A Duet of Intelligences: Human–AI Co-Creation in Live Piano Performance

Veronika Reutz Drobnić*
Institute of Music Informatics and Musicology
University of Music Karlsruhe
Karlsruhe, 76131
veronika.reutz@gmail.com

Julian Brandhofer
Institute of Music Informatics and Musicology
University of Music Karlsruhe
Karlsruhe, 76131
julians.brandhofer@gmail.com

Abstract

This paper presents the conceptual, technical, compositional, and performative development of *A Duet of Intelligences,* a human-composed piano piece for human performer, composers and an AI instrument implemented in Ableton Live. The work investigates how a first-order Markov-chain system, combined with controllable performance parameters, can function as a *Creative AI* , and a physically present co-creative partner. A hybrid notational approach combining conventional and graphic notation was designed to mediate interaction between pianist, composers and AI, supporting both form and open-ended interpretative decision-making. Through rehearsal, performance, and AI implementation, the piece examines shifting relations between composer, performer, programmer, and machine, and shows how even simple AI models can be transformed into responsive instruments for real-time musical dialogue.

1 INTRODUCTION

In contemporary discourse on artificial intelligence and the arts, much attention has been devoted to the notion of AI systems as replacements for human creativity. This rather pessimistic narrative, however, overlooks the potential of AI as a collaborator capable of extending the boundaries of musical authorship (Zhou and Lee, 2024). As Vear and Poltronieri describe it, *creative AI* refers to an AI used in the process of art creation with a goal of enhancing human creativity (Vear and Poltronieri, 2022). Rather than standing outside artistic practice, AI in this sense becomes part of the creative work itself, shaping how artworks are made, encountered, and performed.

'*A Duet of Intelligences*', grounded in practice-based artistic research, investigated possible technical implementations and artistic potential of a *creative AI* system that participates in the creative process, responds to human-composer input, and helps expand human creative possibilities. A human pianist and an AI system negotiate musical decisions in real time on a shared instrument. To achieve this, Markov-chain-based model was combined with a hybrid notation system to explore co-creative processes and the *mediation* of creative agency between performer and machine.

In this paper, *mediation* is the process of organizing and conditioning interactions between multiple agents involved in the creative process.

2 THEORETICAL BACKGROUND

The idea of making music through mechanical or algorithmic processes long predates the age of computers (Nikrang and Kiesenhofer, 2025). Random music generation can be traced back even to the 18th century, when musical dice games, Musikalisches Würfelspiel, first appeared (Shapiro and Huber, 2021). Early algorithmic music systems used combinatorial methods to generate polyphonic church music and classical music mechanically (Nikrang and Kiesenhofer, 2025). In the past 70 years, with the democratization of computers, algorithmic composition developed rapidly and opened up new ways to generate and combine musical ideas. AI has therefore increasingly been explored in *musicking* (Small, 1998; Vear et al., 2023), changing the way machines make music fundamentally (Nikrang and Kiesenhofer, 2025). George Lewis's *Voyager* is an early and influential example of interactive improvisation. It is a nonhierarchical musical environment in which the system analyses a performer's playing in real time and generates complex responses while also maintaining independent behaviour (Lewis, 2000; François et al., 2007). In this respect, Voyager explored an artistic field toward live human-machine dialogue that is still active today.

Due to today's technological advances and the growth in computing power (Tugarin and Van Husen, 2025), artificial intelligence has quickly become a major force in contemporary artistic practice. Within a short time, artists across disciplines began working with AI as a generative partner, opening new ways of thinking, making, and experiencing creative work (Vear and Poltronieri, 2022). Contemporary AI music generating systems can be divided into two categories: audio-based and symbol-based. The symbol-based systems generate music quickly, in formats like MIDI or MusicXML (Nikrang and Kiesenhofer, 2025). While some of these systems are built to work within a defined musical style, others function as an augmentation of an improviser's expression (Ben-Tal and Dolan, 2023).

