*Action, Criticism, and Theory for Music Education*
June 2025. Vol 24 (3): 202–44. doi:10.22176/act24.3.202

# Pedagogy of the Prompt: Music Education, Artificial Intelligence, and Big Tech Magic

adam patrick bell
Western University (Canada)

## Abstract

In this article I examine the generative AI music-making applications of Google called MusicFX DJ and Music AI Sandbox, as well as Udio, an end-to-end generative AI music system created by ex-Google employees. I frame these generative AI music systems designed by Big Tech as part of their neoliberal agenda, which involves influencing music education. Drawing on critical media theorists, I suggest that Big Tech, and in the context of this article, Google, present their generative AI music systems as "magic." I describe how these systems work by examining how algorithmic systems, such as search engines, operate more generally, and then, more specifically, how generative AI music systems function. I aim to make legible the "magic" of generative AI music systems and explain that Google's music-making tools, by design, appear to be agenda-less, but music educators should be wary of this illusion. I forward that Google's "pedagogy of the prompt" may not be explicit to users of their AI music-making applications, but nevertheless it is embedded into them because they are algorithmic systems, and more specifically, generative AI/machine learning-based systems. I suggest that in this present period in which music teachers and learners are using generative AI music systems, knowing how a musical output results from an input, musical or otherwise (e.g., text), is needed.

## Keywords

Music, music education, artificial intelligence, generative artificial intelligence, algorithm, Big Tech, Google, Udio, magic

In April 2024, Dustin Ballard, also known as There I Ruined It (2024), released "The Beach Boys Sing '99 Problems' by Jay-Z." This mashup featured the lyrics of the Jay-Z song from 2003 set to the melody and music of the Beach Boys' "I Get Around" from 1964 with an added dimension: Ballard used a generative artificial intelligence (AI) voice cloning model (Baris 2024) of the Beach Boys so that it sounds like Brian Wilson and his brothers are huddled around a microphone singing Hov's[1] homonymous line, "kiss my whole asshole," amongst other memorable lyrics in harmony. At the end of the video, Ballard included a brief disclaimer and social commentary: "The preceding was a work of parody which comments on the perceived misogynistic lyrical similarities between artists of two different eras." With this statement along with his social media tagline, "I lovingly destroy your favorite songs," it is apparent that Ballard intended these musical mashups to be provocative and based on the responses to this example and others that have followed a similar formula, he succeeded. For example, writing for *Wired* a year earlier about Ballard—in this instance his version of Johnny Cash singing Taylor Swift's "Blank Space" that went viral in the summer of 2023—Knibbs (2023) acknowledged the positive reviews Ballard had garnered, but concluded that this genre is "The most boring version of the future."

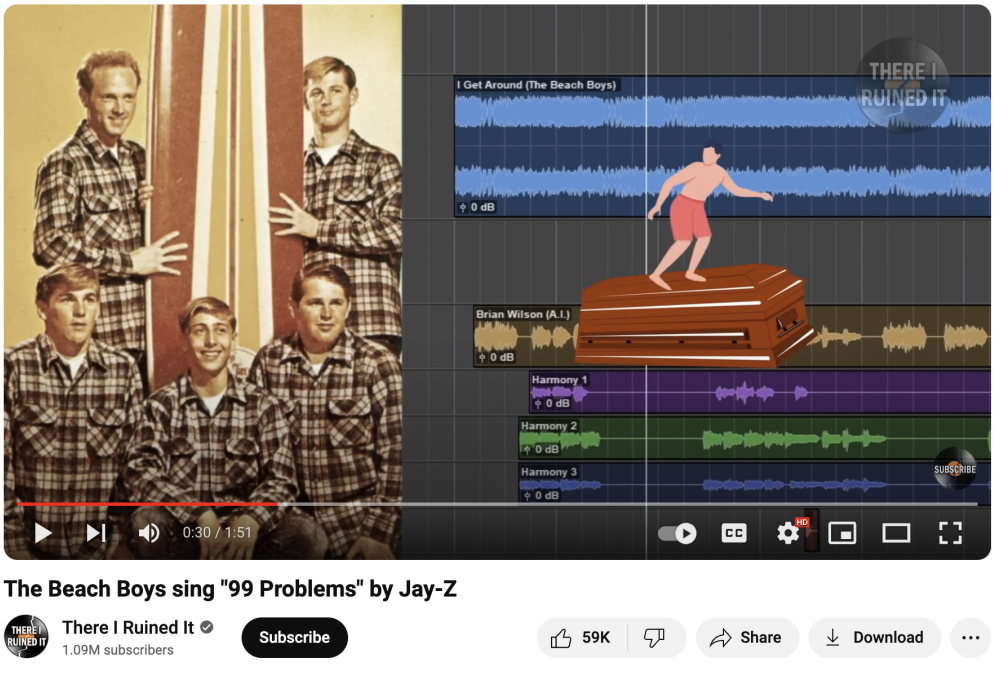


Figure 1: Screenshot of "The Beach Boys Sing '99 Problems' by Jay-Z."

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

Two decades prior to the sonic sorcery of There I Ruined It stoking a range of reactions on TikTok, Instagram, YouTube, and other social media platforms, Brian Burton, also known as Danger Mouse, released *the Grey Album* (2004)—a mashup of Jay-Z's *Black Album* (2003) and the Beatles' "White Album" (1968). Included on this album was a mashup that couched the lyrics and melodies of "99 Problems" within the music of the Beatles' song, "Helter Skelter" (J Scar 2009). In analyzing *the Grey Album*, including its version of "99 Problems," Adams (2015) averred that Danger Mouse's work is a resistive act. Knowing this context, the work of Dustin Ballard/There I Ruined It can be understood as an extension, perhaps an evolution, of the mashup genre. Although hailing from different technological periods, both Ballard and Burton represent innovators of their respective eras. Burton's work demonstrates a careful curation of both the Beatles' and Jay-Z's music, which he relied on to produce novel mashups built on the backbone of precise editing (Fairchild 2014). In comparison, Ballard's work is similar. As he shared on Reddit (Figure 2), precise (and time consuming) editing is a core element in his process, such as using tuning tools like Melodyne (ArmanTNT 2023).

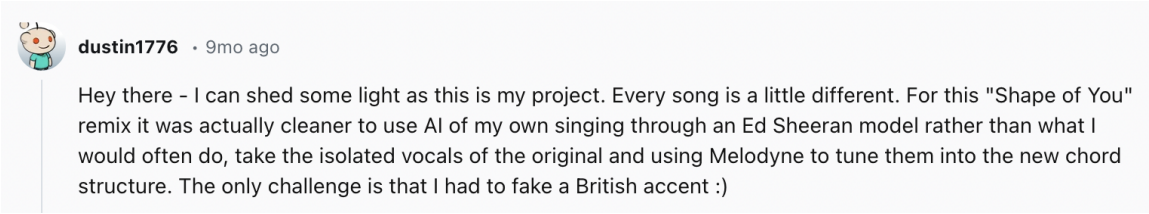


Figure 2: Screenshot of Dustin Ballard's Reddit Comment

And, as he shared in an interview with Rapp (2024), Ballard's process relies on his musicianship: "It would be hard to imagine pulling any of these off if I weren't a musician. Changing a melody to fit a chord structure it was never intended for requires some basic theory and a musician's intuition. Plus, sometimes I'm playing the instruments myself, as well as singing. Even the AI singing is actually my own singing put through a 'filter,' so the inflection, intonation and impersonation still needs to be there first" (para 6).

Comparing the "99 Problems" versions by Burton and Ballard provides a precise point of reference to discern how an AI tool, in this case a voice model, changes the possibilities within music production. Most of what Ballard does is based on

established practices that have been in use for decades. He does not wield a seemingly magical AI wand; instead, he interspersed newer tools with older ones and as result adds to the artistic possibilities within the mashup (and parody) genre. When we pull back the curtain, we see a lot of the same old tricks.

## Magic? Music Making and Artificial Intelligence

Beyond posthumous Johnny Cash covers and the "Taylor Swiftication" (or "ta-AI-lor's Versions") of other artists' songs (Collins and Grierson 2024), such as Japanese artist Yoasobi (AI COVER_JPN 2023), the tech sector has promoted generative AI for music making—or musickAIng (Sturm et al. 2024)—for its promise and possibilities. For example, in a video released by YouTube (2023a), Lyor Cohen, YouTube's "Global Head of Music," provided the following rationale: "Artists are always looking for new ways to get inspired. New ways to create great music and share their work. New ways for their songs to connect with people. Technology is, and has always been, the close collaborator in this. Now AI tools are opening up a new playground for creativity. We're experimenting with new ways to let artists, songwriters, and producers instantly hear their ideas. New ways to use technology to help them create music. Like singing a melody to create a horn section."

In addition to the video referenced by Cohen (YouTube 2023b), the Google DeepMind (2023a) site, "Transforming the future of music creation," features additional experiments, including:

- "Transforming beatboxing into a drum loop" (Google DeepMind 2023b)
- "Transforming singing into an orchestral score" (Google DeepMind 2023c)
- "Transforming chords from a MIDI keyboard into a realistic vocal choir" (Google DeepMind 2023d)

Continuing his narrative, Cohen revealed, "We're also starting to bring these possibilities to creators and fans with a new experiment called Dream Track. Just type in your idea and get a unique soundtrack for your Short." Examples of this text-to-music application include celebrity collaborations with Charlie Puth (YouTube 2023c) and T-Pain (YouTube 2023d). Cohen concluded:

> The potential of AI is incredibly exciting. But with any new technology, we have to approach it responsibly and that's not lost on us. It starts by recognizing that artificial intelligence is meant to amplify human creativity, not replace it. That's why we published a set of principles that include granting appropriate protections and unlocking new opportunities for artists.[2] And we're partnering with

> some of the most creative people on the planet to make sure AI will have a positive impact on the future of music.

Beyond the Googleverse, examples abound of how applications of generative AI are "transforming" music-making practices.[3] Song-generators make bold claims such as "Create original songs in seconds, even if you've never made music before" (Boomy n.d.); "Create any song. Just describe it," (Udio n.d.); "No instrument needed, just imagination. From your mind to music" (Suno n.d.); and "Collaborate with AI to create, customize and release unique music to social media, streaming platforms and more. Designed for modern creators. 100% Royalty-Free" (Loudly n.d.) (Figure 3).

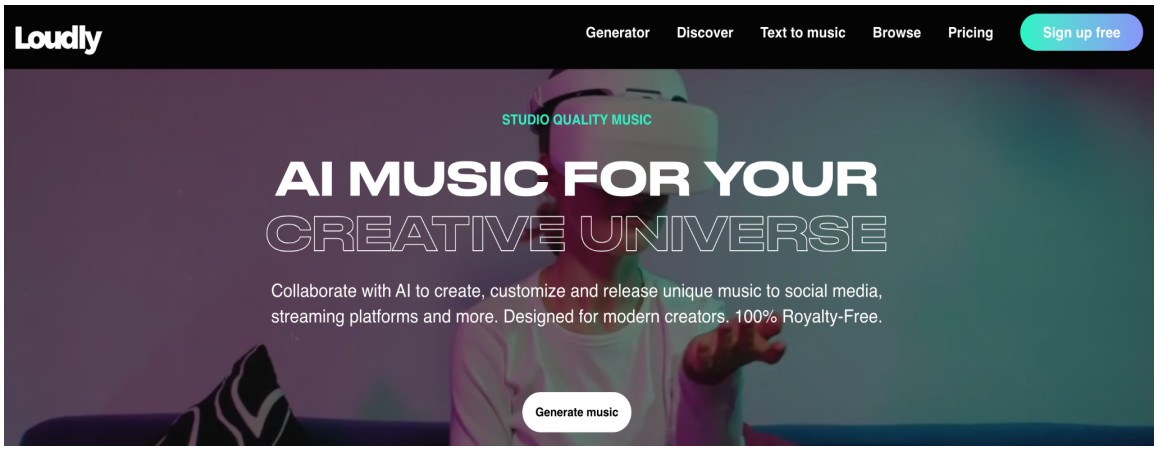


Figure 3: Screenshot of Loudly's homepage.

