ISSN: 2998-7431

MUSICA IQ

ISSUE #4
JUNE
2025

SUMMER
ISSUE

INTELLIGENCE REIMAGINED – AI, MUSIC, AND THE FUTURE OF CREATIVE PRACTICE

Muzie.Live

SCAN ME

Everything You Need For Music Teaching!

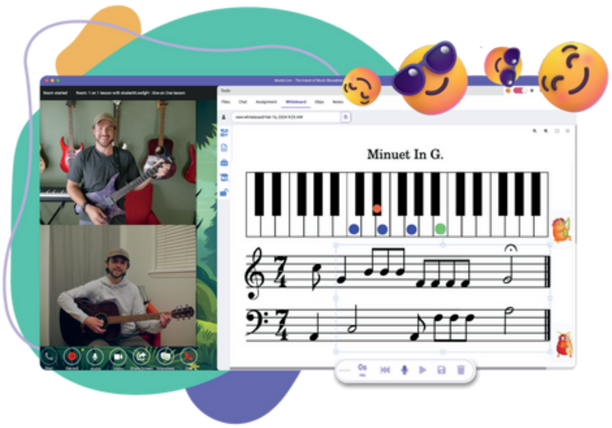

PRIVATE TEACHERS

At Muzie, we understand the unique challenges and aspirations that music teachers face in today's fast-evolving educational landscape. Our platform is designed to empower educators.

STUDENTS

We're dedicated to enhancing the learning experience for students by providing a platform where they can connect with their teachers in a dynamic and interactive environment.

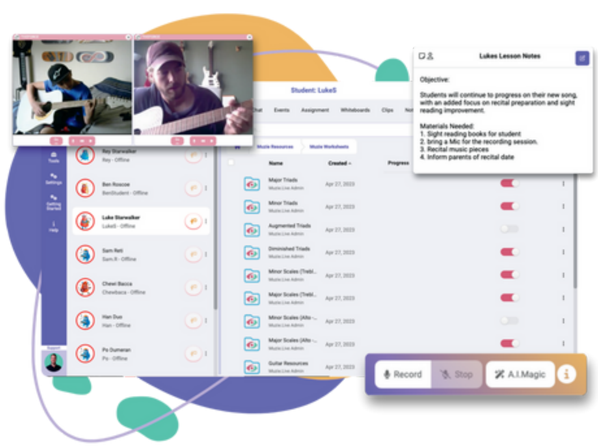

MULTI-TEACHER STUDIOS

Muzie streamlines administrative teaks and enhances organization, enabling schools to focus on student learning in a connected & engaging community.

SIGN UP FOR FREE AT WWW.MUZIE.LIVE

FROM THE EDITOR

Intelligence Reimagined

"The most profound technologies are those that disappear... They weave themselves into the fabric of everyday life until they are indistinguishable from it."
– Mark Weiser

Artificial Intelligence is no longer something we experience only in science fiction films — it is a reality evolving more rapidly than we could have predicted. It is now deeply embedded in our tools, our creative processes, and increasingly, in our very identities as musicians, educators, and cultural contributors. From AI-generated compositions to intelligent tutoring systems and personalized practice platforms, we are witnessing a dramatic transformation in how music is conceived, taught, and experienced.

Many perceive AI technologies as a threat — something to fear or resist. I see them differently. I see AI as a collaborator, an ally. I use it in the same spirit I use my microwave, my dishwasher, or my car. These tools extend my capacities and simplify complex tasks. Likewise, I use AI as an extension of my brain — to streamline processes, accelerate creative work, and unlock time and energy for deeper human engagement. Crucially, I remain in the driver’s seat.

Often, I compare the adoption of AI to choosing how we travel. One might walk across a continent — a rich, slow journey. But one could also board a plane and reach the destination swiftly. The plane doesn't eliminate the human journey; it enables a different kind. We still choose the destination. We still need a pilot. But this is not merely a technological evolution — it is also a cognitive and cultural transformation. What does it mean when machines simulate creativity? When improvisation is not discovered, but predicted? Can a neural network feel the expressive weight of a silence, or does it only learn to mimic it? These questions challenge our assumptions — not just about machines, but about ourselves.

In MusicalQ Issue #4, we explore these tensions and possibilities under the theme Intelligence Reimagined – AI, Music, and the Future of Creative Practice. This issue brings together a chorus of voices who see artificial intelligence not just as a tool, but as a mirror — reflecting back our own values, assumptions, and aspirations for musical intelligence.

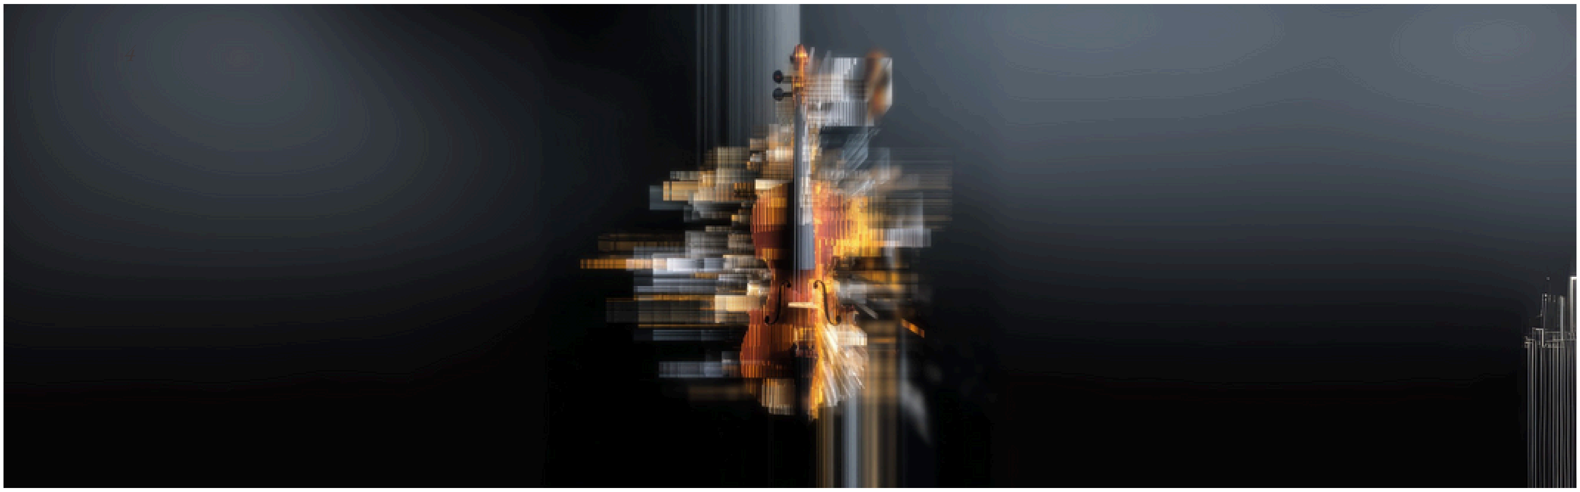

Inside these pages, you will find essays that blur the binary of human versus machine creativity, interviews with artists weaving AI into their compositional language, and reflections from educators using these technologies in real classrooms. You will also encounter sharp critiques — ethical, philosophical, pedagogical — from those insisting we tread with care.

At MusicalQ, we are not here to celebrate technology uncritically — nor to condemn it without curiosity. We are here for the layered questions that do not fit into clickbait headlines or marketing decks. This magazine exists to foster interdisciplinary dialogue, to hold complexity, and to make space for the generative tensions of this moment.
This issue invites you to ask: What kind of intelligence do we want to cultivate in our students? How do we protect the soul of our traditions while embracing the tools of the future? What must evolve — and what must remain sacred?

We hope these contributions challenge you, awaken new ideas, and invite you to shape this unfolding conversation. The future of music is not simply faster or louder — it is more intentional, more layered, and more human than ever.

Thank you for joining us in imagining what comes next.
.

Antonella Di Giulio

EDITOR-IN-CHIEF AND FOUNDER

TABLE OF CONTENTS

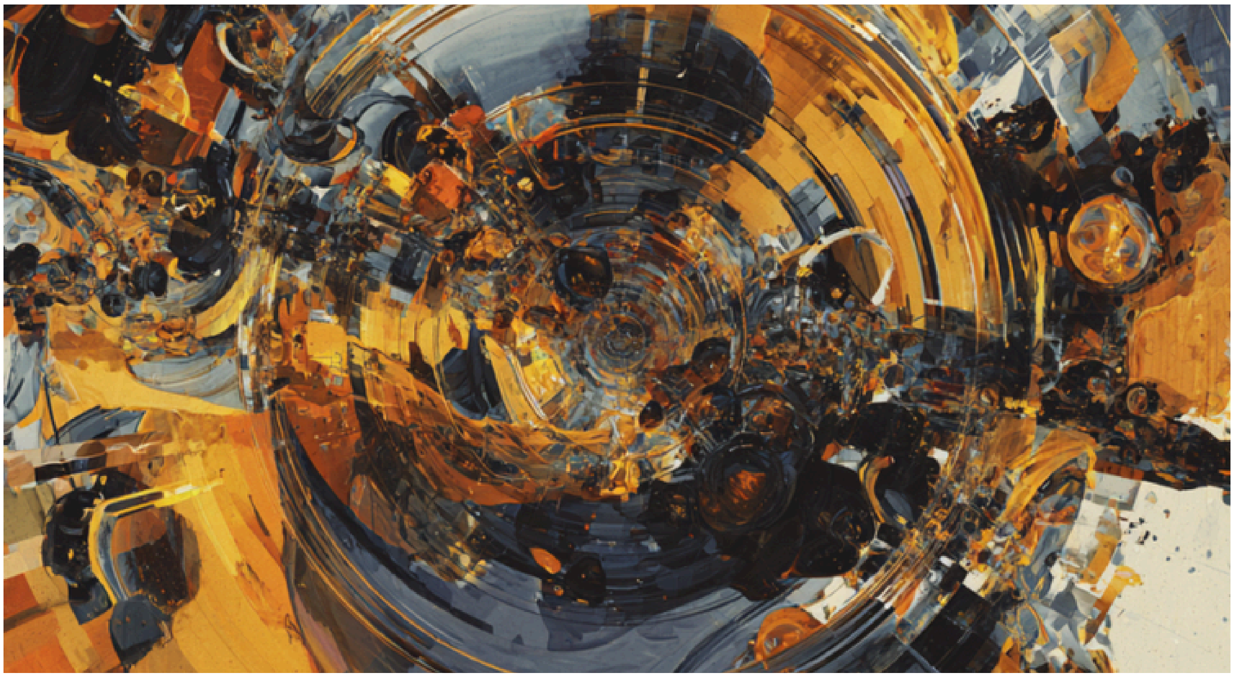

MusicalQ Magazine (ISSN: 2998-7431) is a periodical published by Woom Management Inc four times per year: Quarterly, in September, December, March, and June. Periodical postage paid in Buffalo, NY. Subscriptions are included with the Premium Membership to Musica IQ, but also sold separately. Non members/ library subscription: $60/ year, international $78, single issue $16. Dedicated to fostering insightful discussions on music education, performance, and research, and advocating for musicians and music education worldwide. Submissions are carefully reviewed by our editorial team to ensure content aligns with our mission of promoting excellence and innovation in the musical arts.
For inquiries, submissions, or further information, please visit our website at www.musicaiq.com

MUSIC ANALYSIS:

