced content. Imagine a music-making machine that understands the imagery of lyrics, an image-generating program that connects pictures of baked goods with actual tastes, or a chatbot that knows the warmth of a color. The operations of such models would better approximate human cognitive processes. They would learn and create in a manner much more similar to how humans undertake these tasks.

Chalmers envisions *LLM+s* incorporating sensory data from optical scanners, audio processors, or even tactile sensors to create a richer, more interconnected AI system. This integration would enable the AI to associate sensory experiences with data inputs, like linking the tactile sensation of heat with "warm" colors, or a specific taste with the image of a cookie. Theoretically, an *LLM+* that harnesses datasets across a wide range of media and sensory inputs could parallel human consciousness by learning, thinking, and interacting with the world across different senses and domains, creating memories and associations like a human.

The Haugeland-Chalmers retort further undercuts the Lovelace and Jefferson Objections. Lovelace focused on the limitations of computational systems, suggesting they could only reshuffle existing inputs, while Jefferson emphasized the absence of emotions and experiences in machines. Contrary to these views, Haugeland and Chalmers propose that, through the advancement of abstract programming and the use of increasingly diverse and integrated datasets, computers might indeed mimic human cognitive processes and potentially share a similar inner

life. They argue that these technological improvements could enable AI to perform tasks with a complexity and adaptiveness that resembles human thought, erasing the cognitive gap between machines and humans.

Intentional artistic creativity

But I don't buy it. I think these retorts are so mesmerized by AI's output and engineering that they forget the crucial role the audience plays in realizing artistic expression and how the audience relates to the creators of that expression. Do we value the origins of creative content? Do we consider the internal experiences and emotions of the creators? Is an artist's or composer's identity important to us when we engage with their work?

To my mind, the answer is a clear and resounding "yes." My reasoning focuses on three words: *Intention*, *Art*, and *Creativity*.

I find philosopher Daniel Dennett's concept of the *Intentional Stance* particularly useful in approaching the first part of my formulation. Dennett's theory focuses on the many different ways that humans interpret events in the world. For instance, when we think of a car starting, we make sense of that event in terms of the mechanical processes involved. The car is switched on as a result of the interaction of its ignition system with the rest of the internal combustion engine. As we ponder the car, we might also consider how useful it is as a mode of travel, or we might worry about the machine's contributions to air pollution. However, we do not consider the car's *intention*. Because it's a machine, we don't ask ourselves about the car's motivation for switching on.

In contrast, when we focus on a *person* starting a car, our interpretation shifts. When a human turns an ignition switch, it's not a strictly mechanical event. Rather, it is the result of some human motivation. Perhaps the driver needs to travel somewhere. Perhaps they need to pick up their child at school or get groceries. Or perhaps they're starting the car to run an engine diagnostic. Whatever it may be, they started the car *for a reason*. When we observe human actions, we attribute to them some series of motivations.

This is the Intentional Stance. It's a viewpoint that believes actions are undertaken for some reason and with thoughtful motivation. On the one hand, the Intentional Stance involves empathy. We know what motivates our own actions when we engage a car's ignition, and we project that motivation onto other humans. On the other hand, it depends on viewers' expectations of the sophistication of outside actors. We do not imagine an automobile having consciousness, but we do imagine another human driver having a thoughtful inner life.

We constantly navigate the Intentional Stance in the realms of art and creativity. Consider a subway map displayed on a train platform. We would likely interact with it like we would the mechanics of a car's engine. We engage with it pragmatically, and we use it to figure out which trains to take to our desired destination. If the map is particularly well designed, we might even admire its clarity, usefulness, and aesthetic appeal. However, our appreciation remains anchored to its functionality and efficiency in fulfilling its purpose.

But what if an artist created that same subway map and hung it in a museum? Staring at the map as an installation, our stance would be dramatically different from when we stood on the train platform. We would wonder what the artist meant to convey with their choice of a subway map. We'd consider what the creator was trying to express. Maybe we'd wonder why they chose a particular set of colors or why they used a particular layout. Maybe we'd even scoff at this piece of modern art, thinking how silly and gimmicky it is to hang a subway map in an art gallery. Maybe we'd shrug our shoulders, bewildered at the artist's objectives.

Whatever our reaction, we'd be taking an Intentional Stance toward the map. The platform map and the gallery map may be identical in terms of visual content, but because we know the gallery map was placed there by a human for some reason, we make sense of it very differently than the utilitarian map on the subway wall. We *interpret* the gallery map while we *use* the map on the subway wall.

These different paradigms begin addressing the second of my trio of words, "Art," and these differences can be usefully articulated in the age-old distinction between Artworks and Craftworks.⁹ According to this traditional dichotomy, a Craftwork is characterized by its remarkable construction. It engages its audience through its complexity, sophistication, and the skill used to create it. It elicits pleasure, enjoyment, and occasionally awe. Typically, Craftworks serve some practical purpose or utility. Examples abound: a well-crafted table, a handmade pot, finely wrought jewelry, or an elegantly designed subway map. Each of these items not only performs its intended function effectively but also delights the senses through its craftsmanship.¹⁰

However, Artworks are, in the words of the philosopher Arthur Danto, "interpreted things."¹¹ An Artwork compels an audience to probe deeper into its meaning, to feel emotions, and to ponder some underlying message. It invites the audience to find meaning and expression in its contents. When a painting encourages introspection, a poem makes you cry, or a novel gives you chills, these things become Art. (From now on, I will use capital-A "Art" when referring to this concept, differentiating it from the more general visual "art.")

Just as in the Intentional Stance, Art depends on how humans engage with some given content. The content will change in the mind of the viewer depending on how they are encouraged or inspired to interpret it. In my formulation, the Intentional Stance and Art are intrinsically linked. A piece of content is elevated to the status of an Artwork when it compels its audience to ascribe meaningful interpretation and adopt the Intentional Stance to discern that underlying meaning and motivation.

This definition of Art and Intention can lead to a honed understanding of the last of my word triplet: “Creativity.” In Chapter 2, I described creativity as a process that generates new, useful, and legible content. By integrating Art and Intention in this definition, we can carve out a special case for creative content that encourages an audience to interpret, analyze, empathize, and find meaning. I’ll refer to this sort of content as *intentional artistic creativity*. Works of intentional artistic creativity are those that provoke thought, encourage interpretation, or foster emotional engagement with an audience. The audience believes that the content has some motivation or expressive intention behind it, and they are compelled to interpret it accordingly. Intentional artistic creativity is understandable and useful, but it is also *meaningful*.

By shifting the focus toward the entity that creates the content, the concept of intentional artistic creativity begins to undercut Turing’s retort. In this paradigm, an audience isn’t simply consuming a complex Artwork, shrugging their shoulders, and assuming the entity that created that Artwork must be correspondingly sophisticated. Instead, the audience probes and considers that Artwork, in part, by trying to figure out what sorts of motivations and inner life might be behind the content. Here, we care about who made the content and why it was made. Unlike the Turing Test, in which the identity of the creator is completely removed from an audience’s interpretive calculus, the intentional artistic creativity paradigm requires that Art encourage an audience to probe the identity and intention of its creator to figure out the underlying meaning.

But surely a computer program can exhibit “motivations” and have an “inner life.” After all, the Haugeland-Chalmers retort suggested that sufficiently complex computer engineering could exhibit human-like consciousness. Couldn’t such a sophisticated model exhibit intention and artistry?