Each of these examples demonstrates Sterne and Razlogova's (2019) observation that AI has become a marketing strategy and can obscure the specific technological processes (e.g., machine learning) that take place within an application. Nagy and Neff (2024) affirmed that such cloaking is part of the magic trick performed by the tech sector, which they call "the conjuration of algorithms," and explain as having three purposes: "to conceal the design of technologies, create confusion around the capabilities of technologies, and produce dazzling effects of technologies" (4939). I aim to illustrate how the tech sector's smoke and mirror show creates the illusion of magical music making and that the primary prop to produce the trick is the prompt. On Instagram, music producer Timbaland stated, "I can put something on it I just gotta prompt" (Figure 4), in reference to using AI to finish his uncompleted songs, which demonstrates how the tech sector has succeeded

in convincing him, and perhaps others, that prompting is a future skill of music making.


Figure 4: Screenshot of Timbaland's Instagram video about using prompts to finish his songs

From There I Ruined It, an individual person, to Google, a multinational corporation, the implementation of AI in music making has drawn audiences who, regardless of their reactions to the experience, remain largely unaware of how the illusion is achieved. In my analysis of generative music systems that follows, I use the directions of Sterne and Razlogova (2021), albeit with a particular focus on the implications for pedagogy, given my focus on the music education profession. They

state: "To understand, assess, and intervene in the cultural politics of AI, scholars will need to consider broader questions of how work processes operate, the meanings of the work performed by the AI to its users, the ideologies operating in the interface, what kinds of data can be generated or acquired and the infrastructural conditions within which it must operate" (765).

The majority of my focus throughout this article is on Google and their generative AI music-making applications, MusicFX DJ and AI Music Sandbox. In addition, I also briefly examine an end-to-end generative AI music system called Udio. I begin by framing and situating the involvement of the tech sector in music education as a reflection of their neoliberal values before proceeding to examine how algorithmic systems operate more generally, and then, more specifically, how generative AI music systems work. I aim to make legible the "magic" of generative AI music systems and explain that Google's music-making tools, by design, appear to be agenda-less, but music educators should be wary of this illusion. I forward that Google's "pedagogy of the prompt" may not be explicit to users of their AI music-making applications, but nevertheless it is embedded into them because they are algorithmic systems, and, more specifically, generative AI/machine learning-based systems. I suggest that it would be beneficial for music teachers and learners using generative AI music systems to know how a musical output results from an input, musical or otherwise (e.g., text).

## The Company Curriculum: Neoliberal Music Education and Big Tech

In calling the tech sector's influence on music education "the company curriculum," I am aligning my thinking with the core ideas of Apple (1979/2019) and Giroux (2002), who described how neoliberal practices and policies have shaped educational agendas. Similarly, I conceptualize neoliberalism as a marketplace mentality that pervades all aspects of life, including music teaching and learning. I forward that those in the tech sector with business interests in music education have shifted their general approach from trying to influence formal education systems to bypassing them altogether because they can engage directly with learners through, with, and around their products and services. In short, technology firms

have positioned themselves as music educators. This point has been made previously (Bell 2015; Benedict and O'Leary 2019), but what has changed over the past decade is, (1) the approach taken, which, increasingly, is to position the learner as a "user" of a product or service, and (2) the rate of change, which has accelerated in this current period of generative AI fervor. In making these two points, I am primarily referring to Big Tech (this term typically refers to Alphabet, Amazon, Apple, Meta, and Microsoft) because they are the principal players in generative AI and—with the exception of Apple—are relatively new entrants into the arena of music education; however, in my analysis I also include examples from smaller music technology business entities because historically they have been key contributors in developing the landscape of music education as a marketplace. To contextualize my analysis of the company curriculum in and around music education, I first examine how other researchers have considered neoliberalism's influence and impact on music teaching and learning.

In the introductory editorial for the special issue on neoliberalism and music education for *Action, Criticism and Theory for Music Education*, Goble (2021) provided a primer on the origins and development of neoliberalism, charting how it has changed over time, and concluded that it has "weakened" music education. His assessment is representative of most music education scholarship that grapples with neoliberalism,[4] and leads me to ponder questions such as "How is music education being weakened by neoliberalism?" and "Where is this happening?" Influenced by Foucault, Varkøy (2021) framed neoliberalism as a "political rationality," reasoning that there needs to be a focus on macro-level issues, primarily competition, to understand micro-level issues. The macro/micro distinction is helpful for thinking alongside what previous music education researchers have discussed as it relates to neoliberalism. Some researchers (e.g., Aróstegui 2019; Bates 2021; Horsley and Woodford 2016; Woodford 2019) take a more global approach, albeit one that is clearly grounded in the Global North, discussing how neoliberalism's unwavering commitment to capitalist principles shape societies and their educational institutions. While there may be geopolitical references in these accounts, the authors focus more on the macro perspective, which serves to both illuminate and remind that unfettered capitalism—commerce without borders—aims to be everywhere, and this makes it difficult, if not impossible, to avoid its ever-widening

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

wake. Narrowing the focus to individual nation-states, it becomes clearer how neoliberalism manifests differently in various places. In Costa Rica (Rosabal-Coto 2016), Chile (Angel-Alvorado, Gárate-González, and Quiroga-Fuentesl 2021), India (Avis 2019), and Venezuela (Baker 2014), social inequities, often manifested as classism, can be manufactured in, through, with, and around music education. For example, El Sistema in Venezuela (Baker 2014; Fink 2016) and Costa Rica's similar Sistema Nacional de Educación Musical (Rosabal-Coto 2016) can be construed as mechanisms of control to preserve the neoliberal order.

Elsewhere, in the Anglosphere, including Canada, New Zealand, the United Kingdom, and the United States, the refrain sounds the same, that is, the neoliberalization of music education harms learners (Wright 2019). McPhail and McNeil (2021) explained how in New Zealand there is a blend of neoliberal and post-modern influences on music education, creating systems that vary from locale to locale characterized by a departure from humanist educational aims that have been replaced with instrumental ones. In the UK, Young (2021) detailed how the marketization and privatization of early childhood education has produced a two-tier system that separates poor and affluent families. In the United States, both Bates (2021) and Richerme (2022) identified neoliberal values promoted by the National Association for Music Education, and several other researchers focused on specific characteristics of neoliberalism evident in practices such as competition (Powell 2021), standardization and its sibling, measurement (Louth 2020; Mullen 2019), entrepreneurialism (Canham 2023; Moore 2016; Sadler 2021), and (misguided) multiculturalism (Sánchez-Gatt 2023). Finally, popular music education at multiple levels and in various places has received scrutiny as a site of neoliberalism (Abramo 2022; Benedict 2022; Bylica and Dillon 2024; Hall, Crawford, and Jenkins 2021; Smith 2015).

I situate my analysis amongst this latter group of studies and others like them that focus on specific sites of music education—a micro approach—to explain how neoliberalism operates within a given context; however, my subjects, Big Tech and smaller technology companies, are not typically perceived as sites of music education or organizations with music education agendas, such as Little Kids Rock/Music Will (Abramo 2022; Bylica and Dillon 2024) or Musical Futures (Benedict 2022; Hall, Crawford, and Jenkins 2021). Although some relatively brief critiques

of SmartMusic (Mullen 2019) and Apple (Smith 2015) serve as examples of technology companies approaching music education as a marketplace, this area of music education research that would benefit from more development, following Thompson (2023). Her in-depth research on teacher ambassador programs created by educational technology companies, which she likens to pyramid schemes run by multi-level marketing firms, revealed, "They invest in the teachers who become Ambassadors, and that investment often leads to more buy-in later down the road that can lead to more classrooms using the product, more schools mandating the use of the product, and more school districts using the product, and maybe even county, state, or federal policies that require the use of the product—a business plan" (163).

Although Thompson (2023) made a distinction between ambassador programs and certified educator programs, I conflate them because they serve the same end: making music education a marketplace and having teachers serve as salespeople to students. The corporate branding is explicit as one can become an "Apple Teacher" (Apple 2024), "Google Certified Educator" (Google 2024b) or enroll in the "Soundtrap Educator Certification Program" (Soundtrap 2024). In all cases, these training programs provide product-specific knowledge and skills. They may compete with each other (i.e., Google and Apple) or complement each other, such as Soundtrap promoting Chromebooks, a Google product. The company curriculum has become more overt over the past decade, coinciding with the emergence of Big Tech. Prior to this, there was a time when music technology companies seemed content to produce products and let people figure out how to use them following a do-it-yourself approach (Bell 2017), or have seasoned practitioner-educators facilitate learning experiences, but increasingly, the educator role is being squeezed out, or to use a more fitting music technology analogy, compressed and gated.

As one example of this trend, consider Ableton, a company that produces both music-making software (Live) and hardware (Push) and has been a forerunner in recognizing the potential of education as a market. They offer a free version of their software to classrooms (Ableton 2024a) and for a period of time had an initiative for teachers to apply for free Push units (the current language on their website now states "heavily reduced price"). They have an "Ableton Education Tour," in which "Ableton Certified Trainers" (Ableton 2024b) provide free workshops, and host an

annual conference called "Loop," where one can learn all things Ableton (Ableton 2024c). In this model, Ableton positions itself primarily as an educational resource by teaching educators and providing their software and hardware to schools either free or at a discount. In more recent years, their strategy has changed (or perhaps just expanded), shifting to a model in which they position themselves as the educator. This is apparent with their resources such as "Learning Music" (Ableton 2024d), which provides a series of interactive lessons on topics such as notes, scales, chords, basslines, and melodies, to name a few, all using Ableton interfaces; it is a prepackaging of product and pedagogy. This practice is commonplace in the music production industry. For example, the company iZotope, which makes various music production software products, makes "free" video tutorials, such as "How to Mix if You're Not a Mix Engineer" (iZotope Inc. 2023). The video contains much valuable educational content, but all of the skills conveyed are demonstrated with iZotope products. The company curriculum does not hide its capitalist agenda, as demonstrated by the screenshot of iZotope's "Learn" site (iZotope 2024) in Figure 5. In the menu, "learn" is wedged between "products," "shop," and "deals," and selecting "learn" presents a choice learn by topic or product. iZotope presents the illusion of choice, as both paths arrive at the same destination: learning with proprietary products.