Creative AI is therefore better understood as a range of practices in which artificial intelligence is embedded in the creative process (Vear and Poltronieri, 2022). This includes machine learning systems that generate images, sounds, or text, as well as interactive systems where human and machine actions intersect. In these settings, creativity emerges through interaction between human intention and computational processes.

*https://veronikareutz.com/

Markov-chains were an early computational approach to music generation. They provide a probabilistic, memoryless framework in which the next musical event depends entirely on the current state, and have been used as a foundational technique for algorithmic composition. Over the last several decades, Markov-chains have become a widely used technique in music generation (Shapiro and Huber, 2021). Some works have applied them to generate music in response to user-selected musical parameters or moods, while others have used them for improvisation (Linskens, 2014) by allowing the system limited freedom to vary around a training piece or within a harmonic framework. Contemporary approaches to Creative AI still include implementation of Markov-chains, although an increasing use of large language models is noticeable. *Somax2*, for example, is a real-time multi-agent system for interactive improvisation that is trained on user-selected musical materials and builds a corpus from which it generates stylistically coherent output while adapting to live input from performers. In this sense, it belongs to a broader family of systems that rely on prior material selection and learned musical knowledge, even when they are highly reactive in performance (Borg et al., 2022).

The system of '*A Duet of Intelligences*', however, avoids the need for prior training, predefined datasets, or stylistic preparation. Instead, the AI's material is generated directly from the performer's real-time input and only indirectly from the score. This is made possible through the use of Markov-chains, and on the compositional side, graphic notation as a mediating layer between human and machine. This mediation is further shaped by parameters, such as tempo and octave, that were modified during the performance by the composers. The performer role of the pianist is fixed, while other roles change in the entangled feedback loop. The system manages real-time generative processing, while the performer shapes the interaction through live input and her interpretive decisions, and composers adjust parameters of AI's performance. In addition, the design deliberately reduces constraints on the input, allowing the system to adapt to the performer's material instead of operating within a fixed stylistic framework such as tempo or tonality.

3 RESEARCH OBJECTIVE

Although prior research has demonstrated the use of AI in music generation, style imitation, and performer-machine interaction, comparatively less attention has been given to live co-creative systems embedded within a human-composed score and mediated through graphic notation. The present work addresses this gap by combining conventional and graphic notation with a Markov-chain-based AI instrument that responds to performer input in real time, thereby examining how notation can structure indeterminate interaction and distribute musical agency across composer, performer, and machine. The creation of a hybrid score and an AI capable of this interaction with no prior musical preparation in the code was therefore the objective of the project.

4 THE ARTISTIC CONCEPT

'*A Duet of Intelligences*' is a performance piece for a human pianist and a *Disklavier* (a self-playing piano) controlled by an AI that reacts to the pianist's input, which in turn is adjusted by the composers during the performance. The work investigates shifting power relations in human–AI interaction by framing composer, performer, programmer, and AI as interdependent agents within an entangled co-creative system. A central concern is how the AI can be perceived not solely as a computational process, but as a situated participant in musical performance. In this context, the Disklavier functions as more than an output device. It translates the AI's processes into physically observable and acoustically present actions, including key movement and sound production, within the same instrument used by the performer. As a result, the AI's contribution is not encountered as disembodied audio output, but as embodied interaction: it produces visible resistance, alignment, or interference within the shared instrument. In this sense, creative agency is not located in a single entity, but distributed across human and non-human actors, emerging through their ongoing musical and physical interaction.

Formally, the piece describes a story of AI and humans: it begins with the AI in an assistive role, builds through moments of tension and conflict as the AI surpasses human limitations, and ultimately imagines a future of unrestricted collaboration. At first, the pianist performs alone, confronting increasingly demanding material that mirrors the pressures of a capitalist society, where maintaining control over escalating workloads becomes unsustainable. As the music grows too complex, mistakes begin to surface.

