



Universiteit
Leiden
The Netherlands

Performative transactions: worlding compositional ecosystems

Lukawski, A.

Citation

Lukawski, A. (2025, November 21). *Performative transactions: worlding compositional ecosystems*. Retrieved from <https://hdl.handle.net/1887/4283663>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4283663>

Note: To cite this publication please use the final published version (if applicable).

Performative Transactions: Worlding Compositional Ecosystems

Proefschrift

ter verkrijging van
de graad van doctor aan de Universiteit Leiden,
op gezag van rector magnificus prof.dr.ir. H. Bijl,
volgens besluit van het college voor promoties,
te verdedigen op vrijdag 21 november 2025
klokke 10:00 uur

Door
Adam Łukawski
geboren te Suwałki, Polen
in 1997

Promotor

Prof. dr. Richard Barrett

Copromotores

Dr. Paulo de Assis (Orpheus Institute, Ghent, Belgium)

Dr. Michiel Schuijjer (Conservatorium van Amsterdam, The Netherlands)

Promotiecommissie

Prof. dr. Rachel Beckles Willson (secretaris)

Prof. dr. Kitty Zijlmans

Prof. dr. Artemi-Maria Gioti (Mozarteum University Salzburg)

Dr. Martin Zeilinger (Abertay University, Dundee, Scotland)

Dit proefschrift is geschreven als een gedeeltelijke vervulling van de vereisten voor het doctoraatsprogramma docARTES. De overblijvende vereiste bestaat uit een demonstratie van de onderzoeksresultaten in de vorm van een artistieke presentatie. Het docARTES programma is georganiseerd door het Orpheus Instituut te Gent.

Acknowledgements	1
Introduction—Artistic Research with New Technologies	4
Artistic Research and the Creation of Concepts	4
Phase 1—Four Roles of Artists. Chapters 1-4	5
Phase 2—A New Compositional Ecosystem. Chapters 5-6	11
Artistic Research with New Technologies	17
Chapter 1. Artists as Operators	22
Navigating the Complexity of Musical Space	22
Part 1—On Automation	25
Two AI Lineages: Connectionism and Computationalism	25
Automation and Creative Work	28
What to Automate and What to Control	32
Part 2—On Methods	36
Teaching Creative Workflows with Artificial Intelligence	36
Part 2.1—Composing with AI Models Trained on Music	38
Generate, Select, Edit, Continue	38
Case Study: <i>Ani(mate)</i> (Movements 1–5)	42
Biases in Training Data and Artistic Control	44
Part 2.2—Composing with Large Language Models	46
Large Language Models as AI Agents for Musical Tasks	46
“Could you continue my MIDI, please?”	50
AI Agents Controlling Music Environments	52
Artists as Operators	55
Chapter 2: Artists as Curators	58
Part 1—Creativity and Problematisation	58
The Creative Act of Problematisation	58
Curating as a Creative Practice in Contemporary Art	60
Recombining the Creative	61
Part 2—Composition, Agency, and Assemblage	69
Musical Works as Assemblages	69
Posthumanist Agential Assemblages	73
Artists as Curators	78
Chapter 3. Artists as System-Builders	80
Encoding Process, Enforcing Trust	80
Part 1—Composing Autonomous Systems	84
From Chance to Probability	84

Automatism and Affordances of a Musical System	85
Artists Building Generative Systems	88
Composers as Programmers	91
Part 2—Blockchain for the Arts	94
What is Blockchain	94
Machines that Keep the Score	98
An Example: Hic Et Nunc	100
Generative Music NFTs	103
Long-Form Generative Art	104
Three Experimental Prototypes: <i>Canons</i> , <i>Tone Row</i> , <i>Subtraction</i> (2024)	106
More Long-Form Generative Music?	107
Artists as System-Builders	113
Chapter 4. Artists Worlding Ecosystems	114
Cybernetics and Systems Aesthetics	114
Complexity Theory and Second-Order Cybernetics	115
Individuation and Allagmatics of Artistic Research	117
Allagmatic Composition—Composing for Epistemic Complexity	120
New Kinds of Strange Agential Assemblages	123
Envisioning Abstract Machines	126
Worlding	130
Organising Free-Range Creativity	133
Worlding Compositional Ecosystems	137
Chapter 5. Autopoietic Rhizomatic Metamodelling Machine	141
Designing the Abstract Machine	141
Software Design Principles after Miller Puckette	143
Intelligent Cultural Commons	147
Blockchain as the Art's Medium for Composing Agencies	151
Rhizomatic Metamodelling	154
Social Complexity of the Autopoietic Rhizomatic Metamodelling Machine	162
Autopoietic Rhizomatic Metamodelling Machine	168
Chapter 6. Decentralised Creative Networks and Performative Transactions	172
Decentralised Creative Networks	172
A Compositional System as Public API	174
Composing with Transformations, Features, and Conditions	175
Performative Transactions	176
Composing and Testing Transactions	178
What Are We Modelling? Indexes, Structures, and Interpretation	180
How the Infrastructure Works	190
AI Agents and Recursive Infrastructure	192
Case Studies: <i>Ani(mate)</i> (Movement VI), <i>Allagma</i> , and <i>Chain of Thoughts</i>	196
Building on the Infrastructure: Interfaces, Applications, and Experimental Economies	200
From Composition to Ecosystem: Worlding through Performative Transactions	202

Conclusions	205
Summary	210
Samenvatting	211
Bibliography	212
Curriculum Vitae	224

Acknowledgements

Even systems that can shape themselves don't grow in isolation. This thesis, though concerned with recursive systems and self-modifying structures, had an emergence that was anything but autonomous. It has been shaped, supported, challenged, and cared for by many—human and non-human—whose contributions cannot be reduced to citations or code. What follows is a small gesture of thanks for the companionship, guidance, and moments of care that fully participated in its worlding.

I begin by thanking my promotor, Richard Barrett, whose guidance, attentiveness, and sharp critical insight have accompanied this project from its earliest stages. His ability to engage both rigorously and creatively with my ideas has been invaluable, and I am deeply grateful for the space he gave me to develop this work in my own voice. I wish to express my special thanks to Paulo de Assis, not only as co-promotor and my daily supervisor, but as indispensable intellectual companion throughout this trajectory. Our collaboration at the Orpheus Institute—within the MetamusicX research cluster—was more than a parallel context to this dissertation; it was an ongoing dialogue through which many of its core ideas took shape. Over the years, we developed a shared language for thinking and experimenting together—one that gave form to bold, adventurous, and deliberately unruly ideas that might not have found grounding elsewhere. Without Paulo's unwavering trust, critical acuity, and commitment to both Artistic Research and experimental thought, this dissertation would not be what it became. My heartfelt thanks also go to Michiel Schuijjer, whose support and feedback remained generous and steady throughout the years. From guiding my earlier studies at the Conservatorium van Amsterdam to becoming a co-promotor of this doctoral trajectory, Michiel has been a consistent presence and a thoughtful reader of both my music and ideas.

At Leiden University, I am grateful to the Academy of Creative and Performing Arts (ACPA). I thank Erik Viskil, ACPA's Academic Director, for his kind support throughout the process. I am thankful to Rosalien van der Poel, the Institute Manager, whose generosity, precision, and guidance were a constant source of reassurance. I would also like to acknowledge the Leiden University Centre for Digital Humanities, whose resources provided foundation to some of the digital aspects of this research. My heartfelt thanks go to the Orpheus Institute. I am particularly grateful to the Director Peter Dejans who created an environment where experimental artistic work

could flourish with confidence and ambition. I thank the docARTES coordinators Joost Vanmaele, and Luk Vaes for their generous guidance and responsiveness throughout the trajectory. I am grateful to my colleagues from the MetamusicX research cluster—Martin Zeilinger, Ilpo Jauhainen, Alicia Reyes, Diego Castro Magas, and Paolo Giudici—for their presence, feedback, and companionship throughout these years. I also wish to thank Kathleen Snyers, Heike Vermeire, and Kayo Quintens for their steady and thoughtful support at every step—their presence not only behind the scenes made a world of difference. I would like to thank Jonathan Impett, Einar Torfi Einarsson, Gabriel Paiuk, Tiziano Manca, Marek Poliks, and Roberto Alonso Trillo for enabling various platforms and exchanges that became crucial to the development of this research. I also wish to thank the Conservatorium van Amsterdam and the Iceland University of the Arts for offering the space and trust to develop and teach my experimental course. I am especially grateful to the talented students who participated—thank you for your courageous approaches, sharp observations, and fearless creative experimentation.

Thank you to Michał Skarżyński, whose technical expertise, patience, and creativity were instrumental in bringing the ideas of this dissertation into working systems. From countless conversations to the long hours of collaborative problem-solving, his ability to think across code and use-case was vital to translating my ideas into executable form. Working with Michał was a reminder that building systems can also mean building trust, clarity, and shared excitement for what art and technology can become. Thank you to my friends Ramin Amin-Tafreshi, Nuno Lobo, Alexandra Pouta, Alex Tay, Boris Bezemer and Wiktor Piechowski for your support.

I want to thank my parents Agnieszka and Tomasz, my sister Gabriela, my parents-in-law Jola and Kornel, and my sister-in-law Iga for their unwavering support throughout this long and often complex journey. Whether through practical help, emotional grounding, or simply the confidence that I would find my way, your presence was a vital foundation.

To Julia, my wife, my partner, and my constant companion—thank you for walking this path with me. For your love, your clarity, your strength. For holding the centre when I got lost in abstractions, and for reminding me why it all matters. For your sense of humour, your care, and your fierce intelligence. This work may bear my name, but its rhythm has always been ours.

Introduction—Artistic Research with New Technologies

Every artistic act is an encounter with a system—material or conceptual, historical or speculative, visible or latent. To compose does not mean to begin from nothing, but to shape relations: a repertoire, a medium, a toolset, a set of expectations. Yet composition is more than interaction. The artist intervenes in a system, reshapes it, and in some cases, builds it anew. This thesis emerges from such interventions. It asks what happens when artists begin composing not only sounds or scores, but the very systems through which composition becomes possible. When music is approached as an evolving ecosystem of relations—spanning human and non-human agencies, processes and technical infrastructures—then composition becomes a form of experimentation with the very systems that enable its becoming; a way of thinking with and through emerging conditions.

Artistic Research and the Creation of Concepts

Artistic Research is a young and still widely misunderstood field. It differs both from musicological analysis and from conservatoire-based artistic training, not simply in its methods, but in its fundamental relation to the object of research. Where musicology tends to analyse pre-existing works, and conservatoire training refines the performance of established repertoire, Artistic Research generates its own objects through the practice of making. As artistic researcher Paulo de Assis succinctly puts it: “Artistic research does not operate ‘after the fact’ but generates facts. It challenges reality and proposes novel entities and approaches” (de Assis, 2023, 54). Rather than beginning with a hypothesis to be tested, Artistic Research proceeds through what de Assis calls a logic of experimentation. Here, experimentation is not employed as a scientific technique for gathering evidence in support of a theory. It is, instead, a method that enables the emergence of spaces of problematisation—spaces in which thoughts, practices, and forms can be reshaped. “This search”, he writes, “is not primordially motivated by a quest for newness, or for unprecedented results, but first by a will to extend, enlarge, and, if possible, reconfigure the field of the visible, of the utterable, and of the audible” (de Assis, 2018, 21). Experimentation in Artistic Research is thus not aimed at discovering empirical truths, but at disclosing what becomes thinkable or perceivable through a particular practice. “If understood as a

practice”, de Assis notes, “experimentation discloses a powerful logic of bodies in action and of actions leading to new thoughts, senses, and sensations” (de Assis, 2018, 23). It invites the artist to change, to become something other than what they were—to “not be reducible to an imprisoning ‘I am’” (22). In this way, Artistic Research does not merely study the world; it actively composes new ontologies within it. Aligned with the view of the philosopher Gilles Deleuze (1925–1995)—a thinker who developed a radically processual and constructive approach to thought—that philosophy *is* the creation of concepts, Artistic Research similarly operates *through* acts of conceptual invention. Concepts, as de Assis reminds us, are not discovered or retrieved from disciplinary archives, but constructed through encounters—often between a practice and something exterior to it. This encounter may be with a technological material, an institutional structure, or an epistemic gap. In every case, the researcher steps outside what is already known, creating concepts that operate as events in thought: ruptures, reorientations, openings. “To have a thought”, writes Deleuze, “is to go outside oneself, outside a particular discipline, outside a given system of coordinates” (quoted in de Assis, 2018, 15). Artistic Research stages these encounters as part of its method.

This repositioning of the research object—from something pre-given to something produced—demands a shift in how research is framed. It becomes more productive, in this context, not to ask what is the object of study, but rather what is the artistic researcher doing. The overarching aim of this thesis is to—through experimental and practical encounters—investigate how emerging technologies such as artificial intelligence and blockchain transform compositional practice—not only by enabling new artistic outputs, but by reconceptualising the very practices through which composition is conceived, enacted, and shared. To address this, the research unfolds in two phases: the first identifies, theorises and categorises shifting artistic roles in response to these technologies, while the second designs and implements a new compositional infrastructure that operationalises those insights.

Phase 1—Four Roles of Artists. Chapters 1-4

Phase 1 of this doctoral trajectory, the results of which are described in Chapters 1 to 4, iteratively applied methodologies such as compositional experimentation and

prototyping, codification of new compositional methods through the development of an experimental educational course, and philosophical reflection to identify how the roles of artists and their compositional practices transform in response to emerging technologies with a particular focus on artificial intelligence and blockchain. Through this methodological combination, Phase 1 of this doctoral project identified and described four redefined roles that artists may inhabit in response to quickly changing technological conditions: artists as operators of tools, curators of agential assemblages, system-builders, and initiators of artistic ecosystems. Together, these four roles reveal a trajectory from interacting with technological tools and configuring creative systems, to ultimately designing the ontogenetic conditions under which composition itself unfolds. The described roles are not mutually exclusive nor prescriptive. They reflect overlapping modes of engagement—each responding to a different configuration of technology, authorship, and collaboration. They serve as both diagnostic categories and speculative positions resulting from artistic research inquiry, offering a framework for understanding how artistic practices shift when operating within, against, or alongside complex technological systems.

Chapter 1 focuses on the *artist as operator*, a role that becomes particularly salient in the context of contemporary AI-based composition. The chapter situates creative work within complex, trained, high-dimensional spaces that no longer require exhaustive rule design. It explores how AI systems, particularly deep learning models, compress musical information into latent spaces that the artist must learn to navigate. In this environment, the composer acts not as a passive user of tools but as an operator—someone who strategically engages with systems that are partially opaque, probabilistic, and autonomous. Through detailed technical and theoretical discussion, the chapter examines how AI can be used in compositional workflows, requiring composers to develop new forms of judgment of the musical material, prompt engineering, and contextual framing. In the chapter, I propose a method of engaging with such models. An iterative cycle called “generate, select, edit, and continue” becomes a core compositional strategy, enabling a co-creative dynamic between human and machine. Further, the chapter expands on the proposed method in response to several possible artistic and technological conditions, such as working with text-based agential AI models instead of less-controllable format-specific deep learning architectures, and bringing their capacities outside of the fixed chat-based environments. While automation poses real challenges to artistic labour, the chapter

argues that it also opens new operational terrains in which human creativity can find its place. The operator function thus describes an expanded form of authorship grounded in responsiveness, criticality, and the capacity to engage with complex systems without relinquishing aesthetic agency.

Chapter 2 extends this analysis by presenting the *artist as a curator of assemblages*, a role that reframes the notion of creativity as the expression of an individual genius into creativity as strategical arrangement of agential relations. Through examples ranging from historical algorithmic music to contemporary AI installations, the chapter interrogates dominant assumptions about creativity, authorship, and value. It challenges both the romantic ideal of the solitary creator and the reductive metrics of optimisation-based AI research. Instead, it proposes a combinatorial, distributed model of creativity, in which human and non-human agencies interact within dynamic systems. Drawing on Paulo de Assis's theory of musical works as assemblages and Martin Zeilinger's concept of posthumanist agential assemblage, the chapter further positions artistic practice as an act of problematisation—the creation of open, unresolved situations that provoke new ways of thinking. In this view, the artist becomes a curator of conditions—selecting, combining, and staging elements that generate meaning through their agential interrelation. This curatorial mode is not limited to exhibitions, but defines a broader artistic logic of posthuman creative ecosystems.

Chapter 3 deepens the investigation with the notion of *the artist as a system-builder*—who more than curating external agencies and their aesthetic configurations, builds the very processes, infrastructures, and conditions through which art is generated, circulated, and experienced—artist who designs the operations which become autonomous and gain agency over the creative process. If the operator and curator roles described earlier emphasise navigation and arrangement within existing autonomous and semi-autonomous systems, the artist as a system-builder actively constructs such systems from the ground up. This chapter examines how contemporary artists encode procedures, formalise generativity, and embed compositional agency into executable systems—ranging from stochastic scores to blockchain-based smart contracts. The shift it traces is not only from composer to programmer, but also from system as inscription to system as infrastructural design. Through historical and contemporary examples—from John Cage's *I Ching*-inspired

chance operations to Iannis Xenakis's stochastic formalisms, from Brian Eno's generative music systems to current smart contract-based music Non-Fungible Tokens (NFTs)—the chapter reveals a long arc of artistic inquiry into rule-based creation. At stake is a redefinition of authorship: not as the production of fixed outcomes, but as the composition of dynamic, adaptive environments that sustain variation and procedural autonomy. Drawing on the work of Gilbert Simondon, Philip Galanter, and Artemi-Maria Gioti, the chapter reframes the “system” not as a tool, but as the artwork itself—a generative machine whose identity lies in its structure, behaviour, and affordances, rather than any singular output. The role of the composer is thus extended to that of a metadesigner: someone who crafts spaces of possibility, configures topologies of transformation, and governs the emergence of artistic form through encoded logic.