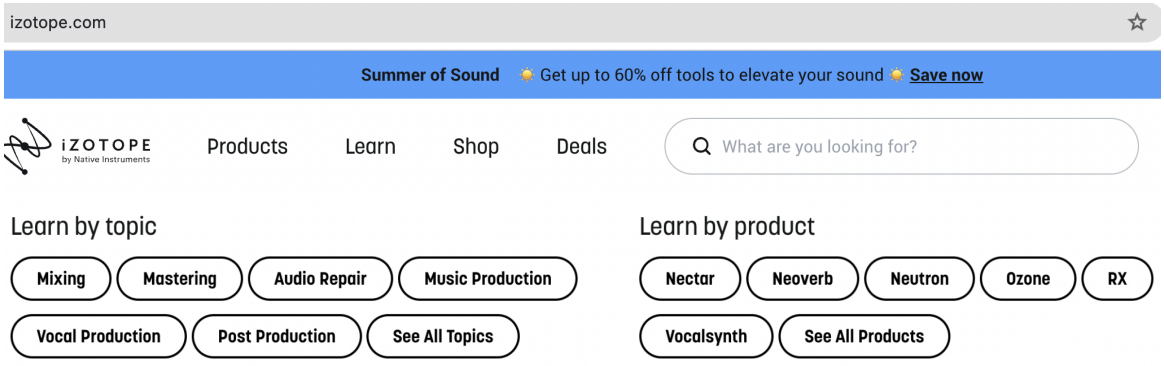


Figure 5: Screenshot of iZotope's "Learn" site

As technology companies position themselves as music educators in this manner, they also position their products and services as sites of music education, including their software, apps, hardware, but also their websites, YouTube channels,

and social media platforms (e.g., TikTok and Instagram). In contrast to the practice of the tech sector attempting to gain access to students-as-customers through school systems in the 2010s, the change (or addition) in course is to leverage products and services as sites of music education to engage directly with learners. Smaller companies such as Ableton may have started this trend, but Big Tech has charged ahead with abandon. Big Tech seems to recognize the (business) value in being involved, and ideally, *integrated* in music education and so position themselves to provide a related product or service. They may not go beyond this step, such as by creating resources for teachers, primarily because there is no need to do so. Whereas a thoughtful teacher would presumably approach an educational context by considering who, where, what, when, and how they will teach, Big Tech tends to append "ever" to these interrogative adverbs (i.e., whoever, wherever, whatever, whenever, however), and this in turn answers the "why" question, because the only thing that counts to them is use—they literally count users to assess their success. The following quote by Wang et al. (2024) sums up the neoliberal values that are embedded in the tech sector: "Due to the fact that music generation systems (1) can provide a far greater volume of content than composers can offer, their cost per minute of music is much lower than that of composers; (2) they can customize and personalize music based on user settings; (3) they can inspire composers and provide assistance; (4) they can be used for music education, among other applications" (6395). To experience the neoliberal education agenda embedded into Big Tech's generative AI systems, let's commence with a Silicon Valley spectacle: Google I/O.

## ~~Voila!~~ Viola! Googling as Music Making

After clamoring out of a colossal coffee mug in what appears to be Joseph's technicolor dreamcoat, improvisational musician Marc Rebillet stumbled around the stage at the pre-show for I/O 2024—Google's developer conference, where they introduce and demonstrate new products (Google 2024a). He uttered the word "Google" five times before asking the audience, "We all ready to do a little Googling? Everyone get out your computers. It's time to have a nice Google. Siri, navigate to google.com." Rebillet proceeded to build a beat using a MIDI keyboard controller, tabletop looper, voice processor, and two laptops. He riffed on the lyric,

"Google's going to wake you up," attempting to rouse the audience from their apparent slumber. Aside from the cascading visuals projected on the screen behind him, it seemed there were no Google technologies involved in Rebillet's routine. Following this first number, Rebillet introduced himself to the audience and with comedic flourishes explained: "I'm here to talk to you about music. And in particular, a very exciting new technology that we've been developing here at google.com—MusicFX DJ. That's right. No more do you need to haul around a case of records. No, the computer will take care of it. MusicFX DJ is a generative artificial intelligence tool that will, from scratch, generate from nothing a series of sounds as it interprets your prompts. I'll demonstrate that to you now, utilizing the technology at my fingertips and utilizing the colossal power of Google's machine—connecting to Google's machine now" (Google 2024a).

The interface of "MusicFX DJ" projected on the screen behind Rebillet was not yet publicly available, but anyone who has used Google or another search engine will recognize the familiar search bar that seemingly invites text entry. Rebillet typed "viola" while explaining to the audience, "We'll start with a viola, how about that? Everyone loves a viola. And if you don't, I don't care. I don't care. You do now. Viola." The sound of a viola began to play and Rebillet reacted: "No one wrote this. Google wrote this" (Google 2024a).

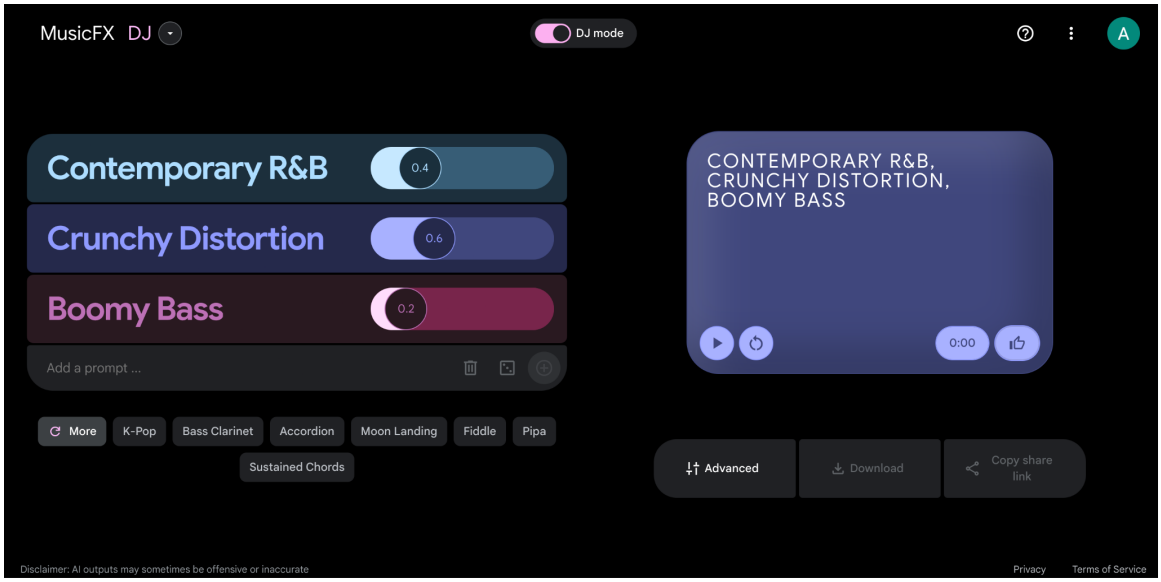


Figure 6: Screenshot of Google's MusicFX DJ User Interface

As depicted in Figure 6, when a word is typed into the interface of MusicFX DJ, a slider appears beside it, which gives the user control over how much or how little of the sound is mixed with the other sounds, akin to how a DJ uses a cross-fader between turntables. As Rebillet continued with his demonstration, he added the prompts "808 hip hop beat" and "Chiptunes" and adjusted the sliders of the three different sounds in search of a mix that appeased him before stopping the playback and stating to the audience: "The machine. The machine is in control now. And that's what it can do." Later, Rebillet facilitated an additional demonstration in which the audience shouted out their preferences from a list of six options presented on the screen (bagpipes, Persian tar, djembe, dulcimer, flamenco guitar, Detroit techno), with the underlying concept remaining the same—that a text-based prompt can generate new music. Put another way, Googling can be a music-making act. Toward the end of the presentation, Rebillet proclaimed, "You know, you're using it. It's helping you. The machine is good" (Google 2024a).

MusicFX DJ is the first example of a Google product that was demonstrated at I/O 2024. The second product, "Music AI Sandbox," was presented in the form of a prepared video (Google DeepMind 2024). Introducing the video, Doug Eck, Senior Research Director of Google DeepMind, boasted:

> I've been working in this space for over 20 years, and this is by far the most exciting year of my career. We're exploring ways of working with artists to expand their creativity with AI. Together with YouTube, we've been building Music AI Sandbox, a suite of professional music AI tools that can create new instrumental sections from scratch, transfer styles between tracks, and more. To help us design and test them, we've been working closely with incredible musicians, songwriters, and producers. Some of them made even entirely new songs in ways that would have not been possible without these tools. Let's hear from some of the artists we've been working with.

The video began with Wyclef Jean, of Fugees fame, in a recording studio, repeating a four-note riff on his guitar, or a slight variation of it, explaining, "I'm going to put this right back into the music AI tool, this same 'boom, boom, boom, boom, boom.' What happens if Haiti meets Brazil?" As Wyclef posed this question aloud, a jump cut to the interface, presumably "Music AI Sandbox," appeared, and someone typed "Brazil" in a search bar under the heading "Describe Your Track" that has the prompt "Include." Following, a button labeled "Transform" is clicked and Wyclef confirmed, "Dude, I have no clue what's about to be spat out. And this is what excites me." Upon hearing the playback, Wyclef appeared pleased with the

results as he stood up and danced a little before the video cut to him in a different scene, providing some context about the significance of this technology: "As a hip-hop producer, we dug in the crates. We playing these vinyls and the part where there's no vocal, we pull it, we sample it, and we create an entire song around that. So right now, we digging in the infinite crate. It's endless." Later in the video, Wyclef affirmed, "The tools are capable of speeding up the process of what's in my head, getting it out. You're able to move light speed with your creativity" (Google DeepMind 2024e).

The second vignette of the video featured the producer Justin Tranter working alongside the artist Blush. The segment opened with Tranter playing what looked and sounded like a Rhodes organ and calling out some chords: "Let's do Gsus2. D major over G. Back to C major over G." Meanwhile, a brief view of the interface of Music AI Sandbox was shown; while not fully legible, it seemed that in the "Include" text box the words "Strummed Acoustic Guitar, chs: Gsus2 C/G" are typed by someone else, presumably an assistant working in the recording studio with Tranter and Blush. In his endorsement for Music AI Sandbox, Tranter explained, "Someone like me, who is not a new-school producer, to be able to hear the things I hear in my head and actually be able to achieve it because AI tools will let me speak the language that I speak as a songwriter, that is very exciting to me." While it is difficult to discern how the songwriting process has progressed (or regressed) since the time of Tranter suggesting some chords as initial prompts, the next scene depicted Tranter and Blush brainstorming prompts to build upon whatever foundation of a song they have made together to this point.

> Tranter: I'd love to get more atmospheric, weirder prompts, like "Paris Fashion Week 1986."
>
> Blush: Can we?
>
> Tranter: Be great.
>
> Blush: We love that, yeah. I want to hear that, for sure.

Briefly, the video showed someone dragging what appears to be an audio file into an input field on the interface, which produced a new musical result, to which Blush reacted, "I love this," and Tranter concurred, "That sound is awesome." The prompt parade proceeded as Tranter announced, "Let's get creative," and lists off a series of text commands: "Sweaty. Gritty. 3:00 AM … No actually, say Berlin in '79." How these prompts are entered is not shown, but soon after the music differs

by featuring a steady disco-esque beat around 126 BPM and Blush singing along to the track with the lyric, "Only time will tell if you're ready." To conclude the vignette, Blush uttered a single word, "Crazy," as an approval of the technologies she had worked with, and Tranter affirmed with "Love it" (Google DeepMind 2024e).

Returning to the concept of the "conjuration of algorithms" (Nagy and Neff 2024), Google's MusicFX DJ and Music AI Sandbox are textbook examples because much is left to the imagination. Neither the act of prompting—one of the few actions these systems allow and the most important one—nor is how it worked are addressed. Instead, users of these applications are expected to accept that "Google wrote this," to quote Rebillet, and to trust "the colossal power of Google's machine" to "interpret" their prompts. This is sleight of hand. Dahlstedt (2021) helpfully explained that with a typical AI system, decisions by multiple humans are involved in design: "The choice of the training set, parameter settings, feedback from the human evaluation that goes into design choices, data representations, tweaks during implementation, changes, decisions about the workflow of the algorithm, how it will interact with its user, and many more" (884). In short, an understanding of algorithmic systems is needed to understand the pedagogy of the prompt.