Once again, I assert that the answer is a resounding “no.” My position is grounded in the essential role that lived human experience plays in intentional artistic creativity. But before I can effectively revisit and reframe the Lovelace and Jefferson Objections through the lens of intentional artistic creativity, we must first delineate the role that bodily, lived human experience plays in musical and other forms of artistic expression.

Associational versus experiential knowledge and the importance of the lived experience to Art

As I write this, my son is in a developmental stage in which he asks a lot of questions about words. Once, when we were talking about a car being parked “under a tree,” he interrupted to ask what it means to be “under.” I gave him a series of descriptions, explaining that when two things are stacked, one will be above and one will be below. “What’s it mean to be ‘stacked?’” he then asked, followed by, “What’s it mean to be ‘below?’” I found myself in an increasingly complex tangle of synonyms and self-referential definitions. Finally, I took two objects from his pile of toys and showed him the toys stacked on top of one another to illustrate what it means to be “below,” “above,” and “under.” He then mimicked the activity, and as he stacked, unstacked, and restacked objects, he proudly narrated the different relative positions of his toys.

My son’s questions highlight the difference between what I call *associational* knowledge and *experiential* knowledge. When I was listing off synonyms, I was giving my son associational knowledge. I was showing him which words are similar to one another and indicating situations in which each might be appropriate. I was listing off the contexts he could *associate* with the words “under,” “above,” and “below.” This sort of knowledge is about meanings, definitions, and usage, and it hinges upon the web of connections between these meanings and how they are used.

In contrast, when my son was stacking toys on top of one another, he developed physical, relational, spatial, and bodily understandings of those words. This experiential learning grounds the definitions of words in the tangible world and our interactions with it. Here, words like “under,” “above,” and “below” derive their meaning from our lived experience with them.

Both approaches are important to learning and communication. I have, for instance, never been to Antarctica. However, I’m confident it’s cold there because I’ve read and heard about the temperature in that part of the world, and the pictures I’ve seen of Antarctica are replete with snow and ice. I’ve *associated* the word “cold” with Antarctica. While I have never *experienced* Antarctic weather, I have experiential knowledge of “cold,” “ice,” and “snow.” I’ve held ice in my hands, bundled myself against the cold, and felt snow against my skin. I can rely on associative knowledge to connect these words, images, and ideas with Antarctica, and ground those associations in my prior experiences.

Humans employ both systems when creating content. Consider an individual who composes a poem about romantic heartbreak despite never having personally endured such an experience. This person might research the topic, studying how other poets have expressed those feelings and

analyzing accounts of what it feels like to be rejected by a loved one. Armed with this knowledge, they can craft a poem that encapsulates the essence of heartbreak. This creation is rooted in associational knowledge. The poet draws on information and emotional descriptions from others to construct narratives and turns of phrase that align with shared cultural understandings of heartbreak. However, if someone gets their heart broken and then expresses those emotions through poetry, they would be using experiential knowledge to write a poem. They choose words and phrases to reflect what they're feeling, and the poet pours their own emotions into the resulting lyrics. They are creating content to express a lived experience.

Experiential knowledge captures lived human meaning, while associational knowledge operates within the realm of how words and symbols convey that meaning. Again, this distinction has been explored by scholars from various disciplines. For instance, the early 20th-century philosopher Maurice Merleau-Ponty argued that all human perception is based in our bodily experiences. To Merleau-Ponty, words, concepts, and ideas all point back to some basic embodied experience, from the chill of low temperatures against one's skin to the neuro-chemical wrenching of heartbreak.¹²

Contemporary cognitive scientist Stevan Harnad arrives at a similar conclusion by examining how meanings are attributed to words, concepts, and images. Harnad's concept of "symbol grounding" occurs when words, images, or references ("symbols") are linked to real-world experiences. Without this grounding, words and images are merely referential, existing in a vacuum where symbols only relate to other symbols, like the piles of synonyms I shuffled in front of my frustrated son. However, when a symbol is connected to an experience—as when my son saw and felt the positions of his toys in front of him—that symbol becomes anchored to a concrete meaning.

Intentional artistic creativity relies on experiential knowledge. For a reader, viewer, or listener to take the Intentional Stance and ask themselves, "What does this content mean?" or "Why did an artist make this?" they need to be convinced that whoever made that content was motivated by some actual idea or experience. If an audience is to view something as an Artwork to be interpreted, they must believe there is something concrete to interpret—the grounded ideas of experiential knowledge—rather than empty self-referential symbols of pure associational knowledge. We don't simply want words, images, and melodies to *seem like* there's intention, understanding, and motivation behind them. We want there to *actually be* some intention, understanding, and motivation behind them. And for there to be meaning behind that intention, the meaning must be grounded in experience. In short, we need to see experiential knowledge inside intentional artistic creativity.

Of course, we don't hold all content to such a high bar. We certainly don't scrutinize every piece of media we consume for its intentional artistry or for some deep experiential meaning, and I will address these cases shortly. Not all content, in other words, is intentional artistic creativity. However, music strongly exemplifies the dynamics of intentional artistic creativity, and very often, it connects to the lived experience of both its creator and its listeners. Therefore, before returning to the larger idea of artistic expression and AI, I want to survey music's particular role in the dynamics of intention, expression, and human experience.

Music as a radically experiential and intentional Art: meaning, social functions, and expressions of authenticity

Music is a deeply emotional art form. As I noted in Chapter 1, when I was a teenager, closeted and struggling with depression, I would end tough days by putting on my headphones and blaring my favorite albums alone in my room. Music became a tool for me to process and cope with my emotions and helped me grapple with my position in a difficult world. A decade later, when my husband and I got married, music was central to some of the most meaningful and memorable parts of the service. The hymns, songs, and instrumentals were vectors in one of the happiest moments of my life. Then last year, when my dad entered hospice and we began planning his funeral, one of our first discussions was about the music we would include. Music would be a focal point for my family's sadness at my dad's passing.

While every art form evokes strong emotional connections, music's connections run particularly deep and are especially dependent on how an audience views a creator's intention and lived experience. This linked reliance on human experience and intention is evident in at least three aspects: the music's meaning, its social role, and the way we associate musical "authenticity" with the identities and biographies of its creators.

A staggering amount has been written about this first category—the ideas and concepts that music communicates to its audience, or in other words, what music *means*.¹³ Despite the multitude of debates and disagreement on the topic, a perennial issue tends to pervade these arguments. Namely, music does not possess clear lexical meaning, at least not as clear as words or objects. You can't pinpoint the concrete *meaning* of a chord or melody as straightforwardly as you can look up the definition of a word in a dictionary or as obviously as you can describe the objective contents of a picture.

The word "cat" and a picture of a member of the species *Felis catus* both have very clear meanings. They both directly indicate a small, furry semi-domesticated creature, and if you or I were to look at that word or

that image, we'd both have the same basic associations. The meaning of the musical climax of your favorite song, on the other hand, is far more amorphous. While it might evoke feelings and emotions, these effects are not specific, they're hard to describe, and you and I might experience very different emotions at that moment. We can't just look up the meaning of some musical gesture in a dictionary.