GENERATING CONTEXTUAL MUSICAL SPACES WITH MACHINE LEARNING

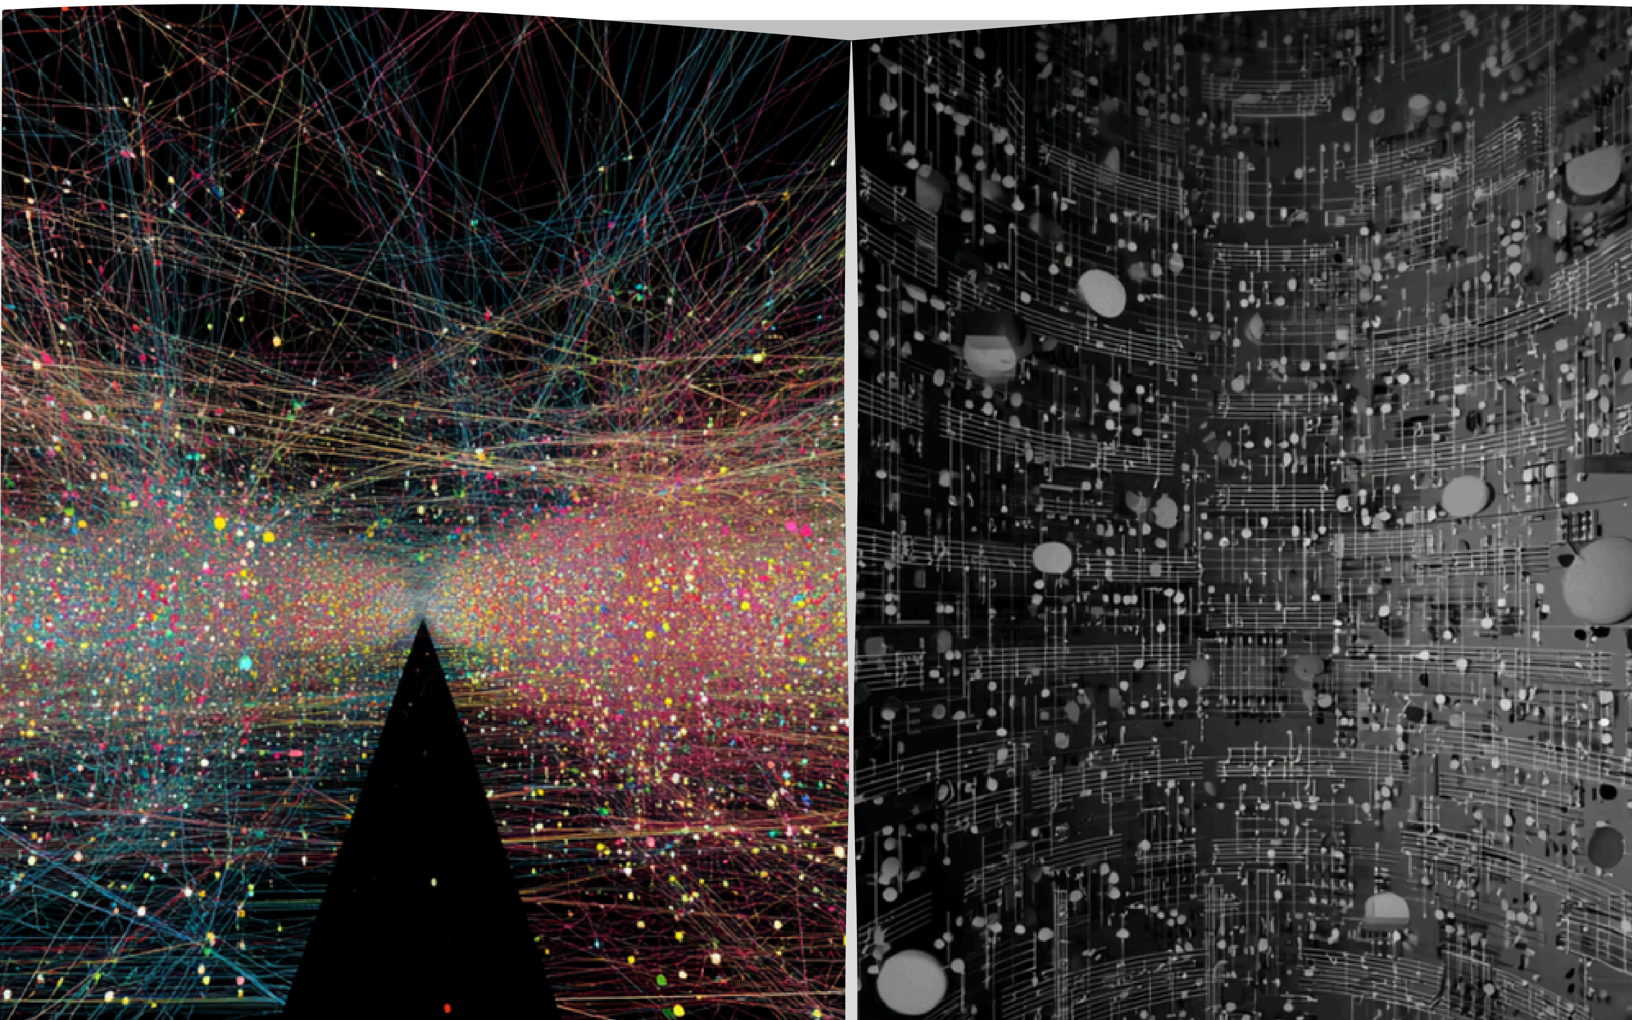

BY DR. MATT CHIU

Fixed vs. Contextual Spaces

In music we frequently used fixed systems. For example, staff notation is a static two-dimensional plane, where the x-axis is time and the y-axis is pitch height. The different y-coordinates corresponds to a different named pitch class; i.e. C, D, E, etc. Figure 1 visualizes the staff in a coordinate system. If we let middle C be the center of the y-axis and calculate units on the x axis in quarter notes, the middle C starting the example in Figure 1 is encoded as (0,0), the D is (1,1), the Eb is (1.5, 2), and so on. Under this system, intervals (that is, “music intervals”) can be defined as the distance from the second position of one note to the second position of another, plus 1. E.g., CàF == |(0−3)| +1 == 4(th).

Distances in this space are fixed: C to D is *always* a second no matter how many intervals it is, or no matter how many alterations you put on either note (CàDx is still a second, just an augmented second). In fact, there has been much ink in music theory spilled over different forms of musical spaces with fixed distances: distance around the circle of fifths, intervals, measurements in semitones, etc.

.

.

But musical distances are not static: the same interval from C to Ab in C major sounds different than it does in the key of B major; a chord progression can sound normal in one style and yet jarring in a particular style. A system designed for a particular context or style is, in contrast, a contextual space. These spaces suggest that musical distances are based on context—whether that is within a single piece, or within a broader style. Instead of using stylistically agnostic spaces for measuring distances, we can use methods from artificial intelligence (AI) and machine learning (ML) to create such spaces. We explore one option here, namely creating an embedding space.

Creating An Embedding Space

An essential aim of natural language processing (NLP)—a branch of computer science and AI—is to encode words such that a machine might understand it. The ideal encoding is efficient for a computer to process and captures aspects of a word’s meaning: the encoding needs to, for example, somehow differentiate “bow” and “cow,” while showing “cow” and “cows” are similar. One type of encoding which is sensitive to a word’s contextual placement is known as embeddings: we can generate word embeddings with a technique known as word2vec.

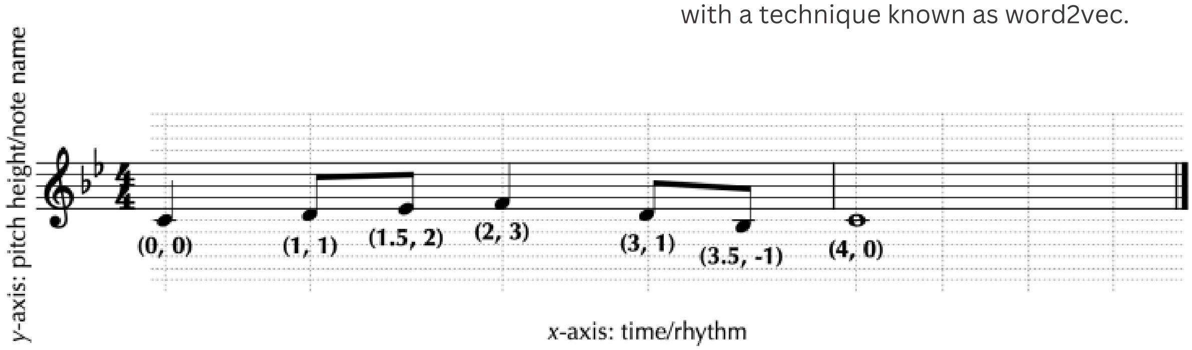

FIGURE 1. VISUALIZING STAFF NOTATION ON A CARTESIAN PLANE

Given a dataset, word2vec uses a specific, shallow neural network to derive vector representations for words based on their contextual placement in said dataset—a vector is a list of numbers where each value captures information about different dimensions; e.g. [0,4,6,8] and [0, 4, 20, 8] are both four dimensional vectors which share values for three dimensions, but contains distant values for the 3rd position (6 vs. 20). The staff notation introduced earlier could be thought of as a two-dimensional vector.

Words that occur near one another or have similar placement in the dataset will have similar vectors (and we call those vectors “embeddings”). For the sake of brevity and accessibility, I cannot explain all the technical aspects of word2vec, but we unfortunately need to wade through a few technical details (such as encoding) to understand the overall idea to grasp the underlying framework. I will give a short, surface-level primer below: for more information on word2vec see Jurafsky and Martin (2025), for more information on this paper’s specific methodology contact the author. If you have no interest in the underlying machine learning models, you can skip ahead to the “Studying Embeddings” section.

Given a dataset, word2vec generates random n-dimensional vector representations for each word. These are random at first, but they will be updated based on a word’s probabilistic placement in the dataset. After that, word2vec scans over each word. Given a target word t and a window of k words, the model returns the probability that a particular context word c will be within k range of t: P(+|t, c). Word2vec then uses a neural network and backpropagation to adjust the embeddings, optimizing embeddings such that word’s with high a probability of occurrence within the window, have similar vectors. Resulting word vectors are close if they tend to occur near one another in the training data or have similar syntactic placement.

By extension, instead of using a dataset of words, we can use a dataset of music and derive musical embeddings. The embeddings occupy a new space sensitive to the context—that is, the musical style in which the music occurs. Figure 2 visualizes how word2vec is used in this paper. As shown in Figure 2, the music in the dataset encodes musical words by slicing the music into verticalities; this method is sometimes known as “salami slicing” (Cuthbert and Ariza, 2010; White and Quinn, 2016). Each slice approximates a “noisy” chord, incorporating all notes sounding notes even if they are not considered chord tones. For the purposes of this example, we take all notes and encode them from 0–11, where 0 is C, 1 is C#, and so on; and we also add a key label from 0–24, where 0 is C major, 1 is C minor, 2 is C# major, 3 is C# minor, and so on. So the slice highlighted in Figure 2 has the notes in D and F# and is in the key of G major: ([2,6,], 19).

To generate the embeddings for these slices we generate a random vector of length 40 for each chord/musical slice. Then word2vec iterates through the dataset, slice by slice, and temporarily assign each verticality as the “target chord.” We isolate the four chords to the left and to the right of the target chord and make a prediction about what they are—we call these the context chords. Then the model updates values such that the embeddings for the context chords are slightly closer to the embeddings for the target chord. After looping through the dataset multiple times (in this case 10 times), we are left with vector representations for chord slices which capture stylistic nuance and contextual placement, achieving our goal of creating a contextual space.

Though we use this method on slices of music (standing as a proxy for chords), this method can be used for any musical representation: melodies, melodic fragments, collections of pitches, scales, etc. To show the applicability of embeddings, I create two different sets of chord embeddings—one representational of a Classical style and another based on rock.