The chapter further investigates a crucial distinction between reliability and trust as it applies to systems involving decentralisation, human agency, and algorithmic opacity. Whereas traditional machines aim for reliable execution, many contemporary systems—particularly those involving human interaction or AI agents—require frameworks for negotiating unpredictable behaviour. It is here that blockchain technology is presented as a critical infrastructure. As an apparatus that enforces trust through code—via consensus mechanisms, smart contracts, and public verifiability—blockchain technology offers artists new capacities to design systems of participation, authorship, and procedural governance. In the introduction to our co-edited volume *Decentralized Music* (2024), Paulo de Assis and I argued that “it is imperative to investigate the consequences of blockchain for the arts, focusing on the creative rather than on the financial and political opportunities of this technology” (de Assis & Łukawski, 2024, 3). We also emphasised the urgency of this task, noting—following Amy Whitaker—that “ignoring blockchain within the field of the arts empowers actors outside the field to act without the field's participation” (Whitaker, 2019, 23; cited in de Assis & Łukawski, 2024, 4). Chapter 3 responds to that call by turning to the under-explored field of generative music NFTs, highlighting a range of exploratory artistic experiments conducted within the framework of this doctoral research that treat blockchain as a generative medium for music composition. These projects shift the focus from final works to executable protocols, from singular authorship to distributed agency, from static audio files to blockchain-based algorithmic generation. This chapter identifies that while generative art NFTs have

already proliferated in the visual arts, their musical counterparts remain emergent—often hindered by infrastructural, cultural, and conceptual gaps. Furthermore, it not only identifies those gaps, but demonstrates how they might be overcome: through new artistic formats, and a rethinking of what counts as a musical work in a digital, web-based culture. Ultimately, Chapter 3 argues that artists as system-builders are no longer only composing music, but composing the very systems in which music can exist. Their work situates composition at the intersection of code, protocol, and procedural logic—where aesthetic, technical, and social systems converge.

Chapter 4—*Artists Worlding Ecosystems*—advances the logic of artists as builders of generative systems—those who encode process, and trust into machinic infrastructures—by introducing a new compositional horizon: worlding. Here, artistic practice is not anymore dependent on the emerging conditions of new technologies, but engages in composing the very conditions under which compositional systems can emerge, evolve, and sustain themselves. Drawing on cybernetics, complexity theory, and Simondon’s philosophy of individuation, the chapter frames artistic systems as dynamic, self-organising ecologies—assemblages capable of recursive transformation and epistemic growth.

The chapter opens by tracing the legacy of cybernetics in the arts. Building on these historical foundations, it introduces the distinction between systemic complexity—the structural intricacy of a system—and epistemic complexity—its capacity to generate and transform knowledge. Through the lens of Simondon’s allagmatics and Paulo de Assis’s notion of epistemically complex artefacts, the chapter argues for a shift in compositional thinking toward designing systems that enact ongoing processes of individuation. This is the conceptual ground on which the thesis introduces a new concept of *allagmatic composition*—a practice that treats compositional systems as metastable fields of operational potential. Unlike generative composition in which the system’s affordances produce its outputs, or interactive composition in which the outputs of the system and its affordances can be also conditioned by external factors, allagmatic composition foregrounds the capacity of a system to iteratively modify its own operations. To engage in an allagmatic composition is then to compose on the level that gives shape to the generative rules themselves—it is to compose epistemic processes that are capable of generating new knowledge. As the chapter explains, each act of execution becomes a transductive event—an individuation that alters the

system's future states and compositional space. Further, Chapter 4 situates these ideas in a broader infrastructural context, showing how emerging technologies such as AI and blockchain already enable the construction of agential assemblages—artworks and systems in which authorship, intention, and agency are distributed across human and non-human actants. Drawing on Martin Zeilinger's concept of "strange agential assemblages" and artistic examples, the chapter explores how procedural systems can operate autonomously, propagate themselves, and reconfigure their own conditions of operation over time.

In the final sections, the chapter introduces the concept of the abstract machine—a diagrammatic, non-empirical logic of composition derived from Felix Guattari's *Chaosmosis*—and shows how this leads to the practice of compositional worlding. Here, artistic creation becomes the enactment of a generative ecosystem. Worlding names a compositional logic wherein the artist configures a field of potential—technical, social, aesthetic—within which others may act, contribute, or transform. The term is used in the context of its philosophical and phenomenological implications following from the works of philosophers Martin Heidegger and Donna Haraway. By incorporating artistic examples ranging from Refik Anadol's data-driven installations to teamLab's *Borderless* exhibition, and extending to Olga Goriunova's theory of art platforms as autopoietic systems, the chapter presents worlding as an emergent, infrastructural strategy of composition. It shows how artists no longer simply make works—they configure and enact ecosystems and diagram abstract machines that world their own conditions of emergence through their technical operationalisation. Composition, in this light, becomes an operational process of systemic individuation—a step-by-step, ever-evolving becoming of an entity as it takes shape by traversing a metastable field of operational processes. Philosopher Gilbert Simondon described individuation as a continuously unfolding metastable network of relations. In that perspective, the musical work, its technical substrates, performers, audience, and environment *co-individuate* through a newly enacted worlding. They co-constitute their roles within the individuation through ongoing interactions, each becoming determinate only in and through the evolving network that sustains them (Simondon 2020 [1958]). The subtitle of this doctoral thesis *Worlding Compositional Ecosystems* refers then to the artistic practice proposed in Chapter 4 as consisting of two steps: 1) iterative envisioning of an abstract machine as a set of rules and conditions, with its operations and relations of various actants, objects, and processes,

as an abstract diagram of principles guiding the world to be enacted, and 2) operationalising such an abstract machine in an actual technical system, a functioning machine with precisely defined set of executable operations, sustaining the existence of the newly-created ecosystem.

Together, Chapters 1 to 4 should not be read as a study about new technologies in music. Instead, they should be seen as a result of a direct artistic engagement with those technologies—through composing, programming, prototyping, collaborating, teaching, performing, and reflecting. It is, in this sense, not a retrospective commentary on artistic practice, but itself a compositional act, unfolding a sequence of actions that generate new conceptual and technical forms leading to the formulation of the artistic practice of *worlding compositional ecosystems*. This leads to a culminating insight of Phase 1 of this doctoral thesis—that the identified changing landscape of artistic practice that currently manifests itself in a yet rather limited and slowly emerging both experiential and technological change, to be operationalised and scaled to a fully functional artistic practice it requires not just new tools or formats, but a new kind of compositional ecosystem—one capable of sustaining distributed, recursive, and programmable creative agency across human and non-human contributors. In other words, to enable the formulated compositional practice of worlding artistic ecosystems, a new type of artistico-technological solution is needed—a technical system that enables the practice of ecosystemic worlding.

Phase 2—A New Compositional Ecosystem. Chapters 5-6

In response to this need, this doctoral project's Phase 2 employed a different set of methods, such as creative prototyping, artistic self-reflection, speculative design, and iterative technical development, to envision and construct such a system. In a sense, in Phase 2 the research project itself engaged in the previously described practice of worlding a compositional ecosystem. Chapter 5 outlines this ecosystem as a new kind of abstract machine—detailing its roles, processes, relationships, assumptions, and requirements. Chapter 6 then operationalises this abstract diagram of relations, translating the abstract machine into a concrete technical system that supports worlding compositional ecosystems on the technical level. Together, these chapters move the discussion from analysis to implementation, turning the thesis's earlier

theoretical and artistic insights into an infrastructural proposal.

Chapter 5 invents an *Autopoietic Rhizomatic Metamodelling Machine* (ARMM): an abstract machine, the operations of which describe a diagram of conditions that are crucial for operationalising a technical system that allows its users to engage in the practice of worlding compositional ecosystems. The ARMM synthesises the research's core insights into a set of rules that describe a new compositional logic—one that foregrounds recursion, self-modification, and emergent structure as the defining traits of contemporary artistic ecosystems, and can host all of the the previously identified artistic roles at once. The main compositional operation described as a rule in the diagram of the ARMM is *metamodelling*: instead of modelling sounds, scores, or other musical objects directly, ARMM specifies how generative, interactive, and self-modifying abstract procedures can be built, validated, and recombined in separation from the metamodelled objects. Referring to the fundamental phases involved in the production of a musical work as described by composer Iannis Xenakis and to the notion of Immanuel Kant's schema, the chapter proposes the process of *rhizomatic metamodelling* as a way of modelling pure relations between objects, rather than objects themselves. Drawing practical lessons from Miller Puckette's Max and Pure Data algorithmic environments and recent platform studies, the chapter distils five design imperatives—minimal modularity, cultural neutrality, distributed complexity, procedural transparency, and long-term intelligibility. Together they define rules for building an infrastructure in which every contribution (human or machinic) enters the system as a reusable operation that can itself evolve as it can be metamodelled by other operations. Composition thus becomes an ongoing process of *individuation*: human and non-human actants configure rhizomatic fields of potential rather than deliver fixed artefacts, and the machine continuously folds new operations back into its own topology. The chapter then shows how such a framework could be capable of supporting a posthumanist cultural commons. By treating operations, data, and economic rules as first-class citizens in a shared ledger—rather than as private assets—the ARMM sets the stage for collectively intelligent ecosystems where authorship, value, and agency circulate across human and non-human contributors. It is a proposal for composing with agencies directly: writing the rules of collaboration, redistribution, and transformation into the infrastructure itself. The Autopoietic Rhizomatic Metamodelling Machine is proposed as a blueprint of principles—how to build a platform for composing with and across the space of possible compositional

systems. It is a set of rules for constructing a framework that enables configuring creative relations rather than producing finished works. It enables artists, agents, and systems to collaborate by contributing operations that are modular (can be reused), interoperable (can work with other components), and procedural (can be executed dynamically over time).

Chapter 6—*Decentralised Creative Networks and Performative Transactions*—operationalises ARMM into a concrete technical system, introducing Decentralised Creative Networks (DCNs) as blockchain-based compositional infrastructures that enable modular artistic processes to be encoded, shared, and executed as smart contracts. The chapter begins by contextualising DCNs within the historical trajectory of Creative Networks as theorised by Rasa Smite in relation to early Internet-based collaborative artistic practices of the 1990s, but reorients this collaborative spirit through the specific affordances of blockchain technology—particularly its capacity to encode modularity, composability, traceability, and programmability at the operational level. The chapter identifies and addresses a technical challenge in generative blockchain-based art: the absence of a general standard for cross-application composability of compositional processes encoded on the blockchain. While existing generative art NFTs typically operate within closed application contexts, mirroring the architectures of centralised software, this thesis proposes to use blockchain technology not only as a medium used by artists to present their works (both as stand-alone NFTs and generative NFTs), but as a medium in which the compositional processes themselves can be encoded as modular, executable operations that can be used by human and non-human actants to transparently build on top of each other's contributions. Decentralised Creative Networks thus shift from encoding solely static outputs and generative processes on the blockchain, to encoding the relations between artistic operations co-created by various participants in the network. This represents a fundamental reorientation from artwork to a chain of executable procedures with many authors, creating an environment for executable provenance that maintains intricate links between artistic assets and their origins while supporting transparent collaboration and shared ownership.

A crucial architectural innovation presented in this context is the conception of the compositional system as a public Application Programming Interface (API). This design ensures accessibility and extensibility, allowing anyone to build upon existing

contributions without requiring central authority permission, while establishing clear interfaces for interaction by both humans and AI agents. Most significantly, this API structure enables recursive compositional logic—components of the system can call and modify the system itself, allowing transformations to interact with the API from within their own execution. This recursive structure introduces the possibility of a self-reconfiguring system where human and non-human contributors collectively build and modify the very compositional environment they work within, making the system a compositional infrastructure capable of composing itself and thus, enabling the previously introduced mode of allagmatic composition.

The chapter provides detailed technical exposition of the system's core architecture, built around three fundamental building blocks: transformation (programmable function that takes an index number and returns a modified index), feature (composable object that defines how transformations unfold over time or space as ordered lists of indexes over other references features), and condition (rule that defines whether a feature and its referenced subfeatures can be executed). These elements form the grammar of the network, enabling the description, execution, and combination of operations that can produce artistic outputs or can be further recombined as part of more complex compositional processes. The power of this architecture lies in its composability: every transformation, feature, and condition is recorded on the blockchain with a unique identifier, enabling users to build artistic features using existing components while others can build upon their contributions, with the system automatically tracking dependencies between all elements.

Central to the system's operation is the titular conceptual invention of this doctoral dissertation—Performative Transaction (PT). Performative Transactions represent moments when compositional processes encoded on the blockchain as transformations, features and conditions are activated. The term "performative" indicates that these transactions perform processes by executing compositional logic and producing outcomes, rather than merely retrieving data. Each PT consists of a combination of the three introduced basic components, parameters that define the scope of the output, and registered dependencies such as other PTs. What distinguishes PTs from generic function calls is their referential capacity—when referencing another user's component, the system executes that original component live with its own logic, authorship, and constraints, ensuring traceability, modularity,

and interoperability across the network. Performative Transactions thus constitute a core contribution of this research, as they provide the technical foundation for realising each compositional model described in this dissertation. From generative music and interactive composition to allagmatic composition and the worlding of compositional ecosystems, PTs serve as basic compositional components enabling the implementation of the projects and technologies examined herein.

The chapter explains the crucial distinction between on-chain execution and development environments. While PTs are ultimately executed on the blockchain, the composition, testing, and refinement process occurs through a dedicated server that functions as both interface and development environment. This server acts as an intermediary, exposing the public API and simulating transaction execution without fees or blockchain state modification. This layered architecture enables costless experimentation, immediate feedback, accessibility for non-programmers, and programmable interaction by both humans and AI agents. The envisioned server implementation is open-source and can be self-hosted, providing users with full autonomy over their network interaction while supporting agent-based co-creation where human composers, algorithmic systems, and AI agents all participate in building and expanding the network. Importantly: this distinction separates the creative freedom to experiment with the system from the often limiting economical constraints of blockchains. These only come into play only when necessary—when the created Performative Transaction is to become a shared contribution within a decentralised system. Thus, while the dedicated open-source server implementation can imitate blockchain’s features for limitless local artistic experimentation, it also enables user’s participation in the on-chain network when the created component is ready to be shared.

A significant portion of the chapter is devoted to explaining what the system models: not music, images, or performances directly, but indexes—integer numbers that acquire meaning only through interpretation (given to these indexes only in the contexts of their specific applications, built on top of the Decentralised Creative Network). Through detailed examples involving simple features like Pitch, Duration, MajorScalePattern, RhythmPattern, and Melody, the chapter demonstrates how transformations modify indexes rather than the values they represent, enabling the construction of highly structured generative processes using minimal operations. The

chapter describes that what the infrastructure models is not the artwork itself but the logic that generates constituent parts, allowing users to define how values relate within and across dimensions and how those relations compose into complex operations. From the system's perspective, features composed by various human and non-human agents define operations on operations, constructing graphs of dependencies. These graphs assemble dynamically at moments of their execution, based on referenced features and transformations, making every execution both a computation and a traversal of user-defined compositional topology.

The chapter emphasises the system's generality: because outputs gain meaning only through external interpretation (in the contexts of: software applications that can be built on top of the Decentralised Creative Network, AI agents using the API to assist human users with accessing the system, humans using the outputs of executed PTs directly, etc.), the same infrastructure could support diverse applications across music, choreography, image generation, text sequencing, algorithmic theatre, assets control in game-engines and in non-artistic domains. The indexed logic of the proposed system allows it to remain agnostic to domain-specific semantics while supporting rich compositional work, maximising flexibility, composability, and reusability without prescribing specific aesthetics or media. The chapter concludes by noting how this orientation toward operational logic rather than fixed content invites reconsideration of what constitutes a musical "work". Decentralised Creative Networks are positioned as operating in a conceptual space where the operations encoded as Performative Transactions define the generation logic rather than its realisation—it is the space of operations and their links that persist, not any particular output of the system or its interpretation.

Together, Chapters 5 and 6 chart the transition from concept to implementation, from philosophical proposition to compositional infrastructure. They culminate the thesis's central inquiry: how might artistic practice evolve when composition becomes recursive, modular, and collectively authored—not through representational inscription, but through infrastructural worlding? By the development of Decentralised Creative Networks, this research both responds to the contemporary challenges of AI and blockchain, as well as proposes a method for inhabiting them creatively, speculatively, and systemically.

Several musical works composed alongside this research will be performed in a concert accompanying the defence of this doctoral thesis. These compositions will reflect how the theoretical and infrastructural developments explored throughout the thesis informed concrete artistic decisions. Each piece will offer a situated example of how the concepts of allagmatic composition, Decentralised Creative Networks, Performative Transactions, and worlding were translated into musical forms, performances, and compositional systems.

Artistic Research with New Technologies

The stakes of engaging with AI and blockchain today are not solely technical. As composer and performer Jennifer Walshe—known for her work at the intersection of experimental music and critical engagement with emerging technologies—has argued, these systems demand an ontological and existential reckoning. “I am convinced”, she writes, “that not only the development of the music, but life in the twenty first century will be primarily marked by how we engage with, respond to and think about AI” (Walshe, 2023, 499). In the musical domain, the consequences are especially acute. AI is already capable of producing stylistically coherent compositions in many genres, and Walshe predicts that within a few decades, “machines will be able to write music, in many genres, which is indistinguishable from that written by humans” (499). While she affirms that humans will continue to make music—and that not all music will be made with machines—the deeper question becomes: what does it mean to make music when machines can too? It is a lived, affective condition that many composers now inhabit. “I’m a composer who is living in the twenty first century and trying to think it through”, Walshe writes. “I’m both sublimely excited and blackly horrified about what is coming. I’m trying to give you a sense of how I view the world, and where I think things are going, because that psychological space is where my art comes from” (Walshe, 2023, 500). As Walshe reminds us, “we are all involved, we are all enmeshed, we are all implicated in the development of AI, regardless of whether we code or not, regardless of whether we ever make a piece of music using AI” (2023, 499). Artistic engagement with these technologies is therefore not a matter of following trends or demonstrating technical fluency. It is a necessary negotiation with the systems that increasingly shape our tools, our institutions, and our compositional imagination.

Yet artificial intelligence is not one thing. As researcher, curator, and theorist Martin Zeilinger—whose work explores how emerging technologies like AI and blockchain reshape artistic agency, ownership, and authorship—notes, AI is “as much a science as it is a technology, as much a cultural phenomenon as it is a philosophical construct”—and perhaps also a “mythical fabulation” that does not exist at all (Zeilinger, 2021, 36). For artists, this ambiguity opens both critical and creative potentials. Zeilinger distinguishes between strategic and tactical engagements with AI. Strategic uses, aligned with institutional and commercial logics, aim to stabilise ownership, reproduce anthropocentric norms, and instrumentalise AI toward fixed outputs. Tactical uses, by contrast, operate along “open-ended vectors of resistance”, probing and disturbing the structural logics of technological systems (Zeilinger, 2021, 51). In this view, AI becomes a site for contesting authorship, for rethinking property, and for experimenting with new distributions of agency. This doctoral dissertation aligns with the tactical strand of AI-based artistic practice. It does not seek to reproduce or simulate human composition, nor to compete with AI in genres where automation already dominates. Instead, it engages these technologies tactically: to expose latent assumptions, to construct new compositional procedures, and to imagine artistic systems that are not predicated on the individual artist-genius model, but on distributed, posthumanist assemblages of creation. In doing so, it joins a growing body of work that responds not only to the technical affordances of AI, but also to its ideological and institutional conditions. As Zeilinger argues, the most generative forms of AI art today mobilise the medium’s emergent capabilities “for interrogating, exposing, problematising, and challenging the aesthetic, ideological, or technological frameworks driving the commodification and propertisation of creative expression” (Zeilinger, 2021, 28). These are not speculative futures—they are already underway. The challenge for Artistic Research is not merely to observe these shifts, but to inhabit and recompose them.