Despite its indeterminate nature, there is no doubt about music's ability to evoke feelings and meanings within listeners. Philosophers often note that musical gestures have recognizable emotional analogs, which Suzanne Langer describes as music's "general forms of feeling." Thinkers like Robert Hatten and Peter Kivy discuss this aspect in terms of metaphor and simile, showing how music can *act like* an emotion or idea. For example, if a melody repeatedly strives but fails to reach a particular note, it can elicit feelings of frustration, disappointment, or sadness, emotions that mirror the melody's failure to achieve its goal. This approach views music as a form of emotional onomatopoeia, where the sounds of music reflect the emotions they evoke. Theorists like Arnie Cox and Janna Saslaw further this notion by suggesting that such a melody gains meaning by triggering our preconscious mind to empathize with its struggle or failure. We imagine what we would be feeling if we were experiencing the same struggle as the melody.

Because musical meaning is amorphous and does not rely on concrete symbols, it directly appeals to our lived experience, becoming rooted in experiential knowledge. When writers like Arnie Cox suggest that we empathize with some musical gesture and imagine ourselves undertaking the same action as some musical phrase, they are suggesting that human experience bestows meaning into the music. When theorists such as Suzanne Langer and Robert Hatten argue that the melodies, harmonies, timbres, and phrases that hit our ears have some foundational similarity with the emotions we feel in our guts, they are also arguing that music gains meaning via human experience.

Despite having varied approaches, all such theories draw a direct connection between musical meaning and lived embodiment, underscoring how music's expressivity relies on the human condition and our collective experiences. These various theories of musical meaning suggest that music's ability to evoke emotion or empathy within us creates direct, immediate, visceral effects in its listeners. It is almost as if music cuts through any layers of associative knowledge altogether to elicit experiences in its audience. And of course, in order for music to express some kind of lived human experience, we need to imagine that experience coming from a musical creator—actively and intentionally—who shares or understands that experience. This is the first of music's deep connections to intentional artistic creativity: its meaning has a direct and raw connection to human experience.

The second crucial connection between music, experience, and intention is its social component. Music gains its potency from being a *social act*. Again, music is not alone among Art in this fact. The Russian philosopher Leo Tolstoy, for example, highlighted Art's ability to forge connections within a community as one of its primary benefits. This sentiment is supported by evolutionary biologists who suggest that artistic activities like sculpture, painting, and music may have supported social bonding in early human societies.¹⁴

However, many scholars argue social connections are especially essential to music. Theorists like Theodor Adorno and Cora Palfy consider music a container for communal values and meanings. Palfy describes music as a “virtual agent,” a persona built from the collective parallels we make between music and human actions. Mariusz Kozak similarly argues that musical concepts like rhythm and meter gain meaning from communal experiences of movement and dance. Furthermore, thinkers like Walter Benjamin, John Molyneux, and Christopher Ariza contend that seeing, feeling, and hearing an actual human perform is a crucial component of musical meaning. Benjamin or Molyneux, for instance, point to the fundamental difference between experiencing live performances and listening to music via a device like a streaming service or radio. Ariza, on the other hand, argues that the embodied process of creating music together—jazz musicians improvising intertwined melodies or a vocalist collaborating with a pianist as they interpret a phrase—is inseparable from musical meaning. In each of these instances, music is tethered to how we experience the world in community with other intentional and expressive humans. This is the second of music's connections to intentional artistic creativity and experiential knowledge. It gains its value, meaning, and content from being shared between humans, and it is used by those humans to create social experiences and community.

Finally, notions of “authenticity” undergird many musical styles. These tend to rely on the identity and background of the human making the music. Scholars like Justin Williams, Keith Negus, and Pete Astor examine how genres like rock and hip-hop are often deemed more authentic when created by individuals from specific racial, socio-economic, or regional groups. For instance, hip-hop and rap are genres deeply rooted in the specific lived experiences of economically and racially oppressed communities. As a result, “authentic” expressions of these genres are often tied to those experiences. Creators from outside these communities may not have the lived experience necessary to access an authentic rap or hip-hop voice. (Consider, for instance, the dissonance some listeners feel between the artist Drake's affluent background and his hip-hop persona.¹⁵)

This interest in a composer's biography can extend to how we interpret individual pieces of music as well. The public's decades-long obsession

with connections between Taylor Swift's love life and her music provides a clear example. We attempt to identify the true ("authentic") meaning behind her songs by voyeuristically scrutinizing the twists and turns of her personal life.¹⁶ We find meaning in trying to identify what aspects of Swift's lived experience she is intentionally embedding into her songs.

The third broad connection between music, experience, and intentional artistic creativity is, then, in the link between human biographies and the meaning and authenticity of the music. Our experience with music is bound to the identity and expression of the artist who made it. Once again, music is tethered to actual, lived human experiences and the intention of the artist.

Revisiting the Lovelace and Jefferson Objections in the face of music and AI

In the rooms of closeted teenagers, wedding celebrations, and parents' funerals, music bolsters the Lovelace and Jefferson Objections while deflating the retorts of Turing, Haugeland, and Chalmers. Meaningful and expressive music asks an audience to take an Intentional Stance. It asks its audience to find their own embodied, lived experience in the music. It invites listeners to empathize with the performer or creator, and it trades on notions of meaning and authenticity that are tied to the artist's lived experience. It acts as a tie to bind us together as a society while connecting to our own individual embodied human experiences.

These dynamics suggest that we won't find music—or any other Art for that matter—truly "creative" if we know it's made by AI. Lovelace argued that we care about the process by which content is created, and the same holds true for the plaintiffs in the many contemporary lawsuits alleging copyright infringement by AI companies. They assert that an algorithmic process is somehow fundamentally different from the way human intuition works. Algorithms merely reorder the source content fed to them, while humans can create something new. Turing attempted to undercut this objection by focusing on the content itself, suggesting that humans can be won over by sufficiently sophisticated computationally generated content and that these impressive machines can even be considered "intelligent."

But, intentional artistic creativity—and especially artistic music—requires an audience to connect and empathize with a creator in order for that Artwork to be understood as true creativity. I maintain that the transformative spark that Lovelace finds lacking in computer algorithms *can be located in the audience's knowledge that a human intentionally produced that content.*

When we want to engage with meaningful Art, we will search for content that we know was made by other fleshy, experienced humans. If we don't know whether a piece of content was made by a human, we will try to find out. And, if we can't be assured that it was, in fact, created by a human, we won't find it artistically creative. Even if an AI produces flawless music that passes Turing Tests with flying colors, we won't take the Intentional Stance toward it unless we see a human behind the creative process. Instead, as soon as we know the content was created by a computer, we'll treat the music more like a car's engine than the human starting the car.

To my mind, these dynamics vindicate the Lovelace Objection. By eliciting a connection to a composer, producer, or performer, the Lovelace Objection suggests that artistic and creative music cannot be generated by a computer, simply because we need to see an intentional human behind the creative process.

But *why* do we need a human behind that process? The answer lies in the dichotomy between associational and experiential knowledge. This is what corroborates the Jefferson Objection and deflates the Haugeland-Chalmers retort. Machine-learned models, even very complex LLMs, create content using only associational knowledge. As we saw in Chapters 2 and 4, these deep learning systems learn the patterns of some datasets by turning them into arbitrary referential symbols, the very epitome of associational knowledge. A neural network knows where to place a word, how to arrange objects in an image, or which note to place in a phrase, because it's learned the appropriate contexts and behaviors of those words, objects, and notes. But it has no access to the experience-based meanings behind these patterns and symbols.