## What's in a Prompt? Assessing Algorithms

A decade ago, Thibeault (2014) asserted that "algorithms … will increasingly be recognized as a key force in music education" (20), and outlined the following five implications for the field: "the shift to an algorithmic wisdom of the crowd, the rise of music as content, the opportunity to engage with the governing of algorithms, the need to understand the aesthetics of algorithms, and the need for resistance to algorithms" (23). I consider alternative approaches to defining an algorithm, but as a starting point, I proceed with an explanation from Fiebrink and Caramiaux (2018): "Creating algorithms for music making can be thought of as defining rules that will subsequently drive the behavior of a machine" (183). Using the example of Google's music-making applications, I take up Thibeault's fourth point (the need to understand the aesthetics of algorithms) in an attempt to demystify the magic trick of the prompt. Google's proposed music-making paradigm is predicated on the prompt, which is understandable given that they are best known for their search engine and conflating prompting with "Googling." At the aforementioned I/O event, Alphabet CEO Sundar Pichai (Alphabet is Google's parent company)

proclaimed: "25 years ago we created Search to help people make sense of the waves of information moving online" (Google 2024c), clearly attempting to apply this same strategy to music making. Presuming people have grown accustomed to typing text into a box, transitioning this user experience (UX) to the user interface (UI) of an application such as MusicFX DJ or Music AI Sandbox is seemingly seamless. Andersson and Sundin (2024) remark that search engines such as Google are "increasingly invisible" (613), citing Bowker and Star (1999, 33): "The easier they are to use, the harder they are to see" (614). How might music educators make the act of prompting (Googling) a more legible process for those who are engaging in prompt-based music making with a generative AI application? I suggest that knowing how information retrieval and search engines work can serve as a starting point for understanding how a user's input (prompt) is mediated within an algorithmic-based system to produce an output, such as music; however, it is important to be mindful that technically, as the following section of this article describes, generative AI models are not simply information retrieval engines (although the design of a system may give that impression to a user). Still, I aver this is a helpful primer for understanding algorithmic operations, which are foundational to the AI applications I examine in this article.

For decades, critical technology and media scholars have closely examined the biases built into computer systems and the applications they support (e.g., internet browsers, search engines), and these ideas remain relevant in the current context of assessing generative AI applications such as Google's MusicFX DJ and Music AI Sandbox. Because such applications are integrated into computing systems, it is important to consider how these broader technological infrastructures and the actions associated with using them in turn influence how AI models are used both technologically and conceptually. Friedman and Nissenbuam (1996) forwarded that biases are often obscured in computer systems and that there are three different types, making them "instruments of injustice" (345):

> 1. Preexisting Bias—Preexisting bias has its roots in social institutions, practices, and attitudes. When computer systems embody biases that exist independently, and usually prior to the creation of the system, then the system exemplifies preexisting bias. Preexisting bias can enter a system either through the explicit and conscious efforts of individuals or institutions, or implicitly and unconsciously, even in spite of the best of intentions.
>
> 2. Technical Bias—Technical bias arises from technical constraints or technical considerations.

> 3. Emergent Bias—Emergent bias arises in a context of use with real users. This bias typically emerges some time after a design is completed, as a result of changing societal knowledge, population, or cultural values. User interfaces are likely to be particularly prone to emergent bias because interfaces by design seek to reflect the capacities, character, and habits of prospective users. Thus, a shift in context of use may well create difficulties for a new set of users. (333)

Building on this analysis of bias in computer systems writ large, Introna and Nissenbaum (2000a) explained how bias in search engines were evident based on their observations of how winners and losers, or in their terms, "systematic prominence" and "systematic invisibility" were produced (54). They affirmed, "search engines mediate" and stressed this reality as a political issue (57). More current research reaffirms that information retrieval remains a process that can produce "filter bubbles"—a limited experience of the internet and social media due to algorithmic sorting (Pariser 2011). Noble's (2018) oft-referenced research on how Google reinforces racist and sexist stereotypes by ranking them highly in their results has helped to draw attention to the biases built into such systems and raise questions about how they work. Furthermore, research specifically on search engine results pages of social media platforms has illuminated their "highly biased" design (Poudel and Weninger 2024), and researchers have raised similar concerns about how search engines promote stereotypes and negative opinions of social groups through auto-complete functions (Liu et al. 2024). In their research on children's perceptions of Google, Girouard-Hallam and Danvitch (2024) observed that children may be able to recognize its "capacities" more than its "limitations" and may not think critically about how information is retrieved or where it came from (10–11). These examples are all drawn from research on information retrieval, but the concepts are also applicable to music, and Holzapfel, Sturm, and Coeckllbergh (2018) have drawn attention to how the biases in a music recommendation system can have negative effects on under-represented cultural groups. While compelling, the narrative that search engines like Google are the sole source of the bias problem, Gillespie (2017) accurately depicted internet interactions as dynamic; information is not passively waiting to be selected by a search engine; producers of information compete for attention, and some are better (or better positioned) than others at playing this game (74). Another factor in this high-stakes game is the quasi-referee—human "quality raters" (Meisner, Duffy, and Ziewitz 2024) and "data cleaners" (Seaver 2021), employed by search engine companies, whose jobs

entail ensuring that search results conform to a company's purported standards.

Although generative AI models built into music applications do not operate in the same way as a search engine, algorithmic biases are still a part of their designs. Whatever sound is produced from a prompt is the result of several human- and machine-made decisions. Furthermore, as MusicFX DJ and Music AI Sandbox evidence, generative AI music systems may be designed to appear as search engines and therefore, at least on a conceptual level, influence the user to approach it as such. The idea, inferred by the interface, which relies on the user's previous experiences of Googling, is that the user is "searching" for a sound. From a technological perspective, it is not possible to know all that goes on within the application due to it being the property of a corporation, but there exists a parallel phenomenon of bias apparent in many, if not most, generative AI models. And, at the very least, the search engine mentality a person is expected to adopt is in itself an unwritten user agreement to accept the limits and biases of a system.

Information retrieval systems were "relatively mature" prior to the introduction of generative AI models such as ChatGPT (Hersh 2024, 1), but this latest development of "GenIR" (Generative Information Retrieval) has not eradicated the aforementioned bias problems; instead, because these models are trained on large sets of data that are not curated, biases persist (Li et al. 2025; Liu et al. 2024). Bender et al. (2021) served a scathing critique of using large datasets for training, driving home the point that increasing the size of the dataset does not correlate with increasing the diversity of perspectives represented within it. For example, they made the case that, "the voices of people most likely to hew to a hegemonic viewpoint are also more likely to be retained. In the case of US and UK English, this means that white supremacist and misogynistic, ageist, etc. views are overrepresented in the training data, not only exceeding their prevalence in the general population but also setting up models trained on these datasets to further amplify biases and harms" (613).

The solution, it seems, would be to understand the algorithms that comprise the architecture of the generative AI models in the systems of MusicFX DJ and Music AI Sandbox, but this is easier said than done for everyone involved, Google included. The forecast is foggy on this front. As a move toward combating the power imbalance plaguing search engines that are in service to the marketplace

and skew the perspectives of most, if not all, internet users, Introna and Nissenbaum (2000b) called for "disclosure of the underlying rules (algorithms) governing indexing, searching, and prioritizing" (181). Their rationale for transparency subscribes to the logic that if it is known how a search engine and its proprietary algorithms work, individuals can identify, and presumably mitigate, the sources of problems. But, as critical media theorists have discussed, the reality is considerably more complex because algorithms are complex, as are the technological systems and sociocultural contexts in which they are embedded. As a result, algorithmic systems are *cloudy*—difficult, arguably impossible, to clearly perceive and understand, even to the people who make them (see Seaver 2019; Wilson 2017), which makes auditing them (Ulloa, Makhortykh, and Urman 2024) and researching them comprehensively an arduous, even impossible, pursuit. This complex reality necessitates a more complex definition of algorithms than the previously cited one by Fiebrink and Caramiaux (2018), which is solely technological. Algorithms are more than just code, they are part of the social world, and to understand the depth and scope of their influence on human activity, expansive research is required to examine how they "play out in practice … mesh into organizations, routines, decision-making and so on" (Beer 2017, 10–11). In this vein of conceptualization, Seaver (2019), who researched music recommendation systems, stated: "'Algorithms,' I propose, are social constructions that we, as outsiders and critics, contribute to. The point of declaring something a construction is to argue that it might be constructed differently" (413). This more expansive conceptualization of algorithms points to a need for understanding how generative AI music applications are constructed to assess the priorities programmed into a system, which I interpret as pedagogical because they shape music learning experiences.

## Where the Magic Happens: The Construction of Generative AI Music-Making Systems

I draw primarily on legal scholars Cooper and Grimmelmann (2024) to explain how generative AI models work because I find their explication adheres to the so-called Goldilocks principle of being "just right," with regard to how much technological terminology and detail they provide. Cooper and Grimmelmann (2024)

posited that the seeming rapid emergence of generative AI models (and the associated heightened attention to them) at the present "moment" is a result of advancements in three areas: model architectures, datasets for training, and computing power (9). Interpolating the line "the files are *in the computer,*" from the comedic film *Zoolander* (2001), with "the training data is *in the model*" (4), Cooper and Grimmelmann (2024) focused on the fact that any given generative AI model is trained on a dataset, which makes it possible to produce an output from a prompt. While Cooper and Grimmelmann (2024) acknowledged the various ways in which generative AI models may differ from each other (e.g., modality, technological architecture, use, purpose, availability), they forwarded that generative AI models share a "common shape," which: (1) involve machine learning models that are made by training on data; (2) produce outputs of the same modality as the training data (that's what makes them "generative"); and, (3) produce outputs that reflect patterns in the training data (6). Intended to mimic the human mind, the *deep learning* that takes place in such a model, as Kaliakatsos-Papakosta et al. (2020) explained, can involve unsupervised learning and may produce abstractions from the dataset that a human cannot (or at least cannot make explicit) (238). In sum, generative AI models "acquire an understanding of patterns and structures within their training data" (Sengar et al. 2024, 4) and can output (or "generate") music that reflects the training data by predicting (e.g., the next note in the melody) or classifying (e.g., identifying the chords corresponding to the melody) (Wang et al. 2024, 6388). Delving deeper, Cooper and Grimmelmann (2024) provided a technical explanation of what happens within generative AI models between the points of a user's initial input (prompt) and the resulting output (text, image, video, audio, etc.):

> At a very high (and very oversimplified) level of abstraction, they [generative AI models] generally consist of *neural networks*: interconnected nodes that can perform computations, and which are organized into layers. The strengths of these connections—the influences that nodes have on another—is what is learned during training. These are called the model *parameters* or *weights*, and they are represented as numbers. To run a generative-AI model on an input—*a prompt* — a computer program takes the prompt and transforms it into a format that can be processed in the model. This typically involves taking the prompt and converting it into *tokens* [smallest units of data] … The transformed, tokenized prompt is passed through the layers of the neural network: the computer program copies the input into the nodes at the first layer of the network, then uses the parameters (i.e., connection strengths) leading out from those nodes to compute the input's

> effects on the nodes in the second layer, and so on, until the last layer has been computed. For example, in large language models (LLMs), this process determines how important each token (i.e., word or part-of-word) is in relation to the entire sequence of tokens that make up the text prompt. At this point, once the prompt has been processed through all of the model's layers, the model will produce an output. (9)

Readers desiring a more detailed and nuanced explanation of how different generative AI models work may wish to consult some of the more recently published literature reviews, such as those by Civit et al. (2022), Sengar et al. (2024), and Wang et al. (2024). One challenge to staying current with these technological processes is their relatively rapid rate of development. For example, Civit et al. (2022) conducted a systematic literature review of AI-based music generation that included 1,391 studies and observed that "most" of the studies used a symbolic model (i.e., scores, lead sheets, MIDI, piano roll), with the alternative of using audio being less popular. However, at present, the landscape has changed considerably, as using audio to train models has become increasingly popular with notable contributions from Big Tech such as Google (Agostinelli et al. 2023a) and Meta (Copet et al. 2024). Google's Music AI Sandbox is similar to their product released in 2023, called MusicLM, and the paper they released to describe it (Agostinellii et al. 2023a) is helpful for understanding how they train their generative AI models. The process is considerably more complex than what I detail here, but a key takeaway is that Google created a dataset called MusicCaps that consists of over 5,500 music-text pairs "with rich text descriptions provided by human experts" to train their model (They also provided publicly available examples. See Agostinelli et al. 2023b). The following is an example of a "rich" description from Google's dataset: "A rising synth is playing an arpeggio with a lot of reverb. It is backed by pads, sub bass line and soft drums. This song is full of synth sounds creating a soothing and adventurous atmosphere. It may be playing at a festival during two songs for a buildup."