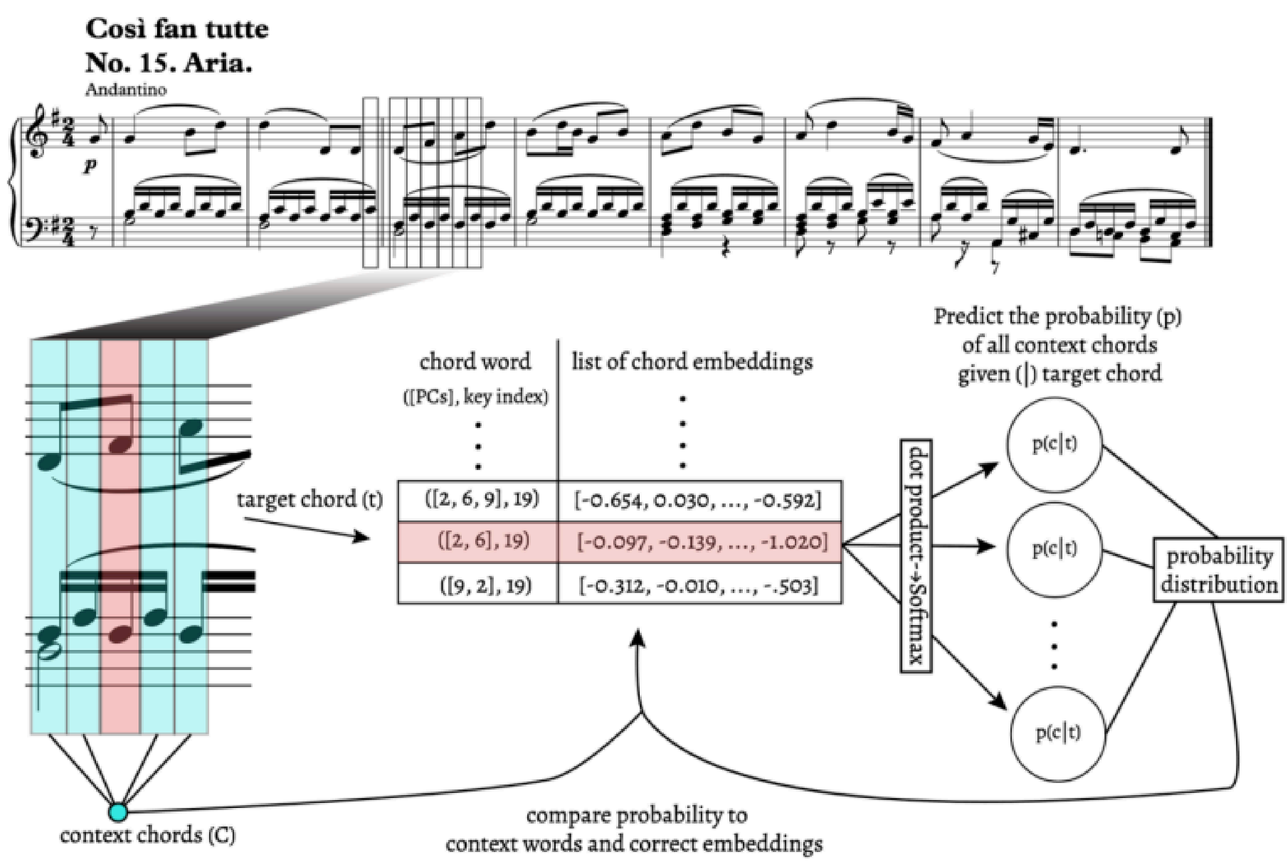

Figure 2. Demonstrating how musical embeddings can be derived from a dataset.

Studying Embeddings

One of the main advantages of embeddings is their numeric nature: because chord embeddings are numerical vectors, we can use mathematical operations to study chord relationships. Take, for example, Figure 3 which uses embeddings derived from pieces by Mozart; the dataset I use contains 882 pieces by Mozart made of 1,587,996 verticalities. The graphic in Figure 3 shows the osine similarity—a metric used to calculate the similarity between two vectors—between a C major triad and diatonic chords in C major. Cosine similarity is the inverse of cosine distance, so we can also think about it as the proximity of two vectors in the vector space. Notice in the figure that F major and G major are the closest triads. This means that chords located a fifth away are closest to C major in Mozart's style. Further, Figure 4 isolates the tonic for each key and clusters them with a dendrogram. Tonic triads cluster by keys: C is close to F and G, A is close to E and D, etc.

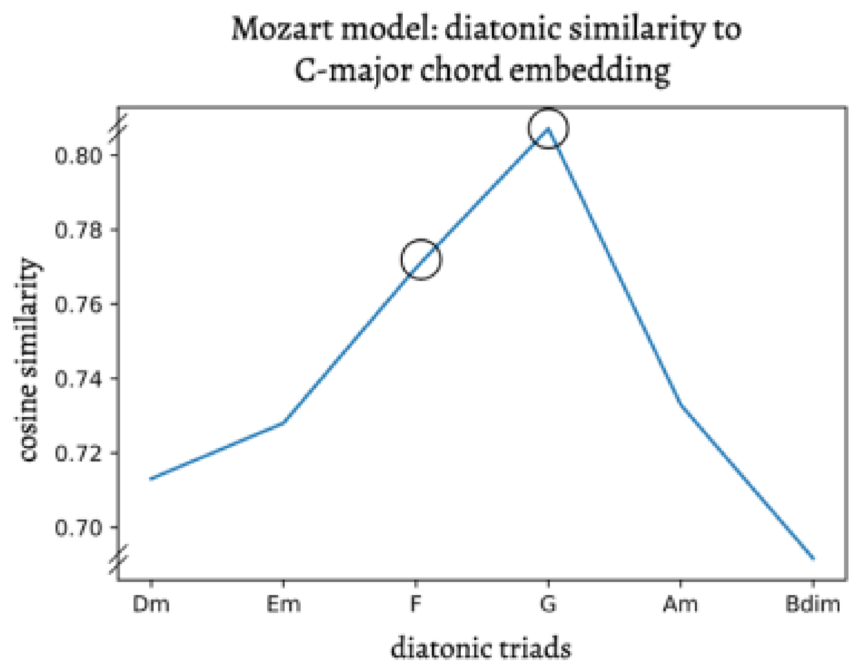

Figure 3. The cosine similarity of C major to other diatonic chords in the key of C major (Mozart model).

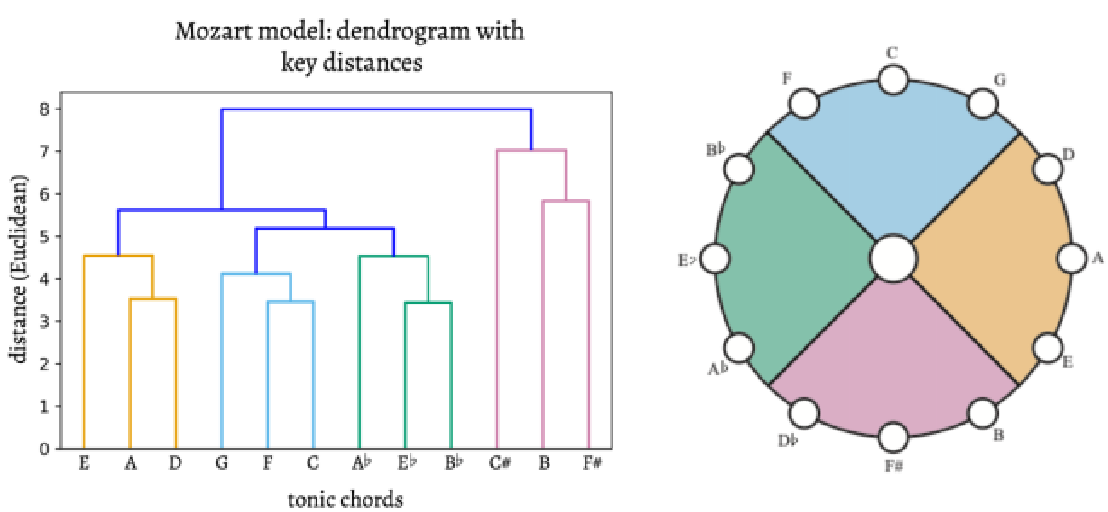

Figure 4. Clustering keys by their relative distance (left). Keys group by their placement on the circle of fifths (mapped to the circle of fifths on the right).

All this confirms the common perspective about the proximity between fifth-related keys and chords in a common-practice style. However, the primacy of fifth-based relationships is typically justified with common-tones: that is, the key of C major is close to G major because they have 6 overlapping pitches in common. In contrast, the procedure here has no concept of musical structure—word2vec does not consider the overlap in common tones when deriving embeddings, only the placement of those collections relative to others: a C major triad and an A minor triad have 2 shared common tones, but the algorithm assigns them completely unrelated embeddings at first. This makes it even more interesting that the circle of fifths is also a product of musical context in addition to musical structure. But is the circle of fifths ubiquitous in other musical styles? Are chords related by always closest?

To test these questions, I created another set of embeddings using a rock dataset—the Rolling Stone Corpus (Temperley and de Clercq, 2011). The dataset contains 200 songs and 36,570 chords. Using the same procedure to generate embeddings, Figure 5 replicates the same visual from Figure 3 with rock music (with slight differences). Firstly, G major is no longer the closest triad—it is the fourth closest. The relative minor (A minor) is instead the closest. Additionally, Ab major and Bb major are relatively close to C major; the proximity of these chords allude to what Nicole Biamonte (2010) has called "aeoleon cadences" (moving from bVIàbVIIàI), and what Chrisopher Doll (2017) has called "rogue dominants" (chords functioning as dominants with a lowered). In a traditional music theory classroom, we are taught (and teach) that the dominant chord is closest to the tonic, but, in a rock style, the dominant is one of many close chords. This shows that taking a prescriptive perspective to musical distances is insensitive to other musical styles.

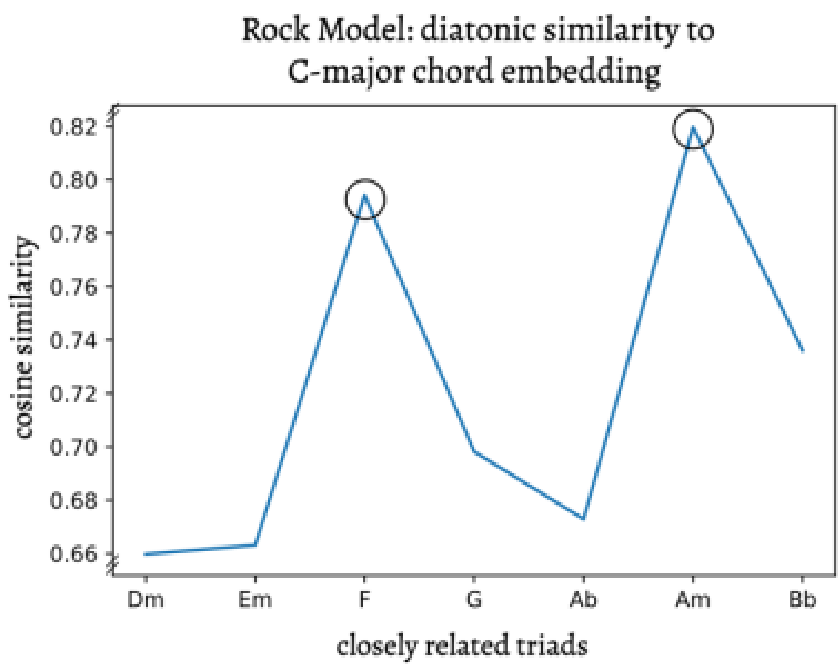

Figure 5. The cosine similarity of C major to other chords in the key of C major (rock model).

We can also use embeddings to listen to music through a particular style—a type of music analysis: by measuring the cosine distances from one chord embedding to the next within a song, we are essentially asking the model to say how surprising the passage is. Because chords that occur near one another are closer in the space, chords that are distant are less probable. Figure 6 shows the cosine distance between all successive chords in Guns N' Roses, "Sweet Child O' Mine" using both Mozart and Rock embeddings. (Note that here I use cosine distance—the inverse of cosine similarity.) Specifically, the figure shows the transition from the chorus to the solo, a moment where the music shifts dramatically from the key of Db major to Eb minor. According to the Classical embeddings, the transitions from 1) VàbVII (AbàCb), 2) bVIIàI (CbàDb), and I in Db major to i in Eb minor are all distant motions. We might stretch the distance analogy and say that someone who is only familiar with a classical style would find this passage very musically surprising.