At its core, an AI does not experience the world. "Solving" this problem by giving the models more data with more complexities—expanding AI's abstract connections using Haugeland's New Fangled AI or reaching across different domains as in Chalmers's LLM+ concept—merely kicks the can down the road. Adding more and more associational knowledge into these systems is not going to flip their understanding into the realm of experience. More associations will never add up to an experience. Just as the potential poet reading all he can about heartbreak won't *experience* that emotion until he's been faced with a devastating personal loss, complex associational systems can only ever *know about* ideas and emotions. They will never *experience* those concepts.

An excellent AI musical model will know where to write a downbeat and may even be able to list all the social, embodied, theoretical, and expressive reasons to write a downbeat. But because it has never jumped or danced or clapped its hands, it will never know what it's like to feel the bodily thud that accompanies a particularly satisfying arrival of a strong beat.

The Jefferson Objection posited that the experience of being human was crucial to Art, and because of the role that experiential knowledge plays in intentional artistic creativity, that rings true to me. When an audience takes an Intentional Stance and actively interprets Art, its members are looking for aspects of their own lived experience. This is particularly relevant in the case of music, where the basis of expression is abstract, embodied meaning and emotion. Here again, for something to become Art, we need to see a human creating it, and this is especially true for music.

These ideas aren't just theoretical. They're supported by observable human behavior. Researchers like Chiara Longoni, Luca Cian, Francisco Tigre Moura, and Charlotte Maw have shown that humans are uncomfortable with AI undertaking more subjective, taste-oriented tasks like recommending clothing or the best recipe for dinner. Further, these studies show that, when we spend our money, we would rather not spend it on something made by AI. Humans also like art less when they believe it's made by AI than when they believe it's by another human. The labs of Lucas Bellaiche, Linwan Wu, Taylor Jing Wen, Nils Köbis, and Daniel Shank have conducted behavioral studies showing over and over again that people like poems, images, and music less when they *believe* they were made by AI, even when they are indistinguishable from human-made content. Participants in these experiments report that the AI doesn't seem to be telling a story, isn't writing from experience, and isn't being truly creative. In other words, these experiments tell us that people view AI as not intentional and devoid of experiential knowledge.¹⁷ On several levels, our guts seem to be telling us that AI-generated content cannot be intentional creative Art.

But what of the more speculative and futuristic side of Chalmers's argument? Wouldn't his idealized LLM+, trained on all sorts of media and integrated with sensory data from a digital body, *experience* the world? By connecting sensory and musical data, couldn't that sort of model *feel* the rising and falling of musical phrases? Couldn't it *experience* a downbeat by mapping the feeling of dancing onto musical beat patterns? Couldn't these complex interconnections lead to consciousness and even creative intention? Couldn't it even experience grief and write music about that feeling?

These are interesting if speculative eventualities, but I remain skeptical. Even if something like this superintelligent experiential consciousness came to pass, its intentions and experiences would be different from ours. It might be very smart and knowledgeable, and it might know a huge amount about the human world,¹⁸ but it would be a silicon-based mind with an inner life very different from our own. It would not be *human*.

And humanity is important. It's hard to imagine a teary-eyed teenager running to their room, throwing on headphones, and listening to music

generated by an algorithm, regardless of the quality of the music or the AI that created it. That computer might know a lot about sadness, and it might have more raw data associated with depression than any human could. But it won't be capable of feeling sad in the way that a human feels sad. A computer can't be bullied. It will never cry.

My dad died shortly after entering hospice. I wrote and arranged several pieces of music for his funeral. Of course, I could have used a musical AI to help with this music. The prospect, however, seemed bizarre, almost disgusting. The idea of using an AI in a religious service celebrating life and mourning death would be antithetical to the essence of such an event. Funeral music facilitates our collective experience of something deeply human, and my dad's family and friends were coming together in empathy and community for such an experience. A musical AI could have a mountain of data about funeral music, and about what sad and mournful music should sound like. Maybe it could put all the notes in the right places. But mourning families do not program funeral music just to hear the right notes played at the right times. Funeral music connects us on a human level as we express human emotions. During my father's funeral, it mattered that humans wrote, arranged, and performed the hymns, anthems, and solos. It mattered that there was a lived experience behind the notes we were hearing and singing.

These dynamics return us to the musical scene in 1984. As the protagonist listens to the laundress inject her humanity into the shallow AI-generated tune, Orwell adds some further descriptions of the surrounding scene. He writes:

But the woman sang so tunefully... He could hear the woman singing and the scrape of her shoes on the flagstones, and the cries of children in the street, and somewhere in the far distance a faint roar of traffic, and yet the room seemed curiously silent.

When the protagonist hears the woman transform the song into something artistic, the world becomes quiet, and his thoughts turn to the footsteps and voices around him. His attention is heightened as he connects the woman's singing to audible expressions of humanity. As the song becomes Art, it sonically magnifies the surrounding human world. When the music switches from plastic AI-generated content to embodied creative expression, the protagonist becomes more intertwined with the human, embodied, social, and visceral world around him.

This expressive, communal, and human experience is exactly what generative AI cannot provide. If musical AI begins to produce flawless content, we would—as the Lovelace Objection notes—find its combinatoric computational process deeply inhuman. We would not take the Intentional

Stance toward that computational algorithm, and we wouldn't find it creative. Even if a complex futuristic model replicates human expression, it will be the result of silicon-based processors rather than emanating from a fleshy, warm, fragile body. As the Jefferson Objection argues, this computational experience will be foundationally different from human experience. It won't know what it's like to cry. Its computational replication of tears would be some manufactured analogy. It won't know what it's like to lose a father. A computer was never born, and it cannot die, at least not in the human sense.

Even if we shovel new and diverse data into future AIs, the pile of information won't replicate a human's actual lived embodied experiences. Given that intentional artistic creativity—and musical creativity, specifically—relies on human experience and community, that distinction will remain important.

AI as wallpaper, Craftwork, and tool

Not every piece of content requires its consumers to plunge into the Intentional Stance and treat the item as an Artwork. Indeed, much of the content we consume in our day-to-day lives is useful, well-made, and enjoyable without being Art. To my mind, well-made AI music can fill these sorts of non-artistic roles.¹⁹ I believe this type of musical content will function in one of three ways: as a background wash, a tool, or a Craftwork.

There's a lot of music in the world we're not meant to pay attention to. We don't consciously listen to music in elevators, bathrooms, hotel lobbies, or five-second commercials on social media. Sometimes we don't even pay much attention to film scores, our exercise playlists, or podcast intro music. So much background music doesn't need to be *good*; it just needs to be *there*. Such enjoyable background sound wash is one of the three ways that I see musical AI being used in the future. When the purpose of music is not to be meaningful but, rather, to fill silence, AI will be a useful tool. Because we don't usually understand this music as capturing or expressing the human experience—it's not meant to be intentional, artistic, or creative—we'll be happy to use AI to fill silence and color the background sonic landscape.

The second role I envision for musical AI is as an aid to human composers. Just as David Cope used his 1980s AI to pull him out of writers' block, musicians can use AI to generate nearly infinite amounts of melodic and harmonic ideas for musical inspiration. Let's give credit where it's due: AI excels at detecting and reconfiguring patterns with a speed and precision beyond human capability. Musicians can leverage this capacity to spark new ideas, accelerate their compositional process, and even solicit

edits to refine their work. In this capacity, AI will function as an aid to human creativity, acting as a tool in service of the final musical product.

