UC Riverside
Conference Papers

Title
AI Music Futures

Permalink
https://escholarship.org/uc/item/7mz4d8t1

Author
Lyon, Eric

Publication Date
2026-06-01

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

AI Music Futures

Eric Lyon
Virginia Tech
ericlyon@vt.edu

Abstract

Artificial intelligence (AI) in music has a long history extending at least to the early experiments of Lejaren Hiller and Leonard Isaacson in the 1950s. In the 2020s the combination of powerful GPUs, large datasets, and pre-trained transformer models has produced a dramatic expansion of AI-based musical practice. This paper explores promising trajectories for AI-based musical co-creativity considering two conflicting realities. First, the ethical and technical challenges of musical AI are not entirely novel. They mirror issues embedded in early digital and pre-computational algorithmic systems. Second, musical AI research and development has undergone a shift in its center of gravity to the private sector as seen in companies like Suno, and Google's Magenta project, due in part to greater corporate computational resources and gigantic datasets that may not always be fully compliant with copyright law. Given the disparity in computational and data resource levels between academia and industry, academic researchers are exploring alternative trajectories, such as establishing ethical guidelines, developing ethically sourced datasets, and creating music from artist-curated datasets. However, a foundational question persists in all these activities: is AI-generated music structurally distinct from music produced without these technologies? This paper nudges that question forward through a combination of historical reflection, musical and technological observations, reports on experiments with AI in music, and speculations about possible futures for AI and music.

Keywords: Artificial intelligence in music; Musical co-creativity; Computer music; Generative AI

1. Introduction

In the 2020s, AI has given computer musicians the gift of destabilization. Since the turn of the 21st century, despite prodigious computer research, innovation, and composition, the basic platforms for computer music creation have remained surprisingly stable. The digital audio workstation (e.g. ProTools), the visual programming language (e.g. Max), the acoustic compiler (e.g. Csound), and the procedural coding platform (e.g. SuperCollider), have all been with us for at least a quarter-century. Underlying all these tools are general purpose programming languages such as C to implement DSP at the sample level. The persistence of platforms such as Csound, Max, Pure Data, and SuperCollider was pointed out early on by Lyon (2002). But in recent years, new forms of AI-based computer music tools have rapidly challenged existing compositional workflows, inviting a radical rethinking of what it means to create computer music.

2. AI Paradigm Onset Speed

One striking aspect of the onset of AI in computer music of the 2020s is how quickly various AI-based tools and procedures moved from fringe research to mainstream adoption, both in the academic computer music world and the commercial music production world. The predecessor research leading to the current explosion of musical usage of AI stretches back many decades.

It has been a respected area of computer music research since the experiments of Hiller and Isaacson in the 1950s (Hiller & Isaacson, 1959). Further attention was given to automated style generation in the work of David Cope in the 1980s (Cope, 1989). Since the early 2010s, machine learning for music (Fiebrink et al., 2009), and transformer-based music AI (Huang, et al., 2018) have been a part of computer music research. But since the early 2020s, an awareness of the potentials of AI in music has increased markedly in both academia and the commercial world. Where the use of AI in music was formerly the domain of specialist researchers, now many non-specialist computer musicians are incorporating AI-based techniques into their work. Even traditional music departments are starting to add courses on music and AI to their curricula. And this surge has occurred in just a few years, creating a palpable sense of progress that computer music has not seen since the 1980s.

3. Party Like It's 1997

William Gibson famously stated that "the future has arrived – it's just not evenly distributed yet" (Rosenberg, 1992). One important future is when AI becomes more proficient than humans at a particular task. It is easy to determine the crossing of this threshold when the task can be objectively evaluated. Such evaluation becomes more challenging when subjectivity is involved, as is the case for work requiring artistic creativity. In either case, there is a threshold between "more proficient than most humans" and "more proficient than all humans." In 1997 the IBM program "Deep Blue" defeated world chess champion Garry Kasparov in a six-game match (Weber, 1997). In that match, Kasparov still won one of the games. The last documented win by a human against a top-performing chess engine occurred in November 2005 when Ruslan Ponomariov defeated Fritz (ChessBase, 2005). Since then, the proficiency threshold for chess has been crossed. However, humans did not stop playing chess. Instead, they integrated chess engines into their training and analysis of the game. The aesthetics of play became important, and distinctions were made between characteristic human play as compared to engine play. Some of the best human players developed bizarre and objectively bad openings as aesthetic gambits, displays of bravado, and to baffle their human opponents (Barden, 2021).