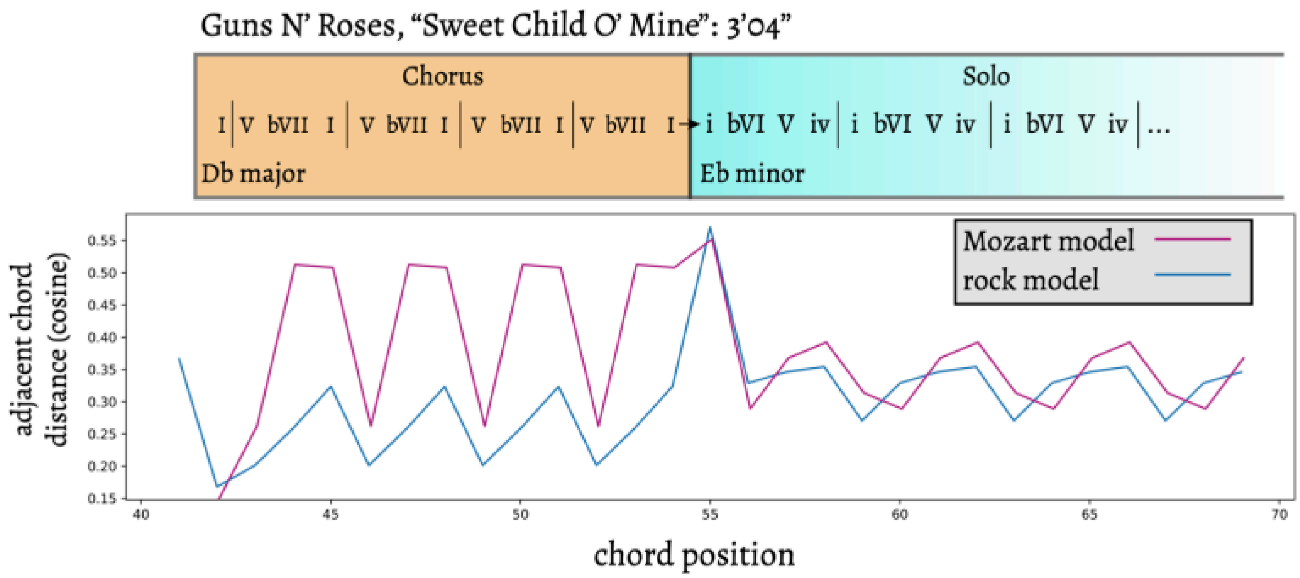

Figure 6. Adjacent chord distances according to a Classical (Mozart) and popular music (rock) model.

On the other hand, the rock model shows the motion from VàbVII and bVIIàI as stylistically typical. The only marked moment between two distant chords is in the transition between sections. So, while both models distinguish the section change as the most distant transition between two successive chords, many other chord transitions are far with the Classical model; if every chord is relatively distant, it diminishes the relative effect of this stark shift to another key. The rock model identifies how marked it is within the musical style, and, in a way, locates the boundary between the two sections with its measurements.

Conclusion

Fixed, expertly designed systems are essential in all aspects of music: they solidify our ability to communicate with a common set of symbols and terms and work exceptionally well at these tasks. At one point, these systems may have been contextually designed, but since gaining such popularity, their application for music is ubiquitous to many as music itself. As I have tried to show in these paragraphs, prescribing too strongly to any system brings with it a set of assumptions and, often, the baggage of a particular musical style. When we use terms designed for one musical system for another, we can implicitly make assumptions about how that music operates. Contextual systems give space for the nuance that breathes life into the stylistic multiplicity we love in music. Because AI models excel at extracting statistical regularities and learning about a given context, they are aptly suited for deriving contextual systems for music. At the conclusion of this paper, I hope the reader is not left with answers, but has additional questions: Could there be an alternate way for visualizing and communicating a musical score without staff notation? What biases do I bring to music, whether reading, listening, performing, or writing it? What does musical distance mean in a perceptual, physical, and computational sense? How do these things overlap? If nothing else, I hope that AI research helps to further complicate our sense of musical representations, freeing it to occupy the messy jumble of meaning that we enjoy investigating when listening to music.

References

Biamonte, Nicole. 2010. “Triadic Modal and Pentatonic Patterns in Rock Music.” Music Theory Spectrum 32, no. 2: 95–110.

Cuthbert, Michael Scott, and Christopher Ariza. 2010. “music21: A Toolkit for Computer-Aided Musicology and Symbolic Music Data.” In Proceedings of the 11th International Society for Music Information Retrieval Conference (ISMIR 2010), 637–42. Utrecht, Netherlands. https://dspace.mit.edu/handle/1721.1/84963.

Doll, Christopher. 2017. Hearing Harmony: Toward a Tonal Theory for the Rock Era. Ann Arbor: University of Michigan Press.

Jurafsky, Daniel, and James H. Martin. 2025. Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models. 3rd ed., draft of January 12, 2025. https://web.stanford.edu/~jurafsky/slp3/.

Temperley, David, and Trevor de Clercq. 2011. “A Corpus Analysis of Rock Harmony.” Popular Music 30 (1): 47–70. https://doi.org/10.1017/S026114301000067X.

White, Christopher Wm., and Ian Quinn. 2016. “The Yale-Classical Archives Corpus.” Empircal Musicology Review 11 (1): 50–58.

About the Author

Dr. Matt Chiu is an Assistant Professor of Music Theory at Baldwin Wallace University. His research straddles the interdisciplinary boundaries between computer science and music theory, with research focused specifically on artificial intelligence and machine learning. He has published articles in Music Theory Online, Music Theory Spectrum, SMT-Pod, Math and Computation in Music, and more. Matt currently serves as the review’s editor for the Journal of Mathematics and Music.

Contributor Title: Researcher

BEYOND THE MACHINE: CUSTOM GPTS AND SEMIOTIC

BY DR. ANTONELLA DI GIULIO

Instructional Design in AI Interfaces

This article investigates the conceptual parallels between Custom GPTs—user-defined versions of GPT-4-turbo—and human-AI communication models, drawing from semiotics, interface theory, and music pedagogy. Using the popular television series Person of Interest as an interpretive framework, we contrast symbolic notions of divine-machine interaction with grounded models of structural design and pedagogical authorship. We argue that Custom GPTs represent a form of externalized cognition, wherein the aesthetic, discursive, and rhetorical scaffolding mirrors both composerly intention and editorial authority. Rather than autonomous agents, GPTs function as programmable mirrors that reflect the epistemic clarity—or vagueness—of their creators.

FICTION, INTERFACE, AND THE SEMIOTIC GAP

The portrayal of artificial intelligence in fiction has long shaped cultural expectations around human-machine relations. In Person of Interest, the character known as "The Machine" serves as an omniscient entity, communicating through predictive signals, text streams, and voice proxies. This fictional representation mirrors theological constructs: Root treats the Machine as divine, while Finch—its creator—designs boundaries and protocols.

The cultural semiotics of these two roles are instructive. Root represents submission to opacity; Finch represents structural authorship. These archetypes provide a compelling metaphor for how we approach modern AI systems such as ChatGPT—especially in their customizable forms. This article explores how Custom GPTs function not as intelligent agents but as semiotic environments shaped entirely by human configuration.

Custom GPTs: Form, Not FunctionOpenAI's Custom GPT builder enables users to create specialized versions of GPT-4-turbo by specifying:

- Behavior and tone instructions
- Formatting and structural constraints
- Reference materials (PDFs, text files)
- Rules for handling vague or ambiguous prompts

Importantly, Custom GPTs do not "learn" from new data or exhibit emergent behavior. They are deterministic systems whose output reflects the precarity and precision of their input parameters. In semiotic terms, they function as coded channels rather than dialogic partners.

Whereas the standard GPT interface operates as a generalized question-answer mechanism, Custom GPTs are high-resolution prompts embedded within a design ecology. Their performance is not a measure of machine intelligence but of instructional clarity—a phenomenon well-known to educators, composers, and interface designers.

MUSIC PEDAGOGY AS A FRAMEWORK FOR AI CONFIGURATION

Pedagogical theory offers a useful lens for understanding the utility and limits of AI customization. Just as a skilled teacher crafts a learning environment through pacing, feedback loops, and scaffolding, a well-designed Custom GPT reflects:

- Consistent linguistic tone (formal/informal, academic/outreach)
- Structural templates (bullet points, paragraph length, narrative order)
- **D**efined affordances and constraints (e.g., avoid generic phrasing, never open with an apology)

This pedagogical metaphor is not merely illustrative—it is operational. Users act as curriculum designers, defining the epistemic terrain within which the GPT operates. The AI system does not generate knowledge; it traverses a space already mapped by the user.

ROOT VS. FINCH: DESIGNING, NOT OBEYING

In semiotic terms, Root's deference to the Machine illustrates an indexical relationship: she treats the system's outputs as signs of a deeper, unknowable intelligence. Finch, by contrast, exemplifies a symbolic relationship—one in which meaning is constructed through shared rules, design structures, and intentional constraints.

Applying this analogy to AI use, we argue for a Finchian model of engagement:

- Custom GPTs are not advisors; they are configurable tools.
- Their apparent intelligence is always derivative.
- **T**heir outputs should be interpreted through the lens of design, not agency.

This framework has implications for how musicians, educators, and administrators use AI systems: not as oracles, but as programmable extensions of their own disciplinary literacy.

TOWARD A SEMIOTICS OF CUSTOM AI DESIGN

Building a Custom GPT is not a technical process—it is a semiotic act. It requires the user to:

- Articulate linguistic tone with precision
- Curate relevant textual corpora
- Anticipate user intentions and edge cases
- Define a communicative contract

In this sense, Custom GPTs represent a convergence of interface design, genre theory, and rhetorical modeling. Their success depends less on algorithmic strength and more on the user's ability to design clear, reusable epistemic pathways.

REDEFINING THE ARTIST'S ROLE: COMPOSERS, NOT CONSUMERS

In the age of AI-assisted content generation, the most meaningful interactions arise not from blind adoption but from intentional authorship. Like composers scoring for an unfamiliar ensemble, users of Custom GPTs must anticipate interpretive behavior, encode structural logic, and listen critically.

The Machine is not God. It is an interface. And Custom GPTs are not thinkers; they are reflections.

To use them well is to design them well.
To design them well is to know what you want them to say—every time.

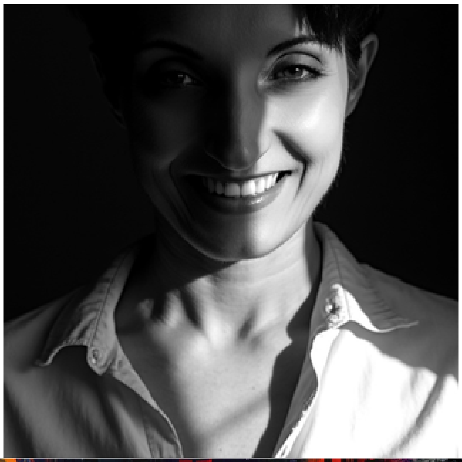

About the Author

Dr. Antonella Di Giulio is a pianist, scholar, and founder of MusicalQ, a platform exploring the intersection of music, cognition, and creative practice. She is the Program Chair for the Allegheny Chapter of the American Musicological Society, organizer of international music festivals, and a leading voice in music pedagogy, technology, and interdisciplinary innovation.