Human-authored text descriptions were used in the training of Google's model to generate music, and following, human listeners were also involved in the process of rating how well the text descriptions matched the music. Additionally, Agostinelli et al. (2023a) demonstrated that melody could be used in addition to text as a prompt. Meta demonstrated similar capabilities with MusicGen (Copet et al.

2024). And, more recently, Wu et al.'s (2024) MusicControl Net demonstrated using melody, dynamics, and rhythm to control music generation. They claim that future controls could potentially be more abstract, citing "emotion and tension" as examples (2702). In sum, there are multiple ways to prompt a music generating system, and at present, and the current research suggests that there may be more on the horizon.

Knowing the possibilities of prompting as music making, what, then, should one expect of a prompt? From a Big Tech perspective, a paper published by Google DeepMind (Cideron et al. 2024) outlined three properties that music generated from a prompt should possess: "adherence to the input text, high acoustic quality (absence of artifacts), and 'musicality' or general pleasantness" (2). Knowing Google's assessment criteria assists music educators in understanding what their applications, like MusicFX DJ and Music AI Sandbox, generate. Consider the case of Marc Rebillet's "Viola" prompt; out of all the possibilities that could be produced, why and how did he get the result that he did? How did a constellation of musical possibilities get condensed into a sonic sample of a few seconds? Returning to Google's criteria for a moment, the sound Rebillet received in response to his prompt sounded like a viola (first criterion), of high acoustic quality (second criterion), and arguably "musical" or "generally pleasant" (third criterion). But as the prompt becomes more complex, so too does assessing its merit using these criteria. In the example with Justin Tranter and Blush, they began with a more specific prompt ("Strummed Acoustic Guitar, chs: Gsus2 C/G") and developed it further by adding prompts such as "Paris Fashion Week 1986," the adjectives "Sweaty" and "Gritty," a time (3:00 AM) and finally "Berlin in '79." How might the prompt of a place, year, and time of day produce a specific sound and meet the criteria? What might be the referents for qualities such as "Sweaty" and "Gritty," and how are their sounds mediated by the other prompts and their permutations?

Regardless of whether Google's assessment criteria are met, ultimately the user must feel satisfied with the output of the system, otherwise they will stop using it. Perhaps Tranter's prompt of "Paris Fashion Week 1986" is specific enough to narrow the scope of sound, but consider as a hypothetical prompt, "New York 1977," which reasonably could be expected to generate punk, disco, or hip-hop. I suspect that in both the Google examples (Rebillet; Tranter and Blush), the system does not produce precisely what either parties want; rather, it produces an output

that resembles what they want or at the very least they are able to understand why the system produced what it did. In both cases, they accept the output because they are simply satisfied that their text prompt produced any music at all. Put another way, it is the magic that is impressive, not the music. Last but not least, it is curious but not coincidental that Google refrains from demonstrating a prompt-based approach that has been frequently used over the past decade on YouTube: the type beat. In this approach, the prompter searches for a soundalike beat of another artist by using the prompt "[artist name] type beat" (see lesleysteele, Morel Jr. and Leight 2018). For Google to promote this approach to prompting would risk the implication that their music-making applications commit copyright infringement. As a result, Google places itself in the awkward position of having to avoid the most direct way to prompt; instead, they promote a prompting approach that resembles the 1960s game show *Password,* in which contestants cannot tell their partners the "password" and instead have to give them clues to correctly guess it.

Revisiting the case of Wyclef Jean, who used a combination of text-based and melody-based prompts—Googling and noodling—recall that he said, "Haiti meets Brazil." How did the system detect his intended prompt about Haiti? Was geolocation used by the system to detect that he was in Haiti? Could it be that there was something about his playing that was distinctly Haitian? Perhaps Wyclef is demonstrating "fixation" and is stuck on the first output the system produces, therefore accepting it as sounding Haitian (Wadinambiarachchi et al. 2024). Wyclef proclaimed to be "digging in the infinite crate," a reference to the LP era of DJing, which is richly described by Schloss (2014). In his research, Schloss discussed the meticulous curation entailed in trying to find the sample that no one else has, or at least doing something with the same sample that no one else does. But Google is not digging in the crates; instead, in an Amazon-esque way they are delivering the crates. The "infinite crate" of Google consumes and catalogs (Dockray and Parker 2023) by subsisting on "hungry listening" (Robinson 2020) and thirsty learning—the latter term I use to refer to the environmental toll that generative AI takes (Crawford 2021; Hogan 2024; Li et al. 2024). Wyclef also marveled at the speed of Music AI Sandbox, using the term "light speed" and "speeding up the process," manifesting yet another neoliberal value—productivity. For a grand finale of examining prompt-based generative AI music systems, I consider the equivalent to the magician's hat-trick, the end-to-end system.

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

## Magic Words? Prompts in End-to-End Music Generation

If saving or reducing time in the act of music making is the aim and/or advantage of using generative AI, then an end-to-end system is the apex achievement because it is the most efficient, and the path to efficiency is prompting. Like a magician proclaiming, "Presto!" prior to pulling a rabbit from a hat, end-to-end music generation systems promote prompts as akin to magic words that produce a complete piece of music. Similar to Sturm et al. (2024), who examined the end-to-end system Boomy, I consider the case of Udio, except I focus on their "how to" resource. Udio is the product of former Google DeepMind researchers, and whereas Google's MusicFX DJ and Music AI Sandbox are presumed to be intuitive by Google and therefore not in need of any accompanying educational resources, Udio offers step-by-step instructions on how to prompt—a prompt pedagogy.

As a first instruction, Udio's site states that making tracks is easy. It provides an example of a prompt, "a jazz song about New York," and it also notes that if the user is not sure what to type, they can click the dice icon to produce a random prompt. Its next instruction introduces an important concept, adding "tags" to the text-based prompt. In their explanation, Udio's creators use the description "a song about summer rain" as the text-based prompt and the words "jazz," "mellow," and "warm" as tags.

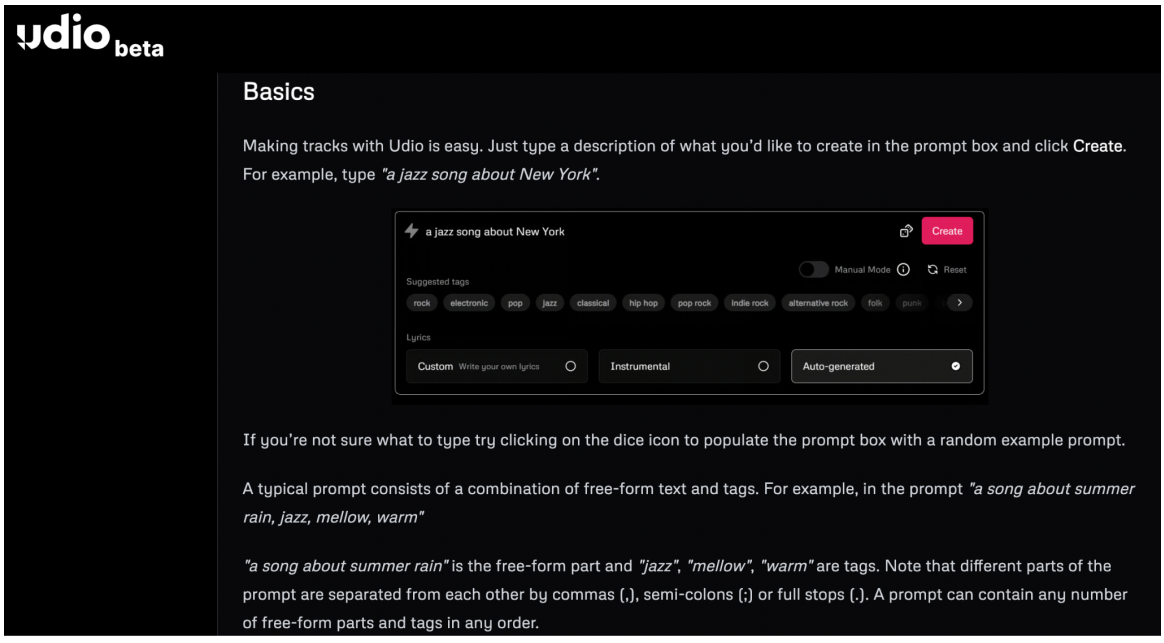


Figure 7: Screenshot of Udio's "Basics" instructions.

In the next instruction, Udio provides prompts for its tags using auto-completions, which is effectively prompting a prompt in progress. The idea is that as a prompt is typed, Udio will suggest or prompt tags to use and the user can choose them or disregard them.

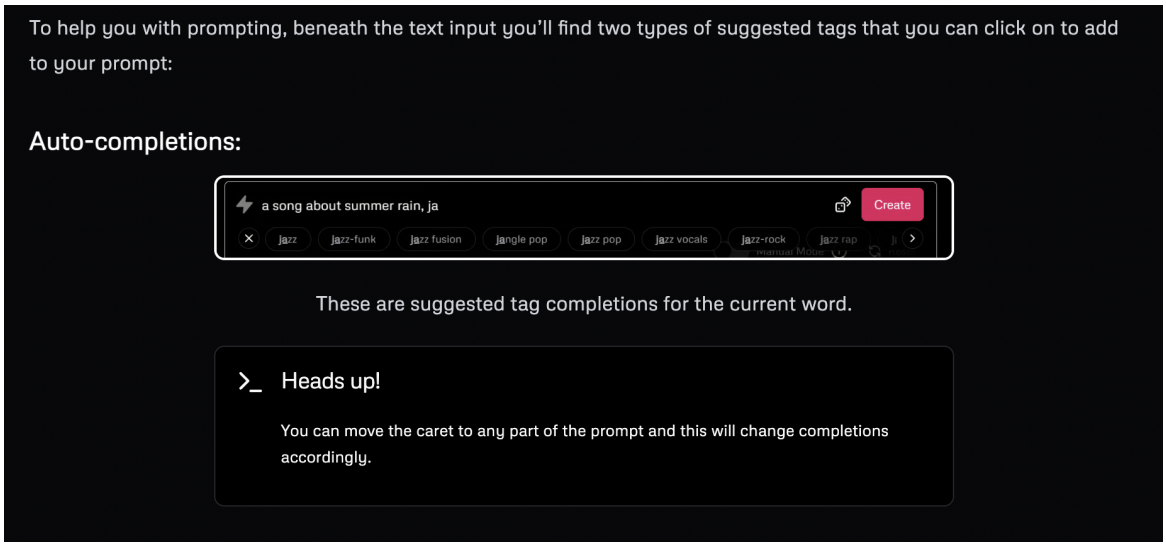


Figure 8: Screenshot of Udio's instructions for using auto-completions

In the following instruction, Udio acknowledges the subjectivity of the prompt, and therefore recommends creating multiple clips with the same prompts. It also states that the prompt box does not reset after "Create" is clicked because Udio's creators recognize that users will likely make more clips with the same or similar prompts to generate the music they want.