Finally, meticulously constructed AI music will be regarded more as impressive Craftworks than Artworks. Recall from our earlier discussions that audiences find Craftworks to be beautiful and impressive, but they don't inspire that audience to seek meaning, expression, or intention within them. In a Craftwork, our focus is less on expressive meaning and more on the pleasure and utility the work provides. We do not engage with a Craftwork through the Intentional Stance, nor do we necessarily find ourselves intrigued by the identity or biography of its creator. While the intricacies of a marble carving might prompt us to ponder a sculptor's expressive intent, we don't usually ask the same questions when admiring the beauty and design of an ornate piece of jewelry. Artworks invite interpretation and meaning. Craftworks are to be appreciated for their design.

Future musical AI will excel at Craftwork. Even now, when I listen to an exceptional piece of AI-generated music, I hear moments of well-made content. Some of the tunes are catchy, and I've even encountered turns of phrase that I find clever. I can find pleasure in the beauty of an AI's chord progressions and delight in the elegant construction of its melodies. Yet, I never find myself pondering the music's meaning, and I won't attempt to interpret an AI's song or use it as a lens for introspection. As long as audiences care about the humans behind musical Artworks, even the finest and most convincing pieces of AI-generated music will be devoid of true human expression. They will, therefore, remain outside the realm of Art. Compelling and impressive AI-generated music will instead be better understood as pieces of Craft.

Some thoughts on the future

I began this chapter by evoking a famously dystopian novel and referencing the bevy of lawsuits that creators are currently bringing against AI companies. I ended by outlining several situations in which AI-generated music will likely succeed in our society. These sentiments were sandwiched by a step-by-step argument about the specific circumstances in which AI will *not* succeed.

I'm a realist. AI will soon be *very good* at making content of all kinds. Will these advances eliminate jobs? Yes. Will humans try to pass off AI music as their own? Certainly. Will fewer kids be inspired to learn to read music and gain fluency on an instrument when they can create impressive content at the press of a button? Likely. Will the next decades be laden with lawsuits about musical ownership surrounding AI models, their outputs, and their training data? Probably. Will younger generations be more

open to finding meaning and expression in musical AI? Maybe.²⁰ Do I wish I could limit or reverse these eventualities? Definitely.

We should engage in pro-music, pro-musician, and pro-creator advocacy. We should pressure regulators and the AI industry to clearly label machine-generated content. We should educate our students about the limits of AI. We should expose our audiences to the beauty of human-made music and live performances. And we should be smart and thoughtful consumers of AI-generated content.

However, AI music will never stamp out human-made Art. Streaming services, records, radio, the printing press—none of these extinguished our collective craving for live performances or newly composed, human-made music, and yet each elicited worry that they would do exactly that.²¹ If Art is about expressing the human experience, generative AI will never be able to create it. It simply cannot capture the full depth of human emotion and the complexities of our experiences. AI lacks the ability to feel and to live through the unique circumstances that shape each composer's—each human's—voice. While AI can mimic patterns and styles of human-made content, it cannot experience a human life. And audiences will continue to care about this, because we care about the people who make the Art we love.

I also think that music will be at the frontline of the discomfort with, and even the resistance to, AI-generated content. For one, as we have seen, musical AI simply develops more slowly than other media. This will give us the opportunity to more thoughtfully reflect on its expansion and evolution. As my previous chapters argued, musical AI is allocated fewer resources for research and development than other areas of AI research, has smaller datasets available to it, is more difficult to represent to a computer, and contains structures and organizations that are difficult for machine learning to identify and learn. However, at some point, musical AI will be capable of truly excellent outputs. When this happens, I believe that music's foundational human components will be put into stark relief. We'll realize that machine-made music lacks a reliable human connection, and we won't like it.

AI will be a tool. It will produce background music. It will even produce beautiful Craftworks. But, as long as we place value on genuine emotional expression and the unique perspectives that only humans can provide, human composition will remain irreplaceable.

Summary

In this chapter, I have argued that audiences do not interpret music generated by AI the same way as they do human-made music. While generative AI may one day create technically perfect music, we will continue to value

the lived experience behind musical expression, a lived experience that will remain inaccessible to AI.

I began by outlining two reasons that we might perceive AI-generated music as unsatisfying: 1) AI programs are not truly creative, and 2) AI does not experience the human world. I connected these arguments to two historical thinkers, Ada Lovelace and George Jefferson, and suggested that their skepticism still applies to contemporary generative AI. In the Lovelace Objection, computer-based intelligence is unable to generate truly novel ideas because it is limited to the information fed into it. The Jefferson Objection, then, argues that artistic expression is based on embodied human experience, and computers will never have access to that experience. I first contrasted these views with arguments from Alan Turing, who believed that, if AI can convincingly mimic human outputs, it should be considered intelligent. I also outlined the thinking of contemporary scholars John Haugeland and David Chalmers, who argue that sufficiently complex computer models might, indeed, have experiences and consciousnesses similar to those of humans.

My own thinking, however, is that human interaction and emotional input are essential for music to be fully appreciated as Artwork. Both aesthetic theory and behavioral data suggest that audiences care deeply about the identities of the artists creating content, and that music is a particularly potent example of this dynamic. Even as AI becomes increasingly human, the very fact that it is *not human* precludes it from creating content that a human audience will interpret as Art.

I ended by suggesting that generative AI will produce music to fulfill three musical roles: as a sonic background filler, a compositional tool to aid human musicians, and a producer of technically proficient but emotionally detached content. Each role leverages AI's strengths in pattern recognition and efficiency while acknowledging its limitations in providing genuine artistic and emotional depth. While AI can support and enhance the process of musical creation, it cannot replace the human elements essential for true artistic expression. The human experience, with its inherent emotions and contextual understandings, remains central to creating music that moves and connects with human audiences. As long as we place value on emotional expression, human experience will remain vital to music as an artform.

Some final thoughts

At the beginning of this book, I outlined five forces that affect musical AI's development. My first chapter explored these five key forces: 1) the *motivations* behind AI development, 2) the *examples* used to train these systems, 3) how music is *represented* to a computer program, 4)

the *structures* a program needs to generate coherent music, and 5) how humans *interpret* and consume AI-generated music. The second chapter addressed the first force. Using some basic market analyses, I showed that musical AI research receives less investment and fewer resources compared to other forms of AI, and that this imbalance is at least partly due to its limited commercial applications.

AI needs vast datasets to learn how to successfully replicate a given medium, and Chapter 3 discussed the scarcity of usable musical datasets. I described the challenges in extracting reliable information from both audio files and musical scores, and I contrasted this with the relatively straightforward ways that information is extracted from text and images. Chapter 4 examined the complexity of music's overlapping elements, which make it difficult for researchers to provide their AIs with digestible musical chunks and cause AIs to stumble when creating new music.

Chapter 5 then analyzed the types of patterns and structures that AI excels at learning, particularly *nested determined proximities*, sequences of nearby, frequently occurring events that follow some overarching logic. I reviewed music generated by some current musical AIs and demonstrated how these systems favor nested determined proximities. I contrasted these characteristics with the more varied, dispersed, and unpredictable structures of human-made music.