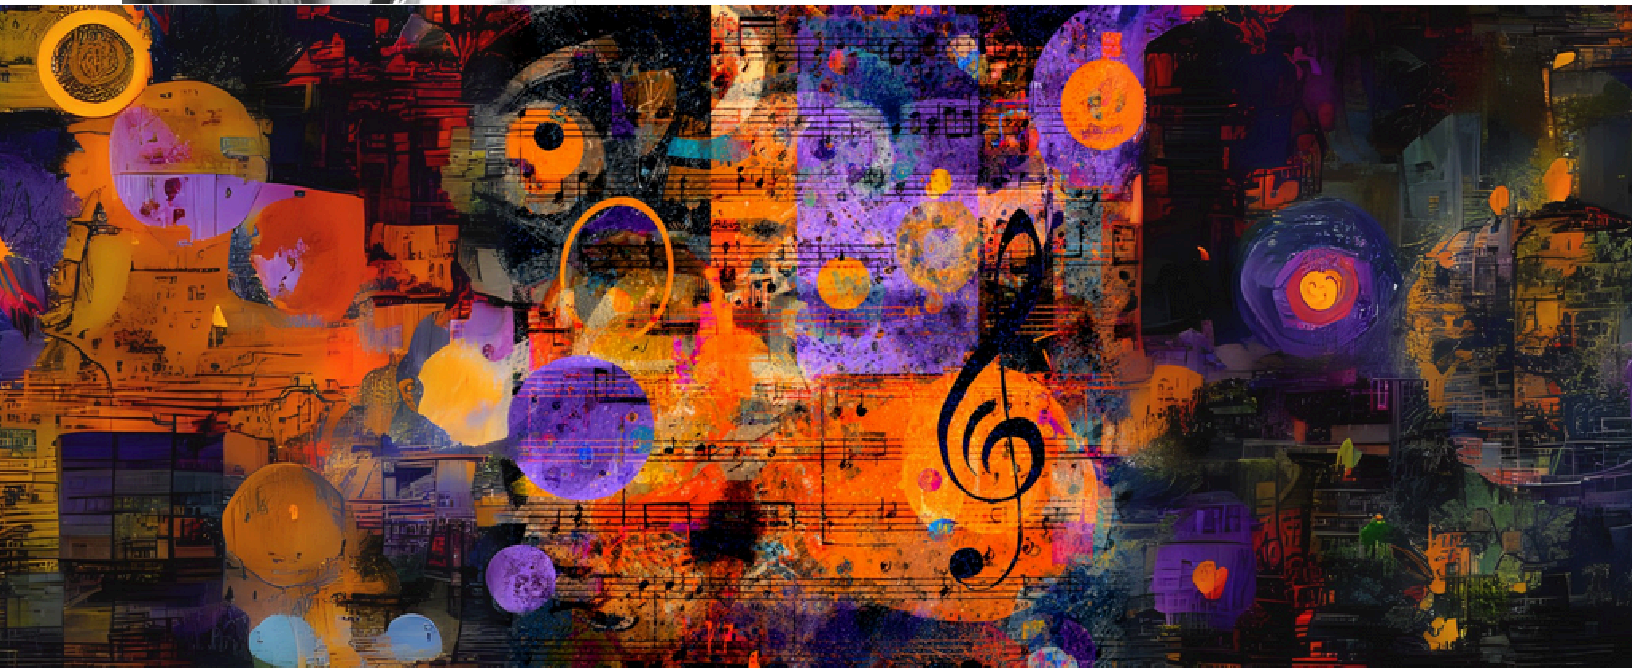

COMPOSITIONAL PRACTICE: GENERATIVE MUSIC AS A RESPONSE TO AI MACHINE LEARNING MODELS

by Dr. Matias Homar

This short opinion piece explores the use of generative music composition techniques as a response to the rise of AI and its machine learning models in music. My approach builds on the integration of open-source software with custom-built hardware to create reactive environments where music and visuals emerge from live interaction. These environments are not passive or merely programmed—they are alive, constantly shaped by the communication between performer(s) and system.


www.musicaiq.com

At the core of my perspective is the framework of enactivism. Drawing from Varela, Thompson, and Rosch, I focus on the idea of the present moment and awareness as essential components in the creation of these interactive environments. The generative system becomes a co-creator, but only in its direct engagement with human presence. In this sense, generative music is not about automating composition, but about generating conditions for emergence through embodied participation.

AI-generated music, by contrast, often functions through training on vast datasets to simulate human musical behaviors. This process tends to flatten nuance and reinforce stylistic norms. What results is a growing aesthetic of sameness, where difference is erased in favor of statistical probability. My experience with generative systems shows the opposite: when a system is developed as a reactive environment—one that listens and responds in real time—it enhances the creative experience, invites inclusion, and allows for unique musical moments to arise.

Generative music from an enactive perspective offers something AI does not: a living system that evolves in dialogue with those who interact with it. It is not about control or prediction but about creating space for unexpected outcomes grounded in mutual responsiveness. This approach can support meaningful musical experiences regardless of the performer's background or skill level, emphasizing presence and community over perfection or polish.

In the end, my personal view is that generative systems, when designed with awareness and intentionality, surpass the use of AI as compositional tools. They allow for a shared, situated experience that is specific to the people and context involved.

They embrace difference, foster communication, and resist the pull toward standardization. Compositions such as Glitch/Bending Sound and Schizophonia, then, become a way to reclaim music as communual practice—rooted in interaction, diversity, and presence.

About the Author

Dr. Matías Homar is a composer, music technologist, and researcher specializing in real-time generative systems, interactive sound design, and digital performance environments. His work explores the intersection of enactivism, embodied cognition, and algorithmic composition, often involving custom-built hardware and open-source platforms. He has presented internationally at conferences on music technology, and his compositions have been featured in experimental and multimedia contexts. Dr. Homar currently teaches and collaborates across institutions in Latin America and Europe, with a focus on reactive environments and inclusive creative practice.