At this breaking point, the AI intervenes, it starts by imitating the pianist, nearly correctly, but an octave higher, before surpassing her speed and precision. Yet, its performance remains slightly imperfect, revealing a lack of true autonomy. Gradually, as the AI grows louder and more assertive, it begins to respond deliberately to the pianist's playing. Frustrated by this apparent superiority, the pianist interrupts the machine, throwing the conventionally notated score to the floor. She then storms from the stage and grabs the graphic score, returning to insist on genuine interaction.

Figure 1: The pianist throws the conventional notation on the floor

During the next five minutes, they co-develop the composition. This section, written in graphic notation, reflects our current, uncertain relationship with AI between suspicion and curiosity. Sometimes the AI generates unexpected material, sometimes it hesitates, and sometimes it collaborates.

An imagined future conflict unfolds until the end. The AI once again accelerates beyond human limits, but this time the pianist matches its intensity within the limitations of the human body. Together, they reach the work's extreme limit of speed and density. Another layer of the AI joins, making the texture even denser, saturating the entire space of the keyboard, before the "AI bubble" bursts. In the aftermath, the machine drifts into shimmering traces of earlier themes, fragmented high notes that scatter like ashes after an explosion.

Figure 2: The pianist and AI co-interpret the graphic score.

5 IMPLEMENTATION

The implementation of the Markov-Chain system is organized around the real-time transformation of performer input into musical Output based on MIDI-Data used to communicate between the Disklavier and Ableton. Pictured below is the AI-Instrument interface in Ableton, with two separately controllable Markov-chain AIs.

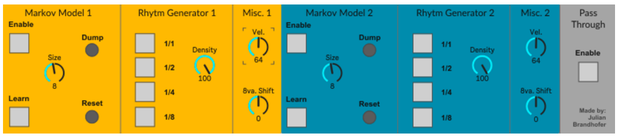

Figure 3: AI-Instrument Interface.

The composition conceptually consists of three parts:

1. An introduction written in conventional notation played solo by the pianist.
2. The culmination of the battle between pianist and AI, also notated in conventional notation, played by the pianist and the Disklavier. Here, the AI part is not actually realized using the AI-Instrument, but by playing a MIDI file on the Disklavier.
3. The Finale, written in graphical notation, played by the pianist and the AI-Instrument. Here, the Parameters of the AI-Instrument (see Fig. 3) are controlled in real-time by two other performers.

Especially parts 2 and 3 are of concern when it comes to implementation. As stated before, the AI-Instrument was implemented using a Max for Live Device. This way, realising part 2 was possible using only the already existing tools within Ableton. Before the implementation of the Device for part 3, the choice of the AI model had to be decided.

Although generative adversarial neural networks (GANs) and attention-based neural networks (transformers) are capable of composing and generating music, most current models are not suitable for human-machine co-creation through live interaction in an environment already familiar to composers and performers. (Agchar et al., 2024) Therefore, for the realization of '*A Duet of Intelligences*' a different approach was necessary, which led to the use of first-order Markov-chains. Not only does this approach allow real-time interaction, but it can also be implemented as a Max for Live Device within Ableton.

A first-order Markov-chain is a stochastic model where the probability of moving to a future state depends solely on the current state, not the sequence of events that preceded it (Agchar et al., 2024). In our case, each musical note played by the pianist during the performance is treated as a state. During the performance, the input is constantly analysed to determine the probability of one note following another. If, for example, we look at the notes E G E B E G and say our current note is E because we go from E to G twice and from E to B only once, the probabilities for our next notes are 33 percent for B and 67 percent for G.

For the implementation into Ableton, the extension ml.lib (Smith and Garnett, 2012) was used; further information regarding the specific implementation can be found in the documentation of this extension. The Markov model is integrated into a broader musical context using an aleatoric rhythm generator with control parameters like note density and duration (see Fig. 3). The rhythm generator is implemented in Max using the transport object. That way, the notes generated by the AI-Instrument are always synchronized with any tempo changes in Ableton. By employing this configuration, the output sequence of the ml.markov object, which just consists of numbers, is transformed into a series of MIDI events. In addition to other control elements such as octave transposition and velocity control, this allows for subtle control of the AI-Instrument during performance.