The Suno platform (Suno, n.d.), released in 2023, can compose entire pop songs in just a few seconds. The quality of the output is debatable, but over multiple runs, humans can select their favorite generated songs, augmenting Suno's speed with human taste. Given that Suno was largely trained on whatever music could be scraped from the Internet (Koebler, 2024), it is possible that the system is a closed loop that can only produce variants of the music used for its training. Furthermore, we have not yet heard songs from Suno that approach the transcendent qualities of great commercial songs such as "Ain't No Mountain High Enough" or "Hey Jude." But these comparisons land us firmly in subjectivity where we would struggle to define what makes one song "transcendent" or "great" and another song not. For creative musicians, it is still the year 1996, but 1997 might be coming soon.

One indicator to keep an ear out for is the "jack move," a concept developed by William Gibson that refers to unconventional, lateral, improvisational, and completely unexpected moves. Quoting Gibson (1996), "what it took, sometimes, was just your basic jack move, something nobody, maybe even you, was expecting." An example of a musical jack move takes place in the Choral of J.S. Bach's Cantata BWV 90 (Bach, 1723). In a setting of the text "verleihein selges Stündelein" ("grant a peaceful little hour") Bach introduces a severely chromatic deceptive cadence that implies a resolution to F-major (the relative major of the D-minor key of the chorale) but instead shockingly resolves to D-flat major. This is pure text painting - the

"little hour" is a metaphor for the hour of death. This was also the kind of maneuver that in 1706 won Bach this official reprimand from the Arnstadt Consistory: "we charge him with having hitherto made many strange variations in the chorale and mixed many strange tones into it, so that the congregation has been confounded by it" (David & Mendel, 1998, p. 46). Bach is now celebrated for his musical jack moves that previously attracted censure.

Generative music AI tools can make jack moves too. When working on a piece for E-flat clarinet and piano, I prompted Suno to generate some music for that combination. None of it was useful for my compositional needs, but something strange happened in one of the runs. I used the prompt, "Atonal modernist chamber piece for E-flat clarinet and piano, slow Andante with jagged cross-rhythms, abrupt tempo turns, and shifting meters; the clarinet traces brittle angular cells from set class 0,1,4,5,6 while avoiding note repetition, piano answers with clipped clusters, registral leaps, and asymmetrical attacks. Sparse openings, fractured surges, hard cut transitions, dry close-mic mix with biting clarity., flat, clarinet." Despite the instrumentation clearly being stated as E-flat clarinet and piano, Suno introduces a drum set into the composition at 3'46" and it works beautifully, in a track that Suno entitled "Splintered E-flat Study" (Lyon, 2026). The introduction of the drum set violated the instrumental parameters of my prompt, producing better-than-expected results – an AI-generated musical jack move.

While Suno is mostly known as a song generator, it can produce other kinds of music. To explore the possibilities for generating textures of greater interest to me as an experimental computer musician, I prompted Gemini with the following: "Give me a prompt that will totally confuse Suno." It responded with this: "1930s field recording wax cylinder, extreme data corruption, maximum bitcrush noise, silent gaps, erratic 300bpm microtonal pacing, zero melody, sudden explosive dynamic shifts." The resulting track is called "Wax Cylinder Fracture" (Lyon, 2026), a noisy, non-metric, textural composition that was quite different from what I had grown to expect from Suno.

Both examples suggest that while generative AI music tools might not be able to equal the best in human creativity, they can still surprise us with materials that can then be put to use by humans in a co-creative working relationship with AI.