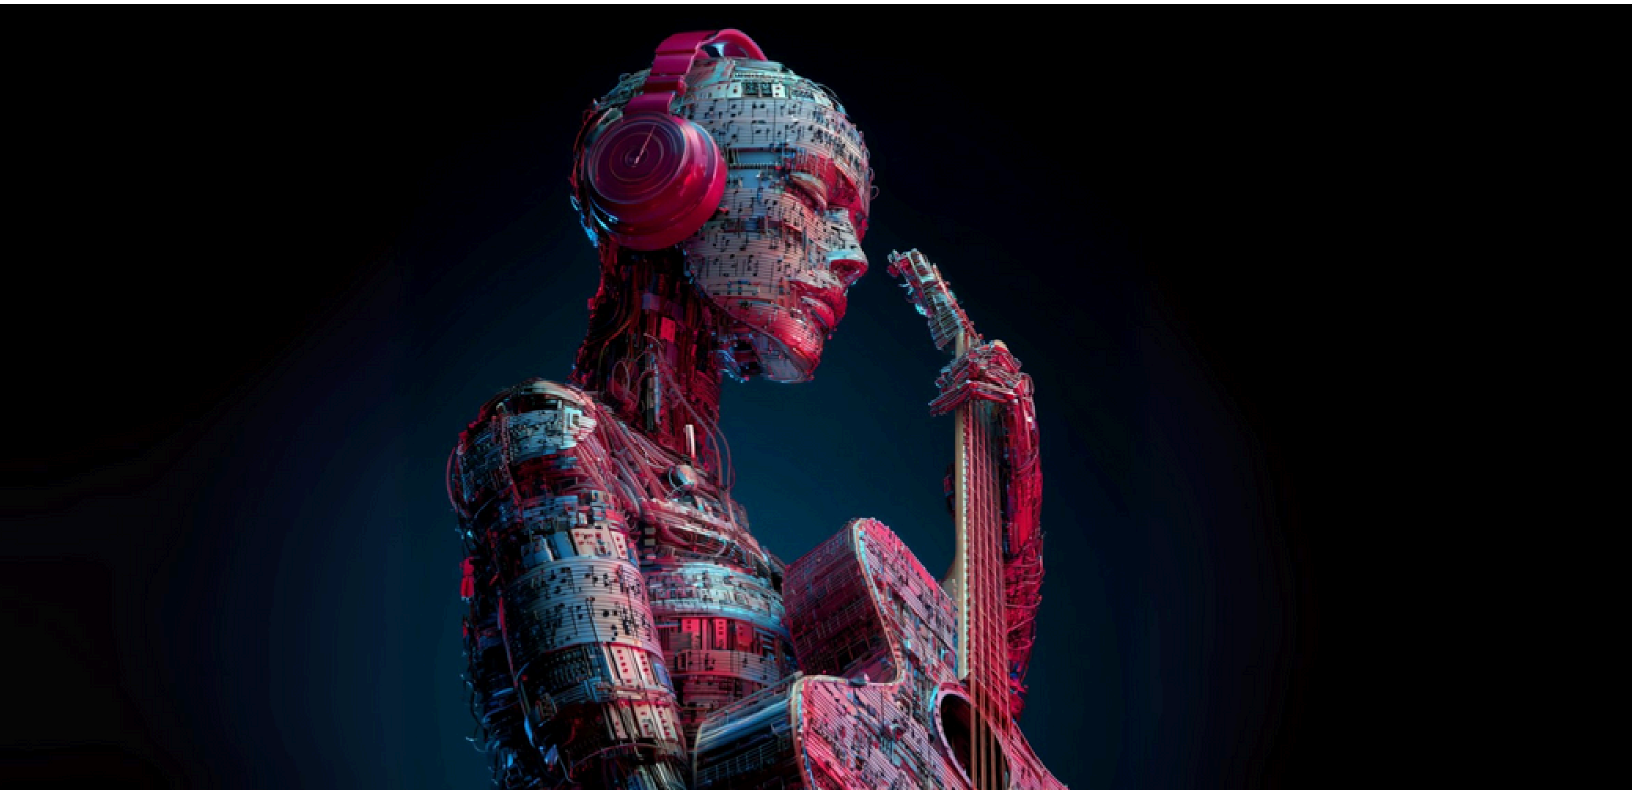

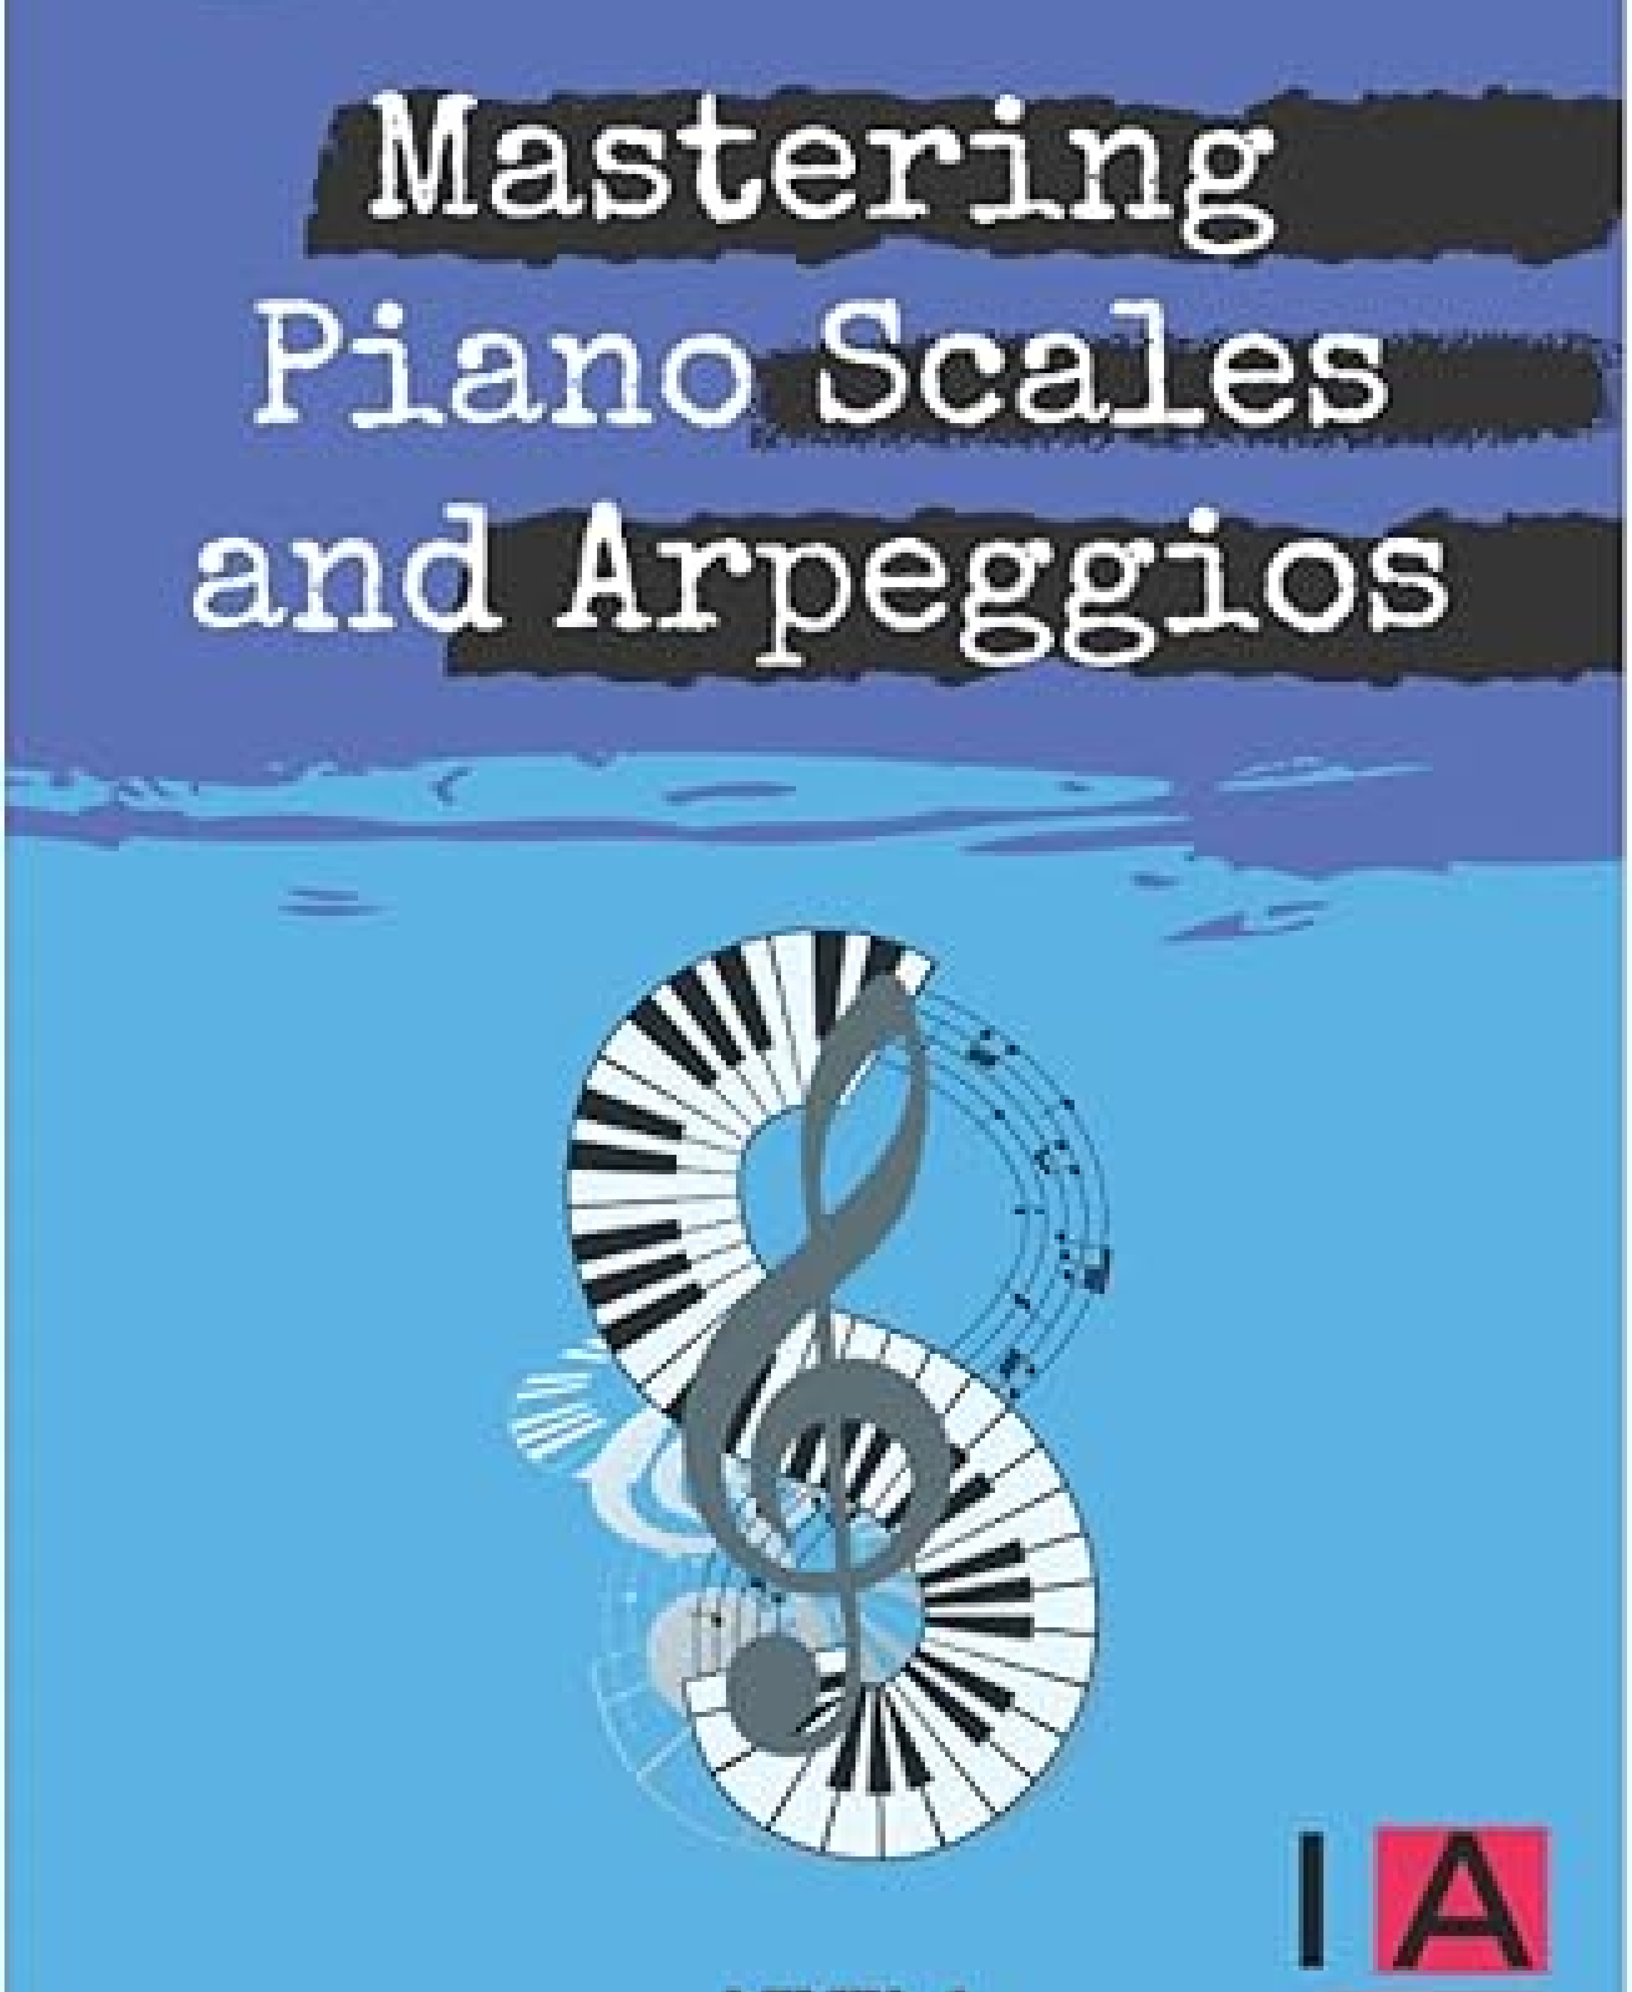
Mastering
Piano Scales
and Arpeggios

BOOKSHELF

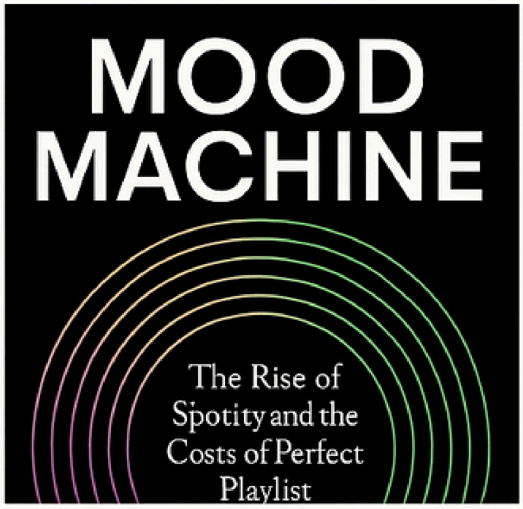

Mood Machine

The Rise of Spotify and the Costs of the Perfect Playlist

Liz Pelly

A scathing analysis of ardigrothmic listening and streaming culture, Pelly questions what's lost when music Becomes background, optimized for mood instead of meaning.

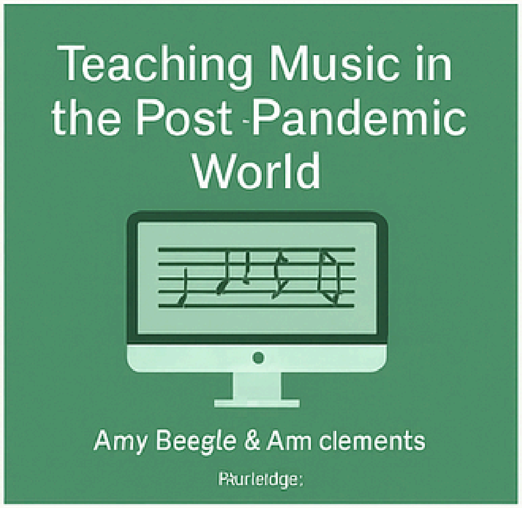

Teaching Music in the Post-Pandemic World

Amy Beegle & Ann Clements

Routledge

This collection explores how remote instruction, digital tools, and pandemic disruptions have reshaped music teac-hing: Topics include accessibility, hyb-rid learning, and trauma-informed pod-agogy.

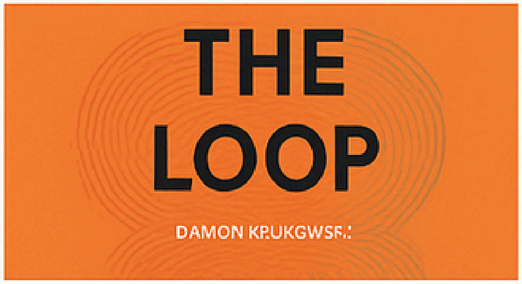

The Loop: How Technology is Creating a New Musical Language

Damon Krukowski

Musician and writer Damon Krukowski explores how digital tools—loops, samples, and streaming interfaces—have transformed the grammar of music itself. With deep insight and sharp critique, he asks: What happens to musical meaning when we design music for infinite repeatability? And how does platform logic reshape what artists create?

Digital Creativity in Music Education

Jude Brereton & Liz Dobson

Routledge

This essential volume bridges theory and practice in digital music pedagogy. It provides educators with real-world case studies on integrating creative coding, DAWs, loop-based composition, and digital improvisation in the classroom. An ideal reference for rethinking 21st-century music instruction across age groups and abilities.