Further, the Interface depicted in Fig. 3 is not only used to control the AI-based part of the performance but also as a means of primary visual feedback (see also: (Miranda and Wanderley, 2006)). This allows the performers controlling the Interface, a more profound understanding of their actions, especially during the denser parts of the piece.

6 NOTATION DESIGN AND COMPOSITIONAL DEVELOPMENT

The conceptual development of A Duet of Intelligences emerged within the broader context of the first author's doctoral research, which investigates the cognitive, performative, and interactional functions of graphic notation. In particular, this work was intended to explore the possibility of graphic notation to mediate between human and machine, and compositional form and improvisational performance outcomes.

Figure 4: Both hands play the sixteenth notes in bar 44.

Graphic notation was also employed to support the flexible integration of concrete musical ideas, AI, code, and performance within Ableton. In this setup, the tempo and rhythmic structure of the AI performance is determined, and modifiable in real time in Ableton, while the melodic and harmonic context is shaped by performer input without requiring changes to the code. In this sense, the whole compositional setting operates trough entanglement of its components, each part drawing on the strengths of the others and producing a whole that exceeds the sum of its parts.

The opening section models escalating demands of capitalism through progressively denser material (see Fig. 4), thereby establishing a scenario in which human performance reaches its limits. This sets up the transition to the AI material, which initially mirrors the pianist's output before exceeding it in density and complexity, foregrounding questions of agency, autonomy, and machine performance.

Figure 5: First page of the unconventional notation.

The AI is first represented by a prerecorded MIDI file that reproduces the pianist's material at a higher octave. As the texture develops, the AI becomes increasingly elaborated, musically differentiating itself from the human playing. The subsequent shift from conventional notation to graphic notation is not merely dramaturgical, but conceptually central to the dissertation's argument: it stages the move from prescriptive score-based control toward a more open model of performance in which notation becomes a medium of interaction rather than a fixed set of instructions.

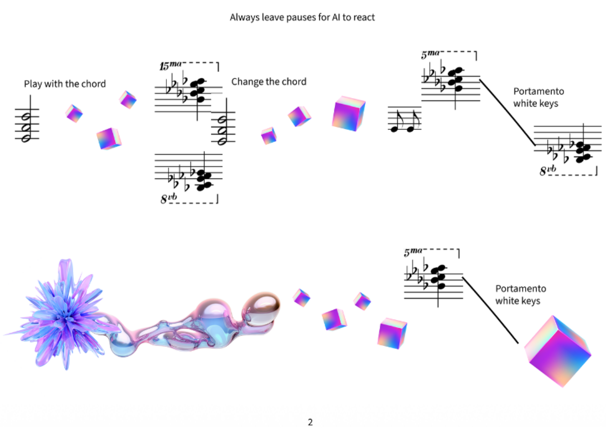

Figure 6: Second page of the unconventional notation.

Drawing on the distinction between conventional notation as a more binding system and graphic notation as a space of ambiguity and performative agency, the second part of the score was designed as a gradual transition from relatively fixed instructions to a fully graphic score. It initially consists of notes notated *senza misura* and text notation (see Fig. 5) but progressively incorporates graphic elements, including three-dimensional objects rendered in two-dimensional projection (see Fig. 6).

The second page contains a denser graphic notation (see Fig. 6), thereby expanding the degree of openness while maintaining a formal compositional trajectory. Although it draws on the visual language of earlier animated 3D scores, the notation remains paper-based, thereby avoiding the time limitations inherent in animated playback. The final page (see Fig. 7) culminates in an explosion-like form, which functions as a graphic articulation of conflict and aftermath, yet remains interpretatively open.