Finally, the current chapter connected these engineering problems back to how we enjoy and value music. In the last several pages, I argued that even if AI produces flawless music, it won't resonate with people as much as human-made music. These issues then connect back to the broader problems that began this book. If the human element is a crucial reason that audiences value music, why should we expect users and financiers to invest in musical AI? The lack of motivation is compounded by the lack of value that listeners place on music created by AIs.

At the end of the classic 1939 movie *The Wizard of Oz*, the mechanical Tin Man famously receives a "heart" from the titular wizard. The Tin Man had been constructed with nothing in his chest, which he believed made him unable to feel or experience emotion. The wizard then presents the Tin Man with a plastic trinket in the shape of a heart with the words, "A heart is not judged by how much you love, but by how much you are loved by others."

This sentiment represents an issue underpinning this entire book, and especially this final chapter. We can worry about market forces, and how computers represent music. We can dive deeply into how machine learning models work, and scrutinize the datasets they require. We can wring our hands all day about whether an AI can truly be conscious, or if some computer is capable of feeling. We can theorize about different types of knowledge, and we can make analogies between neural networks and the

brain all day long. But, if we do not value music made by a computer as much as human-made compositions—as I believe we do not and will not—there will be limited motivation to develop musical AI, compile large datasets, parse that data, and train it to learn sophisticated meaningful musical structures. We won't believe that musical AI has a "heart," and that belief will undercut all other forces that might otherwise support the technology's development.

If that turns out to be the case, I'll be very happy.

Notes

- 1 See Menabrea (1843); Stein (1984); Hollings, Martin and Rice (2018); and McCully (2019).
- 2 Lovelace's criticism was itself influenced by contemporary British Romantic aesthetic theory. English thinkers like John Keats (themselves influenced by continental philosophers like Immanuel Kant and Georg Hegel) argued that original creativity arises from the artist's connection to some set of ethereal larger truths, and the artist transforms their relationship with those truths into a new and unique work of art (Stein 1984; Bonds 2006). Because of her Romantic sensibilities, Lovelace would doubtlessly argue that a computer could not generate poetry of the quality created by a human. More important to my logic, her criticism is also about the authorial process. To Lovelace, a true artwork does not involve cutting and pasting from past artworks, but from an artist's personal relationship to the creative process.
- 3 Notably, this distinction is nowhere near new and has many precedents in theories of creativity in general. To Plato, the lowest form of Art involved copying while the highest was some representation of a larger ideal. Several millennia later, Samuel Coleridge Taylor would make a distinction between "Fancy," which takes items in the memory and reorders them, and "Imagination" which dissolves these associations and creates something truly new. Immanuel Kant considered someone a "genius" when they could synthesize something truly new from their intuition (Bonds 2006).
- 4 Other theorists have made similar arguments. In 1980, the philosopher John Searle criticized supporters of AI through his "Chinese Room" analogy. Imagine a man in a locked room with a library of Chinese phrases organized by some coding system. Intermittently, the man receives a letter with a sequence of codes on it, and he walks through the library withdrawing the corresponding Chinese phrases. He meticulously copies the phrases onto a paper, stuffs the result into an envelope, and returns the letter. In the parable, the man can *produce* Chinese, but he does not *know* Chinese. Searle suggests that this is the way AI works. When computers learn to generate an appropriate series of symbols in an appropriate context, they don't actually know about the content they're generating any more than the man in the room knows Chinese.
- 5 This sort of approach to AI would come to be known as "Strong AI." In Strong AI, we "don't need to know how the brain works to know how the mind works" (Searle 1980). Rather, we just need to analyze the behavior of humans and AIs to understand their intelligence. This contrasts with a "Weak AI" stance that argues that computational systems should be evaluated based on their outputs and utility. To a theorist of Weak AI—and I must admit that I

- count myself a member of this club—an AI can act *like* a human and can even *model* a human thought process, but it cannot actually *be the same as* human intelligence.
- 6 By calling these the “Lovelace Objection” and “Jefferson Objection,” I’m giving homage to Turing’s original paper, in which he describes “Lady Lovelace’s Objection” and directly cites Jefferson’s argument.
 - 7 See Satel and Lilienfeld (2013).
 - 8 Proust (1913).
 - 9 It’s likely that debates over the definition of Art began as soon as the paint dried on the earliest cave drawings. The contours of the debate about “What is Art,” how that conversation has changed over time, and the different notions of art between cultures fill piles of books. In *The Transfiguration of the Commonplace* (1981), for example, Arthur Danto suggests that part of what makes something an artwork is that it offers the possibility—if not the necessity—of interpretation. Importantly, Markowitz (1994) argues that the distinction between what is art and what is not is based on the race and class of both the viewer and the creator. I must admit that I hold a particularly Romantic/modernist default in my own thinking, inasmuch as I center expressive meaning and authorial intention in my discussions and definitions of Art.
 - 10 See Markowitz (1994), Danto (1981), and Dissanayake (2015).
 - 11 Danto (1981), 135.
 - 12 While I’ve singled out Merleau-Ponty, ideas of embodied cognition weave their way through many academic discourses. Indeed, much intersectional feminism and Black feminism locates issues of identity and politics in how bodies exist in the power structures of the world around them (Richardson 2013, Allen 2011, Davis 2001, Laqueur 1990), and many theories of the mind rely on bodily metaphors (Saslaw 1996). In the past several decades, music studies in particular have seen increased focus on the role that the body plays in musical experience (Kozak 2019, De Souza 2017, Le Guin 2005, Palfy 2022, Cox 2016). Additionally, several AI researchers have explicitly focused on why physical neurons are different from the silicon-based processing chips of computers (Block 1995, 2009).
 - 13 For a few reviews of these discussions, see Agawu (1991), Cook (2001), and Palfy (2022).
 - 14 See Piilonen (2024) and Tomlinson (2015) for sophisticated and subtle descriptions of music’s role in theories of evolution.
 - 15 See Pope (2016) and Isai (2022).
 - 16 See, for instance, Spanos (2016) or Greene (2022). This sort of “authenticity” extends to classical music as well. Studies have shown that people like classical music more when they know that it was written by someone respected and famous (Fischinger *et al.* 2018, White *et al.* 2022), that it was made by people whose biography is similar to their own (Greenberg *et al.* 2021), and that it is relatable to their own experience (Pelowski *et al.* 2017). For a review of how the framing around a piece of music influences our assessment, see Leder and Pelowshi (2022).
 - 17 Bigman and Gray (2018).
 - 18 My argument here interacts a bit with the concept of AI superintelligence. A model would be “superintelligent” if it can undertake some reasoning or investigation faster and with more sophistication than could a human, and even could identify some solution using logic that humans can’t comprehend (Bostrom, 2014). For instance, a superintelligent AI might be able to make pharmaceutical breakthroughs by understanding interactions between organic and bodily chemistry that extend far beyond the capacity of human

- researchers. This type of superintelligence would be impressive, useful, and even scary—but, its knowledge would still be associational. My hypothetical pharmaceutical breakthrough would come from the superintelligent AI reading and studying inhuman amounts of biological, and chemical research about the human body, not because the AI *has* a body.
- 19 To use a term from Chapter 5, this music would achieve the status of an “adult” model.
 - 20 Indeed, Deirdre Loughridge’s historical analysis of how humans relate to “human-like” mechanized sounds shows that the relationship between humans and machines has changed over time. It’s a definite possibility that future generations will have very different understandings of mechanized expression and creativity than those I am describing here.
 - 21 See Blair (2003) and Fischer (2010).