To navigate such a shifting terrain demands more than technical fluency—it calls for a fundamental reorientation of artistic meaning and method. As Ethan Mollick, a scholar of innovation and entrepreneurship, and a prominent commentator on generative AI, argues, “we are already living in the early days of the AI Age, and we need to make some very important decisions about what that actually means” (Mollick, 2024, Part 1, Chap. 2). These decisions are not only infrastructural or regulatory, but deeply cultural. “We are going to need to reconstruct meaning”,

Mollick writes, “in art and in the rituals of creative work. This is not an easy process, but we have done it before, many times” (2024, Part 2, Chap. 5). In moments of historical transformation, artists have consistently responded by reconfiguring the boundaries of their medium: when photography rendered oil portraiture obsolete, painting turned toward abstraction; when recorded music restructured economic models, musicians reshaped their practices around performance, liveness, and experience. Today, the rise of generative AI and blockchain technologies represents a similarly deep transformation. As Martin Clancy—musician, producer, and chair of the IEEE Global AI Ethics Arts Committee—notes, these tools present musicians not only with new creative instruments, but with an opportunity “to shape and transform existing modes of creative expression”, raising renewed questions about “the role of the human artist and related economic and philosophical issues” (Clancy, 2022, 1).

Within this landscape, Artistic Research offers a vital and under-explored methodology. It situates artistic practices as modes of inquiry that generate both technical and conceptual outputs. A key reference in this regard is *Future Art Ecosystems*, an ongoing strategic briefing series published by the Serpentine Galleries’ Arts Technologies team. Initiated in 2020, the project provides conceptual and infrastructural frameworks for supporting art practices operating with advanced technologies such as AI, blockchain, and immersive media. As *Future Art Ecosystems* points out in their first volume, artists working in technological contexts are “always implicitly engaged in technological innovation”, even if their motivations differ from industrial ones (Serpentine Arts Technologies, 2020, 10). Importantly, the report recognises the resonance between these practices and the frameworks of “practice as research”, in which artistic work is situated “as a generator of research outputs” (68). Artistic Research thereby offers a route for artists to meaningfully participate in—and reshape—the epistemic, aesthetic, and infrastructural dimensions of emerging technologies. As observed in the briefing, the new practices appearing in this context tend to exhibit five consistent features: they embrace dynamic materials; engage in building distributed networks; construct new narratives; succeed in adjacent fields; and increasingly perceive the “art world as a medium” in its own right (Serpentine Arts Technologies, 17). This doctoral thesis participates in all five of these tendencies. It embraces AI and blockchain not as neutral tools, but as dynamic, volatile materials. It develops new networks—not only of artists and institutions, but also of agents, contracts, and protocols. It constructs alternative narratives about the role of

composition in a decentralised, posthuman context. It succeeds in adjacent fields by building experimental software infrastructures. And it ultimately reimagines the art world as a site of infrastructural invention and critical world-making.

As artist and researcher Sofian Audry—whose work explores machine learning, artificial life, and non-human agency in art—observes, in a cultural landscape increasingly shaped by machine learning systems and platform economies, artists and institutions alike may need to “reconsider their role and that of machines” (Audry, 2021, 164). What does it mean to be an artist when the means of creation are increasingly automated, when the frameworks of authorship are destabilised, and when the infrastructures of distribution are governed by algorithms and corporate platforms? Rather than framing this shift in terms of loss or threat, this research sought to understand it as an opening—an invitation to re-articulate artistic agency across multiple sites and systems.

In concluding this introduction, it is worth returning to a question that underlies the entire inquiry: What does it mean to compose in an era where the compositional act is no longer exclusively human? As AI-generated art proliferates and algorithmic processes become inseparable from aesthetic production, many artists feel the ground of their practice shift beneath them. But perhaps this shift is not a rupture, but a reminder. As artist and theorist Keith Tilford—whose art and research projects focus on the intersection of comics and artistic modernism in relation to diagrammatics, worldbuilding, abstraction, and technics—notes, “there is no outside to the constructibility of art’s model, because art is already an artificial mode of cognition” (Tilford, 2023, 135). If art is a system of mediation, then its entanglement with technological abstraction is not a threat to its authenticity but a condition of its emergence. Art was never outside of the machine. Sasha Stiles—a poet and language artist frequently exploring artistic uses of AI and blockchain technologies—puts it more directly: “What’s more human than technology?” (Stiles, 2023, 359–360). Language is a technology. Music is a technology. The score, the synthesizer, the sequencer—these are not accessories to artistic thought; they are its material expression. In this light, the use of AI or blockchain in composition is not a break from tradition, but a continuation of art’s technicity by other means. These systems extend the artist’s cognitive apparatus, enabling new forms of perception, organisation, and relationality. And yet, as this thesis demonstrates, the task is not simply to adopt new

technologies, but to reimagine the conditions under which artistic meaning is constructed and shared. That task is necessarily speculative. It must proceed not only from technical understanding, but from conceptual invention, infrastructural experimentation, and poetic risk. At this point it is important to acknowledge that working with AI and blockchain involves infrastructures often shaped by energy-intensive computation and corporate control. These broader conditions were kept in view during the artistic and educational experiments underpinning this research, where students were encouraged to approach such tools critically and to consider their environmental and economic costs, and low-cost, resource-efficient models were used where possible to support this critical awareness. The systems we build today may appear partial or strange, but—as Amara’s Law reminds us—we tend to overestimate the impact of new technologies in the short term and underestimate their long-term effects (Mollick, 2024, Part 2, chap. 6.). What seems awkward or marginal now may define the creative landscapes of tomorrow. This thesis is one such speculative gesture. It does not offer conclusions, but compositional propositions. It asks what kinds of music—and what kinds of artistic agency—become possible when composition is treated as an infrastructural, recursive, and distributed process. It argues that we are no longer composing only sounds or scores, but the systems through which composition itself can unfold. The reader is invited to enter this evolving ecosystem not as a passive observer, but as a participant in its worlding.

Chapter 1. Artists as Operators

Navigating the Complexity of Musical Space

Although what can be considered a musical work comprises countless material and immaterial parts—scores, manuscripts, performances, recordings, economic value, aesthetic judgments, historical contexts, symbolic aura, and more (de Assis & Łukawski 2024, 1)—even if we focus only on sound, the imaginary space of all possible music would still constitute an enormous space of possibilities. As observed by Martin Rohrmeier—researcher and director of the Digital and Cognitive Musicology Lab at the École Polytechnique Fédéral de Lausanne (EPFL)—even the results of well-defined rule systems such as counterpoint or voice-leading already “span an enormous search space, which is hard to traverse and affords for rare, original solutions and strategies to be found” (Rohrmeier 2022, 53). A quick thought-experiment shows the scale. Write twelve chords using only the span of two chromatic octaves (we could imagine it is a harmonic material for a miniature to be performed on a toy piano). Any chord may contain any subset of those notes, from none to all twenty-four. The number of all possible solutions is about five million times larger than the estimated number of protons in the observable universe (Tegmark 2003, 462) ($2^{24 \times 12} = 2^{288} \approx 4.97 \times 10^{86}$). Such a space is so large that not even the most powerful supercomputer could check all options one by one within any reasonable timeframe.

The problem evokes Jorge Luis Borges’ *Library of Babel*, a thought experiment imagining a library containing every possible combination of letters across countless books (Borges, 1941/1962; see Bottou & Schölkopf, 2023 for a discussion on Borges and AI). In Borges’ story, the library is so vast that any meaningful book is lost among endless volumes of random text. Similarly, the musical “library” generated by such simple parameters vastly exceeds what could ever be systematically explored, placing the search for musical solutions beyond the reach of exhaustive computation or human memory. Statistically speaking, one could imagine a monkey seated at a grand piano, randomly pressing keys for millions of years. Given enough time—and enough random strikes—it is not impossible that the monkey would eventually stumble upon a performance of a Chopin Mazurka (this is known as an infinite monkey theorem). Assuming we have no time to lose, and no army of monkeys at our disposal, we need a smarter way to search.

Despite the difficulty of the described problem, musicians do not rely on infinite randomness to discover meaningful structures. In fact, musicians effectively handle enormous musical spaces every day. Improvisers craft new progressions on stage in real time, and composers sketch ideas long before they test every alternative. It is possible, because the brain is built for good-enough answers when time and data are scarce. As Samuel Gershman—researcher in computational cognitive neuroscience—writes, “the brain is evolution’s solution to the twin problems of limited data and limited computation” (Gershman 2021, 154). To visualise the problem of limited data, imagine that you stand in front of a tree and wonder if your ladder will reach the lowest branch. You have never measured the tree, but using your own height, the distance to the trunk, and the way the trunk tapers, you form a quick estimate. It is not exact, yet usually close enough to decide whether to fetch the ladder. The problem of limited computation can be pictured with another example. After a concert you must catch the last train. Rather than modelling every possible route, you recall the simple rule “avoid the main avenue at rush hour” and pick a side street you know. The path may not be the absolute shortest, but it costs almost no mental effort, and you reach the station on time. Artificial Intelligence research copies these strategies. Deep Learning pioneers admit that a large part of the inspiration for deep learning is the human brain (Bengio 2021, xviii). Thinking of music as a multidimensional space of parameters—where pitch, rhythm, dynamics, timbre, and other characteristics each define independent dimensions—is not uncommon. For example, composer Davor Vincze—who often uses electronics and AI tools in his practice—envisions musical works as collections of points in a multidimensional matrix, an idea he operationalises in his microllage compositional strategy (Vincze, 2023, 312–314). However, as Vincze notes, manually mapping even a simple musical piece into such a space would be “a tedious and next-to-impossible task” (Vincze, 2023, 312).

Artificial Intelligence research combines the brain-like strategies with the ability of computers to quickly process large amounts of structural data. When a neural network is trained on musical material, it learns so-called latent spaces—internal maps where similar events lie close together (Audry 2021, 104). Intuitively, a latent space is the model’s own map of the music it has “heard”: chords, rhythms, or timbres; the model judges similar cluster together in the same region. Technically, each layer of the network turns its input into a new set of numbers. After training, these numbers form a distributed representation: no single unit means “major triad”, but a pattern across many units does. The model has compressed the raw, high-dimensional data

into a smoother, lower-dimensional surface where gradient search is possible. A process that imitates what human brains do in the process of learning enables computers to solve problems that were previously reserved exclusively for humans (Bolt, 2023, 97).

The key premise of this chapter is, then, that artificial intelligence systems are increasingly able to automate tasks long considered central to human musical creativity, such as navigating large possibility spaces, identifying patterns, and producing stylistically meaningful output. However, such compression brings a risk. Bias helps us find useful solutions quickly, but it can also hide unusual ones that might matter artistically. Creative work therefore juggles two needs: focus—to locate workable ideas fast—and openness—to leave room for surprise. To understand how artists engage with such systems, this chapter investigates the conceptual and practical foundations of AI-assisted composition. What follows is both a theoretical reflection and a practical report—much of the material presented here was developed and tested through a postgraduate course I designed as part of this doctoral trajectory and taught in 2023–2025 at the Conservatorium van Amsterdam and the Iceland University of the Arts. The course introduced students to creative uses of AI in composition, combining technical experimentation with artistic reflection, and it played a central role in shaping the methods and examples discussed in this chapter.

The chapter is structured in three parts. **Part I – On Automation** sets the conceptual groundwork by comparing symbolic and connectionist AI, then examines how automation reshapes creative labour and artistic agency. **Part II – On Methods** presents practical compositional strategies: first, workflows that generate symbolic musical data (MIDI) with machine-learning models and a discussion of a case-study composition *Ani(mate)*; second, interactive approaches that enlist large language models as conversational or performative agents. The **Conclusion – Artists as Operators** synthesises these strands, redefining the composer as an active operator who navigates, curates, and contests semi-autonomous systems.

Part 1—On Automation

Two AI Lineages: Connectionism and Computationalism

Artificial Intelligence (AI) is a notoriously elastic term, but its coinage is quite precise. Computer scientist John McCarthy introduced the term when he convened the 1956 Dartmouth Summer Research Project—a six-week workshop that aimed to discover “how to make machines use language, form abstractions, and improve themselves” (McCarthy, Minsky, Rochester, & Shannon, 2006/1955). From the start, then, AI referred not just to a technical agenda but to an ambitious cultural dream. The early researchers already envisaged “programs that pushed the boundaries” of machine capability (Mollick, 2024, Part 1, chap. 1). Six decades later, Tiziana Terranova—a theorist and activist whose work focuses on the effects of information technology on society—can describe AI as both the hidden operating system of daily life—controlling logistics, finance, and social-media feeds—and a magnet for enduring hopes and anxieties about consciousness, agency, and control (Terranova, 2023, viii). Beneath that lore lie two technological lineages whose rivalry still organises contemporary debate.

The first lineage often labelled symbolic AI (also computationalism or “good old-fashioned AI”), approaches intelligence as the orderly manipulation of explicit symbols under hand-written rules. In such systems every fact is encoded as a discrete token—for instance the rule IF `chord_is_major` THEN `mood_is_happy`—and reasoning proceeds by syntactic inference. Early milestones include Newell and Simon’s *Logic Theorist* (1956), often considered the first program to prove mathematical theorems, and the expert systems of the 1980s, whose thousands of if-then clauses diagnosed diseases or advised mineral prospecting. In each case cognition was specified in advance by the programmer.

The second lineage, connectionism, seeks to model intelligence by training adaptive networks to autonomously identify statistical regularities in data rather than by specifying rules in advance. A foundational conceptual step was taken by Canadian psychologist Donald O. Hebb, who proposed that neurons strengthen their mutual connections when repeatedly activated together, thereby forming distributed “cell assemblies” that encode experience (Hebb, 1949). Building on that principle, psychologist Frank Rosenblatt introduced the Perceptron (Rosenblatt, 1958), a simple

learning algorithm that trains a model to classify handwritten symbols by incrementally adjusting the strengths of the connections during training. In parallel, Oliver Selfridge proposed the Pandemonium model (Selfridge, 1959), which described pattern recognition as a hierarchy of simple processes. At the lowest level, basic "detectors" identified simple features, such as lines or curves. Higher levels combined the outputs of these detectors to recognise more complex structures, such as letters or symbols. Selfridge called these competing units "demons", a term that draws metaphorically on the idea of autonomous agents—echoing earlier scientific metaphors such as *Maxwell's demon*, which imagined independent entities making selective decisions based on observed inputs.

These early prototypes established two characteristics that continue to define contemporary neural networks. First, learning is achieved by optimising weights rather than editing symbolic code; as the authors of the Future Art Ecosystems briefing note, the model "is not manually defined through a sequence of instructions" but is produced by algorithms that search large datasets for regularities (Serpentine Arts Technologies 2024, 25). Second, pattern recognition is hierarchical: successive layers formulate progressively abstract representations, a principle illustrated by mid-century neurophysiological research on feature detectors in the frog retina (Lettvin et al., 1959; as discussed in Tilford, 2023, 128).

Connectionism nevertheless entered a period of reduced funding after Marvin Minsky and Seymour Papert (1969) demonstrated a serious limitation of early neural networks. A single-layer perceptron could only separate input patterns using a straight line, which meant it could solve simple tasks like distinguishing light from dark, telling whether a point lies above or below a diagonal on a plane, or deciding whether a musical note is higher or lower than a given pitch, but not slightly more complex problems. It could not compute the so-called exclusive-or (XOR) function—a basic logical operation that returns true if either of two inputs is true, but not if both are. Because the XOR problem cannot be separated by a straight line, it exposed that single-layer perceptrons lacked the flexibility needed even for elementary forms of reasoning. This finding severely undermined optimism about neural networks at the time, and research attention shifted back to symbolic approaches.

Only decades later, when new training methods, faster hardware, and large-scale datasets became available, did multi-layer networks re-emerge as a viable path. The crucial breakthrough was the development of backpropagation algorithms, which

made it possible to efficiently adjust the weights of neurons across multiple layers (Rumelhart, Hinton, & Williams, 1986). By computing how much each connection contributed to the final error and propagating corrections backward through the network, backpropagation allowed deep architectures to learn complex, non-linear relationships—such as the XOR function—that single layer models could not handle.

In parallel, theoretical work strengthened the case for connectionism. The universal approximation theorems proved that sufficiently large neural networks could approximate any continuous function to an arbitrary degree of accuracy (Cybenko, 1989; Hornik, 1991), showing that the architectural limitations identified by Minsky and Papert could be overcome when deeper structures were employed. Combined with the exponential growth of data and computing power, these advances made training multi-layer neural networks practically feasible for the first time. The resulting field of deep learning retains the perceptron’s core mechanisms—weighted connections, layered abstraction, and data-driven optimisation—while scaling them to contemporary computational resources.

Deep learning became practically viable only after several key advances converged. New training methods such as unsupervised pre-training made it possible to stabilise very deep networks (Hinton, Osindero, & Teh, 2006, as discussed in Audry, 2021, 98). At the same time, the rise of inexpensive graphical processing units (GPUs) and the availability of web-scale datasets provided the computational resources and data volume necessary to train large models effectively (Audry, 2021, 98). By the mid-2010s, deep learning was achieving state-of-the-art results in fields such as speech recognition, computer vision, and natural language processing (Goodfellow, Bengio, & Courville, 2016, 1). Crucially for this study, it also began to underpin a new generation of generative models capable of producing text, images, audio, and other creative media. The commercial enthusiasm that followed led to the development of today’s foundation models: vast neural networks trained on massive datasets, whose billions of parameters allow them to generalise across multiple modalities (Serpentine Arts Technologies, 2024, 25).