4. Secret Agents

Multi-agent agentic AI is a system with multiple specialized AIs that operates autonomously to achieve broad goals with minimal human supervision. This is relevant because human music production is highly agentic. Simply composing a new piece of music often involves working on many separate but integrated tasks such as devising formal structures, designing generative systems, generating sonic material, sifting through that material and choosing promising materials, assembling the materials, spatializing them, designing live performance interfaces, and writing piece-specific generative code. Further agentic tasks might involve booking performance dates, finding musical performers, scheduling rehearsals, applying to grants to support the work, booking travel, and so on. Wouldn't it be nice if AI could do all of that for us? But this question raises a typical form of AI anxiety – if AI can do everything that we did in a previously human creative domain, what is left for us to do now?

In a strikingly prescient passage, Karl Marx states,

> In no way does the machine appear as the individual worker's means of labour. Its distinguishing characteristic is not in the least, as with the means of labour, to

> transmit the worker's activity to the object; this activity, rather, is posited in such a way that it merely transmits the machine's work, the machine's action, on to the raw material -- supervises it and guards against interruptions. Not as with the instrument, which the worker animates and makes into his organ with his skill and strength, and whose handling therefore depends on his virtuosity. Rather, it is the machine which possesses skill and strength in place of the worker, is itself the virtuoso, with a soul of its own in the mechanical laws acting through it; and it consumes coal, oil etc. (matières instrumentales), just as the worker consumes food, to keep up its perpetual motion. The worker's activity, reduced to a mere abstraction of activity, is determined and regulated on all sides by the movement of the machinery, and not the opposite. (Marx, 1858/1993, p. 692).

As agentic AI enters musical workflows, computer musicians will be challenged to consider what parts of the creative workflow we want to keep to ourselves, and if we can afford to do so.

5. Speed and Shipwrecks

The adoption of generative music AI into both the academic computer music world and the commercial music world has proceeded with remarkable speed. We may therefore wish to consider the views of Paul Virilio who coined the term Dromology to describe the systematic study of speed. Virilio was greatly interested in the cultural impacts for a society under the influence of speed. He argued that "progress has pushed our hyper-anticipatory and predictive society toward a simple *culture of chance*, a contract on the aleatory" (Virilio, 1979/2009) This culture of chance has long been part of the computer music community that works with algorithmically generated music. The AI and music explosion of the last few years brings this culture of chance fully mainstream. Virilio also suggests that "A technology detached from socio-economic or cultural preconceptions, desires to become the metaphor of the world, while envisioning itself as a revolution of consciousness – finally replacing the pseudo-state of rational wakefulness with an artificial condition of paradoxical wakefulness, while furnishing people with an assistance-becomes-subliminal" (Virilio, 1995). This underscores our AI anxiety that we will not only lose our agency but even lose our awareness that we have lost it. Virilio's anthropomorphizing of technology in his statement becomes doubly ironic.

Virilio also observed that "when you invent the ship you also invent the shipwreck" (Virilio, 1999). There is no shortage of AI shipwrecks in music. AI could flood the market with too much music. The Suno platform alone generates about seven million songs per day (Stassen, 2024). The averaging function of AI could institutionalize conformity. Some AI models have been trained by copyrighted music, with ongoing intellectual property rights lawsuits (*UMG Recordings, Inc. v. Suno, Inc.*, 2024). By becoming pervasive, AI-generated music could elevate artistic mediocrity. AI could reinforce the most statistically popular music at the expense of rarer, stranger work. By treating music as data, AI strips context such that music used as training data becomes no longer of a given time, place, and community. Its aura as described by Walter Benjamin is surgically removed. Perfection can become standardized while losing the imperfections that make music special and give it personality.

To the extent that we computer musicians perceive these shipwrecks as problems, it is our responsibility to counteract them with alternative co-creative models, ethically sourced AI models, and advocacy for alternatives to AI-generated slop. But we must also consider the other side of the shipwreck question. Not all ships get wrecked, and once you invent the ship, you

can travel further and faster than when you're limited to swimming. Can AI-assisted "further and faster" lead us to previously unreachable shores? This tantalizing question is also part of our contemporary AI and music experience.

6. AI Music Tools

Amid the current AI rollout, it's worth considering a few practical questions. Is AI for music extending our capabilities in new ways? If so, how specifically? Are we now capable of creating a new kind of music that would not have been possible prior to the introduction of these tools? Are ontological shifts in computer music creation emerging due to these tools? Will our self-conceptions of ourselves as composers also drift somewhere new? Let's address these questions by considering the AI tools that have come into common use in recent years.