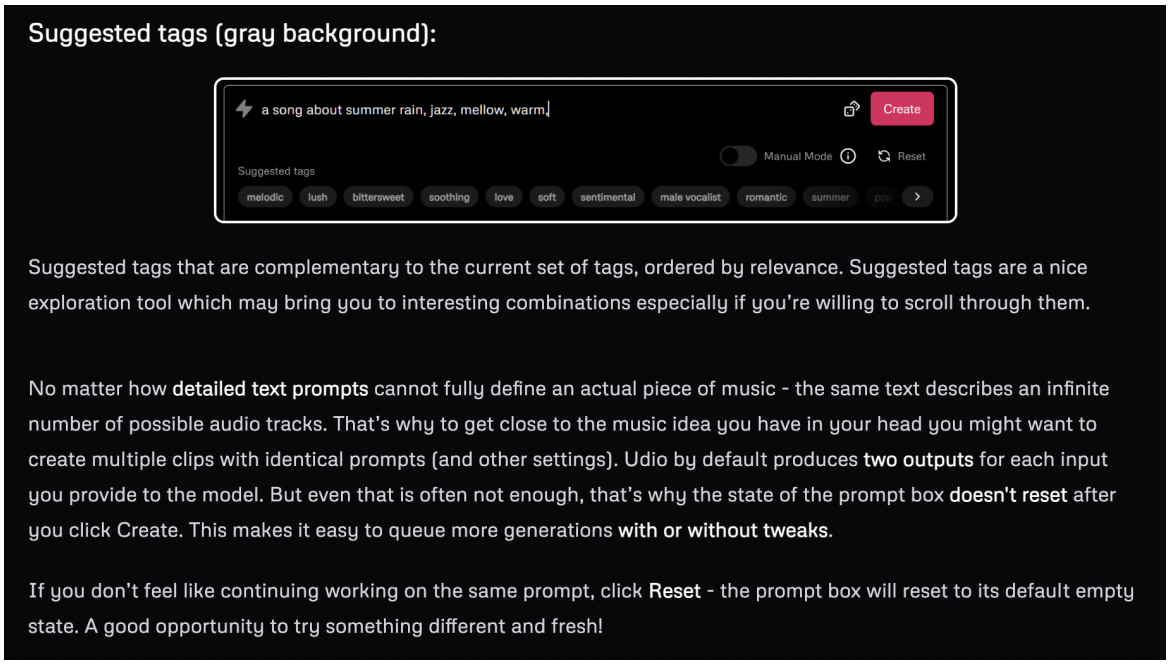


Figure 9: Screenshot of Udio's instructions for using "Tags"

Once "Create" is clicked, Udio produces two 32-second segments, and the user must decide which segment to develop further. This presents a false dilemma, as Udio can produce more options but only provides two at a time. After a user picks an option, they can add sections before and/or after it to create a longer song.

In summary, Udio provides the following formula to make a 90-second song (verbatim):

- Generate the middle (main) part by using regular creation mode. This is the "meat" of your track, the most exciting section.
- Enter Extension mode for the section you just created.
- Select Add Intro option in the Extension Placement area and click Extend. This will create a build up to your main section. The resulting two-section track is now 1 minute long.
- Enter Extension mode for the 1-minute track.
- Select Add Outro and click Extend. That's it - now you've created a piece of music that has a proper beginning and conclusion. And it's all yours!

Beyond creating a song, Udio has a remixing option that functions on a single variance slider between the poles of "similar" and "different." Finally, it is notable that Udio uses "inpainting" as a corrective action to "regenerate" a part of song—

akin to macro editing the song (editing all parts at once) as opposed to micro editing a part of a song (one part at a time).

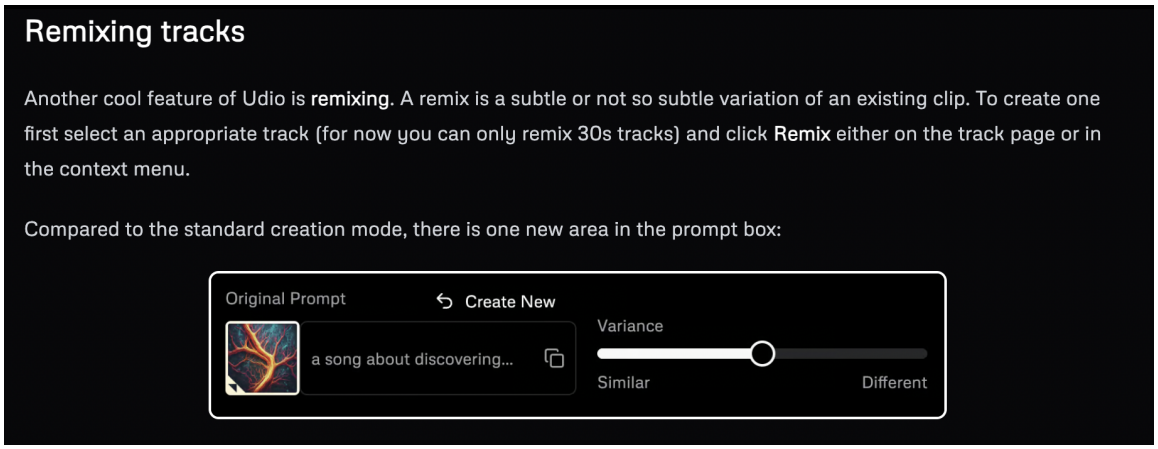


Figure 10: Screenshot of Udio's instructions for remixing tracks

End-to-end generative music systems such as Udio demonstrate how generative AI is a neoliberal genre because they are product-over-process to the point of eliminating as much of the process—the acts of making music—as possible. This is the function of prompting in such systems, and therefore having instructions on how to prompt is a pedagogy that is antithetical to music education. Generative AI music systems are neoliberalism's babies that never sleep, suck up water, and spew out semblances of songs. Compared to Google's MusicFX DJ and Music AI Sandbox, Udio is considerably more overt with guiding its users in how to prompt to achieve optimal outputs. Whereas Google presumes "Googling" will transfer to the music making domain, no explanation needed, Udio opts for a more prescriptive approach. On the surface, Google's and Udio's products are different, but beneath their veneers the same neoliberal ideology influences their designs and the actions of their users. They typify the "conjuration of algorithms" by concealing their designs, creating confusion around what their technologies can (and can't) do, and focusing on the seemingly magical effects of their technologies (Nagy and Neff 2024).

## Sentiments of Sousa

Admittedly, my assessment of generative AI music systems, especially end-to-end systems, bears resemblance to the critiques forwarded by John Philip Sousa in 1906 about new music technologies at that time. Remarking on "The Menace of Mechanical Music," such as player pianos and recorded music, Sousa, the American king of marching band, wrote that they were a substitute for human skill, intelligence, and soul. Player pianos still exist, but perhaps their biggest influence at present is in the piano roll interface of most digital audio workstations. Meanwhile, recorded music continues to be a thriving practice, and as Thibeault (2022) aptly observed, in time, Sousa changed his tune about recording and participated in the practice. I recognize that my discussion of generative AI music systems is limited to prompt-based approaches to create music, and that there are other applications that could potentially be a boon to music making such as musical expression for novices (Louie et al. 2020), musical co-creation (Dahlstedt 2021), and mixing with digital audio workstations (Deruty et al. 2022; Tsiros and Palladini 2020). Furthermore, as someone who frequently attends disability arts events, I have experienced impactful implementations of AI in music making, such as Canadian artist Dyllan Lambert Monroe using an AI voice clone model as part of his wheelchair dance performance to critique ableism in the queer community, and UK-based artist Kris Halpin performing his music with MiMU gloves, which use machine learning.[5] Such experiences give me optimism that perhaps through alternative influences to Big Tech, such as the Distributed AI Research Institute, which is "rooted in the belief that AI is not inevitable, its harms are preventable, and when its production and deployment include diverse perspectives and deliberate processes it can be beneficial" (DAIR n.d.), the profession of music education can find an ethical way forward with AI.

Returning to Big Tech and their generative AI music-making systems, in the present context at least, I am afraid that I am short on optimism because this is not a case of balancing pros and cons—far from it. This is a case in which a neoliberal extractivist system is increasing inequity worldwide while exacerbating our planet's climate crisis (Hogan 2024; Valdivia 2024). Big Tech's generative AI music systems present users with limited control of music making in the form of a prompt, all the while promoting neoliberal values of efficiency and productivity and offering in return products such as unlimited "Gloops" (the term coined by

Marc Rebillet for Google-made loops). Dear music educators, I trust that you will not be duped by these Big Tech neoliberal magic tricks and that you will strive to help those who you teach understand how generative AI music systems work. This will require our profession to continually examine the pedagogy of the prompt and the neoliberal logics and systems that uphold it.

## About the Author

adam patrick bell is Canada Research Chair in Music, Inclusion and Accessibility and associate professor of music education at Western University, Canada. He is the author of *Dawn of the DAW* (Oxford 2018), editor of *The Music Technology Cookbook* (Oxford 2020), and editor of *Canadian Music Educator* and *Journal of Music, Technology & Education*.

## References

Ableton. 2024a. Offers for schools. https://www.ableton.com/en/classroom/offers/

Ableton. 2024b. Become a certified trainer. https://www.ableton.com/en/education/certification-program/

Ableton. 2024c. Loop. https://www.ableton.com/en/blog/categories/loop/

Ableton. 2024d. Learning music. https://learningmusic.ableton.com/

Abramo, Joseph Michael. 2022. *Re*-sistance in the age of cooption by the political right. *Philosophy of Music Education Review* 30 (2): 119–25. https://muse.jhu.edu/article/867640

Adams, Kyle. 2015. What did Danger Mouse do? *The Grey Album* and musical composition in configurable culture. *Music Theory Spectrum* 37 (1): 7–24.

Agostinelli, Andrea, Timo I. Denk, Zalán Borsos, Jesse Engel, Mauro Verzetti, Antoine Caillon, Qingqing Huang, et al. 2023a. MusicLM: Generating music from text. arXiv. http://arxiv.org/abs/2301.11325

Agostinelli, Andrea et al. 2023b. MusicLM: Generating music from text. https://google-research.github.io/seanet/musiclm/examples/

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

AI Cover_JPN. 2023. YOASOBI - アイドル/Oshi no Ko OP (Taylor Swift Ai Cover). https://www.youtube.com/watch?v=cOOo-nWLPKY

Apple. 2024. Your Apple teacher journey. https://education.apple.com/learning-center/T006360A

Apple, Michael W. 2019/1979. *Ideology and curriculum*. 4th ed. Routledge.

Andersson, Cecilia, and Olof Sundin. 2024. The elusive search engine: How search engine use is reflected in survey reports. *Journal of the Association for Information Science and Technology* 75 (5): 613–24. https://doi.org/10.1002/asi.24819

Angel-Alvarado, Rolando, Bayron Gárate-González, and Isabel Quiroga-Fuentes. 2021. Insurrection in Chile: The effects of neoliberalism from a music education perspective. *Action, Criticism, and Theory for Music Education* 20 (3): 108–31. https://doi.org/10.22176/act20.3.108

ArmanTNT. 2023, September. How does "There I Ruined It" do this? *Reddit*. https://www.reddit.com/r/FL_Studio/comments/16zhxd7/how_does_there_i_ruined_it_do_this/

Aróstegui, José Luis. 2020. Implications of neoliberalism and knowledge economy for music education. *Music Education Research* 22 (1): 42–53. https://doi.org/10.1080/14613808.2019.1703923

Avis, Rupert. 2019. Bi-musical curricula and "abyssal thinking": The case of KM Music Conservatory, India. *Action, Criticism, and Theory for Music Education* 18 (3): 26–55. https://doi.org/10.22176/act18.3.26

Baker, Geoffrey. 2014. *El Sistema: Orchestrating Venezuela's youth*. Oxford University Press.

Baris, Antonios. 2024. AI covers: Legal notes on audio mining and voice cloning. *Journal of Intellectual Property Law & Practice* 19 (7): 571–76. https://doi.org/10.1093/jiplp/jpae029

Bates, Vincent. 2021. Music education, neoliberal social reproduction, and play. *Action, Criticism, and Theory for Music Education* 20 (3): 82–107. https://doi.org/10.22176/act20.3.82

Beer, David. 2017. The social power of algorithms. *Information, Communication & Society* 20 (1): 1–13. https://doi.org/10.1080/1369118X.2016.1216147

Bell, Adam P. 2015. Can we afford these affordances? GarageBand and the double-edge sword of the digital audio workstation. *Action, Criticism, and Theory for Music Education* 14 (1): 44–65. https://act.maydaygroup.org/articles/Bell14_1.pdf

Bell, Adam P. 2017. DIY recreational recording as music making. In *The Oxford handbook of music making and leisure*, edited by Roger Mantie and Gareth Dylan Smith, 81–98. Oxford University Press.