References and Further Reading, Chapter 6

- Agawu, K. V. 1991. *Playing with Signs: A Semiotic Interpretation of Classic Music*. Princeton: Princeton University Press.
- Ariza, C. 2009. “The Interrogator as Critic: The Turing Test and the Evaluation of Generative Music Systems.” *Computer Music Journal* 33(2): 48–70.
- Allen, J. S. 2011. *¡Venceremos!: The Erotics of Black Self-Making in Cuba*. Durham: Duke University Press.
- Bellaïche, L., R. Shahi, M. H. Turpin, et al. 2023. “Humans Versus AI: Whether and Why We Prefer Human-Created Compared to AI-Created Artwork.” *Cognitive Research: Principles and Implications* 8(1): 42. <https://doi.org/10.1186/s41235-023-00499-6>.
- Benjamin, W. 1935 (1968). “The Work of Art in the Age of Mechanical Reproduction.” In *Illuminations: Essays and Reflections*, edited by Hannah Arendt, translated by Harry Zohn, 217–251. New York: Schocken Books.
- Bigman, Y., and K. Gray. 2018. “People Are Averse to Machines Making Moral Decisions.” *Cognition* 181. <https://doi.org/10.1016/j.cognition.2018.08.003>.
- Blair, A. 2003. “Reading Strategies for Coping with Information Overload ca. 1550–1700.” *Journal of the History of Ideas* 64(1): 11–28. <https://doi.org/10.2307/3654293>.
- Block, N. 1995. “The Mind as the Software of the Brain.” In *Thinking: An Invitation to Cognitive Science*, edited by Edward E. Smith and Daniel N. Osherson, vol. 3, 377–425, 2nd ed. Cambridge, MA: MIT Press.
- Block, N. 2009. The mind as the software of the brain. In S. Schneider (Ed.), *Science fiction and philosophy: From time travel to superintelligence*: 170–185. Wiley-Blackwell.
- Bonds, M. E. 2006. *Music as Thought: Listening to the Symphony in the Age of Beethoven*. Princeton: Princeton University Press.
- Bostrom, N. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Braguinski, N. 2022. *Mathematical Music: From Antiquity to Music AI*. New York: Routledge.
- Chalmers, D. J. 2022. *Reality+: Virtual Worlds and the Problems of Philosophy*. New York: W. W. Norton.
- Chalmers, D. J. 2022. “Could a Large Language Model be Conscious?” Talk Given at NeurIPS. <https://philpapers.org/archive/CHACAL-3.pdf>.
- Cook, N. 2001. “Theorizing Musical Meaning.” *Music Theory Spectrum* 23(2): 170–195.

- Cox, A. 2016. *Music and Embodied Cognition: Listening, Moving, Feeling, and Thinking*. Bloomington: Indiana University Press.
- Danto, A. 1981. *The Transfiguration of the Commonplace: A Philosophy of Art*. Cambridge, MA: Harvard University Press.
- Davis, A. 2001. "Black Women and Music: A Historical Legacy of Struggle." In *Black Feminist Cultural Criticism*, edited by Jacqueline Bobo, 217–232. Malden, MA: Blackwell Publishing.
- Dennett, D. C. 1971. "Intentional Systems." *The Journal of Philosophy* 68(4): 87–106.
- Dennett, D. C. 1981. "True Believers: The Intentional Strategy and Why It Works." In *Scientific Explanation*, edited by A. F. Heath. Oxford: Oxford University Press, 150–167.
- Dennett, D. C. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
- De Souza, J. 2017. *Music at Hand: Instruments, Bodies, and Cognition*. New York: Oxford University Press.
- Dissanayake, E. 2015. *What Is Art For?* Seattle: University of Washington Press.
- Fischer, S. 2010. "A History of Media Technology Scares, from the Printing Press to Facebook." *Slate*, February 15, 2010. <https://slate.com/technology/2010/02/a-history-of-media-technology-scars-from-the-printing-press-to-facebook.html>.
- Fischinger, T., M. Kaufmann, and W. Schlotz. 2018. "If It's Mozart, It Must Be Good? The Influence of Textual Information and Age on Musical Appreciation." *Psychology of Music* 48(4): 579–597. <https://doi.org/10.1177/0305735618812216>.
- Gray, H. M., K. Gray, and D. M. Wegner. 2007. "Dimensions of Mind Perception." *Science* 315(5812): 619. <https://doi.org/10.1126/science.1134475>.
- Greenberg, D. M., Kosinski, M., Stillwell, D. J., Monteiro, B. L., Levitin, D. J., & Rentfrow, P. J. 2021. The song is you: Preferences for musical attribute dimensions reflect personality. *Social Psychological and Personality Science*, 12(3), 300–308.
- Greene, J. 2022. "Taylor Swift's Dating History: How Her Past Boyfriends Wind Up in Songs." *Today.com*. October 10, 2022, updated April 18, 2024, posted at <https://www.today.com/popculture/music/taylor-swift-boyfriends-dating-history-rcna50748>.
- Harnad, S. 1990. "The Symbol Grounding Problem." *Physica D* 42: 335–346.
- Haugeland, J. 1996. "What Is Mind Design?" In *Mind Design: Philosophy, Psychology, and Artificial Intelligence*, edited by John Haugeland, 1–16. Cambridge, MA: MIT Press.
- Hollings, C., U. Martin, and A. Rice. 2018. *Ada Lovelace: The Making of a Computer Scientist*. Oxford: Bodleian Libraries.
- Hong, Joo-Wha, Katrin Fischer, Yul Ha, and Yilei Zeng. 2022. "Human, I Wrote a Song for You: An Experiment Testing the Influence of Machines' Attributes on the AI-Composed Music Evaluation." *Computers in Human Behavior* 131: 107239. <https://doi.org/10.1016/j.chb.2022.107239>.
- Huang, C. A., C. Hawthorne, A. Roberts, M. Dinculescu, J. Wexler, L. Hong, and J. Howcroft. 2019. "The Bach Doodle: Approachable Music Composition with Machine Learning at Scale." In Proceedings of the 20th International Society for Music Information Retrieval Conference, 100–107. Delft, The Netherlands: ISMIR. <https://arxiv.org/abs/1907.06637>
- Isai, V. 2022. "When Drake Is on Your Course Syllabus." *New York Times*. www.nytimes.com/2022/01/22/world/canada/drake-course-toronto-ryerson-university.html.