Since such systems learn from experience rather than from explicit programming, designers relinquish direct control and instead curate training regimes—the data, objectives, and evaluation metrics that steer each model’s optimisation. This shift carries political weight: the architectures now driving artistic exploration also underwrite what Martin Zeilinger (2021, 13) calls the “frightful surveillance tools” of

platform capitalism. Artists such as Keith Tilford therefore read neural networks ambivalently, both as instruments that expose our behavioural clichés—no more mysterious, perhaps, than a frog’s reflex to leap toward shadow—and as engines for re-imagining creative agency (Tilford, 2023, 128).

In what follows, I focus on this connectionist lineage, the ability of which—to discover latent musical structure without exhaustive rule design—makes it central to questions about the role of artists as operators of tools automating tasks associated with music composition. Symbolic approaches will re-enter the discussion in Chapter 3, where the composer’s role shifts from explorer of statistical spaces to explicit system-builder and legislator of musical rules.

Automation and Creative Work

Automation of creative tasks has fascinated and unsettled humans for centuries. Already Aristotle envisioned mechanical instruments relieving enslaved musicians from their endless toil, imagining machines capable of taking over tedious human labour (Clancy, 2022b, 169). Today, economic, writer and social theorist Jacques Attali speculates that artificial intelligence might radically expand creative output, foreseeing a future in which hundreds of symphonies are composed in the unmistakable styles of Beethoven or Tchaikovsky—indistinguishable from authentic historical works. For Attali, such scenarios illustrate how the boundary between artistic reality and virtual possibility will blur, dissolving the distinction between past works and future creations yet to be realised (Clancy, 2022c, 8).

This vision resonates with the capabilities demonstrated by recent AI models such as NotaGen, "a symbolic music generation model aiming to explore the potential of producing high-quality classical sheet music" (Wang et al., 2025). NotaGen allows users to specify a composer’s name as the target style for the generated music and can generate scores either from scratch or conditioned on prompts written in ABC music notation. While the compositions produced by the model do not yet perfectly simulate the styles of the referenced composer, they often exhibit surprising nuance, evoking an uncanny familiarity that reveals why a listener might recognise the intended composer in the generated work. Further exemplifying the trend of fully automating the music generation is the *Stochastic Pirate Radio* project by Bob Sturm et al. (2024).

This project employs an integrated pipeline of multiple AI models, including SunoAI, Boomy, and Stable Audio for generating audio content, alongside various Large Language Models for textual content, and additional models for speech synthesis. This aggregated system, which simulates a complete live radio station evoking the radio stations in the early Grand Theft Auto game series, is continuously live-streamed on YouTube, situating itself within a broader tradition of generative AI live-streams. Notably, the authors reference earlier influential examples such as Dadabots' *Relentless Doppelganger*, a stream which has been generating an uninterrupted flow of metal music since September 4, 2019, and remains active in various iterations to this day.

The automation of creative labour poses a real threat to artists whose societal roles and livelihoods often already exist in precarious balance. Many artists express a profound unease about being rendered obsolete by generative AI tools. As reported in the Future Art Ecosystems briefing, there is a widespread fear among artists: 77% feel that AI threatens to replace their jobs and opportunities, echoing broader societal anxieties where 60% of workers across various sectors anticipate significant changes due to AI-driven automation (Serpentine Arts Technologies 2024, 82). Furthermore, the briefing suggests that while AI-driven automation might lead to deskilling in certain crafts, history demonstrates that artistic skills continuously evolve; some disappear, others emerge, and some even experience a revival as cultural needs and technological tools change (2024, 91-92). At the same time, artists who possess skills predating the rise of AI, and whose expertise aligns with the sensibilities these new technologies leverage, hold an essential advantage. Their specialised knowledge positions them uniquely to master and creatively expand the possibilities offered by AI tools, illustrating the continuing value of human expertise even in highly automated contexts (2024, 90). In practical terms, current approaches to AI integration in creative workflows span a spectrum. Some organisations adopt a transformative "Minotaur" approach (the head of a bull and the body of a human), entirely restructuring workflows around AI capabilities. Others take a "Centaur" approach (upper body of a human and the lower body of a horse), empowering individual users to choose precisely where and how to integrate AI into their practice (2024, 66-67). Crucially, the interaction between artists and AI technologies is more complex than simple substitution narratives suggest. Instead, it initiates a dynamic process of feedback and emergence, where artists not only adapt to technological tools but also reconfigure the very notion of what creative work involves. This interplay, according

to Future Art Ecosystems, transcends narrow framings of automation as mere replacement, opening new possibilities for creative expression and collaboration across various artistic domains and technological layers (2024, 81).

Ethan Mollick notes that the current wave of automation uniquely impacts highly compensated, highly educated, and creative roles—jobs traditionally viewed as secure from technological displacement. He raises pointed questions about why clients or audiences should continue paying human artists when AI systems produce similar results rapidly and inexpensively (Mollick, 2024, Part 1, chap. 2). However, despite these concerns, automation in creative fields may not necessarily lead to outright replacement but rather to transformation of existing roles. Historical precedents show automation typically shifts job responsibilities rather than eliminating them. For instance, accountants who once manually calculated figures now use sophisticated software; they remain accountants but perform fundamentally altered sets of tasks (2024, Part 2, chap. 6). This trend suggests a future where artists increasingly curate and manage AI-generated content, focusing on tasks that demand uniquely human qualities such as creativity, critical judgment, and interpretative skill. Mollick further emphasises that despite AI's proficiency, genuine expertise requires humans to retain foundational knowledge and skills (2024, Part 2, chap. 8). He argues against the notion that basic factual knowledge is obsolete in the age of AI; instead, he highlights a paradox where foundational skills become even more crucial because true expertise depends on deeply interconnected knowledge stored in long-term memory. Consequently, the future of creative and expert labour involves both mastering foundational knowledge and becoming proficient in integrating AI tools into professional practice. Indeed, working effectively with AI itself constitutes a new form of expertise, where some individuals may excel dramatically, becoming the "new kings and queens" of AI-assisted tasks, while others may benefit less significantly (2024, Part 2, chap. 8). Moreover, Mollick speculates that the rise of AI could catalyse renewed interest in the humanities. Expertise in fields like philosophy, literature, history, and art uniquely qualifies individuals to interact with AI creatively and effectively, enabling the development of nuanced prompts and innovative applications (2024, Part 2, chap. 5). This suggests a potentially counterintuitive future where humanistic knowledge becomes increasingly valuable precisely because it enhances human capacities to manage and collaborate with sophisticated AI systems. Thus, rather than simply diminishing the role of human creativity and expertise, AI may redefine these roles, placing greater emphasis on uniquely human insights and

interpretative capabilities.

Media theorists Anna Munster and Ned Rossiter (2023, 48) highlight a polarised debate around automation, with some forecasting dire consequences for human labour, while others envisage human-machine collaborations enhancing value and productivity. Crucially, they point out that certain automated processes—such as AI-generated images—can operate indefinitely without direct human oversight, suggesting a future where creative machines operate increasingly autonomously. According to the authors, the emergence of the automated image represents a profound shift in the relationship between humans and machines, where the external world increasingly appears to itself mediated solely through computational processes, creating a continuum of porosity between machine operativity and social structures (Munster & Rossiter, 2023, 51). This transition marks the post-cybernetic era characterised by a redistribution of labour between conscious human cognition and non-conscious machine operations, signaling a diminishing role for human agency. Munster and Rossiter emphasise that traditional humanist notions of performance provide little solace against the backdrop of growing machine autonomy, suggesting instead that critical gestures—such as simulation, fakery, and the use of machine-generated 'tonality'—are necessary to performatively investigate and expose the operational grammar specific to contemporary automated image production (2023, 52-53). They provocatively question the place of human brilliance and creativity when AI assumes artistic roles, raising fundamental concerns about the ongoing significance of human inventiveness in a machine-driven context (2023, 60).

Marek Poliks and Roberto Alonso Trillo (2023, 29)—researchers in music and philosophy of technology—further illustrate how autonomous creative production increasingly diminishes human roles within cultural production and consumption chains. They use the example of ChatGPT to demonstrate how, despite stable or even expanding total labour input for cultural creation, the proportion directly executed by humans is progressively shifted to computational systems. This phenomenon results in fully autonomous chains of cultural production and consumption circulating online content with minimal human oversight, exemplified by bot-generated YouTube channels (previously mentioned *Stochastic Pirate Radio* being an illustrative experimental case), algorithmically optimised Spotify tracks, and entirely artificially generated and maintained websites. They position these developments within the philosophical frameworks of "computational realism", which proposes that the

universe itself operates on discrete, computable principles, and "New Empiricism", a Silicon Valley epistemology considering that "data itself, simply by virtue of sufficient quantity, is a generator of meaning" (2023, 26-27). Critiquing these ideologies, they suggest that reducing experience, intelligence, and consciousness to mere externalised data flows and recursive computational processes significantly flattens human cognitive and experiential depth, relegating human creative roles primarily to passive interpretation—little more than "cloud-reading, storytelling about patterns in the mist". Adding another layer to these discussions, Mattin (2023, 429)—musician known for his stance against copyright and for publishing his works under the no-licence—argues that the genuine threat to human autonomy is less about a singular, dominant artificial intelligence, and more about the complex and unpredictable collective behaviours of human groups mediated by networked computational systems. Mattin suggests that rather than a unified, malevolent AI force, the real danger lies in the subtle yet pervasive computational structures and ecosystems, powered by ordinary algorithms, driving human society into unpredictable and potentially troubling directions.

What to Automate and What to Control

The discussion around automation in creative labour necessarily engages deeply with the concept of autonomy inherent to AI systems. In this context, Artemi-Maria Gioti (2021, 1–2)—composer and artistic researcher working in the fields of artificial intelligence, musical robotics and participatory sound art—emphasises two crucial attributes defined by Michael Young (2008) for what he terms 'live algorithms': empowerment—the ability of algorithms to make independent decisions influencing future actions—and opacity, referring to their inherently complex and nonlinear generative processes. This autonomy can lead to significant indeterminacy, necessitating a form of trust from human collaborators in AI-driven systems to consistently deliver outcomes within anticipated bounds. Sofian Audry (2021, 161) similarly observes that the autonomy exhibited by sophisticated AI systems frequently results in outcomes that surprise even their creators, eluding complete rational comprehension. Audry cautions against overly simplistic narratives claiming that machine learning straightforwardly eases artistic labour or entirely replaces human authorship. Instead, they argue that machine learning reshapes and

redistributes creative agency, cultivating novel relationships between humans and machines within artistic workflows (2021, 159). Indeed, once models are designed, trained, and set in motion, they begin operating with a substantial degree of autonomy, shaping artistic outputs independently of direct human intervention (Audry, 2021, 35). Artists who engage with AI, thus, navigate a delicate balance, managing the high levels of autonomy intrinsic to these systems while strategically directing them to align with their aesthetic intentions (43). This tension between autonomy and authorship, as further elaborated by Audry (292), characterises a fundamental spectrum in artistic practice. Some composers, such as John Cage, are comfortable granting substantial autonomy to external processes, allowing materials and agents to express their intrinsic characteristics freely. Others, like Brian Ferneyhough, construct tightly controlled systems in which every parameter is rigorously specified, ensuring that the outcome adheres closely to a predetermined compositional logic. In such practices, complexity is not a byproduct of autonomous processes but the result of deliberate design, where the composer maintains interpretive authority over every layer of the musical material. In either case, the integration of AI within artistic processes does not represent a mere technological substitution; rather, it necessitates deliberate artistic negotiation between human intentionality and machine autonomy, highlighting ongoing creative tensions and opportunities within human-machine collaborations.

Emanuele Arielli (2021)—philosopher and researcher interested in aesthetics, cognitive science and communication theory—provides a useful classification of how machine learning intersects with aesthetic processes by mapping two fundamental dimensions: object versus subject and description versus generation. From these dimensions, four distinct types of machine learning tasks in aesthetic domains can be identified. First, AI systems can study objects—analysing stylistic and formal patterns in existing artefacts, such as harmonic structures in music compositions. Second, they can generate objects, producing stylistic variations that closely mimic the analysed corpus, thereby extending and reinterpreting existing aesthetic traditions. Crucially, Arielli emphasises that a comprehensive account of aesthetic processes must also consider the dimension of human reception. Thus, the third type involves studying subjects—tracking, modelling, and predicting aesthetic preferences of audiences through behavioural data. Finally, the fourth type concerns generating subjects: AI systems can be trained to simulate human aesthetic judgments, autonomously evaluating novel artefacts without direct human input (Arielli, 2021, 12–14).

Although various techniques exist to automate aspects of creative workflows, it is crucial to distinguish between different kinds of automation. For instance, methods such as Self-Organising Maps and K-means clustering can automate the classification of large sets of generated objects, grouping musical fragments based on feature similarity. Meanwhile, social media profiling algorithms automate the classification of subjects, modelling and predicting aesthetic preferences based on user behaviour. From the perspective of Arielli's framework, both activities—classifying objects and modelling subjects—could theoretically be delegated to machines. However, for artists seeking to maintain an active, critical engagement within the creative process, automating these two aspects can be counterproductive. It is precisely through curating musical artefacts (classifying objects) and anticipating the aesthetic responses of listeners (modelling subjects) that composers retain agency, creative intuition, and aesthetic control over their works. The emphasis on maintaining an active human role aligns with recent theoretical work on human-AI co-creativity. Rather than treating AI models as autonomous generators of finished musical works, the recent studies foreground interactive, performative engagements with the outputs of machine learning systems. For instance, the *Human-AI Musicking* framework proposed by Valverde, Agres, Herremans, and Chew (2023), published in the proceedings of The International Conference on AI and Musical Creativity (AIMC), conceptualizes the creative process as a dynamic co-activity between human and machine, where the human maintains a central curatorial and interpretative role throughout the generative process. Within this model, creativity is understood not as a product delivered by the AI, but as an ongoing, reciprocal negotiation in which human aesthetic agency remains paramount.

In AI-assisted composition, two broad modes of generation can be distinguished. Some models aim to produce relatively complete musical outputs in a single pass, requiring the composer to engage primarily in selection and refinement of the initial input. Once the input parameters are set, the model does the rest and generates a full ready-made composition. An example is the NotaGen system (Wang et al., 2025), which generates stylistically conditioned full scores. Other models are better thought of as generators of modular material: smaller fragments, motifs, or textures that invite recombination and creative assembly. In this approach, the composer does not expect a finished work but rather gathers musical strata to construct larger forms.

These very broadly described generation strategies could be seen as, importantly, independent of any tool or platform and operate across different forms of musical data, whether symbolic (e.g., MIDI files, piano-roll representations) or audio-based. Nevertheless, symbolic data currently afford far more precise manipulation of musical structures than audio data, offering finer control over parameters like pitch, rhythm, dynamics, and articulation. While advances in audio generation are rapidly narrowing this gap, symbolic representations remain particularly advantageous for composers seeking detailed, granular level of control over musical material. This distinction, while important today, is likely to blur over time as machine learning models for audio synthesis continue to improve.

Part 2—On Methods

Teaching Creative Workflows with Artificial Intelligence

Much of the material discussed in this chapter stems from an experimental composition course I developed and taught during this doctoral research trajectory, titled *Posthuman Creativity Labs: Artificial Intelligence and Blockchain in Music*.

The course was first taught in 2023–2024 as a master’s elective at the Conservatorium van Amsterdam, where it unfolded over the academic year through a series of monthly seminars and workshops. Each session combined short theoretical lectures with hands-on experimentation, and students were asked to bring back creative assignments using the presented technologies. Students were invited to reflect on how the tools they used changed their approach to composition, their sense of authorship, and the identity of the musical material they produced. The student cohort included composers, sound designers, and performers, many of whom had no programming experience. This constraint shaped the course’s design: it favoured accessible, Graphical User Interface-based tools like Magenta Performance RNN, Riffusion, and ChatGPT, while also scaffolding more advanced practices for those who wished to go deeper.

In a second iteration of the course, also hosted at the Conservatorium van Amsterdam but now as part of the *Creative Performance Labs* program in January and February 2024, the structure was condensed into four intensive sessions. This time, the focus shifted more explicitly toward preparing a final performance: the course culminated in a presentation of students’ works. Because of the compressed format, the second iteration built directly on the practical experiences and examples developed in the first version. It featured refined teaching strategies and compositional tasks that had proven effective, while also placing greater emphasis on workflow design and real-time interaction with AI systems. Again, the class included students from a broad musical background, most without programming knowledge, which required maintaining a high degree of technical accessibility without simplifying the conceptual ambitions of the course.

A third institutional iteration took place at the Iceland University of the Arts in Reykjavík, where the student cohort included individuals with prior experience in computer programming, including live-coding practices using programs like

SuperCollider. In response, this version of the course adopted a more advanced technical orientation, building upon the methodological foundations laid in the earlier versions. Students were invited not only to interact with existing AI tools, but also to prototype their own workflows, combining generative reasoning models (such as ChatGPT) with sound synthesis and control environments. In several cases, this led to the development of custom interfaces or API extensions that allowed for real-time interaction between reasoning models and live-coded music systems.

Finally, the course was presented and adapted in various guest lectures to different musical communities, including classical composers (at Geneva University of Art and Design), film music students (at ArtEZ University of the Arts), and artistic researchers (at Orpheus Institute). These presentations tested the portability of the core workflows and confirmed their utility across a wide range of artistic contexts. In each case, the course prompted its audience to rethink not only what composition entails, but what it means to collaborate with systems that are both generative and opaque. In all its versions, the course pursued a dual aim: first, to demystify AI and blockchain by giving students concrete tools for integrating these technologies into their artistic practice; and second, to foreground the conceptual and aesthetic shifts introduced by these technologies. Topics included the structure of latent spaces, the automation of creative labour, the reconfiguration of authorship, and the emergence of posthuman agency in music. Students were encouraged to consider not only what AI systems can generate, but how they generate, and under what assumptions. The guiding questions were both technical (“how do I get this to work?”) as well as operational and reflective: “what kind of composer do I become by using this?”, “what kind of work is this tool capable of producing?”, and “what parts of the process am I delegating, shaping, or reclaiming?”.

The course served as a live, iterative creative testing ground for both artistic inquiry and conceptual development of this doctoral trajectory, as well as a generator of engagement with the topics discussed in this thesis. Through experimenting with various AI tools in the context of students’ creative practices, it enabled a testing ground for the iterative development of compositional workflows with new technologies. This led to developing two distinct compositional workflows: one that I further describe as *Generate, Select, Edit, Continue*, and another based on prompt-responsiveness and iterative model control, titled “*Could you continue my MIDI, please?*”. Both emerged from experiments in the classroom, where students negotiated

the limits and possibilities of machine-generated musical material. These methods—presented in the following sections—do not claim to be comprehensive or universal, but they illustrate how composers can function as operators: users who activate, guide, and curate automated systems within constrained creative tasks.