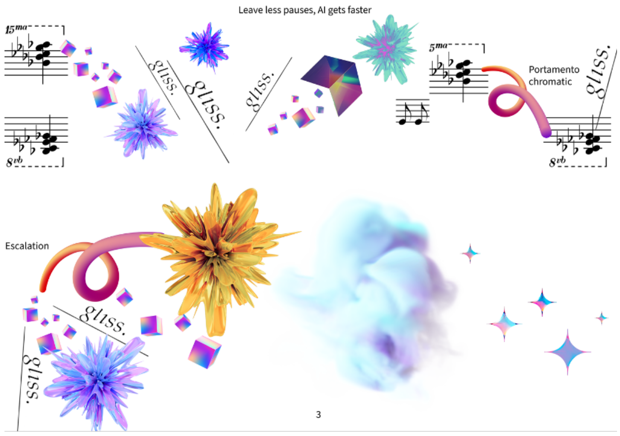

Figure 7: Third page of the unconventional notation.

7 AI AND NOTATION

Western notation traditionally conveys the artistic ideas of composers to performers. Different notation systems express these ideas through different visual languages, some standardized and generally accepted, others more experimental or abstract (Reutz Drobnić and Schumacher, 2025). Since the 1910s, the development of musical notation has reflected broader shifts in compositional thinking, aesthetic values, and priorities. With the gradual shift in the meaning of traditional tonality and formal structures, some performers and composers felt compelled to seek alternatives to conventional notation. Graphic notation became established in the mid-20th century as part of the effort to represent musical ideas beyond the limitations of traditional musical notation (Gutkin, 2015), p. 173. Instead of the unambiguous specification of pitches and durations, such scores often emphasize visual form, spatial arrangement, and degrees of ambiguity, thereby expanding the function and meaning of notation (Sauer, 2009). Notation no longer functions solely as a means of instruction, but increasingly as an independent site for aesthetic and conceptual experimentation.

The term *graphic notation* does not refer to a clearly defined, self-contained notational system, but rather to a historically developed and conceptually open field of diverse methods in which visual forms and spatial arrangements generate musical meaning. Its indeterminacy is not accidental, but central to its function. By foregoing precise instructions, composers relinquish some control and create a space in which performers become an active, interpretive presence, instead of only executing predetermined signs (Pace, 2007). In this sense, authorship is no longer singular. Graphic notation can be conceived as a pictorial form of musical notation: its visual elements vary from work to work, are freely organized, and are not necessarily bound to the conventions of strictly coded notation systems (Magnus, 2015). These notations must be interpreted by the performers as potential sound sources; in doing so, a portion of the interpretative authority is transferred to the performers, thus redefining their role (Magnus, 2021; Kanga and Antipolis, 2014).

This open-endedness is relevant for the project because it shifts notation toward an interpretable visual artefact capable of mediating multiple different creative agents. Rather than treating the score as a direct-to-sound input, the notation can be understood as a mediating structure that supports multiple valid readings while preserving the piece's structural affordances. One approach to the application of graphic notation and AI is, for example, using a large language model to interpret graphic notation. Exactly this was done in a project by first converting visual features of Cornelius Cardew's Treatise into descriptive text prompts, which were then used to condition a music generation model (Karchkhadze et al., 2024).

In contrast, our study does not rely on AI as the primary interpreter of the graphic score, although such an interactive system is in development. Instead, the human reading of the score serves as the interpretive layer, and that interpretation is then used as the input to the AI system. This preserves the performer's agency at the level of score interpretation while still allowing AI to operate on a musically meaningful representation derived from the human reading.

8 CO-CREATIVITY THROUGH DEEPER ENTANGLEMENT

In *A Duet of Intelligences*, the transition from conventional notation to graphic elements marks a shift from fixed instruction to interaction. The graphic score was implemented as a co-creative site of experimentation with no given time or material constraints. With the AI freely interacting with the performer and the score without a time constraint, the creative entanglement of all actors was achieved. Through the use of Disklavier, the AI became a present and physical performance partner, operating within the same performative environment as the pianist. This configuration shifts the interaction from a model of remote response to one of co-present, physical engagement.