6.1. Stem Separation Based on Deep Learning

This tool is an extension of earlier source separation algorithms. In 2019, Spleeter was introduced, an AI tool based on TensorFlow. In 2021 the more user-friendly tool RipX was introduced by Hit'n'Mix. In 2024, AI-based stem splitting was introduced into Apple's Logic Pro, and in 2025, stem splitting was introduced into Ableton Live. So, it's only in recent years that AI-based stem splitting has been widely accessible. The technique has been widely adopted. It is a useful tool for such applications as spatialization, studying details in a dense electronic mix, or recreating a song in a DAW with a different orchestration. None of these applications is transformational, but AI-based stem splitting is a good example of a new capability provided by AI that has quickly been put to work.

6.2. Voice Cloning

Voice cloning is a special case of speech synthesis, an invention that dates back at least as far as Homer Dudley's vocoder (Dudley et al., 1939). Earlier examples of text-to-speech voice cloning relied on concatenative synthesis. AI-based voice cloning starts with the WaveNet system introduced by Google in 2016 (van den Oord et al., 2016). Proof of concept for voice cloning was provided by Lyrebird AI in 2017 (Lomas 2017) and an accessible platform was introduced by ElevenLabs in 2023. The quality of AI-based speech synthesis and voice cloning is a major improvement on previous methods. That said, it is not a paradigm shift from a user perspective, though the high quality of this engineering feat may encourage more users to adopt voice cloning for creative purposes.

6.3. Timbre Transfer

Pre-AI timbre transfer used techniques such as concatenative synthesis (Schwartz, 2000) and spectral reanimation (Lyon 2003). The Google Magenta project released the Differentiable Digital Signal Processing approach to timbre transfer in 2020. In 2021, IRCAM released the Realtime Audio Variational autoencoder (RAVE). All these tools have proved useful in various musical projects. It's not yet clear if they will lead to fundamentally different sonic or musical results compared with earlier techniques. However, there has been promising work by Holly Herndon and Matt Dryhurst (2019), with the development and musical use of their neural net voice model, Spawn. There has been similarly promising work by Alexander Schubert working

with RAVE, and creating rich, timbrally hybrid compositions based on sound corpuses recorded by himself and his collaborating musicians. It is significant that both Herndon and Schubert base their work on bespoke sound corpuses which then form the basis for individualistic exploration of the corpuses with modern AI techniques. Earlier techniques such as concatenative synthesis and spectral reanimation were also capable of certain kinds of timbre transfer. The difference with the more modern AI-based tools seems to be in the experimental workflow now possible in exploring the learned latent space of a model trained on bespoke sonic data. This aspect of creative computer music is worthy of careful attention in the coming years.

6.4. Audio-to-MIDI

Audio-to-MIDI is the least revolutionary of the AI-based tools discussed here, as there is a long history of the concept. Initial work in the 1970s focused on pitch detection. The 1980s-1990s saw the advent of Pitch-to-MIDI hardware. The 1990s brought waveform-to-MIDI software using conventional DSP techniques such as spectral analysis, onset detection, heuristic note tracking, etc. Music Information Retrieval researchers began relying on machine learning in the 2010s. Google's Onsets and Frames provided significant improvement in automated transcription (Hawthorne et al., 2018). In the 2020s, AI brings audio-to-MIDI to a more general purpose, instrument-agnostic transcription. However, well before those refinements, useful audio-to-MIDI was available as a mainstream DAW feature, such as in Ableton Live by 2013. Notably, Ableton has not yet augmented this capability with AI. The Spotify tool Basic Pitch, introduced in 2022, does use AI and can produce more accurate results, but it's not clear that there is a transformational benefit from this tool. More promisingly, audio-to-MIDI provides fluent bidirectional movement between audio and symbolic representations of a sound, which could be useful for further AI-based multi-perspectival manipulation of sound. As an analogy, consider the benefits of Fourier processing which, by providing fluid switching between spectral and time-domain representations of sound has provided endless useful tools since the 1980s. This is not intended as a technical analogy since, among other discrepancies, the Fourier transform does not lose information when switching between the time-domain and frequency domain; audio-to-MIDI either loses or adds information when switching domains. The point is the rich creative opportunities afforded by representational domain switching for both transformational techniques.