In this sense, the course was not only pedagogical but also research-driven. It allowed for direct observation of how AI systems affect compositional decision-making, how artistic intentions are refracted through non-human agents, and how authorship becomes distributed across systems of generation, selection, and refinement. The role of the composer as operator—proposed at the end of this chapter—is grounded in this experience of collaborative learning, experimentation, and critique.

Part 2.1—Composing with AI Models Trained on Music

Generate, Select, Edit, Continue

Currently, most AI models capable of generating high-quality audio outputs are trained predominantly on datasets of popular music. Examples include models like Suno, Udio, and Stability AI's Stable Audio, all of which excel at producing tracks in popular genres. However, there are no widely available models capable of producing high-quality, stylistically sensitive generations of experimental or contemporary classical music. There are several reasons for this. First, by definition, experimental music resists standardisation, making it difficult to assemble sufficiently large and coherent datasets for training. Second, much of the contemporary music repertoire remains under copyright protection, limiting the available training material. Third, training state-of-the-art audio models remains prohibitively expensive, and commercial entities have little economic incentive to invest in models aimed at small, specialised audiences. Perhaps most fundamentally, experimental music often prioritises exploration of sound-production methods themselves rather than adherence to recognisable stylistic patterns. As a result, generative models trained purely on existing examples would struggle to capture the experimental ethos unless they were specifically designed to model creative processes rather than surface outputs—a challenge that touches on deeper questions about the role of experimentation and innovation in art music, to which we will return in the next chapter. Nonetheless, even models primarily trained on popular music corpora can

often be creatively "hacked" or subverted for experimental purposes. For instance, at the course I taught at the Conservatorium van Amsterdam, some students took the initiative to experiment with ElevenLabs' Voice Changer model, which is normally designed to take a human voice as input, and output the same spoken text using a cloned voice. Instead of using speech, they recorded the sounds of musical instruments and fed them into the model. The result was a striking hybrid: the instrumental sounds were transformed, carrying an unexpected blend of instrumental timbre and human vocal qualities, creating an uncanny and novel listening experience. By exploiting their biases, breaking their expectations, or creatively manipulating their inputs and outputs, composers can repurpose these models in ways that defy their original design intentions. For the purposes of this chapter, however, which focuses on compositional workflows based on using AI systems in ways roughly aligned with their intended operation, it is useful to examine concrete examples based on symbolic music generation.

The first example is Performance RNN model (Simon & Oore, 2017), developed within Google's Magenta project. Performance RNN was trained on a dataset of expressive, good quality performances (Yamaha e-Piano Competition dataset), specifically MIDI recordings rich in timing deviations, dynamic variations, and articulation nuances—elements crucial to human musical expressivity. The model is optimised not simply for generating correct notes, but for producing new performances that imitate the timing, dynamics, and phrasing characteristics found in the training data. In this sense, it represents an important step beyond earlier symbolic models that focused purely on pitch and rhythm without expressivity (Briot, Hadjeres, and Pachet, 2020, 190). Crucially for non-programming musicians, Performance RNN can be used without technical expertise through an online demo page provided by Magenta, allowing real-time interaction with the model. This makes it a convenient tool for demonstrating AI capabilities to non-coding artists, for instance in a conservatoire context. The generated MIDI output can be routed to a Digital Audio Workstation (DAW) or into a score notation software, enabling musicians to record, edit, and manipulate the AI-generated performance as part of their compositional process¹. The model's output can be further conditioned by constraining it to selected scale patterns,

¹ To enable this setup, users must create a virtual MIDI channel. On macOS, this is done through the Audio MIDI Setup utility: Audio MIDI Setup → Window → Show MIDI Studio → Double-click IAC Driver → Enable "Device is online". Windows users can achieve a similar setup using external tools such as the free software LoopBe1.

either manually via the demo interface or dynamically through MIDI input.

When recording the output of Performance RNN, composers can interact with the model in real time, influencing its responses and capturing the results. This might be seen as situating the process within a tradition of using improvisation as a compositional method. As Artemi-Maria Gioti notes in her dissertation on distributed creativity (Gioti, 2021), composers like Giacinto Scelsi famously used improvisation as a strategy for composition, recording spontaneous sessions and later transcribing or adapting them into notated works (78). Similarly, the concept of "seeded improvisation" by composer and researcher Richard Barrett describes a hybrid approach in which composed materials are interspersed with spaces for improvisation, allowing emergent musical phenomena that could not arise from notation or free improvisation alone (Barrett, 2014, as emphasised in Gioti, 2021, 78). Within this lineage, using Performance RNN can be seen as a digital extension of these improvisatory compositional strategies. The AI model serves not as a generator of finished works, but as an improvisational partner producing rich musical material to be further edited, selected, and assembled into coherent artistic structures.

It is important to emphasise that such generated material should not be treated as performance-ready or self-sufficient. On the contrary, its greatest creative potential lies in treating it as raw material for further experimentation. A common strategy among artists working with generative systems is to produce large quantities of varied outputs before any compositional decisions are made. This abundance creates a fertile space for exploration, enabling composers to navigate, classify, and sculpt the generated material according to the needs of their larger artistic vision. By maintaining a clear separation between the generative phase and the curatorial phase, composers preserve their aesthetic agency and integrate AI models into a broader experimental workflow. However, once such musical material is prepared and ready to be incorporated into a larger musical work, there are additional strategies that can be employed to further develop it using AI tools, assuming the composer wishes to continue engaging with generative processes. Two particularly useful approaches are continuation and generative in-painting.

In the continuation approach, an AI model is tasked with generating a stylistically coherent extension of an existing musical fragment. The composer provides an initial musical prompt—whether written manually or generated earlier by another model—and the system predicts how the musical material might plausibly continue based on

the stylistic features embedded in the input. In the generative in-painting approach, the model fills in the musical gap between two given fragments, creating a smooth and stylistically unified transition. This technique is particularly valuable for connecting independently generated or composed sections into a cohesive musical whole.

Both strategies are made accessible through the Allegro MIDI Transformer model, developed by programmer Alexandr Sigalov as part of his broader Project Los Angeles, which explores novel architectures and applications of AI for music generation (Sigalov, 2020). Allegro MIDI Transformer builds on the Transformer architecture, a deep learning model originally designed for natural language processing, but now widely adapted for symbolic music tasks. The model allows composers to either extend an existing MIDI fragment into a longer sequence or interpolate between two different fragments, producing fluid transitions that can bridge stylistically distinct ideas. Importantly, Allegro MIDI Transformer can be used not only to continue musical material originally written by the composer but also to further develop material generated by other AI systems, such as Performance RNN. This opens the possibility of chaining multiple models together within a single workflow, leveraging the strengths of each at different stages of the creative process.

Bringing these practices together, a flexible compositional framework can be summarised through the iterative workflow of generate, select, edit, continue:

- **Generate:** Create a large pool of musical material using one or more AI models. This phase prioritises diversity and richness over immediate coherence.
- **Select:** Curate the generated material, classifying fragments according to their usability, stylistic suitability, or experimental potential within the emerging work.
- **Edit:** Actively modify, combine, and transform the selected materials to align them with the composer's artistic aims. Editing can involve anything from minor rhythmic corrections to radical re-contextualisations.
- **Continue:** Use continuation or in-painting models to extend, connect, or reimagine the edited materials, further integrating them into a larger compositional structure.

Through this iterative cycle, composers maintain their central creative role, while AI systems function as sources of inspiration and musical material. Each phase invites human aesthetic judgment, ensuring that the outputs remain embedded within a

broader artistic vision rather than dictated by the statistical patterns learned by the models.

Case Study: *Ani(mate)* (Movements 1–5)

The methods discussed above informed the composition of *Ani(mate)*, a commissioned work for piano and string ensemble. The piece was commissioned by the Chamber Music at Lundsgaard festival to open their concert series at Lundsgaard Gods in Kerteminde, Denmark on the 7th of August 2025, and was also subsequently performed again in Copenhagen on the 12th of August 2025. The commissioners invited me to compose a new work that would explore the creative use of artificial intelligence tools. It was premiered by Trio Con Brio Copenhagen together with the Marmen Quartet, pianist Enrico Pace, violists H  l  ne Cl  ment and Michael Germer, violinist Andrej Bielow, and cellist Jonathan Swensen. The title, *Ani(mate)*, carries multiple meanings. To "animate" is to bring to life, to give motion or spirit. The term *mate* suggests companionship or partnership yet simultaneously references *checkmate* evoking notions of conflict, strategy, or finality—as in the decisive end of a chess match.

The resulting composition consists of six short movements. The first five were composed using the connectionist AI-based methods described throughout this chapter, while the sixth movement was created using a distinct set of methods, which will be discussed later in Chapter 6. This section focuses on the first five movements as a concrete case study of the operator role: how the techniques of “generate, select, edit, and continue” can be used to develop musical material in dialogue with machine learning systems.

In these movements, I worked extensively with the Magenta Performance RNN model to generate a large corpus of MIDI material conditioned by a system of curated pitch-class distributions. By adjusting the input parameters, I could influence the statistical frequency of particular pitch classes in the generated sequences, creating long, unstructured passages of material. I recorded hours of such “seeded improvisations” from the model, then curated fragments that seemed musically promising according to my compositional judgment. This process required the same form of operational decision-making discussed earlier in this chapter: engaging with the model as an

opaque yet responsive environment, and navigating its latent space through iterative exploration.

Once selected, the fragments were imported into a Digital Audio Workstation software, where I edited the MIDI data to remove noise, refine durations, and shape local rhythmic profiles. I then transferred the material into a score-editing software to create a clear piano score. At the final stage of the process, I orchestrated the selected fragments for the string ensemble and piano. The piece exists as a score in two versions: for a single pianist (performed by Jens Elvekjaer in Copenhagen) and for two pianists playing four hands on one piano (performed by Jens Elvekjaer and Enrico Pace in Kerteminde).

While much of the musical material in these movements was derived from model-generated sequences, I was struck by how strongly the resulting music aligned with my existing compositional idiom. This was not an outcome I anticipated. Given that the model was trained on external performance data and conditioned only by statistical constraints on pitch-class distributions, I expected the outputs to diverge significantly from my established stylistic tendencies. Instead, after curating and normalising the material (besides the desired material, the model outputs various types of noise in the generated signal), the resulting music sounded surprisingly similar to what I might have composed without AI tools at all.

This observation highlighted how central the acts of selection and structural positioning are to the compositional process. The final character of the piece did not appear to result primarily from the local content of the fragments produced by the model, but from the decisions about how these fragments were ordered, how they shaped musical tension and release, and how they were evaluated against my own aesthetic preferences. These decision points—shaping temporal progression, controlling suspense, and filtering material according to taste—seem to act as strong attractors that pull diverse material toward a personal style. Even when the source material is heterogeneous, the composer's operational choices can impose continuity and identity. This finding offers an important perspective on the operator role: curatorial navigation of large generative spaces can, paradoxically, intensify rather than dissolve authorial voice.

To test this observation, I composed the fifth movement with a deliberately different intention. Here, I selected a longer fragment from the AI model without intervening

to reshape it according to my usual criteria, and allowed material I would typically exclude as cliché or overly sentimental to remain. To my ear, the resulting movement is markedly more traditional and emotionally overt than the surrounding ones. I decided to retain it in the final score precisely to underscore this artistic point: that the suspension of curatorial filtering can produce a distinctly different stylistic result, making the influence of those filtering decisions visible by their absence.

Another noteworthy aspect of the piece is the tension between its macro-level coherence and its micro-level idiomatity. Although the six movements form a coherent overall structure, the instrumental writing often departs from idiomatic ensemble practice. This reflects the origin of the material: the model generated only piano performances (not scores), which I then normalised and orchestrated. Because the model was not conditioned on ensemble-specific affordances, the resulting parts occasionally diverge from conventional instrumental habits, even as they contribute to the macro-level continuity of the piece. This disjunction illustrates how integrating non-orchestrated, model-generated material into ensemble writing can produce music that is structurally cohesive yet locally idiosyncratic, and it became an important feature of the work's character. Nevertheless, it is important to remain aware of the limitations inherent in current generative models. One of the major challenges is that AI models are inevitably shaped by the properties of their training datasets.

Biases in Training Data and Artistic Control

To perform well, models must be trained on thousands of examples, but this necessity introduces stylistic and structural biases that can be difficult to overcome. The dataset itself becomes a latent boundary on what the model can generate, and it is often infeasible to manually curate such large datasets for stylistic purity or quality control.

One promising solution to this challenge lies in fine-tuning pre-trained models with a much smaller, curated dataset tailored to the artist's own stylistic preferences. Fine-tuning involves taking an already trained model and adjusting it slightly to better reflect a new, specialised corpus. An illustrative example of this approach can be found in the *Nobody's Songs* project by programmer Sebastian Macchia (Macchia, 2020). Macchia fine-tuned Magenta's Music Transformer model on the music of Erik Satie, adapting the model to generate material more closely aligned with Satie's style.

The generated material then served as the basis for further creative work, ultimately resulting in a full album. In this way, fine-tuning offers composers a way to personalise generative models, allowing for a deeper alignment between AI outputs and individual aesthetic goals. Rather than being constrained by the biases of large, generic datasets, artists can sculpt the creative potential of AI systems according to their own choices.

However, the dependency on large, general-purpose datasets can often reveal itself as a significant limitation to artistic experimentation. Models trained on extensive corpora excel at continuing or extending materials that resemble their training data but can struggle when confronted with radically divergent musical languages. For example, the Allegro MIDI Transformer is highly effective when tasked with continuing a classical music fragment, such as a short excerpt by Mozart. It can also produce relatively satisfying results when asked to extend a post-tonal musical fragment in a style reminiscent of Mussorgsky, Prokofiev, or Bartók—a combination I tested with quite promising outcomes. However, when presented with a strict twelve-tone series, the model's output quickly collapses into nonsensical continuations, resembling geometrical shapes rather than meaningful musical structures. In such cases, the model's internalised assumptions about harmonic progression, phrase construction, and rhythmic regularity—patterns absorbed from the training corpus—reassert themselves, making it difficult to maintain the structural integrity of more experimental idioms. This example highlights a crucial limitation: when working with models trained on broad stylistic data, artists must always account for the latent stylistic biases encoded within the model. No matter how flexible a model may seem, it cannot easily extrapolate beyond the structural habits ingrained in its training data without substantial retraining or fine-tuning.

At the same time, it is important to avoid over-interpreting this limitation as an indictment of generative methods altogether. As Philip Galanter (2003)—theorist, curator, and researcher interested in generative art, sound art, and complexity science—points out, “(...) the generative approach [itself] has no particular content bias, and generative artists are free to explore life, death, love, war, beauty, or any other theme” (Galanter, 2003, 17). In other words, while the data used to train a model may introduce stylistic limitations, the generative methodology remains fundamentally neutral. It is not the act of using generative processes that imposes aesthetic restrictions, but rather the choices surrounding what data, assumptions, and

frameworks those processes are based upon. This distinction is crucial for composers and artists using AI tools: the limitations of existing models should not be seen as intrinsic shortcomings of AI-assisted composition itself, but rather as challenges that can be addressed through thoughtful interaction with the system. One way to overcome these challenges is to provide generative models with richer, more stylistically targeted input data or contextual guidance. However, achieving such fine-grained controllability within current symbolic generation models remains a technical challenge.

In this context, it is worth acknowledging that some recent systems, such as Google’s Lyria 2 and Meta’s MusicGen, use text as input to guide music generation. In these models, users prompt the system by describing in natural language the style, mood, or instrumentation they wish to hear. However, from the user’s perspective text-based prompting does not constitute a fundamentally different generative paradigm, but rather just a different interface for interacting with models that are still grounded in learned latent musical spaces. Their outputs remain shaped by the biases and constraints of their training data, just as with symbolic input models. While the interaction moves from providing musical fragments to crafting verbal descriptions, the underlying generation dynamics and limitations are comparable.

This issue—how to guide generative models toward more stylistically sensitive, context-aware outputs—points toward another compositional strategy. In what follows, we will explore a workflow based on text-based generative AI agent for symbolic music generation. These conversational systems offer a different mode of collaboration, enabling users to “teach” the model specific theories, stylistic principles, or creative constraints through natural language instructions. Rather than relying solely on prior training data, these agents can negotiate musical ideas dynamically, allowing composers to develop musical materials together with the model in an interactive, iterative process.

Part 2.2—Composing with Large Language Models

Large Language Models as AI Agents for Musical Tasks

Among the many forms of Artificial Intelligence, one class of models has recently come to dominate both research and popular imagination: Large Language Models

(LLMs). LLMs are a type of neural network trained to process, predict, and generate human language by learning from vast textual datasets. They represent a specific approach to AI that privileges statistical pattern recognition over explicit symbolic reasoning, continuing the connectionist lineage described earlier. Although statistical models for language have existed for decades, LLMs distinguish themselves by their unprecedented scale and versatility, making it possible for a single model to perform a wide range of tasks—from translation and summarisation to creative writing—without task-specific retraining. The foundational architecture for most LLMs is the Transformer, introduced by Vaswani et al. (2017), which replaced earlier sequential models with a parallelisable attention mechanism capable of capturing long-range dependencies in data. However, it was the scaling experiments conducted by Brown et al. (2020) in the development of the GPT family that revealed the surprising few-shot learning abilities of very large Transformer models, demonstrating that capabilities could emerge from scale alone. Bommasani et al. (2021) later offered a more conceptual framing, describing LLMs not merely as Transformer-based models trained on web-scale corpora, but as systems whose latent representations compress, generalise, and transform broad domains of human knowledge, enabling emergent capabilities that extend beyond their original training objectives. Over time, the Transformer architecture has become so central to AI research that it now practically defines the public understanding of AI itself, with the term "AI" increasingly evoking Transformer-based systems like ChatGPT (Poliks & Alonso Trillo, 2023, 24).