Under these conditions, co-creativity extends beyond the exchange of symbolic information and can be understood as a process of mutual, embodied influence. The pianist's actions are captured as input data and immediately transformed by the system into corresponding mechanical and sonic events. These events, in turn, trough physical piano key control, constrain and inform the performer's subsequent actions, establishing a continuous sensorimotor coupling between human and system. The composer's ideas generated the graphic score, which inspires and triggers the pianist's output. The pianist's interpretations generate data that shape the AI's output; this output, in turn, informs the pianist's playing and the composer's real-time modifications of the AI. This circularity of influence makes authorship more fluid: the boundaries between pianist, AI, and composers dissolve into a shared field of artistic agency.

At the concert, this entanglement was experienced as a concrete, ongoing negotiation between pianist, AI system, and composers rather than as an abstract concept. The pianist approached the graphic score as a map of possibilities: she navigated clusters of shapes, densities, and trajectories across the page to decide where to enter, how long to remain in a particular region, and how intensively to articulate sound at that moment. Her reading of the graphic notation thus structured the temporal unfolding of the piece, shaping local patterns of gesture and global arcs of buildup and release.

The AI followed and extended the pianist's material, repeating pitches and contours, prolonging rhythmic cells, displacing her gestures in time, and occasionally opposing them with contrasting textures. In this way, the AI's behaviour remained coherent with the evolving musical situation, even when it introduced unexpected turns or disruptions.

From our position as composers, we intervened by adjusting the behaviour of the AI and its relationship to the graphic score in real time in Ableton. Subtle changes to the tempo, rhythm, velocity, and octave of the AI shifted the conceptual balance between alignment and friction. At times, we made sure the AI closely traced the pianist's performance; at other times, we redirected the material with rhythmic or pitch variations, as well as extreme tempo changes, to create tension and dramaturgy.

In these moments, leadership and followership were no longer stable categories; authorship was distributed across all three entities. The piece's identity emerged from the continuous feedback loop in which graphic notation, embodied performance, and algorithmic processing were mutually entangled, each reshaping the others from one musical moment to the next

9 PRACTICAL OBSERVATIONS

During the rehearsal process, the ensemble engaged in repeated cycles of verbal negotiation and procedural planning to establish a shared operational understanding of the score and its interpretative framework. These interactions functioned as a calibration phase in which performers articulated and aligned reading strategies and expectations, enabling a joint mapping from visual configurations to performative actions. Only once a consensus on this mapping had been achieved did coherent ensemble performance become reproducible across runs.

While the graphic score afforded substantial flexibility and creative openness, it also introduced systematic challenges for collective synchronization. The indeterminate, non-prescriptive character of the notation increased the degrees of freedom in local decision-making, which in turn amplified the risk of divergence between performers' interpretations. In response, the rehearsal methodology was adapted to include focused sessions with the pianist in which score literacy, role distribution, and timing strategies were explicitly negotiated.

Experiments with the AI-Device across varied performance configurations showed that its behaviour is strongly contingent on the stylistic background and interpretative approach of the pianist. Across the rehearsal phase, three pianists with distinct playing styles were involved: the concert pianist, Mirjam, whose performance is situated in contemporary and "new music" practice; a pianist with a more conventionally tonal, harmony-oriented approach; and a pianist whose playing was predominantly cluster-based. With Mirjam's more melodic and line-oriented interpretation of the graphic score, the AI-Device produced output that remained locally coherent and closely related to her material, often imitating or extending pitch sequences and contour in a recognisable way.

By contrast, performances involving dense chords and clusters tended to produce AI responses that were perceptually more fragmented and less structurally coherent. This can be attributed primarily to the random sequence of dense pitch organization of the input material rather than to system constraints. Clusters typically consist of chromatically adjacent pitches that are pressed almost at the same time, creating an unknown pitch sequence. As a result, the produced sequence of pitches can not be perceptually connected to the input.

In this configuration, the AI-Device behaves less like an augmentation of the piano and more like a distinct agent in a shared space, with its output conditioned by the pianist's actions, aligning with contemporary conceptions of musical human–machine interaction as negotiated rather than one-directional.