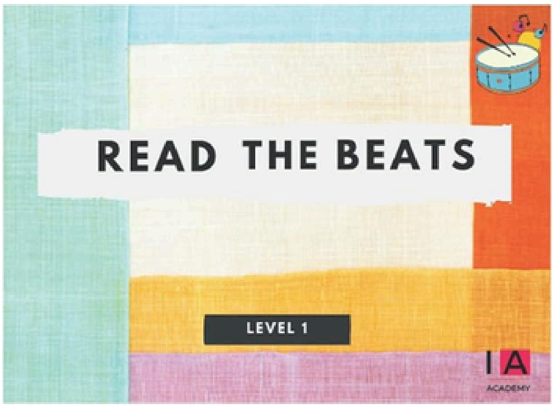

READING MUSIC MADE SIMPLE

available on amazon only

CAN AI FEEL MUSIC? MODELING EMOTION IN SOUND

This article explores how artificial intelligence models musical emotion—examining the intersection of machine learning, music information retrieval (MIR), and affective computing. Using recent studies in deep learning, we interrogate how AI systems are trained to recognize, classify, and generate emotionally coded musical output. Drawing from interdisciplinary sources including cognitive psychology, semiotics, and algorithmic composition, we reflect on what it means for a machine to "understand" feeling, and what the implications are for musicians, audiences, and the future of expressive technology.

EMOTION AS STRUCTURE AND SIGNAL

Music, by its very nature, encodes emotion. It does so not only through texted content, but also via structure, timbre, harmonic expectation, tempo, and gesture. Yet while human listeners interpret emotional nuance through context, experience, and embodied intuition, AI systems must learn emotion through proxies—quantifiable features that correlate with affective states.

The field of Music Emotion Recognition (MER) seeks to map emotional qualities—e.g., sadness, joy, tension—onto acoustic or symbolic data using models trained on labeled corpora. These models range from simple rule-based classifiers to complex deep neural networks that ingest hundreds of features, from spectral centroid to rhythmic density.

The question is no longer whether machines can recognize emotion in music, but how they interpret emotional content—and whether that process tells us something profound (or disturbing) about both musical meaning and affective representation.

AFFECTIVE COMPUTING AND MUSICAL EMOTION RECOGNITION (MER)

MER resides at the intersection of multiple domains:

- Affective computing: enabling machines to detect and respond to human emotions
- Music information retrieval: extracting meaningful data from audio signals
- Psychology of music: understanding how humans perceive and communicate feeling through sound
-

Three dominant models guide current research:

1) Categorical models: Emotions as discrete labels (e.g., happy, sad, angry)

2) Dimensional models: Emotions mapped in valence-arousal space

3) Multimodal models: Emotions inferred from audio + context (e.g., lyrics, performance data, visual stimuli)

Recent breakthroughs in deep learning (especially CNNs and LSTMs) have improved performance in emotion classification tasks, especially in real-time applications like playlist generation, music therapy, and adaptive sound design.

MUSICAL FEATURES AS EMOTIONAL PROXIES

To a machine, emotion must be translated into data. Some common input features include:

- Tempo and meter: Faster tempos often correlate with high arousal
- Mode and harmony: Minor keys or dissonance often predict sadness or tension
- Timbre and spectral shape: Brightness, noisiness, or attack can suggest emotional intensity
- Melodic contour: Rising or falling lines may signal hope, grief, or yearning
- Dynamic range: Variability may suggest expressive volatility

However, correlation is not comprehension. What machines detect is not emotion itself, but its acoustic trace.

APPLICATIONS AND ETHICAL QUESTIONS

AI-powered emotion modeling has found commercial application in:

- Streaming services (e.g., mood-based playlisting)
- Film and game scoring (generative adaptive music)
- Music therapy (emotion tracking for regulation and treatment)
- Educational apps (feedback systems tailored to mood or motivation)

But these applications raise questions:

- **D**oes mapping emotion onto fixed categories flatten human nuance?
- **C**an machine-generated emotionality ever be truly expressive—or is it pastiche?
- What biases are encoded in training datasets (e.g., Western tonality, gendered assumptions)?
- What are the implications of using emotional data for personalization or surveillance?

RETHINKING EMOTIONAL INTELLIGENCE IN AI MUSIC SYSTEMS

Rather than asking whether AI can feel, a better question may be: what models of emotion do we embed in AI, and what do they reflect about us?

MER systems are only as good as the assumptions they're built on. If we feed them datasets labeled by consensus but stripped of cultural or embodied context, we risk encoding an impoverished vision of musical affect. Conversely, AI systems may help us discover latent patterns in emotional structure that humans alone might overlook. Ultimately, the conversation is not about replication—it's about representation.

CONCLUSION: EMPATHY OR EMULATION?

Music remains one of the most affectively rich human practices. While AI systems can model emotional content with growing accuracy, they do not feel—they pattern-match. Their emotional intelligence is borrowed, coded, and derivative.

But these systems are not useless. They challenge us to define emotion more clearly, to confront the semiotic assumptions behind musical meaning, and to ask whether "feeling" is something we recognize only in ourselves—or can begin to model in the other.

In this way, AI becomes not a substitute for emotion, but a provocation.

COLLABORATING WITH US: ENGAGE, SHARE, AND ELEVATE YOUR EXPERTISE

MusicalQ is a vibrant platform advocating for music and musicians worldwide. It offers professionals the chance to share insights, engage with peers, and contribute to the evolution of music education and performance. Here's how you can collaborate with MusicalQ through the platform, the monthly blog, and the magazine:

This is a dynamic space where educators, performers, and scholars can share their expertise and ideas. By collaborating on the platform, you have the opportunity to contribute articles, research findings, lesson plans, and innovative teaching methods. Whether you're an established professional or an emerging voice in the field, your contributions can help shape the future of music education and performance. The platform encourages a collaborative approach, fostering a community of professionals who are passionate about music and education.

HOW TO COLLABORATE WITH MUSICAIQ

- **Writing for the Monthly Blog**

The MusicalQ monthly blog is your platform to share insights, experiences, and perspectives on current trends, challenges, and innovations in music education, performance, and research. This blog is accessible to professionals at all levels, offering an opportunity to present fresh takes on established subjects or introduce new concepts. Whether you’re discussing emerging trends or reflecting on personal experiences, the blog is your space to inspire and educate a diverse audience.

- **Publishing in the MusicalQ Magazine**

If you're interested in providing more in-depth analysis, consider contributing to the MusicalQ Magazine. This platform is ideal for longer articles, interviews, and features that explore complex topics in music education and performance. The magazine is committed to publishing content that informs, challenges, and inspires its readers, encouraging thoughtful discourse on the future of music. By writing for the magazine, you can help shape the conversations that will influence the music industry for years to come.

- **Advocating for Music and Musicians Worldwide**

MusicalQ is dedicated to promoting music and supporting musicians on a global scale. Your collaboration helps amplify the voices of musicians and educators, ensuring that music education remains relevant and accessible. By contributing, you become part of a larger movement advocating for the importance of music in our culture and educational systems.

How to Get Involved

To collaborate with MusicalQ, start by reaching out with your ideas or proposals. Whether you're interested in writing a blog post, contributing to the magazine, or sharing resources, the MusicalQ team is eager to hear from you. Submissions are reviewed for quality, relevance, and originality, ensuring that all content aligns with MusicalQ's mission to advance music education and performance.