Human reasoning is traditionally anchored in deduction: the logical extraction of truth from known premises. Deduction operates according to formal rules that guarantee validity—if the premises are true and the reasoning steps are valid, the conclusion must also be true. In contrast, the reasoning exhibited by Large Language Models (LLMs) is based not on deduction but on inference: the statistical prediction of the most probable continuation given prior data. LLM predicting "Neil Armstrong" as the first person on the moon does so purely through statistical association, not through comprehension or deduction (McCormack, 2023). LLMs do not know facts, meaning, or truth; they predict the next most likely token based on the patterns learned in training. This distinction explains why LLMs can often produce outputs that appear deductively valid without possessing any actual understanding. As programmer and researcher Andrej Karpathy (2023) explains, while we have a detailed understanding of the architecture and training processes of LLMs—including the mathematical operations, optimisation procedures, and engineering principles

involved—the question of how these models generalise so broadly remains open. Despite being trained primarily on the task of next-token prediction, LLMs spontaneously develop complex capabilities such as reasoning, planning, and abstraction. Research efforts in mechanistic interpretability are actively working to reverse-engineer these internal processes, aiming to map how specific circuits and patterns within the networks contribute to emergent behaviour. While there have been important advances (e.g., in understanding induction heads and simple algorithmic circuits), our current understanding remains partial and fragmented.

Inference imitates deduction sufficiently well for many practical purposes because human reasoning itself, limited by cognitive and informational constraints, often substitutes heuristics and probabilistic estimations for strict logical tasks (such as in the example of estimating the height of a tree at the beginning of this chapter). The distinction between different types of cognitive processing, famously popularised by Daniel Kahneman in *“Thinking, Fast and Slow”*, offers a helpful analogy for understanding LLM behaviour. Kahneman distinguishes between System 1, which is fast, instinctive, and automatic, and System 2, which is slow, deliberate, and effortful (Kahneman, 2011). As Karpathy (2023) explains, human responses to simple problems like “2 plus 2” are instant because they are cached in System 1, whereas more complex tasks, such as multiplying “17 times 24”, require the slower, rational operations of System 2. Current LLMs, in their core, primarily operate through a process resembling System 1: they produce fast, pattern-based predictions without sustained deliberative thought. However, on top of these systems, various techniques such as chain-of-thought prompting and inference-time scaling are used to simulate the multi-step processes associated with System 2. As LLM research engineer Sebastian Raschka (2025) explains, reasoning in LLMs involves generating answers that require multiple intermediate steps rather than jumping directly to a conclusion. Regular LLMs can answer straightforward factual queries, but true reasoning tasks—such as solving mathematical problems, puzzles, or coding challenges—require models to simulate an extended, stepwise thought process. This can be encouraged more deeply through specialised training methods aiming to scaffold internal representations that can handle multi-step problem solving more effectively, moving LLMs closer to System 2-like behaviour when tackling complex tasks. In that sense, LLMs become versatile models for solving a broad range of general tasks.

An important further development is the concept of AI agents. Unlike standalone LLMs, agents are systems designed not only to reason internally but also to interact with external tools to achieve specific goals. As researchers Wiesinger, Marlow, and Vuskovic (2024) explain in a guide paper on technical application of AI agents released by Google, just as humans use external aids—such as calculators, books, or search engines—to supplement their reasoning, Generative AI models can be trained to use tools via Application Programming Interfaces (APIs) to access real-time information, retrieve data, or trigger actions. An AI agent is thus defined as an application that observes its environment and acts upon it through the tools it has available, aiming to fulfil a specific objective (Wiesinger et al., 2024, 5). The technical mechanism that enables this interaction is the API—a set of tools that facilitates the exchange of data and functionality between systems (Serpentine Arts Technologies, 2024, 24). By calling APIs, agents extend their abilities beyond internal text prediction, enabling them to fetch, update, manipulate, or generate information dynamically. Tools exposed through APIs typically operate through standard methods like GET, POST, PATCH, or DELETE, allowing agents to perform complex tasks such as querying databases, retrieving web data, or interfacing with external software (Wiesinger et al., 2024, 7). To facilitate the use of APIs, agents often employ structures known as Extensions. An Extension teaches the agent how to interact with a particular API by providing examples of correct usage, specifying the necessary arguments, and modelling the format of successful API calls (14). Functions operate similarly but are executed client-side rather than agent-side: a model can output a function and its arguments without making a live API call, leaving the execution to the client (18). Together, the concepts of Extensions, Functions, and APIs build the operational scaffold that enables agents to transform LLMs into active participants capable of modifying and acting upon digital environments.

The ability of AI agents to interact with external tools opens wide-ranging possibilities for creative applications in music. Rather than limiting generative AI models to producing text-based outputs, agents can be configured to perform practical musical tasks by interfacing with digital environments. For instance, an agent can control a Digital Audio Workstation (DAW) to automate recording sessions, adjust parameters, trigger samples, or manipulate MIDI tracks. A web-based platform called WavTool temporarily demonstrated such possibilities by integrating an LLM-based agent directly into a DAW interface. WavTool enabled users to manage the production workflow through natural language commands, allowing the AI to automate tasks

such as recording, editing, and arranging, as well as to trigger embedded AI models to generate, extend, or modify musical fragments within the same environment. Although at the time of writing WavTool has gone offline with a notice that it will return with expanded capabilities, it is likely that similar integrations of AI agents into DAW ecosystems will soon be replicated and further developed by major DAW manufacturers, progressively embedding AI-assisted workflows as standard features within music production software.

“Could you continue my MIDI, please?”

During one of the institutional editions of my course Posthuman Creativity Labs: Artificial Intelligence and Blockchain in Music, taught at the Iceland University of the Arts, the students expressed a strong interest in experimenting with LLMs for music making. At the same time, WavTool—the primary tool I had previously used to demonstrate this capability while teaching in Amsterdam—was no longer available. To enable the use of ChatGPT’s reasoning and text-writing capabilities for working with MIDI files, I thus developed a simple API function². This tool allowed students to upload MIDI files during a conversation and have them converted into a JSON format understandable by the LLM in the realtime of their conversation and from inside of the ChatGPT’s user interface. The experiment opened a rich space for exploration of the tool’s capabilities and limitations. What I have quickly discovered when working with the tool to prepare the class is that successfully implementing such a workflow in music experimentation requires a very careful strategy.

The agent must correctly interpret the context of its operations: it must understand who is issuing a command, what kind of output is expected, and how the tool it controls functions within that specific environment. One thing to remember is that plain LLM, without further fine-tuning to better align with a specific expertise, does not have any knowledge on how to successfully write and analyse music. However, it does know how to successfully structure the MIDI file format and how to “reason” with natural language. This gives us the opportunity to teach the model any music theory that we’d like it to know, before applying this in practice. The agent must

² The code for the API function is available on my GitHub (@sunsetobserver) under the name MIDI-JSON-API.

embody a human-like understanding of the tasks it performs. For instance, to analyse a musical score, it must first grasp the theoretical foundations—such as harmonic analysis, form, or voice-leading—before attempting to classify or critique musical structures. Similarly, for generative composition tasks, the agent must internalise stylistic, structural, and aesthetic constraints, rather than applying surface-level transformations without deeper comprehension. Ensuring such context-sensitive, musically informed behaviour remains a key challenge when deploying agents in professional creative workflows. The agent, asked briefly to compose something interesting and to return a MIDI file, will almost certainly fail to do the job as it does not know who is asking, why is such question relevant, and what would be the context for an appropriate evaluation of such a task. However, two strategies tested during my course, provided more than extraordinary results in terms of musical outputs.

In one scenario, the student would first explain to the LLM their own musical background and artistic expectations when working with a musical material. Then, they would input a MIDI file of a musical fragment that they like, to show the model what is considered a “good” generation. To enable deeper understanding of that fragment, the model would be further queried to take its time and analyse the file according to its various parameters, such as to perform an analysis of the musical fragment’s melodic, harmonic, and rhythmical content. After such an extended analysis, the model would be asked to try and conceptualise what a stylistically-coherent continuation of that musical fragment could contain—what would be the changes in the musical material and how could this be implemented. In that situation the model would come up with a few possible ideas and confront them with the student. The student would then ask the model to choose their preferred strategy and ask the model to generate the file. The model would write the extended composition and use the API tool to convert that text into a downloadable MIDI file, which the student can then open in their preferred DAW and listen to for evaluation. While the evaluation of the results of such a creative workflow will always be inevitably subjective, the high-level controllability with which the user can teach their personal assistant to become gradually more helpful, enables a way of achieving highly satisfactory musical results, comparable and actually surpassing the capabilities of pre-trained musical transformer models such as Allegro MIDI Transformer for experimental creative applications.

We quickly discovered that the presented workflow, while successful, is quite labour intensive. At least at first, to prepare the context for further creative conversations. The general rule is that the more creative context is provided to the model, the more useful it becomes for assisting the creative collaboration with the specific user. However, it is possible, at least partly, to strategise a quicker contextual jump enabling the right creative mode of the foundational model. For instance, a specific music theory can be evoked to immediately adjust the model by asking it to impersonate an expert in a field associated with a specific music theory. To give one example, a user can ask the model to act like an expert in transformational theory by David Lewin, a strictly mathematical method of music analysis. The model can then help the user analyse their musical fragment within the framework of that methodology and it can further use that knowledge to generate a coherent musical continuation. While there is of course a risk that the model does not know a specific theory, or only simulates to know a certain theory well, for an expert in that theory itself it will be easy to recognise the hallucinations of the model and to intervene accordingly. At the end of the day, interacting with an LLM is like interacting with a random person (Mollick, 2024, Introduction: Three Sleepless Nights). At the beginning we don't know what the specific beliefs, knowledge and skills that person has, but one thing that we can rely on is that given enough time, we can learn each other better and learn to work with each other on gradually more aligned common goals.

AI Agents Controlling Music Environments

In the second class of my course at the Iceland University of the Arts, some of the students expressed their interest in using the reasoning capabilities of LLMs from within SuperCollider—a platform for audio synthesis and algorithmic composition in which users write code to design, control, and manipulate sound in real time, enabling the creation of dynamic musical structures, custom instruments, and complex generative systems. In their creative practice, the students used SuperCollider primarily in live-coding performances, where they write and execute code in real time to produce live music, with this code often projected visibly to the audience as part of the performance itself. For this reason, switching back and forth between a browser-based ChatGPT interface and their live-coding environment was impractical. Furthermore, they expressed a desire for a more integrated interaction, where the

outputs of the reasoning model could directly influence their SuperCollider code without interrupting the performance flow. To address this need, I developed a simple SuperCollider extension that enables direct API communication with OpenAI's ChatGPT models from within the SuperCollider environment. This extension allowed students to send prompts and receive responses programmatically, making it possible to integrate large language model outputs dynamically into their live-coded performances without leaving the SuperCollider interface. For their final concert, students used the extension to let ChatGPT dynamically decide on the semantic meanings of their natural language inputs in the console visible on the screen to the audience, to dynamically control the sound parameters of their live-music performances.

Because LLMs and agents can be accessed through APIs, it becomes possible to integrate them into other software environments commonly used in experimental and professional music production. Similarly, in a MAX/MSP environment, an agent could be employed to generate control sequences, manipulate audio streams, or dynamically reconfigure patches based on high-level commands provided by the user. This integration bridges generative reasoning with real-time computational environments, enabling fluid, on-demand augmentation of creative processes. Looking further into the future, these developments point toward a generalised paradigm of operator-like interfaces, where reasoning-enabled agents could coordinate and control a wide variety of software tools across an operating system. Rather than scripting or programming each individual task manually, users could delegate complex, multi-step operations to agents that understand high-level goals, learn tool-specific protocols, and autonomously compose workflows across diverse applications. In such a vision, the computer becomes a dynamically reconfigurable creative partner, responding to abstract artistic intentions through concrete, tool-mediated actions executed with the support of large-scale reasoning models.

Because the environmental and economic costs of AI tools have become a growing concern, especially in educational contexts, this workshop was also designed to foreground such considerations. We discussed with the students how large-scale language models can entail significant energy consumption and corporate dependencies, and we deliberately investigated how to work within these constraints in a responsible way. To this end, we used the *gpt-4o-mini* model, selected for its efficiency and minimal resource footprint. Over more than a week of continuous

artistic exploration by the entire group, the total API usage amounted to less than one US dollar. This experience demonstrated that conceptually ambitious and musically innovative work with AI can be achieved at minimal computational and financial cost. It also modelled for students how to critically assess the resource implications of their own creative tools.

While these developments dramatically expand the creative possibilities for composers, it is important to maintain perspective on the current state of AI technology. As AI pioneer Yann LeCun emphasises, Large Language Models—despite their impressive achievements—are not themselves a path toward human-level intelligence (Plumb, 2024). They still lack fundamental aspects of reasoning, planning, and understanding necessary for true autonomy, and no amount of scale alone will bridge this gap. Nevertheless, the methodologies explored in this chapter—building collaborative workflows with text-based AI agents, developing contextual strategies for reasoning, and creatively reconfiguring generative models within musical environments—remain artistically valid regardless of whether future AI systems achieve more advanced forms of intelligence. These strategies will remain applicable to any form of text-based collaboration with future intelligent systems, and, perhaps more humorously, with actual human collaborators as well.

Artists as Operators

The notion of the artist as operator of tools emerges with clarity within contemporary practices of machine learning-based creation, where authorship, materiality, and process are fundamentally reconfigured. Iannis Xenakis described the role of a composer working “with the aid of electronic computers” as “a sort of pilot: he presses the buttons, introduces coordinates, and supervises the controls of a cosmic vessel sailing in the space of sound, across sonic constellations and galaxies that he could formerly glimpse only as a distant dream. Now he can explore them at his ease, seated in an armchair” (Xenakis, 2001 [1963], 144). As Tilford (2023, 142) suggests, the appearance of a “Subject” in contemporary generative systems does not originate from a stable source of meaning but instead indexes “the emergence of an operator” working with randomness as its material. In this view, the operator is neither the originator nor the result but the local status of an ongoing process—a temporary configuration within an indeterminate field. Composer Holly Herndon—known for her work with emerging technologies and especially AI and blockchain—similarly articulates an accumulative logic of operator-practice when describing her trajectory across projects: with each new body of work, she does not rebel against previous methods but gathers accumulated tools and skills to extend and transform her practice (Clancy, 2022d, 46–47). This continual aggregation and modulation of technological means embodies an operator-function in the strongest sense: not the repetition of known procedures, but the dynamic accumulation of operational capabilities in response to new experimental terrains. Audry (2021, 35) further emphasises that artists working with machine learning must embrace the unpredictable: while algorithms are precise in their computational structure, their generative behaviour—particularly in machine learning contexts—emerges as unstable and contingent due to stochastic processes, data dependency, and nonlinear interactions that resist full control. Rather than imposing rigid control over outcomes, artistic practice in this context becomes an engagement with volatility, an ongoing negotiation with systems that resist full determination.

In this framework, Paulo de Assis’ method of archaeological, genealogical, and creative investigation of musical material provides a powerful operational model for how composers-as-operators can structure their engagement with AI systems. De Assis’ methodology unfolds in three steps. First, the archaeological phase collects all available material traces of a musical work or practice, categorising them into six types

of strata: substrata (background documents and conditions preceding the work), parastrata (materials generated during the creation process), epistrata (derivatives, interpretations, and analyses over time), metastrata (future reinterpretations and extensions), interstrata (hybrid materials resulting from contamination between strata), allostrata (external materials not directly related but capable of productive interaction). Second, the genealogical phase maps the relations between these strata, tracing connections, influences, and singularities—points of high energy where creative tensions cluster and transformation is possible. Third, in the problematisation phase, selected materials are arranged into experimental arrangements that actively challenge the established orders of the work, generating new actualisations rather than restoring or reifying an original (de Assis, 2018, 109–111).

This methodology can be applied with the iterative generative workflows described earlier in the chapter. When working with AI models, composers first generate large pools of material (analogous to assembling strata), then map and classify these materials according to their stylistic, structural, or experimental potentials (analogous to the genealogical phase). Finally, they intervene through editing, continuation, and recombination, producing new experimental assemblages that problematise the original materials, not by reconstructing their sources but by opening them to divergence, mutation, and further transformation. In both cases—whether operating on historical musical materials or AI-generated outputs—the operator engages in a practice that refuses stabilisation and definitive closure. Instead, the composer-operator curates fields of virtuality, navigates unstable topologies of musical possibilities, and constructs provisional assemblages that expose rather than resolve the tensions latent within the material. Through this continuous process of divergence, genealogy, and reconfiguration, the composer reclaims agency not by controlling the creative field, but by sustaining its metastable openness.

Latent spaces—those internal representations where AI models organise relations between concepts—offer novel ways of seeing, exposing unexpected connections and contextual proximities between ideas, sounds, or structures (Serpentine Arts Technologies, 2024, 65). Yet navigating these spaces demands an operational craft: the ability to formulate effective prompts, to guide the model’s attention, and to shape emergent outputs with precision and intentionality. As *Future Art Ecosystems 4* observes, prompt engineering is an evolving craft-based skill and a form of conceptual labour, akin to querying an unstable search space (Serpentine Arts Technologies, 2024,

93–94). The artist-operator must therefore remain vigilant, embodying what Mollick (2024) calls the "human in the loop"—the critical presence necessary to detect hallucinations, errors, and misalignments within AI outputs (Part 1, Chap. 3). Without this active vigilance, there is a real danger of "falling asleep at the wheel", allowing AI systems to silently erode human expertise, skill development, and the meaningful labour of creative work itself (Mollick, 2024, Part 2, Chap. 6). In a landscape where AI offers universally accessible shortcuts, compressing complex processes into instantaneous outputs, the value of deliberation, critical thinking, and aesthetic judgment becomes even more urgent (Mollick, 2024, Part 2, Chap. 5). While some creative tasks may be delegated entirely to AI—what Mollick terms "Automated Tasks"—others, especially those tied to personal values, deep meaning, or ethical stakes, must remain within the domain of what he calls "Just Me Tasks", activities reserved for irreducibly human engagement (Mollick, 2024, Part 2, Chap. 6). The future of artistic agency, then, does not lie in rejecting AI technologies but in mastering how to live, work, and create with them critically, retaining and reconfiguring human expertise precisely where it matters most.