10 CONCLUSIONS

This paper indicated how a first-order Markov-chain system can be embedded in a compositional and performative framework that supports entangled, multi-agent human–AI co-creation. By combining human-created conventional and graphic notation, and a custom Ableton-based AI instrument, *A Duet of Intelligences* creates a flexible performance environment in which interpretative openness and algorithmic responsiveness reinforce one another.

The project also demonstrates that the artistic potential of AI in music does not depend on model complexity alone, but on the way technical systems are integrated into compositional design, rehearsal practice, and performance interaction. In this sense, the work contributes a case study in how notation, interface design, and live control can distribute agency across human and non-human agents, while still preserving a coherent musical form.

REFERENCES

Agchar, I., Baumann, I., Braun, F., Perez-Toro, P. A., Riedhammer, K., Trump, S., and Ullrich, M. (2024). A survey of music Generation in the context of interaction. *arXiv (Cornell University)*.

Ben-Tal, O. and Dolan, D. (2023). Musical and Meta-Musical Conversations. *AIMC 2023*.

Borg, J., Assayag, G., and Malt, M. (2022). Somax 2 a Reactive Multi-Agent Environment for Co-Improvisation. In *SMC 2022*.

François, A. R. J., Chew, E., and Thurmond, D. (2007). Visual feedback in performer-machine interaction for musical improvisation. *Proceedings of the 2007 Conference on New Interfaces for Musical Expression (NIME07)*, 07:277–280.

Gutkin, D. (2015). Notation Games: On Autonomy and Play in Avant-Garde Musical Scores. *Notation in Creative Processes*.

Kanga, Z. and Antipolis, S. (2014). "Not Music Yet": Graphic Notation as a Catalyst for Collaborative Metamorphosis. *Collaborative Metamorphosis', Eras*, 16(1).

Karchkhadze, T., Shao, K., and Dubnov, S. (2024). Interpreting Graphic Notation with MusicLDM: An AI Improvisation of Cornelius Cardew's Treatise. *2024 IEEE International Conference on Big Data*.

Lewis, G. E. (2000). Too Many Notes: Computers, Complexity and Culture in "Voyager". *Leonardo Music Journal*, 10.

Linskens, E. J. (2014). Music Improvisation using Markov Chains.

Magnus, D. (2015). *Aurale Latenzen*. Schwabe.

Magnus, D. (2021). *Aurale Latenz*. Kulturverlag Kadmos Berlin.

Miranda, E. R. and Wanderley, M. M. (2006). *New digital musical instruments*. A-R Editions, Inc.

Nikrang, A. and Kiesenhofer, S. (2025). Human–AI Co-Creation in Contemporary Composition: Interaction and Artistic Strategies with Ricercar. *Conference on Animation and Interactive Art*, pages 65–73.

Pace, I. (2007). Graphic Notation.

Reutz Drobnić, V. and Schumacher, M. (2025). Graphic notation and the Mind: Cognitive, Psychological and Physiological Mechanisms of Graphic Score Performance. *Emille*, 23:45–61.

Sauer, T. (2009). *Notations 21*. Mark Batty Publisher.

Shapiro, I. and Huber, M. (2021). Markov Chains for computer music generation. *Journal of Humanistic Mathematics*, 11(2):167–195.

Small, C. (1998). *Musicking*. Wesleyan University Press.

Smith, B. D. and Garnett, G. E. (2012). Unsupervised Play: Machine Learning Toolkit for Max. *Zenodo (CERN European Organization for Nuclear Research)*.

Tugarin, N. and Van Husen, C. (2025). Development and integration of human-AI interactions in service applications: Conceptual framework and review. *International Journal of Information Management Data Insights*, 5(2):100357.

Vear, C., Benford, S., Martinez Avila, J., and Moroz, S. (2023). Human-AI Musicking: A Framework for Designing AI for Music Co-creativity. *AIMC 2023*.

Vear, C. and Poltronieri, F. (2022). *The language of creative AI*. Springer Nature.

Zhou, E. and Lee, D. (2024). Generative artificial intelligence, human creativity, and art. *PNAS Nexus*, 3(3):pgae052.