7. Is Any of This Legal?

One side-effect of the speed of implementation of AI in the corporate world, given the need for vast amounts of data to get optimal results, is that the build-out and release of public-facing AI music products has moved faster than the law, with a corporate assertion that training of neural networks on commercial intellectual property (IP) is inherently fair use (OpenAI, 2019). This theory has not yet been adjudicated in the US and remains a vulnerability for some of the AI music companies. Copyright infringement lawsuits against major online AI-based generative music composition sites such as Suno and Udio, include plaintiffs such as Universal Music Group, Sony Music Entertainment, EMI, GEMA, and a class action suit from thousands of independent songwriters. The unresolved critical legal questions around training models on copyrighted information underscore that we are in a transitional period. Definitive legal decisions may create a new normal regarding how AI music models are produced and utilized.

Two earlier legal battles illustrate how new technologies can be tamed and incorporated into existing industry systems. Napster was started in 1999 by Shawn Fanning and Sean Parker to facilitate file sharing MP3s over the Internet. The threat to the music industry was obvious, and the landmark copyright case "A&M Records, Inc. v. Napster, Inc." put a permanent end to Napster, opening the way for legal digital marketplaces such as Apple's iTunes (*A&M Records, Inc. v. Napster, Inc.*, 2001). Another case illustrates a more culturally complex situation centered around sampling and hip hop. Digital sampling quickly became prevalent in the hip hop community in the late 1980s relying on recent technologies like the E-mu SP-1200 and Akai MPC. The landmark lawsuit in *Grand Upright Music, Ltd. v. Warner Bros. Records, Inc.* (1991) brought in the wake of Biz Markie's sampling of Gilbert O'Sullivan coincided with a decline in the average number of samples per hip hop song as shown in (Sewell 2014) and further documented in a larger corpus analysis by McLeish (2023). While this correlation does not prove causation, McLeish also discussed a statement from Chuck D of Public Enemy that the sampling lawsuits of the time forced changes to Public Enemy's use of samples by 1991.

While there are multiple perspectives on whether "Grand Upright Music" more colloquially titled by the judge, "Thou Shalt Not Steal" was a good decision, it's worth entertaining the thought experiment of what effects a similar decision on musical intellectual property might have had in 1491. Given the intense intertextuality of much European classical music, pressure would have been put on such practices as parody mass, shared chord progressions, freely used hymn tunes, variations on themes by other composers, rearrangement of existing compositions by other composers, or simply the added pressure of composers suing each other because say, a four-bar passage in a Handel fugue sounded too similar to a passage in a fugue by J.S. Bach. It seems quite possible that European classical music might have evolved quite differently under those constraints and offers the tantalizing question of how hip hop production might have continued to evolve if not shackled to having to clear every sample in a track.

8. Beyond Copyright

There are aspects of Suno's rollout that might have implications for the future of music as a marketable object. For their basic plan, Suno retains ownership over all songs created by users of the site. This is an interesting departure since traditionally, the manufacturers of creativity tools do not claim ownership of work created with those tools. Manufacturers of music paper make no IP claim to music written on their paper, nor do DAW vendors make IP claims on music created using their systems. The US Copyright office has determined that works generated entirely by AI, without sufficient human authorship, cannot be copyrighted. However, Suno asserts its ownership based on contract law. By using their service, you agree to the terms of their contract. Suno also explicitly states that they do not own a copyright to their output, and importantly, make no warranty that any copyright will vest in their output. This arrangement might expose Suno users to future copyright litigation, but it might also open the door to a future in which copyright becomes irrelevant.