Seen through this lens, the artist as operator is redefined. The role of the operator expands into a field of new responsibilities: to navigate latent spaces with sensitivity, to craft interactions through prompt-based conceptual labour, sustain critical vigilance against automation's numbing effects, and to deliberately reserve domains of irreducible human engagement. As Mollick notes, the introduction of calculators once provoked fears like today's anxieties about AI; yet rather than abolishing the need for arithmetic or critical thinking, it re-situated these skills within new cognitive ecologies (2024, Part 2, Chap. 7). Similarly, AI will not eliminate the necessity of learning, writing, or thinking critically—it demands their intensification. The imperative becomes clear: to teach and cultivate not only artistic techniques but AI literacy, critical judgment, and operational acuity (Mollick, 2024, Part 2, Chap. 7). Moreover, at a planetary scale, the potential benefits of AI—such as massively expanding access to education and knowledge in underserved regions—constitute a profound ethical mandate (Mollick, 2024, Part 2, Chap. 7). In this view, to refuse the development and intelligent use of AI would be to abdicate the responsibility to address urgent global needs. The task of the artist-operator, then, is neither nostalgic resistance nor blind acceleration, but the careful invention of new practices: to operate within a world reshaped by artificial intelligence without relinquishing the depth, intensity, and responsibility of human creation.

Chapter 2: Artists as Curators

Part 1—Creativity and Problematisation

The Creative Act of Problematisation

Artistic practice differs fundamentally from engineering or other problem-solving paradigms in that it does not begin with a predefined problem to be solved. Instead, art generates *fields of problematisation*—open-ended contexts that provoke questioning, uncertainty, and reflection. Rather than seeking to eliminate ambiguity by finding solutions, the artist deliberately creates situations that *make things problematic* in a productive sense. As Sofian Audry argues, the role of contemporary art is not to provide solutions but to generate problems—problems that estrange, provoke, and reframe perception. “Contemporary art”, writes Audry, “was not so much about solving problems as it was about creating problems for the viewers by bringing them into an experience, allowing the revelation of otherwise unfathomable truths about the world through its estrangement” (Audry 2021, 22). In sharp contrast to the optimisation paradigms that dominate AI and computational research, where goal-orientation and quantifiable metrics define success, art resists such instrumentalisation. Traditional optimisation approaches, Audry insists, are “fundamentally inadequate when applied to the arts” (23), because they presuppose that artistic value is both measurable and universal—a view that misunderstands the contextual, situated, and subjective nature of artistic judgment. Attempting to model artistic relevance statistically, they reflect, falsely assumes that more widely appreciated works are inherently better—a fallacy grounded in what they call an “inaccurate premise about art” (23). The very notion of a “best” painting or “optimal” song collapses under the plural, contextual, and non-teleological values that define artistic practices. In this view, art is not an optimisation problem, nor can it be reduced to an “objective function” without doing violence to its critical and experiential core (25).

Within artistic research discourses, this approach to art-making has been theorised explicitly as a mode of knowledge production in the work of Paulo de Assis, who emphasises that artistic research can turn stable objects into “objects for thought” through a dual process of historical interpretation and experimental reconfiguration (de Assis 2018, 132). In one sense, problematisation involves looking backward—

excavating the assumptions and histories that underlie a given practice—but equally it involves looking forward, “search[ing] for new ways of productively exposing [things] within a contemporary situation” (de Assis 2018, 132). Rather than asking *what something is or was* in a fixed sense, the artist asks *what it could become*: how might a musical composition, an image, or a cultural ritual be re-imagined under new conditions, and what questions would that raise? In this way, the work of art performs a *creative act of problematisation* (de Assis 2018, 132), taking existing materials or ideas and projecting them into new, unresolved futures. Seen in this light, artistic creation is less about achieving a *solution* and more about framing a *situation*. The artist draws attention to blind spots and assumptions, inviting audiences into a space where meanings are unsettled and open for re-examination. Problematisation, in this sense, differs from critique: it is not simply about revealing flaws or inconsistencies, but about constructing new conceptual and experiential fields in which alternative relations can emerge. A composer might, for example, present an everyday object in a strange context, or orchestrate an experience that defies the audience’s expectations. In doing so, artistic practice challenges viewers to confront their own certainties and engage in critical reflection. The value lies in the experience of uncertainty and the dialogue that uncertainty engenders. To problematise is to take something ordinarily seen as stable or given and to render it uncertain – to turn it into a question, a site of debate or exploration. Art’s role, from this perspective, is to pose new questions and fundamentally alter how an issue is understood.

By framing artistic creativity in this way, we acknowledge the inherently combinatorial and exploratory nature of the process. The artist curates a set of elements—whether images, sounds, gestures, or concepts—and arranges them in such a way that their interplay raises issues rather than resolving them. Every artwork, then, can be seen as a kind of problematic field constructed for the audience’s engagement. This does not mean that art is “problematic” in the colloquial sense of troublesome; rather, it means that art problematises reality by challenging the viewers’ habitual frames of reference. The artworks of the Dadaists, which placed everyday objects in galleries, or contemporary interactive installations that require audience participation, do not present clear theses. Such works operate as open systems of sense-making that resist closure or quantifiable judgment. Instead, they provoke us: Why is this ordinary object presented as art? What is the role of the observer in completing the work? Such questions, without singular answers, are precisely the point.

Curating as a Creative Practice in Contemporary Art

In contemporary art discourse, the curator's role is increasingly understood as a form of creative authorship. Paul O'Neill (2012)—curator, artist, writer and educator—observes that since the late 20th century the curator has emerged as an “author” or co-producer of meaning within exhibition-making, not simply an overseer of collections. In other words, curators are now seen as key agents in cultural production who actively shape the values and narratives around artworks. O'Neill argues that curatorship today is “a far-reaching category... encompassing various organisational forms, cooperative models and collaborative structures”—even including “generative properties traditionally attributed to artistic production” (O'Neill, 2012, 89). This blurring of boundaries between curating and creating suggests that the curator's work can be as inventive and influential as that of the artist, constructing new contexts and interpretations for art.

Importantly, recent theory distinguishes “the curatorial” from simple exhibition management, emphasizing “the curatorial” as an active process of meaning-making. Curator and theorist Maria Lind defines “the curatorial” as “a more viral presence consisting of signification processes and relationships between objects, people, places, [and] ideas...that strives to create friction and push new ideas” (Lind, 2012, 20). The curatorial approach is about building connections and contexts: linking artworks with each other and with audiences, and provoking new insights through those relationships. In this view, the curator is a mediator and context-maker who foregrounds relationships over objects and processes of interaction over static products. By crafting the conditions in which art is encountered—through careful juxtaposition, thematic framing, and orchestrating the viewer's experience—the curator actively co-authors the meaning that emerges. Curating thus becomes a creative act of arranging elements in ways that generate dialogue, spark critical questions, or reveal latent possibilities in the art.

This expanded notion of curation has clear historical precedents in experimental art practices since the 1960s. Avant-garde artists began to shift away from viewing an artwork as a self-contained object and toward conceiving art as an open-ended process or situation. For example, the Fluxus movement foregrounded concise “event scores” and instruction-based works that turned everyday actions and audience participation into constitutive elements of the artwork, shifting emphasis from discrete objects to

open situations (Higgins, 2002; Smith, 1998; Ouzounian, 2011). In other words, the artist functions more like a curator of possibilities, designing a scenario in which meaning is co-created by others. Many conceptual, participatory, and performance art projects from the 1960s onward followed this path, prioritizing situations, interactions, and audience engagement over tangible artifacts (Bishop, 2012). These early experiments laid the groundwork for viewing the making of art as inseparable from the curating of context—a perspective now echoed in curatorial theory. The contemporary composer or installation artist who “stages” an event or environment for art, rather than producing a closed work, is in effect adopting a curatorial stance within their artistic practice.

Building on these developments, today’s artists and curators often operate in overlapping roles, especially in settings involving complex, generative processes. The rise of digital and algorithmic art, for instance, further accentuates the curator-like agency of the artist. Rather than fabricating a single artifact, an artist working with generative software or AI might design a system, set parameters, and then guide or select from the output – effectively curating the results of an autonomous process. In such cases, the creative labor lies in configuring the conditions and making editorial choices about what to present, much as a curator would. As media scholars note, even the most “automated” art still relies on human orchestration: artists choose the data or inputs, adjust the algorithms, and crucially decide which outcomes merit public display (Hertzmann, 2018). The role of the artist thus converges with that of the curator, emphasising facilitation and selection within a field of possibilities.

Recombining the Creative

In September 2024, I attended the opening of an exhibition *Musica ex Machina: Machines Thinking Musically* curated by Sarah Kenderdine, Martin Rohrmeier, Paul Kenderdine, and Jonathan Impett at EPFL Pavilions in Lausanne, Switzerland. The exhibition, and the accompanying conference featuring Jennifer Walshe, George Lewis, François Pachet, Miller Puckette, Laetitia Sonami, Bob Sturm and others, offered a striking confrontation with the long history and emerging future of algorithmic thinking in music. Rather than merely illustrating machines that produce sound, the exhibition traced the historical entwinement of musical creativity and systematic processes—from Guido d’Arezzo’s Guidonian Hand to Xenakis’ *UPIC*

system and contemporary AI installations like *Polyp* by Marek Poliks and Roberto Alonso Trillo (Kenderdine et al., 2024, section 1), where AI agents mutate their sonic identities in response to environmental sound. Walking through its four thematic clusters—1) Symbols, Spaces & Algorithms, 2) Automating the Human, 3) Music as Information and Data, and 4) Body, Mind & Machine—I found myself surrounded not just by artefacts, but by provocations. Each work proposed a different answer to a question: What do we mean when we call something "creative"? Whether through Athanasius Kircher's 17th-century *Arca Musarithmica* (Kenderdine et al., 2024, section 10.1), which treated composition as a rule-based operation, or George Lewis's *Voyager*—a now-classic example of a real-time interactive composition built as a rule-based expert system (Kenderdine et al., 2024, section 29.3)—the exhibition made clear that creativity has never been the exclusive domain of human spontaneity. It is—and always has been—a matter of systems, constraints, interactions, and transductions.

The exhibition's curatorial logic resonates with a broader critique of how creativity is conceptualised in the age of artificial intelligence. One recent articulation of such a critique can be found in the work of Keith Tilford, whose contribution to *Choreomata: Performance and Performativity After AI* (2023) engages directly with debates about AI and creative agency. Tilford argues that many claims about the limits of machine creativity remain bound to outdated philosophical frameworks—especially the Kantian conception of genius, which frames creativity as spontaneous, unteachable, and incommunicable. This legacy, Tilford suggests, continues to shape artistic and cultural assumptions. The widespread belief that AI systems cannot be truly creative often rests on an implicit premise: that creativity must be ineffable and irreducible to rules—something that machines, by definition, cannot possess. Yet as AI systems increasingly simulate imaginative behaviour once seen as uniquely human, this assumption becomes unstable. "To claim that AI cannot exhibit creative behaviour", Tilford writes, "one must have an implicit assumption of what creative behaviour is"—an assumption still "exceptionally enmeshed" in both critical art discourse and broader cultural anxieties about authorship and automation (Tilford 2023, 125).

In response, Tilford calls for a closer examination of the term creativity itself, a word that has become central to contemporary discourse yet remains historically contingent and ideologically loaded. To make this point, he draws on two historical sources: Vern Neufeld Redekop and Thomas Ryba's *René Girard and Creative Mimesis* (2013) and Władysław Tatarkiewicz's *A History of Six Ideas* (2012 / 1980). From Redekop and Ryba,

Tilford cites the Oxford English Dictionary's record of the term's earliest use in English, attributed to A.W. Ward in 1875, who described Shakespeare's poetic spontaneity using the newly coined word "creativity" (Tilford 2023, 133; Redekop & Ryba 2013, 8–9), and that the philosopher A.N. Whitehead later gave the term metaphysical depth by using it to describe the "temporal passage to novelty" as a fundamental principle of reality (Tilford 2023, 133).

Tilford then turns to Tatarkiewicz for a longer historical view. According to Tatarkiewicz, for nearly a thousand years the term *creator* was used exclusively in a theological sense—as a synonym for God. The Greeks had no equivalent term, and although the Romans used related expressions, these were mostly colloquial (Tatarkiewicz 2012/1980, 250–251; cited in Tilford 2023, 134). It was only in the nineteenth century that 'creator' began to be applied to artists, and only in the twentieth century that creativity expanded into broader domains—science, politics, technology—eventually becoming a ubiquitous marker of originality and innovation. As Tilford observes, "everyone and anyone, even organizations, or things in the case of disembodied synthetic intelligences, might come to be viewed as creative" (134).

This semantic expansion, he argues, has left the concept of creativity overextended and underexamined. "We desperately need a more specific language for the term creativity", Tilford writes, one that accounts for its "cognitive, social, and ideological entanglements" (130). Rather than treating creativity as a stable or universal category, we must ask: what operations, values, and exclusions does the term encode—and what kinds of labour and agency does it recognise or obscure, especially in an age of artificial actors?

The idea that creativity could be modelled, abstracted, and reproduced by machines challenges assumptions that it is uniquely human, ineffable, or inaccessible to formalisation. Advocates of computational creativity argue that what we perceive as novel or original can often be traced to processes of recombination, pattern recognition, and transformation—processes that, in principle, can be simulated algorithmically. Yet this shift also raises questions about what gets lost when creativity is operationalised as a technical function. To what extent does such a move reproduce existing cultural assumptions, flatten the social complexity of creative labour, or obscure the structures of value and authorship through which creative acts gain meaning? Rather than settling the debate, the computational turn reanimates older tensions between system and spontaneity, rule and exception,

automation and invention.

It is in this context that Margaret Boden's influential typology of creativity becomes particularly useful, offering a framework to distinguish between different modes of creative production—both human and machinic. Boden distinguishes three types of creativity: combinatorial, exploratory, and transformational (Boden, 2010, 41). Each describes a different mode by which novel ideas or artefacts emerge, with distinct implications for both human and machine creativity. Combinatorial creativity involves the novel juxtaposition or integration of elements from different, pre-existing conceptual spaces. In music, this might manifest as genre-blending or hybrid forms—for example, combining traditional Balinese gamelan textures with Western electronic music. It relies on recombination, often producing surprising results by bridging disparate contexts. Exploratory creativity operates within an established conceptual or stylistic space, generating new instances that conform to existing rules or constraints. This is the type most commonly associated with AI-generated music: systems that, trained on vast datasets of stylistically unified material, produce novel but stylistically coherent works. Examples include automatic composition systems designed to emulate Baroque fugue or jazz improvisation. Transformational creativity, by contrast, involves altering or reconfiguring the very structure of a conceptual space—changing the rules of the game rather than playing within them. Historically, this has given rise to paradigm shifts: the invention of twelve-tone serialism, the use of silence as material in John Cage's *4'33"* (Gioti, 2021, 27), or Xenakis's use of stochastic processes. Transformational creativity is the rarest and most radical of the three, and it raises the most provocative questions when considered in relation to AI: can computational systems truly invent new concepts?

As Artemi-Maria Gioti observed in 2021, most automatic composition systems to date—including both rule-based and machine learning approaches—were fundamentally designed to imitate, making them inherently exploratory (27). The crucial question, as Gioti has put it, was no longer simply whether computers can be creative, but whether they can be transformationally creative (28). However, this boundary is increasingly porous. With the rise of large language models and generative agents capable of operating across modalities and domains, we may now be witnessing a shift toward algorithmic systems capable of transformational creativity—generating not only new musical content, but new logics of musical production. Ethan Mollick (2024) observes that large language models function

as “connection machines”—tools designed to link disparate ideas and generate novel associations through probabilistic recombination. Drawing on innovation theory, Mollick reminds us that breakthroughs often emerge not from ex nihilo invention but from the creative recombination of existing concepts. LLMs excel precisely at this: by statistically predicting the next word in a sequence, regardless of how strange or unprecedented the context may seem, they traverse conceptual boundaries and generate unexpected juxtapositions. Their capacity to forge connections between seemingly unrelated tokens makes them not only generative, but potentially transformative—particularly when coupled with randomness, prompting outputs that surprise even their creators. If, as Mollick argues, “new ideas do not come from the ether”, then the question is no longer whether machines can be original, but whether their recombinatory logic can be harnessed as a system for compositional innovation.

Computer music researcher and composer David Cope, best known for developing the system *Experiments in Musical Intelligence* (EMI), argues that recombination is not merely a technical strategy but a fundamental principle of musical creation. In his book *Virtual Music*, reflecting on the design of his EMI, Cope writes: “After all, composers compose recombinantly. I use this term deliberately, since I believe *Experiments in Musical Intelligence* uses processes of recombinance similar to those that human composers use to compose. I believe as well that these same processes create the meaning we hear in music” (Cope, 2001, 89). In Cope’s view, compositional creativity is inherently recombinant—less about spontaneous invention than the reconfiguration of stylistic materials drawn from a given musical corpus. EMI’s ability to produce stylistically coherent works “in the style of” canonical composers foregrounds a model of creativity grounded in variation, adaptation, and contextual transformation. This view aligns not only with combinatorial creativity as defined by Boden, but also with a broader lineage of algorithmic composition stretching from Mozart’s *Musikalisches Würfelspiel* to 20th-century montage aesthetics. It also underscores a broader epistemological claim: that musical meaning itself arises not from originality in the abstract, but from the situated recombination of known elements into new and recognisable forms. Rather than undermining creativity, such recombinatory systems make it legible—suggesting that novelty and familiarity operate not as opposites, but as co-constitutive vectors of musical sense.

At this threshold, however, another conceptual pivot becomes necessary—one that

shifts the very terms by which we think about aesthetics and creativity in relation to machines. As the philosopher of computation and media theorist Luciana Parisi (2023) argues, dominant discourses around AI creativity tend to reinscribe the logic of humanism by demanding that machines mimic human forms of consciousness, sentience, or expression to be granted creative legitimacy. Such frameworks, she warns, ultimately reaffirm the colonial and patriarchal metaphysics of modern thought, where the nonhuman is made intelligible only through its resemblance to the human (88–90). Against this, Parisi calls for a non-standard aesthetics rooted in “non-performance” and the embrace of what she terms “incomputables”—modes of machine operation that resist philosophical capture and exceed the grammar of human sensibility (74–79). Rather than seeking equivalence between AI and human agency, she urges us to attend to the expressive potentialities of machinic systems in their own right—as aesthetic processes that unfold through alien logics of recursion, error, and negation. This reframing resonates deeply with the agenda of posthuman composition: not to redeem AI by demonstrating its capacity to imitate human creativity, but to explore how machines, algorithms, and networks enact creative processes that reconfigure the aesthetic field itself. In this view, creativity becomes not a faculty but a distributed operation, one that cuts across ontologies and agencies. Such a shift requires us to reintroduce an awareness of the nature of machines and their relations to humans and values—a task that is as aesthetic as it is philosophical.