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? 🦜. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 610–623. ACM. https://doi.org/10.1145/3442188.3445922

Benedict, Cathy. 2022. Resisting neoliberal subjectivities: Friendship groups in popular music. *Philosophy of Music Education Review* 30 (2): 132–44. https://doi.org/10.2979/philmusieducrevi.30.2.04

Benedict, Cathy, and Jared O'Leary. 2019. Reconceptualizing "music making": Music technology and freedom in the age of neoliberalism. *Action, Criticism, and Theory for Music Education* 18 (1): 26–43. https://doi.org/10.22176/act18.1.26

Boomy. n.d. Boomy. https://boomy.com

Bowker, Georffrey C., and Susan Leigh Star. 1999. *Sorting things out: Classification and its consequences*. The MIT Press.

Bylica, Kelly, and Jonathan Edan Dillon. 2024. Problematizing prescriptive methods in popular music education: Neo-liberalism, care and Little Kids Rock/Music Will. *Journal of Popular Music Education* 8 (1): 9–26. https://doi.org/10.1386/jpme_00121_1

Canham, Nicole. 2023. Living with liminality: Reconceptualising music careers education and research. *Research Studies in Music Education* 45 (1): 3–19. https://doi.org/10.1177/1321103X221144583

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

Cideron, Geoffrey, Sertan Girgin, Mauro Verzetti, Damien Vincent, Matej Kastelic, Zalán Borsos, Brian McWilliams, Victor Ungureanu, Olivier Bachem, Olivier Pietquin, Matthieu Geist, Léonard Hussenot, Neil Zeghidour, and Andrea Agostinelli. 2024. MusicRL: Aligning music generation to human preferences. In *Proceedings of the 41st International Conference on Machine Learning (ICML'24)* 235, Article 358: 8968–8984. https://dl.acm.org/doi/10.5555/3692070.3692428

Civit, Miguel, Javier Civit-Masot, Francisco Cuadrado, and Maria J. Escalona. 2022. A systematic review of artificial intelligence-based music generation: Scope, applications, and future trends. *Expert Systems with Applications* 209: 118190. https://doi.org/10.1016/j.eswa.2022.118190

Collins, Nick, and Mick Grierson. 2024. Avoiding an AI-imposed Taylor's version of all music history. *arXiv*. https://arxiv.org/abs/2402.14589

Cooper, A Feder, and James Grimmelmann. 2024. The files are in the computer: On copyright, memorization, and generative AI." *arXiv*. https://arxiv.org/abs/2404.12590

Copet, Jade, Felix Kreuk, Itai Gat, Tal Remez, David Kant, Gabriel Synnaeve, Yossi Adi, and Alexandre Défossez. 2024. Simple and controllable music generation. *arXiv*. http://arxiv.org/abs/2306.05284

Crawford, Kate. 2021. *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Dahlstedt, Palle. 2021. Musicking with algorithms: Thoughts on artificial intelligence, creativity, and agency. In *Handbook of artificial intelligence for music*, edited by Eduardo Reck Miranda, 873–914. Springer International Publishing. https://doi.org/10.1007/978-3-030-72116-9_31

Deruty, Emmanuel, Maarten Grachten, Stefan Lattner, Javier Nistal, and Cyran Aouameur. 2022. On the development and practice of AI technology for contemporary popular music production. *Transactions of the International Society for Music Information Retrieval* 5 (1): 35. https://doi.org/10.5334/tismir.100

Distributed AI Research Institute. n.d. About us. https://www.dair-institute.org/about/

Fairchild, Charles. 2014. *Danger Mouse's The Grey Album*. Bloomsbury.

Fiebrink, Rebecca, and Baptiste Caramiaux. 2018. The machine learning algorithm as creative musical tool. In *The Oxford Handbook of Algorithmic Music*, edited by Roger T. Dean and Alex McLean, 181–208. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190226992.013.23

Fink, Robert. 2016. Resurrection Symphony: El Sistema as ideology in Venezuela and Los Angeles. *Action, Criticism, and Theory for Music Education* 15 (1): 33–57. act.maydaygroup.org/articles /Fink15_1.pdf

Friedman, Batya, and Helen Nissenbaum. Bias in computer systems. *ACM Transactions of Information Systems* 14 (3): 330–47.

Gillespie, Tarleton. 2017. Algorithmically recognizable: Santorum's Google problem, and Google's Santorum problem. *Information, Communication & Society* 20 (1): 63–80. https://doi.org/10.1080/1369118X.2016.1199721

Girouard-Hallam, Lauren N., and Judith H. Danovitch. 2024. How does Google get its information?: Children's judgements about Google Search. *British Journal of Developmental Psychology* 42 (3): 334–47. https://doi.org/10.1111/bjdp.12487

Giroux, Henry. 2002. Neoliberalism, corporate culture, and the promise of higher education: The university as a democratic public sphere. *Harvard Educational Review* 72 (4): 425–64. https://doi.org/10.17763/haer.72.4.0515nr62324n71p1

Goble, J. Scott. n.d. Neoliberalism and music education: An introduction. *Action, Criticism, and Theory for Music Education* 20 (3): 1–18. https://doi.org/10.22176/act20.3.1

Google. 2024a. Google I/O 2024: Pre-show. https://www.youtube.com/watch?v=wwk1QIDswcQ

Google. 2024b. Distinguish yourself in your classroom and career with certifications from Google for Education. https://edu.google.com/intl/ALL_ca/for-educators/certification-programs/product-expertise/?modal_active=none

Google. 2024c. Google I/O 2024 Keynote: Sundar Pichai opening remarks. https://www.youtube.com/watch?v=uFroTufv6es

Google DeepMind. 2023a. The future of music creation. https://deepmind.google/discover/blog/transforming-the-future-of-music-creation/

Google DeepMind. 2023b. Introducing music AI tools - Transforming beatboxing into a drum loop. https://www.youtube.com/watch?v=30rCswhotL0&

Google DeepMind. 2023c. Introducing music AI tool - Transforming singing into an orchestral score. https://www.youtube.com/watch?v=aC8I2YvL6Uo

Google DeepMind. 2023d. Introducing music AI tools - Transforming chords from a MIDI keyboard into a realistic vocal choir. https://www.youtube.com/watch?v=49T-TYNx_0A

Google DeepMind 2024.The ultimate collaborator in music | Music AI Sandbox. https://www.youtube.com/watch?v=-dPqc7l2zu8

Hall, Claire, Renée Crawford, and Louise Jenkins. 2021. Questioning convergences between neoliberal policies, politics and informal music pedagogy in Australia. In *The Routledge handbook to sociology of music education*, edited by Ruth Wright, 121–35. Routledge.

Hersh, William. 2024. Search still matters: Information retrieval in the era of generative AI. *Journal of the American Medical Informatics Association* 31 (9): 2159–61. https://doi.org/10.1093/jamia/ocae014

Hogan, Mél. 2024. Artificial intelligence is a hot mess. In *Training the Archive*, edited by Inke Arns, Eva Birkenstock, Dominik Bönisch, and Francis Hunger, 33–54. Verlag der Buchhandlung Walther und Franz König. https://trainingthearchive.ludwigforum.de/en/book/

Holzapfel, Andre, Bob L. Sturm, and Mark Coeckelbergh. 2018. Ethical dimensions of music information retrieval technology. *Transactions of the International Society for Music Information Retrieval* 1 (1): 44–55. https://doi.org/10.5334/tismir.13

Horsley, Stephanie, and Paul Woodford. 2016. Facing the music: Pursuing social justice through music education in a neoliberal world. In *The Oxford handbook of social justice in music education*, edited by Cathy Benedict, Patrick Schmidt, Gary Spruce, and Paul Woodford, 62–77. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199356157.013.6

Introna, Lucas D., and Helen Nissenbaum. 2000a. Defining the web: The politics of search engines. *Computer* 33 (1): 54–62. https://doi.org/10.1109/2.816269

Introna, Luca D., and Helen Nissenbaum. 2000b. Shaping the web: Why the politics of search engines matters. *The Information Society* 16 (3): 169–85. https://doi:10.1080/01972240050133634

iZotope, Inc. 2023. How to mix if you're not a mix engineer. https://www.youtube.com/watch?v=Bfh2E6qNF2o

iZotope. 2024. Learn. https://www.izotope.com/en/learn.html/

J Scar. 2009. Danger Mouse - 99 problems. https://www.youtube.com/watch?v=DtfEYuZhDjo

Jackson, Liz, and Rua Williams. 2024, April 4. How disabled people get exploited to build the technology of war. *New Republic*. https://newrepublic.com/article/179391/wheelchair-warfare-pipeline-disability-technology

Kaliakatsos-Papakostas, Maximos, Andreas Floros, and Michael N. Vrahatis. 2020. Artificial intelligence methods for music generation: A review and future perspectives. In *Nature-Inspired Computation and Swarm Intelligence*, edited by Xin-She Yang, 217–45. Elsevier. https://doi.org/10.1016/B978-0-12-819714-1.00024-5

Knibbs, Kate. 2023, November 6. Johnny Cash's Taylor Swift cover predicts the boring future of AI music. *Wired*. https://www.wired.com/story/johnny-cash-artificial-intelligence-taylor-swift-barbie-girl-covers-music-future/

lesleysteele, Jacques Morel Jr., and Elias Leight. 2018, March 9. How "type beats" have changed hip-hop production. *Genius*. https://genius.com/a/how-type-beats-have-changed-hip-hop-production

Li, Pengfei, Jianyi Yang, Adam Wierman, and Shaolei Ren. 2024. Towards environmentally equitable AI via geographical load balancing. In *The 15th ACM International Conference on Future and Sustainable Energy Systems*, 291–307. ACM. https://doi.org/10.1145/3632775.3661938

Li, Xiaoxi, Jiajie Jin, Yujia Zhou, Yuyao Zhang, Peitian Zhang, Yutao Zhu, and Zhicheng Dou. 2025. From matching to generation: A survey on generative information retrieval. *ACM Transaction on Information Systems*. https://doi.org/10.1145/3722552

Liu, Geng, Pietro Pinoli, Stefano Ceri, and Francesco Pierri. 2024. A comparison of online search engine autocompletion in Google and Baidu. *arXiv*. http://arxiv.org/abs/2405.01917

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

Loudly. n.d. Loudly. http://loudly.com

Louie, Ryan, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J. Cai. 2020. Novice-AI music co-creation via AI-steering tools for deep generative models. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. ACM. https://doi.org/10.1145/3313831.3376739

Louth, J. Paul. 2020. Emphasis and suggestion versus musical taxidermy: Neoliberal contradictions, music education, and the knowledge economy. *Philosophy of Music Education Review* 28 (1): 88. https://doi.org/10.2979/philmusieducrevi.28.1.06

McPhail, Graham, and Jeff McNeill. 2021. Music education and the neoliberal turn in Aotearoa New Zealand. *Action, Criticism, and Theory for Music Education* 20 (3): 44–81. https://doi.org/10.22176/act20.3.44

Meisner, Colten, Brooke Erin Duffy, and Malte Ziewitz. 2024. The labor of search engine evaluation: Making algorithms more human or humans more algorithmic? *New Media & Society* 26 (2): 1018–33. https://doi.org/10.1177/14614448211063860

Moore, Andrea. 2016. Neoliberalism and the musical entrepreneur. *Journal of the Society for American Music* 10 (1): 33–53. https://doi.org/10.1017/S175219631500053X

Mullen, Jess. 2019. Music education for some: Music standards at the nexus of neoliberal reforms and neoconservative values. *Action, Criticism, and Theory for Music Education* 18 (1): 44–67. https://doi.org/10.22176/act18.1.44

Nagy, Peter, and Gina Neff. 2024. Conjuring algorithms: Understanding the tech industry as stage magicians. *New Media & Society* 26 (9): 4938–54. https://doi.org/10.1177/14614448241251789

Noble, Safiya Umoja. 2018. *Algorithms of oppression: How search engines reinforce racism*. New York University Press.