Suno's contracts have not yet been tested in a US court. There is potentially tension between the US Copyright Office assertion that AI-generated works are in the public domain, and contracts that restrict access to so-defined public domain works. But, again, this tension has not yet been tested in an actual US court case. A Munich court has ruled against Suno in a copyright case brought by GEMA, with Suno currently appealing the ruling (Reuters, 2026). Meanwhile,

Suno continues to move forward. As of 2025, Suno has a deal with Warner Music Group to legally use their catalog to train Suno's models, and a similar deal with BMG as of 2026.

Here, we speculate that Suno might be expanding from a business-to-consumer model into a business-to-business model as well. If so, a key strategy will be to build a walled garden within which users interact with Suno-generated content and then lease that service to other platforms. This could redefine Suno as a gamified, interactive entertainment center. It would be the YouTube-ification of AI-based music creation. The successful business model would be software as a service, based on Suno's music infrastructure. If this speculation is correct, the new model will not need copyrights to protect the music generated by the system. According to Suno, users generate seven million songs every day. With this much volume, the value of each individual song as a product effectively drops to zero. The value of the service is rather in flow and engagement, monetized through subscriptions, services, copyright-free music, and (eventually) advertisements. While important aspects of Suno remain untested – the legality of its contracts, the legality of its music models, and the ability to generate a profit – the company has demonstrated a significant engagement level with two million paid subscribers (Millman, 2026). Therefore, even in its inchoate state, Suno should be considered as a possible future for music and AI, against which we can compare other futures that we may wish to invite into existence.

9. Acknowledgements

This paper explores some ideas from my keynote lecture "How to Compose AI Music That Isn't Mid" presented at the Co-Creativity in Music, Sound, and AI conference. I would like to thank the Organizing Committee for the opportunity to share my current thoughts on AI in music at the conference, and Conference Chair Paulo C. Chagas for the invitation to dig deeper into the implications of some of the ideas I presented at the conference, in this paper. My notes for the talk are online (Lyon, 2026) and cover more ground than could be addressed in an academic paper. My current views lie somewhere between the talk and the paper.

10. References

A&M Records, Inc. v. Napster, Inc., 239 F.3d 1004 (9th Cir. 2001).

Bach, J. S. (1723). *Es reißet euch ein schrecklich Ende, BWV 90*

Barden, L. (2021, March 18). Double bongcloud: Why grandmasters are playing the worst move in chess. The Guardian https://www.theguardian.com/sport/2021/mar/18/bongcloud-meme-opening-carlsen-nakamura

ChessBase. (2005, November 22). Bilbao – the humans strike back. https://en.chessbase.com/post/bilbao-the-humans-strike-back

Cope, D. (1989). Experiments in musical intelligence (EMI): Non-linear linguistic-based composition. *Journal of New Music Research*, *18*(1-2), 117–139.

David, H. T., & Mendel, A. (Eds.). (1998). *The new Bach reader: A life of Johann Sebastian Bach in letters and documents* (C. Wolff, Rev.). W. W. Norton & Company.

Dudley, H., Riesz, R. R., & Watkins, S. S. A. (1939). A synthetic speaker. *Journal of the Franklin Institute*, *227*(6), 739–764.

Fiebrink, R., Trueman, D., & Cook, P. R. (2009). A meta-instrument for interactive, on-the-fly machine learning. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME'09)* (pp. 280–285).

Gibson, W. (1996). *Idoru*. G.P. Putnam's Sons.

Grand Upright Music, Ltd. v. Warner Bros. Records, Inc., 780 F. Supp. 182 (S.D.N.Y. 1991).

Hawthorne, C., Elsen, E., Song, J., Roberts, A., Simon, I., Raffel, C., Engel, J., Oore, S., & Eck, D. (2018). Onsets and frames: Dual-objective piano transcription. *arXiv*. arxiv.org

Herndon, H., & Dryhurst, M. (2019). *Proto*. 4AD.

Hiller, L. A., Jr., & Isaacson, L. M. (1959). Illiac suite for string quartet. In *Experimental music: Composition with an electronic computer* (pp. 182–197). McGraw-Hill.

Huang, C.-Z. A., Vaswani, A., Uszkoreit, J., Shazeer, N., Simon, I., Hawthorne, C., Dai, A. M., Hoffman, M. D., Dinculescu, M., & Eck, D. (2018). *Music Transformer: Generating music with long-term structure*. arXiv.