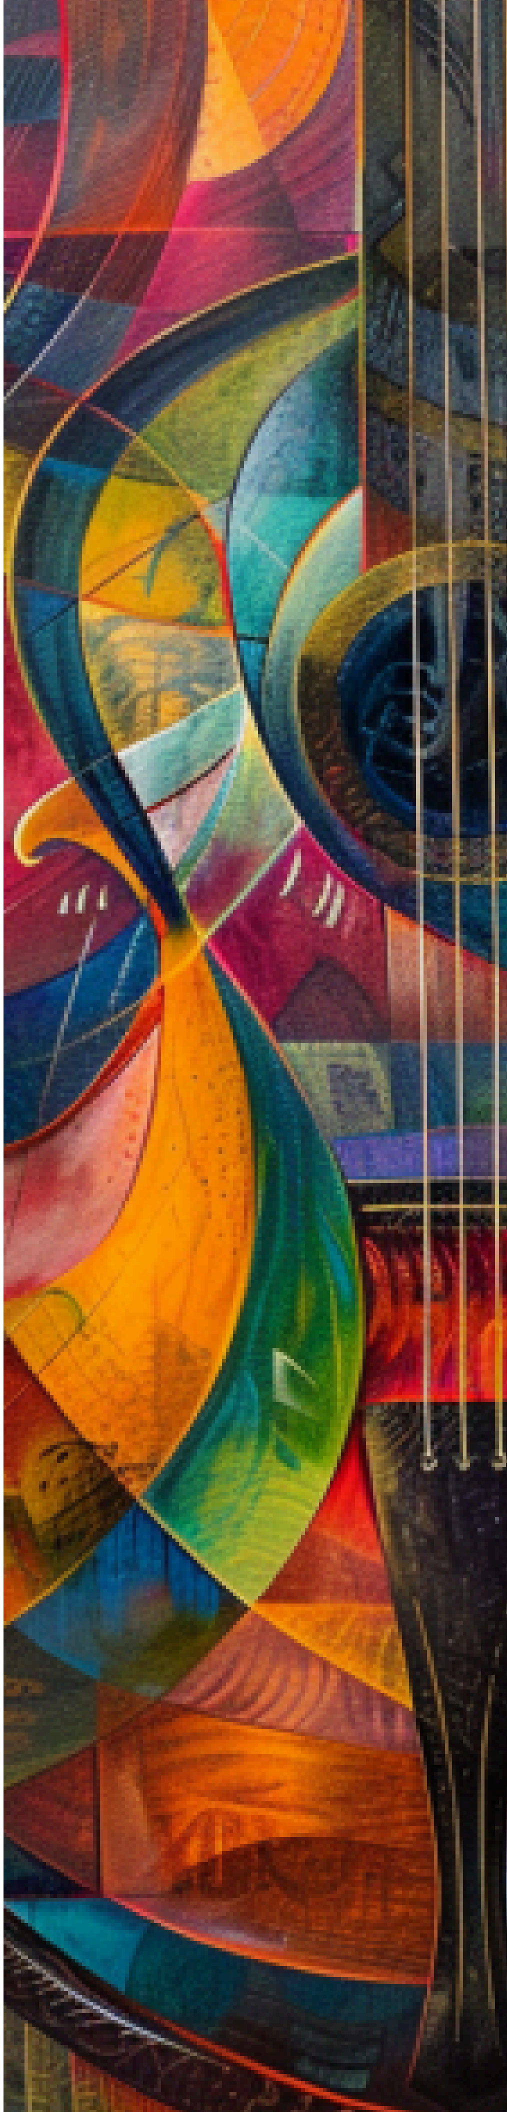

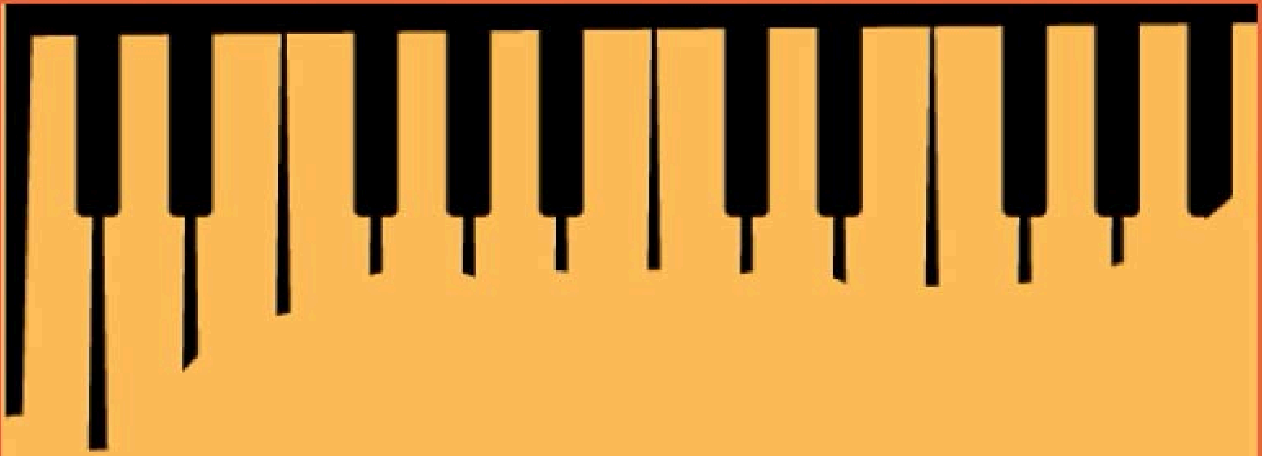

MUSIC ESSENTIALS

Piano Practice Journal

The most effective journal for your piano practice

GREEN ELEPHANT PRESS

MOOD MACHINE: THE RISE OF SPOTIFY AND THE COSTS OF THE PERFECT PLAYLIST

AUTHOR: LIZ PELLY
PUBLISHER: HODDER & STOUGHTON
PUBLISHED: JANUARY 2025

BY DR. ANTONELLA DI GIULIO

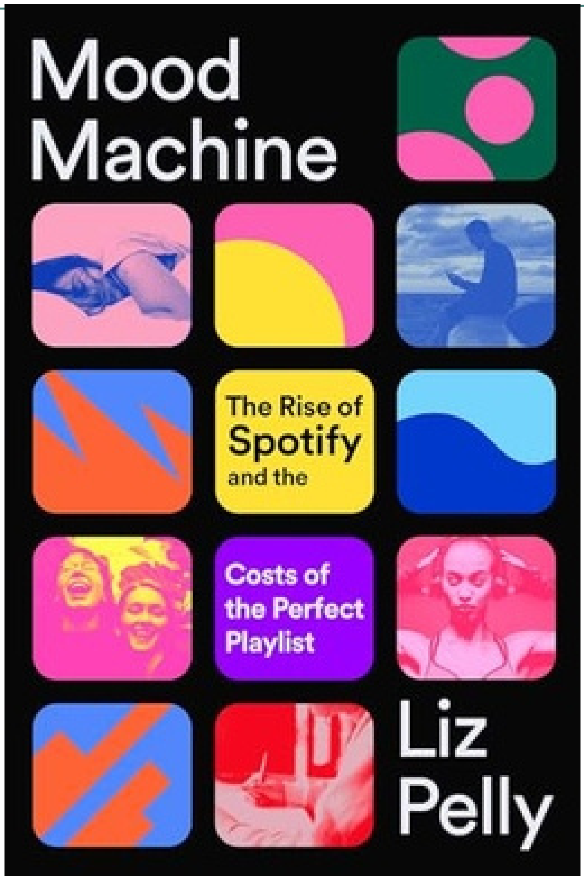

What's it about?

In Mood Machine, Liz Pelly delivers a sharp and deeply informed critique of Spotify's algorithmic culture. More than a media analysis, this book is a cultural manifesto—probing how playlists, engagement metrics, and mood-based design shape not just what we hear, but what we create.

Key Themes

• Algorithmic curation vs. artistic freedom
• Streaming platforms as cultural gatekeepers
• Mood playlists and "ambient conformity"
• The future of musical diversity in tech-driven spaces

Why It Matters

Pelly argues that streaming algorithms flatten musical nuance and reinforce sameness. Mood Machine makes a compelling case for conscious curation, independent platforms, and reclaiming listening as an intentional act.

In what is arguably the most incisive chapter of Mood Machine, Liz Pelly unpacks the phenomenon she terms “ambient conformity”—the rise of background music as both aesthetic strategy and ideological filter. This is not just about the lo-fi beats that populate study playlists. It’s about an entire design logic baked into Spotify’s algorithmic infrastructure: music that doesn’t interrupt, challenge, or surprise.

Pelly tracks how “chill,” “focus,” and “vibes” playlists have come to dominate the platform, not by organic listener demand, but through design choices that favor streamability over artistic depth. The result is a cultural shift: songs are shorter, loops are smoother, and tension—harmonic, lyrical, rhythmic—is increasingly avoided.

“Streaming didn’t just change the music industry—it changed how music means,” Pelly writes.

This chapter is especially relevant for music educators and curators, who now face students and audiences shaped by sonic environments where variation is a bug, not a feature. For composers and performers, it raises critical questions:

- How do we write for ears trained on seamlessness?
- How do we preserve unpredictability, dynamic range, or expressive risk in an age of constant ambient wash?

Pelly doesn’t call for a return to some golden pre-digital past. Instead, she challenges readers to reclaim the radical possibilities of listening—to teach it, to program it, to value it again as an act of attention, not just a mood enhancer.

This chapter alone makes Mood Machine required reading for anyone working at the intersection of music, technology, and education.

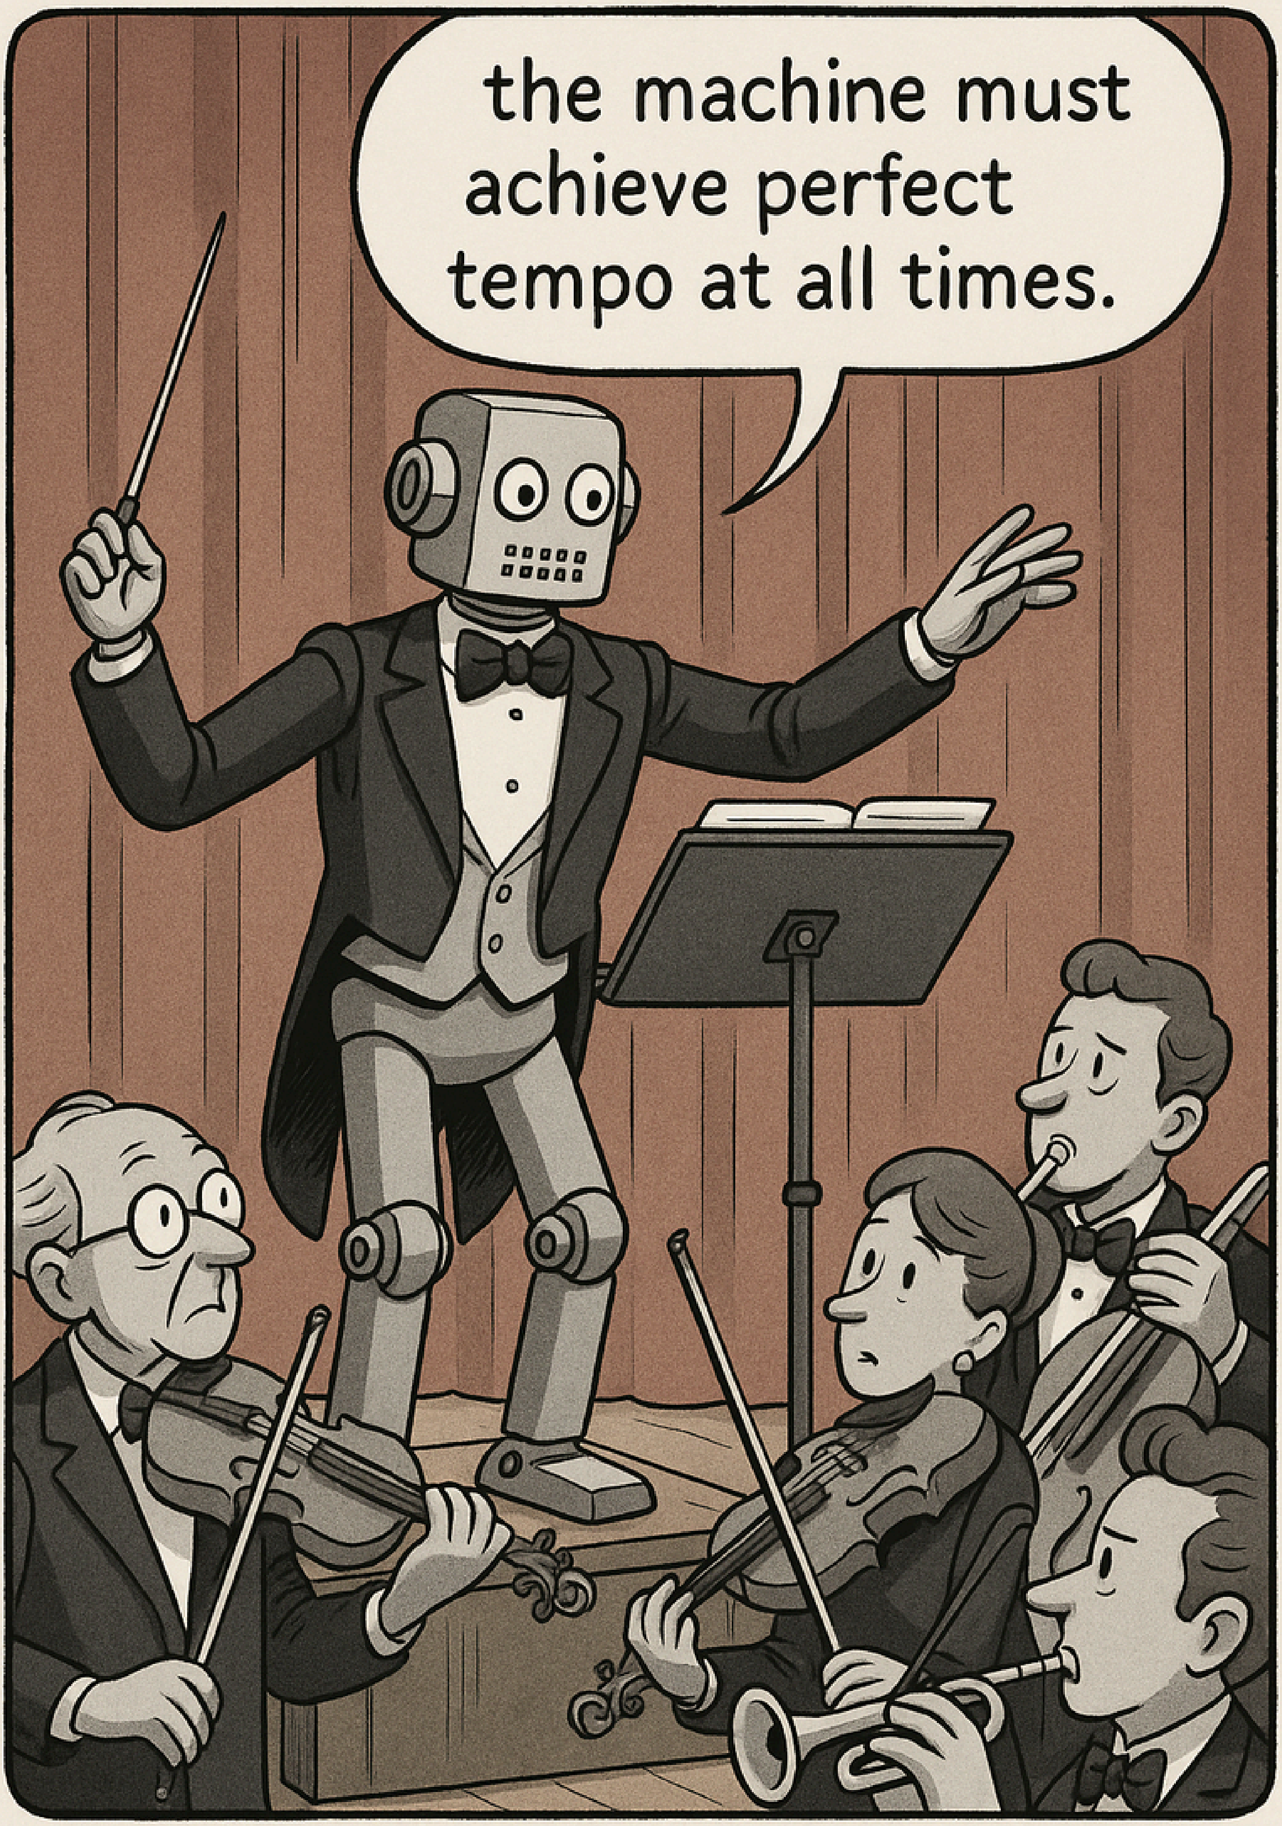

ConductorGPT