Pariser, Eli. 2011. *The filter bubble: What the internet is hiding from you*. Viking.

Parker, James E. K., and Sean Dockray. 2023. "All possible sounds": Speech, music, and the emergence of machine listening. *Sound Studies* 9 (2): 253–81. https://doi.org/10.1080/20551940.2023.2195057

Perrine, William M. 2017. The perils of repressive tolerance in music education curriculum. *Action, Criticism, and Theory for Music Education* 16 (2): 6–38. https://doi.org/10.22176/act16.2.6.

Poudel, Amrit, and Tim Weninger. 2024. Navigating the post-API dilemma. In *Proceedings of the ACM Web Conference 2024 (WWW '24)*: 2476–84. https://doi.org/10.1145/3589334.3645503

Powell, Sean. 2021. Competition, ideology, and the one-dimensional music program. *Action, Criticism, and Theory for Music Education* 20 (3): 19–43. https://doi.org/10.22176/act20.3.19

Rapp, Alison. 2024, August 31. Meet the man ruining your favorite rock songs with mashups. *Loudwire*. https://loudwire.com/interview-there-i-ruined-it-song-mashup/

Richerme, Lauren Kapalka. 2022. The hidden neoliberalism of CASEL's social emotional learning framework: Concerns for equity. *Bulletin of the Council for Research in Music Education* 232: 7–25. https://doi.org/10.5406/21627223.232.01

Robinson, Dylan. 2020. *Hungry listening: Resonant theory of Indigenous sound studies*. University of Minnesota Press.

Rosabal-Coto, Guillermo. 2016. Costa Rica's SINEM: A perspective from post-colonial institutional ethnography. *Action, Criticism, and Theory for Music Education* 15 (1): 154–87. act.maydaygroup.org/articles/Rosabal-Coto15_1.pdf

Sadler, Katherine. 2021. A critical assessment of entrepreneurship in music higher education. *Action, Criticism, and Theory for Music Education* 20 (3): 132–55. https://doi.org/10.22176/act20.3.132

Sánchez-Gatt, Lorenzo. 2023. Divining an Afrofuturist music education. *Action, Criticism, and Theory for Music Education* 22 (4): 131–58. https://doi.org/10.22176/act22.4.131

Schloss, Joseph G. 2014. *Making beats: The art of sample-based hip-hop*. Wesleyan University Press.

Seaver, Nick. 2019. Knowing algorithms. In *digitalSTS: A field guide for science & technology studies*, edited by Janet Vertesi and David Ribes, 412–22. Princeton University Press. https://doi.org/10.1515/9780691190600-028

Seaver, Nick. 2021. Care and scale: Decorrelative ethics in algorithmic recommendation. *Cultural Anthropology* 36 (3): 509–37. https://doi.org/10.14506/ca36.3.11

Seaver, Nick. 2022. *Computing taste: Algorithms and the makers of music recommendation*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226822969.001.0001.

Sengar, Sandeep Singh, Affan Bin Hasan, Sanjay Kumar, and Fiona Carroll. 2024. Generative artificial intelligence: A systematic review and applications. *Multimedia Tools and Applications*. https://doi.org/10.1007/s11042-024-20016-1

Smith, Gareth Dylan. 2015. Neoliberalism and symbolic violence in higher music education. In *Giving voice to democracy in music education*, edited by Lisa C. DeLorenzo, 65–84. Routledge.

Soundtrap. 2024. Soundtrap & YPG to launch revamped educator certification program. https://www.soundtrap.com/content/news/soundtrap-ypg-cert-program

Sousa, John Philip. 1906. The menace of mechanical music. https://ocw.mit.edu/courses/21m-380-music-and-technology-contemporary-history-and-aesthetics-fall-2009/18ab3aba9fe7aa1502a55cd049333659_MIT21M_380F09_read02_sousa.pdf

Sterne, Jonathan, and Elena Razlogova. 2019. Machine learning in context, or learning from LANDR: Artificial intelligence and the platformization of music mastering. *Social Media + Society* 5 (2). https://doi.org/10.1177/2056305119847525

Sterne, Jonathan, and Elena Razlogova. 2021. Tuning sound for infrastructures: Artificial intelligence, automation, and the cultural politics of audio mastering. *Cultural Studies* 35 (4–5): 750–70. https://doi.org/10.1080/09502386.2021.1895247

Sturm, Bob L. T., Ken Déguernel, Rujing Stacy Huang, André Holzapfel, Oliver Bown, Nick Collins, Jonathan Sterne, Lora Cros Vila, Luca Casini, David Alberto Cabrera, Eric A Drott, and Oded Ben-Tal. 2024. MusAIcology: AI music and the need for a new kind of music studies. https://doi.org/10.31235/osf.io/9pz4x

Suno. n.d. About. https://suno.com/about

There I Ruined It. 2024. The Beach Boys sing "99 Problems" by Jay-Z. https://www.youtube.com/watch?v=Z5EGZ145K_w

Thibeault, Matthew D. 2014. Algorithms and the future of music education: A response to Shuler. *Arts Education Policy Review* 115 (1): 19–25. https://doi.org/10.1080/10632913.2014.847355

Thibeault, Matthew D. 2022. John Philip Sousa's historic resistance to technology in music learning. *Journal of Historical Research in Music Education* 44 (1): 5–25. https://doi.org/10.1177/15366006211033966

Thompson, Kali. 2023. Teacher ambassador programs as pyramid schemes: A neoliberal capitalist tale of the exploitation of women educators' labor within educational technology companies. PhD diss., University of Georgia. https://esploro.libs.uga.edu/esploro/outputs/doctoral/TEACHER-AMBASSADOR-PROGRAMS-AS-PYRAMID-SCHEMES/9949558925902959

Tsiros, Augoustinos, and Alessandro Palladini. 2020. Towards a human-centric design framework for AI assisted music production. Proceedings of the International Conference on New Interfaces for Musical Expression. https://doi.org/10.5281/zenodo.4813436

Udio. 2024. Udio. https://udio.com

Ulloa, Roberto, Mykola Makhortykh, and Aleksandra Urman. 2024. Scaling up search engine audits: Practical insights for algorithm auditing. *Journal of Information Science* 50 (2): 404–19. https://doi.org/10.1177/01655515221093029

Valdivia, Ana. 2024. The *supply chain capitalism of AI*: A call to (re)think algorithmic harms and resistance through environmental lens. *Information, Communication & Society*, 1–17. https://doi.org/10.1080/1369118X.2024.2420021

Varkøy, Øivind. 2021. Neoliberalism as political rationality: A call for heretics. In *The Routledge handbook to sociology of music education*, edited by Ruth Wright, 247–57. Routledge.

Wadinambiarachchi, Samangi, Ryan M. Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The effects of generative AI on design fixation and divergent thinking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–18. ACM. https://doi.org/10.1145/3613904.3642919

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

Wang, Lei, Ziyi Zhao, Hanwei Liu, Junwei Pang, Yi Qin, and Qidi Wu. 2024. A review of intelligent music generation systems. *Neural Computing and Applications* 36 (12): 6381–401. https://doi.org/10.1007/s00521-024-09418-2

Willson, Michele. 2017. Algorithms (and the) everyday. *Information, Communication & Society* 20 (1): 137–50. https://doi.org/10.1080/1369118X.2016.1200645

Woodford, Paul G. 2019. *Music education in an age of virtuality and post-truth*. Routledge. https://doi.org/10.4324/9780429451775

Wright, Ruth. 2019. Envisioning real utopias in music education: Prospects, possibilities and impediments. *Music Education Research* 21 (3): 217–27. https://doi.org/10.1080/14613808.2018.1484439

Wu, Shih-Lun, Chris Donahue, Shinji Watanabe, and Nicholas J. Bryan. 2024. Music ControlNet: Multiple time-varying controls for music generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 32: 2692–2703. https://doi.org/10.1109/TASLP.2024.3399026

Yin, Zongyu, Federico Reuben, Susan Stepney, and Tom Collins. 2023. Deep learning's shallow gains: A comparative evaluation of algorithms for automatic music generation. *Machine Learning* 112 (5): 1785–822. https://doi.org/10.1007/s10994-023-06309-w

Young, Susan. 2021. The branded product and the funded project: The impact of economic rationality on the practices and pedagogy of music education in the early childhood sector. *British Journal of Music Education* 38 (2): 107–18. https://doi.org/10.1017/S0265051720000248

YouTube. 2023a. An early look at the possibilities as we experiment with AI and music. https://www.youtube.com/watch?v=rrk1t_h2iSQ

YouTube. 2023b. Introducing music AI tools—Louis Bell, producer/songwriter, builds a track with just a hum. https://www.youtube.com/watch?v=aeYU3_7IKCM

YouTube. 2023c. Introducing Dream Track - an experiment on YouTube Shorts - featuring Charlie Puth. https://www.youtube.com/watch?v=1gjuHUy0IMM

YouTube. 2023d. Introducing Dream Track - an experiment on YouTube Shorts - featuring T-Pain. https://www.youtube.com/watch?v=F4Ngs4mZ5cY

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201

# Appendix – List of Technologies Discussed

| Company/Corporation | Product or Service | Intended Use(s) | Intended User(s) |
|---|---|---|---|
| Ableton | Live (Software) | Software-based music production | Electronic music producers of all levels of experience |
| Ableton | Push (MIDI Controller) | Hardware-based music production (using Ableton Live) | Electronic music producers of all levels of experience |
| Ableton | Learning Music (Website) | "Learn the basics of music making" | Anyone, but especially those with no prior experience |
| Boomy | Boomy (Website) | Text-based prompting to create a song | Anyone, but especially those who are inexperienced in making music and/or want to generate income by distributing their songs through streaming services |
| Google | Music FX and Music FX DJ (Website) | Text-based and music-based prompting to create music excerpts that can be used in songwriting and DJing | Anyone, but especially music producers |

| | | | |
|---|---|---|---|
| iZotope | Learn (Website) | Learn how to use iZotope products and/or music production skills | Music producers of all levels of experience |
| Loudly | Loudly (Website) | Text-based prompting to create a song | Anyone, but especially those who want to generate income by distributing their songs through streaming services and/or create royalty-free music for "digital projects in seconds." |
| Meta | MusicGen (Website and application available through LimeWire) | Text-based and music-based prompting to create music excerpts | Anyone who wants to engage in "music composition." |
| Suno | Suno (Website) | Text-based prompting to create a song | Anyone who wants to create a song |
| Udio | Udio (Website) | Text-based prompting to create a song | Anyone who wants to create a song |

## Notes

[1] https://www.thebookofhov.com/8-hov-did-that

[2] The following are Google's stated principles: 1. Be socially beneficial; 2. Avoid creating or reinforcing unfair bias; 3. Be built and tested for safety; 4. Be accountable to people; 5. Incorporate privacy design principles; 6. Uphold high standards of scientific excellence; 7. Be made available for uses that accord with these principles.

[3] A list of the technologies discussed throughout this article are provided in an Appendix.

[4] An exception may be evident in the writing of Perrine (2017), who does not argue explicitly in favor of neoliberalism but instead critiques those who critique it.

[5] While it is beyond the scope of this article, see Jackson and Williams (2024) regarding their discussion of how disabled people have been exploited in design research for the development of war technologies.

bell, adam patrick. 2025. Pedagogy of the prompt: Music education, artificial intelligence, and big tech magic. *Action, Criticism, and Theory for Music Education* 24 (3): 202–245. https://doi.org/10.22176/act24.3.201