Koebler, J. (2024, August 1). *AI music generator Suno admits it was trained on "essentially all music files on the internet"*. 404 Media. https://www.404media.co/ai-music-generator-suno-admits-it-was-trained-on-essentially-all-music-files-on-the-internet/

Lomas, N. (2017, April 25). Lyrebird is a voice mimic for the fake news era. TechCrunch. https://techcrunch.com/2017/04/25/lyrebird-is-a-voice-mimic-for-the-fake-news-era/

Lyon, E. (2002). Dartmouth Symposium on the Future of Computer Music Software: A Panel Discussion. *Computer Music Journal*, *26*(4), 13–30.

Lyon, E. (2003). Spectral reanimation. *Proceedings of the Ninth Biennial Symposium on Arts and Technology* (pp. 103–105). Connecticut College.

Lyon, E. (2026). *Splintered E-flat Study* [Song generated by Suno AI]. Suno. https://suno.com/s/i7Rg0IOLeH2iTcaK

Lyon, E. (2026). *Wax Cylinder Fracture*. Suno. https://suno.com/song/02448cb6-5714-4409-bfa0-e7699dc5cc3c

Lyon, E. (2026). *How to compose AI music that isn't mid*. Google Slides. https://docs.google.com/presentation/d/1u5zVO_9zjpt2gW1CAa71x8KaKjLGy-QX/present

Marx, K. (1993). *Grundrisse: Foundations of the critique of political economy* (M. Nicolaus, Trans.). Penguin Books.

McLeish, C. E. A. (2023). Hip-hop sampling aesthetics and the legacy of Grand Upright v. Warner. Popular Music, 42(1), 79–103.

Millman, E. (2026, February 26). *AI music platform Suno says it has surpassed 2 million paid subscribers*. The Hollywood Reporter. https://www.hollywoodreporter.com/music/music-industry-news/suno-ai-music-reaches-two-million-subscribers-1236516241/

OpenAI. (2019, October 30). *Comment on patent, trademark, or copyright issues involved in artificial intelligence*. United States Patent and Trademark Office. https://cdn.openai.com/policy-

submissions/OpenAI%20Comments%20on%20Intellectual%20Property%20Protection%20for%20Artificial%20Intelligence%20Innovation.pdf

Reuters. (2026, July 31). German court rules AI music firm Suno broke copyright rules. https://www.reuters.com/world/german-court-rules-ai-music-firm-suno-broke-copyright-rules-2026-07-31/

Rosenberg, S. (1992, April 19). Cyberpunk: The primary subculture of the late 20th century. *San Francisco Examiner*, p. E1.

Schwarz, D. (2000, December 7–9). *A system for data-driven concatenative sound synthesis*. International Conference on Digital Audio Effects (DAFx-00), Verona, Italy.

Sewell, A. (2014). How copyright affected the musical style and critical reception of sample-based hip-hop. *Journal of Popular Music Studies*, 26(2–3), 295–320.

Stassen, M. (2024, April 17). *AI music app Suno claims its users are generating 7m tracks per day – enough to replicate Spotify's entire catalog every 2 weeks*. Music Business Worldwide.

Suno. (n.d.). *Suno: AI music generator*. https://suno.com/

UMG Recordings, Inc. v. Suno, Inc., No. 1:24-cv-11611 (D. Mass. filed June 24, 2024). https://www.riaa.com/wp-content/uploads/2024/06/Suno-complaint-file-stamped20.pdf

van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., & Kavukcuoglu, K. (2016). *WaveNet: A generative model for raw audio*. arXiv. https://arxiv.org/abs/1609.03499

Virilio, P. (1995). *The art of the motor* (J. Rose, Trans.). University of Minnesota Press.

Virilio, P. (1999). *Politics of the very worst* (P. Petit, Interviewer; S. Lotringer, Ed.; M. Cavitch, Trans.). Semiotext(e).

Virilio, P. (2009). *The aesthetics of disappearance* (P. Camiller, Trans.). Semiotext(e).

Weber, B. (1997, May 12). Swift end to a long chess match leaves Kasparov looking blue. *The New York Times*, p. A1.