- Jefferson, G. 1949. "The Mind of Mechanical Man." *British Medical Journal* 1(4616): 1105–1110. <https://doi.org/10.1136/bmj.1.4616.1105>. PMID: 18153422; PMCID: PMC2050428.
- Kivy, P. 1980. *The Corded Shell: Reflections on Musical Expression*. Princeton: Princeton University Press.
- Köbis, N., and L. D. Mossink. 2021. "Artificial Intelligence Versus Maya Angelou: Experimental Evidence That People Cannot Differentiate AI-Generated from Human-Written Poetry." *Computers in Human Behavior* 114. <https://doi.org/10.1016/j.chb.2020.106553>. (<https://www.sciencedirect.com/science/article/pii/S0747563220303034>)
- Kozak, M. 2019. *Enacting Musical Time: The Bodily Experience of New Music*. New York: Oxford University Press.
- Kozak, M. 2021. "Feeling Meter: Kinesthetic Knowledge and the Case of Recent Progressive Metal." *Journal of Music Theory* 65(2): 185–237. <https://doi.org/10.1215/00222909-9143190>.
- Kramer, L. 2002. *Musical Meaning: Toward a Critical History*. Berkeley: University of California Press, 2002.
- Langer, S. 1942. *Philosophy in a New Key*. Cambridge: Harvard University Press.
- Laqueur, T. 1990. *Making Sex: Body and Gender from the Greeks to Freud*. Cambridge, MA: Harvard University Press.
- Leder, H., and M. Pelowski, 2022. "Empirical Aesthetics: Context, Extra Information, and Framing." In *The Oxford Handbook of Empirical Aesthetics*, edited by Marcos Nadal, and Oshin Vartanian. Oxford: Oxford University Press. Online edition: <https://doi.org/10.1093/oxfordhb/9780198824350.013.43>.
- Le Guin, E. 2005. *Boccherini's Body: An Essay in Carnal Musicology*. Berkeley: University of California Press.
- Long, R. 2022. "Key Questions about Artificial Sentience: An Opinionated Guide." *Experience Machines* (Substack). <https://experiencemachines.substack.com/p/key-questions-about-artificial-sentience>.
- Longoni, C., and L. Cian. 2020. "Artificial Intelligence in Utilitarian vs. Hedonic Contexts: The 'Word-of-Machine' Effect." *Journal of Marketing* 86(1): 91–108. <https://doi.org/10.1177/0022242920957347>.
- Lonsdale, A. J., and A. C. North. 2009. "Musical Taste and Ingroup Favouritism." *Group Processes & Intergroup Relations* 12(3): 319–327. <https://doi.org/10.1177/1368430209102842>.
- Loughridge, D. 2023. *Sounding Human: An Expansive Analysis of the Relationship between Human and Machine in Music*. Chicago: University of Chicago Press.
- Markowitz, S. J. 1994. "The Distinction between Art and Craft." *Journal of Aesthetic Education* 28(1): 55–70. <https://doi.org/10.2307/3333159>.
- McCully, E. A. 2019. *Dreaming in Code: Ada Byron Lovelace, Computer Pioneer*. Somerville, MA: Candlewick.
- Menabrea, L. F. 1843. "Sketch of the Analytical Engine Invented by Charles Babbage, Esq." In *Scientific Memoirs, Selections from the Transactions of Foreign Academies and Learned Societies and from Foreign Journals*, edited by Richard Taylor, F.S.A., vol. 3, article 29, 666–731. London: Richard and John E. Taylor. Originally published in *Bibliothèque Universelle de Genève*, no. 82, October 1842.
- Merleau-Ponty, M. 1945. *Phénoménologie de la perception*. Paris: Gallimard.
- Molyneux, J. 2020. *The Dialectics of Art*. London: Haymarket.
- Orwell, G. 1949. 1984. London: Secker and Warburg.

- Ostermann, F., I. Vatolkin, and G. Rudolph. 2021. "Evaluating Creativity in Automatic Reactive Accompaniment of Jazz Improvisation." *Transactions of the International Society for Music Information Retrieval* 4(1): 210–222. <https://doi.org/10.5334/tismir.90>.
- Palfy, C. S. 2022. *Musical Agency and the Social Listener*. New York: Routledge.
- Pelowski, M., P. S. Markey, M. Forster, G. Gerger, and H. Leder. 2017. "Move Me, Astonish Me... Delight My Eyes and Brain: The Vienna Integrated Model of Top-Down and Bottom-Up Processes in Art Perception (VIMAP) and Corresponding Affective, Evaluative, and Neurophysiological Correlates." *Physics of Life Reviews* 21: 80–125. <https://doi.org/10.1016/j.plrev.2017.02.003>.
- Piilonen, M. 2024. *Theorizing Music Evolution: Darwin, Spencer, and the Limits of the Human*. Oxford: Oxford University Press.
- Pope, A. 2016. "Musical Artists Capitalizing on Hybrid Identities: A Case Study of Drake the 'Authentic' 'Black' 'Canadian' 'Rapper.'" *Stream* 9: 3–22.
- Proust, M. 1913 (1992). *In Search of Lost Time*. Translated by C. K. Scott Moncrieff and Terence Kilmartin. Revised by D. J. Enright. 6 vols. New York: Modern Library.
- Richardson, M. 2013. *The Queer Limit of Black Memory: Black Lesbian Literature and Irresolution*. Columbus: Ohio University Press.
- Sarmiento, P., L. Jackson and B. Mathieu. (2024). Between the AI and Me: Analysing Listeners' Perspectives on AI- and Human-Composed Progressive Metal Music. 10.48550/arXiv.2407.21615.
- Saslaw, J. 1996. "Forces, Containers, and Paths: The Role of Body-Derived Image Schemes in the Conceptualization of Music." *Journal of Music Theory* 40(2): 217–243.
- Satel, S., and S. O. Lilienfeld. 2013. *Brainwashed: The Seductive Appeal of Mindless Neuroscience*. New York: Basic Books.
- Searle, J. 1980. *Minds, Brains, and Programs*. Cambridge: Cambridge University Press.
- Shank, D. B., C. Stefanik, C. Stuhlsatz, K. Kacirek, and A. M. Belfi. 2023. "AI Composer Bias: Listeners Like Music Less When They Think It Was Composed by an AI." *Journal of Experimental Psychology: Applied* 29(3): 676–692. <https://doi.org/10.1037/xap0000447>.
- Spanos, B. 2016. "Ex-Factor: Taylor Swift's Best Songs About Former Boyfriends." *Rolling Stone*. <https://www.rollingstone.com/music/music-lists/ex-factor-taylor-swifts-best-songs-about-former-boyfriends-19644/>.
- Stein, D. K. 1984. "Lady Lovelace's Notes: Technical Text and Cultural Context." *Victorian Studies* 28(1): 33–67.
- Tigre Moura, F., and C. Maw. 2021. "Artificial Intelligence Became Beethoven: How Do Listeners and Music Professionals Perceive Artificially Composed Music?" *Journal of Consumer Marketing* 38(2): 137–146. <https://doi.org/10.1108/JCM-02-2020-3671>.
- Tolstoy, L. 1897 (1995). *What Is Art?* Translated by Richard Pevear and Larissa Volokhonsky. London: Penguin.
- Tomlinson, G. 2015. *A Million Years of Music: The Emergence of Human Modernity*. New York: Zone.
- White, C. Wm., M. O'Harra, K. Coker, I. Fuentes, S. Franciosa. 2022. "Exploring Implicit and Explicit Bias in Music." Paper Presented at the *Society for Music Perception and Cognition*, Portland, Oregon.
- Williams, J. A. 2011. "Historicizing the Breakbeat: Hip-Hop's Origins and Authenticity." *Lied Und Populäre Kultur* 56: 133–167.

- Wu, L., and T. J. Wen. 2021. "Understanding AI Advertising from the Consumer Perspective." *Journal of Advertising Research* 61(2): 133–146. <https://doi.org/10.2501/JAR-2021-004>.
- Wu, Y., Y. Mou, Z. Li, and K. Xu. 2020. "Investigating American and Chinese Subjects' Explicit and Implicit Perceptions of AI-Generated Artistic Work." *Computers in Human Behavior* 104: Article 106186. <https://doi.org/10.1016/j.chb.2019.106186>.