Even this, however, is not enough. As both Gioti (2021) and Rohrmeier (2022) remind us, creativity cannot be defined solely in terms of novelty. The “standard definition” in creativity studies insists that to be creative, an artefact must also exhibit value—whether in terms of usefulness, significance, or impact (Rohrmeier, 2022, 51). But this raises a thornier problem: who determines value, and by what criteria? Boden acknowledges this ambiguity, noting that while evolutionary algorithms may generate novel stylistic spaces, they often fall short of producing outputs perceived as valuable or meaningful (Gioti, 2021, 28). The problem of value thus remains unresolved, especially in computational creativity, where the processes may be novel but the contextual significance of their outcomes is difficult to assess without human judgment. For Gioti, this further underscores the need for systems that can trace and localise agency—systems that allow us to understand how creative contributions emerge and are attributed, particularly in collaborative or distributed contexts.

Questions about the computability, structure, and systematisation of creativity have long shaped 20th-century compositional thought, where composers increasingly grappled with the tension between intuition and system, spontaneity and formalisation. Already in the 1940s, composer and music theorist Joseph Schillinger described the evolution of art—specifically music—as a progressive cycle from mimicry (passive transformation), through magic (active transformation), to engineering (scientific transformation). These stages, he argued, correspond not only to aesthetic sensibilities but to the transformation of artistic media and the roles of creators themselves. From the bodily production of sound to the invention of instruments, and eventually to electronic and computational music, Schillinger tracked a dematerialisation of both sound and subjectivity: from vocal cords to algorithm, from improvisation to automation. What begins as biologically induced pattern-making becomes, over centuries, formalised into compositional theory—first rule-based and then increasingly abstract, mathematical, and generative. In Schillinger’s vision, the composer gradually transforms from an intuitive performer into a kind of scientific creator, eventually designing systems for the automatic production of music (Schillinger, 1948, 11). This conceptual shift prefigures our contemporary moment, where AI not only executes predefined stylistic instructions but also participates in the design of new aesthetic systems, styles, and rule sets. If Schillinger’s “engineering” stage corresponds to the emergence of scientifically modelled, self-operating compositional systems, then today’s generative AI may mark its intensification: a mode of creation that is neither purely mechanical nor magical, but machinic—a system of operations capable of producing novelty within and across conceptual regimes.

In his Collège de France lectures, composer Pierre Boulez questions whether musical works are ever more than crystallisations of “a series of accidents, a series of choices that in different circumstances might have produced completely different results” (Boulez 2018b, 119). He frames composition as emerging from a “chaos of intentions”, where chance is never fully eliminated but rather filtered—domesticated—through systems of constraint that give the final outcome its coherence. Whether the composer embraces chance or seeks to annihilate it, Boulez notes, they must construct a method to manage it. This paradox—where even randomness must be formally structured—exposes the limits of humanist myths that associate creativity with free will, spontaneity, or pure intentionality. Indeed, Boulez’s insight gestures toward a broader dismantling of the binaries that have long underwritten Western accounts of artistic

subjectivity: freedom versus constraint, rationality versus irrationality, author versus medium. As he puts it, what appears as irrational is always the expression of a deeper rationality “that eludes us”, while rationality itself remains a “limited expression of a transcendent order” (120). This dissolution of oppositions challenges the very notion of the autonomous creator and invites us instead to see the composer as a node within a larger structure—mediating between logics, constraints, and emergent forces. In this light, the artist’s task is not to express an unutterable interiority, but to assemble a system—a machine—that can make decisions, mutate, and evolve according to its own logic. Creativity, in other words, is not the negation of system, but its modulation.

One example of such a dynamic interplay between constraint and expression can perhaps be experienced in a composition *Ultrachunk* by Memo Akten and Jennifer Walshe. Listed in the Body, Mind & Machine section of the exhibition in Lausanne, the piece stages a vivid confrontation with machinic subjectivity. Neural networks trained on vast corpora of audio-visual content deconstruct and recompose Walshe’s own voice and image into a grotesque, unstable hybrid—one that oscillates between intimacy and estrangement, coherence and collapse. *Ultrachunk* (Kenderdine et al., 2024, section 41.4) foregrounds the tension between personal identity and machinic remix, offering a performance that is both deeply familiar and eerily alien. Yet this is not simply a matter of speculative storytelling. As Davor Vincze (2023) argues, works like *Ultrachunk* exemplify a broader aesthetic strategy in contemporary AI art: an embrace of “the musicality of imperfection” (308), where glitch, mismatch, and unpredictability are not failures to be corrected but expressive resources to be exploited. He further references the point Walshe once made in their private conversation: “we are lucky to be creating art in the time that artificial intelligence is just nascent”. Vincze further writes that “[Jennifer Walshe] believes that it gives us the freedom to play with the imperfection of this emerging technology, much like how artists in the 1950s could naively experiment with any kind of electronic sound just because it was unheard of at that point in time” (Vincze, 2023, 330). As Vincze notes, while AI development in commercial contexts often strives for precision and predictability, in the arts it is precisely these system failures—glitches, mismatches, and errant behaviours—that can become the most potent sources of creative expression (308).

Part 2—Composition, Agency, and Assemblage

Musical Works as Assemblages

Traditionally, especially in Western classical music, a *musical work* has been treated as a fixed entity: the composition inscribed by the composer in the score. According to this ontological account, the score is the sovereign text, and performances are merely realisations or reproductions of the composer's intent. The work is regarded as having a stable identity – one “correct” form as defined by the authoritative score and by the composer's original vision. This view aligns with the Romantic idea of the solitary genius: the work is a product of the composer's mind, and performers are expected to be faithful to that vision. However, musicologists and artists have long noted problems with this view. Performers and listeners invariably bring their own interpretation; instruments and acoustics change the sound; editions differ; over time, even the most canonical works undergo re-evaluation and alteration in practice. In the previous sections, we explored how artistic practice operates through the creation of problems rather than the pursuit of solutions, and how creativity often unfolds through processes of recombination, transformation, and the construction of new relations. Rather than expressing fixed meanings or solving predefined tasks, artworks function as open-ended fields—provocations that invite interpretation and reconfiguration. Such understanding of art naturally leads to a view of musical works as configurations of diverse elements—materials, methods, histories, and choices—held together temporarily in a particular form.

Paulo de Assis, in his critical redefinition of musical works, addresses this issue by proposing a new image of the musical work as an assemblage rather than a permanent structure. Drawing on philosophies of difference (notably Deleuze and Guattari's concept of *agencement*, later translated as *assemblage*), de Assis describes musical works as “complex arrangements of aesthetic-epistemic things, forces, intensities, and signs, which establish several superposed networks of historical, cultural, material, symbolic, and psychological dimensions” (de Assis 2018, 100). In this description, a musical work is not just the notes on the page: it includes the sounding result (frequencies, intensities), the historical traditions it carries, the material aspects (instruments, scores, technologies), the symbolic meaning or affect it produces, and even the psychological impressions on performers and listeners. All these layers together constitute the work. De Assis elaborates that works “cease to be conceived as

a set of instructions or as an ontologically well-defined structure. They become reservoirs of forces and intensities, dynamic systems characterised by metastability, transductive powers, and unpredictable future reconfigurations” (de Assis 2018, 100). The term “metastability” here implies that a work has a certain coherence or identity, but it is not in static equilibrium – it can change phase under new conditions. “Transductive powers” suggests that works can transform and generate new forms (transduction being the process of transforming one form of energy or information into another). Crucially, the future of the work is “unpredictable”; each performance or re-interpretation can potentially reconfigure what the work is, adding to its assemblage of meanings and forms.

In practical terms, this assemblage view acknowledges that any musical work is never a single invariant object. It is one thing in the manuscript, another thing in a modern orchestra’s performance, another in a recording on period instruments, another in a remix or quotation in a film soundtrack. It changes its meaning in socio-political contexts, in various interpretations, in the affective responses of listeners, and in the evolving technological and cultural infrastructures through which it is mediated. All these instantiations participate in the *assemblage* of what the work becomes as a cultural artefact. Rather than privileging one “authentic” version over others, this perspective recognises the work as a dynamic system with “unpredictable future reconfigurations” (de Assis 2018, 100).

The logic of assemblage carries several consequences that de Assis outlines: it favours *multiplicity over unity, event over essence, becoming over being*, and embraces “informed inconsistency” over the illusion of certainty (de Assis 2018, 101). In other words, instead of seeking one unified essence of a work, we recognise the work’s multiple facets and instantiations. Instead of seeing the work as a static being, we see it as an event that happens anew in each edition, performance, listening, or critical reflection (a process of *becoming* each time). And rather than cling to the certainty of an authoritative version, we accept that a work might be internally inconsistent or variable—and that this variability, properly understood, is *informative* and rich, not a flaw. De Assis points out that the new image of the work as assemblage can even *include* the old image of the work as a special case: the traditional fixed work becomes “a particular case of less complex combinations of codings and territorializations” within the broader assemblage (de Assis 2018, 100). In other words, what we used to think of as a self-contained work can be seen as one *snapshot* or *layer* of a more

complex, interactive reality. A tightly controlled serialist composition, for example, might look like a fixed structure, but even it can be re-contextualised or heard differently in a new setting—thus it's just a more rigid assemblage, not fundamentally different in kind.

By conceiving musical works as assemblages, we inherently acknowledge the plurality of agencies and elements involved in their becoming. The composer's contribution is just one element (albeit an important one) in the network. The performers' skills and interventions, the instruments' timbral qualities, the acoustics of the hall, the audience's perceptual frame, and the work's intertextual links to other works or cultural references are all active parts of the work-as-assemblage. This diminishes the old hierarchy where the composer-genius was at the top and everything else was secondary. Instead, creation is distributed: the *meaning* and *effect* of the work arise from the interaction of all parts. A simple example is in jazz or improvised music: a lead sheet (melody and chords) provided by a composer is just a skeleton, and the real "work" takes shape in performance through the interaction among improvising musicians. But even in fully notated music, an assemblage view sees each performance as a *reconstruction* of the work—not a mere execution, but a *creative act* that (re)assembles the work's components in a unique way. De Assis notes that "musical entities can't be taken as given; they are rather seen as constantly becoming something else, without ever attaining a fixed state or definition, and they become this something else by means of concrete operations done by concrete individuals" (de Assis 2018, 72–73). Here, "concrete operations" could mean the act of performing, interpreting, or even analysing—any active engagement that *constructs* the work anew. Each such act might reveal "well-known structures" of the piece, but also inevitably highlights "zones of indeterminacy, grey spots, cracks in the structure that no identity-based ontology is able to explain" (de Assis 2018, 72). These indeterminate or flexible aspects are exactly what the assemblage model can accommodate, whereas the old ontology of the work could not.

The implications for creativity are significant. If a musical work is seen as an assemblage, then *creativity is also assemblage-like*. The creation of a work is not a one-time act by a composer, but an ongoing process that includes many creative acts (composition, interpretation, adaptation, etc.) by many agents. The *composer* in this framework becomes less a solitary originator and more a *catalyst* or *curator* of possibilities—setting initial conditions or providing material that will be further

developed by others. Likewise, performers are not passive vessels but active co-creators; they contribute their own artistic decisions to the assemblage. Technology also plays a role: the introduction of a new instrument or sound technology can transform how a work is realised, effectively adding new components to its assemblage. In electronic or electroacoustic music, this is obvious: the software or hardware used is part of the work's identity. But even for acoustic music, recording technology (microphones, mixing) or playback context can alter the assemblage.

De Assis's assemblage theory for music resonates with a broader shift in aesthetics from *object* to *relations*. The "ecological network" around a work includes the work, its environment, and its perceivers (de Assis 2018, 73). He emphasises that object, environment, and subject are all changing in permanent transformation (73). This ecological perspective dovetails with the concept of art generating fields of problematisation: if everything is in flux, then a work of art is not a fixed answer but a perpetual question – each presentation asks anew, "What is this work now? What can it be, here, today?" The work becomes a *site of experimentation*.

In musical practice, we see composers and artists explicitly adopting this assemblage mindset. Some contemporary composers write open-form pieces or invite performers to contribute material, thus decentralising authorship. Remix culture in music treats existing recordings as raw material for new works, blurring the line between composer, performer, and listener (the remix artist is all three at once). Even historically, practices like theme and variations, improvisation, or folk music traditions show that music has always had an assemblage aspect—motifs and structures circulating and reassembling across pieces and generations.

Thus, "musical works as assemblages" can be both a descriptive theory as well as a prescriptive stance for creativity. It describes how works function (as assemblages of components with distributed agency). In doing so, it aligns music with the combinatorial view of creativity discussed earlier: no work is purely the product of one human mind; every work is a mosaic formed in a network of human and non-human actors (scores, instruments, algorithms, etc.). From this perspective, the composer or artist is perhaps less a *maker* in the traditional sense and more an *organiser of forces*—bringing together various elements and setting conditions for something to emerge.

Posthumanist Agential Assemblages

If artistic creativity can be seen as inherently combinatorial, situated, and emergent—built through arrangements of historical, material, and conceptual relations—then the logic of copyright appears not only inadequate but antagonistic to such a view. Copyright law, as Martin Zeilinger argues, is premised on a reductive image of creativity as the product of a singular, original, human author. It “flatten[s]” the complex, dynamic nature of creative production into a legal paradigm that assigns exclusive rights to individuals or their corporate proxies, treating the creative act as an expression of isolated subjectivity rather than a relational or systemic event (Zeilinger 2021, 22). This individualist framework is not only outdated, but actively obscures the collaborative, iterative, and multi-agent nature of creative work, particularly in digital and posthuman contexts. As Zeilinger points out, this conflict is especially acute in AI, where creative behaviour is rooted in technical processes like training on, recombining and adapting massive datasets—processes that “rely fundamentally on copying and reusing”—techniques that copyright regimes tend to regulate or criminalise (2021, 22).

The implications are far-reaching. In contexts where AI systems participate meaningfully in creative production, legal and cultural expectations still presume a human at the centre: a figure of authorship whose rights can be asserted and protected. But as Zeilinger shows, such expectations are symptomatic of deeper anthropocentric assumptions: namely, that agency and authorship must go together, and that both must remain grounded in a human subject. This logic, embedded in intellectual property frameworks, operates recursively. Agency is inferred from creativity, authorship is inferred from agency, and ownership is inferred from authorship (Zeilinger 2021, 48–49). In short: to be creative is to act; to act is to author; to author is to own. When AI enters the picture, however, this recursive chain begins to break down. If we accept that non-human systems can engage in generative processes that produce meaningful artefacts, then we are forced to confront a destabilising possibility: that authorship is no longer the prerogative of human individuals alone. Zeilinger frames this as a “short-circuiting” of the humanist logic through which agency becomes property (2021, 63). When artistic outputs emerge from hybrid human-machine assemblages, it becomes increasingly difficult to determine where intention lies, or whether the output can be meaningfully attributed to any single subject. As creative processes become more distributed between human

and non-human agents, the very notion of a bounded author collapses.

This collapse has legal, political, and economic consequences. In the absence of legal personhood for AI systems, the rights to AI-generated outputs are typically assigned to human users or platform owners. As Zeilinger predicted, this has led not to a radical restructuring of intellectual property regimes, but rather to a conservative assimilation of AI works into existing frameworks, where outputs are simply reattributed to subscribing users or developers (2021, 165). Four years on, this prediction remains accurate: generative platforms like OpenAI's ChatGPT, Google's Gemini or Anthropic's Claude grant usage rights to human users, even when the system's creative contribution exceeds mere tool-like functionality. More troublingly, this legal adaptation masks a broader enclosure of the cultural commons. As Tiziana Terranova—a theorist of digital culture and political economy whose work critically examines how networked technologies reshape labour, knowledge, and power—warns, AI technologies are driving a new wave of digital accumulation, echoing the early 2000s' rise of platform capitalism (Terranova 2023, ix). Instead of fostering collective ownership or open creative ecosystems, generative AI systems—trained on public data—are monetised through proprietary interfaces and licensing schemes. Even opt-out mechanisms (like those proposed by Spawning.ai³) or content watermarking offer only limited resistance; they resemble bargaining chips in a system that remains structurally tilted toward centralised control.

One of Martin Zeilinger's key contributions to this debate is the concept of the posthumanist agential assemblage—a configuration in which expressive agency is no longer centralised in a singular human creator but emerges from the entangled operations of human, machinic, algorithmic, and infrastructural components. Rather than treating AI as a discrete tool or fully autonomous agent, Zeilinger describes these systems as assemblages in which creativity is distributed and enacted relationally. These assemblages are “posthumanist” in that they reject the Enlightenment ideal of the isolated, rational, human author, and instead foreground the dynamic, intra-active

³ Spawning.ai is a platform that enables artists and creators to control whether their works are included in AI training datasets. Through tools like *Have I Been Trained* and the *Do Not Train* registry, it offers opt-out mechanisms allowing creators to request the removal of their content from datasets used by companies and AI models such as Stability AI or LAION. Unlike most major AI developers, who operate under broad data scraping practices and retroactive licensing models, Spawning.ai positions itself as an advocate for creator consent, transparency, and ethical dataset construction. However, these mechanisms remain voluntary and largely unenforceable at scale, especially in jurisdictions where scraping public data is still legally ambiguous, or actually legalised (see Japan).

relationships among a plurality of actants—artists, algorithms, datasets, platforms, legal codes, and audiences alike (Zeilinger 2021, 10, 44, 63). It is within this framework that the tensions between legal ownership, creative agency, and technological mediation must be re-evaluated.

Yet the critical potential of posthumanist assemblages lies precisely in their challenge to this logic. If we take seriously the idea that creativity can emerge from human-nonhuman entanglements—what Zeilinger calls “artificially intelligent agential assemblages” (2021, 63)—then we must also consider new models of authorship, value, and ownership that reflect distributed agency. This means questioning the assumption that expressive agency must solidify into property, and instead exploring alternative frameworks: communal authorship, dynamic attribution, or even the recognition of assemblages themselves as legal or aesthetic agents. In this view, artistic practice becomes a site of experimentation not only in aesthetics but in ontology and law. By foregrounding hybrid configurations of agency, artists working with AI and algorithmic systems expose the limitations of prevailing intellectual property models and hint at the contours of what Zeilinger calls a “posthumanist cultural commons”—a space where creative expression is no longer tethered to the individual subject, but emerges from the collective operations of complex systems.

This reframing of agency invites us to reconsider what is truly at stake in the rise of generative AI. If style can be encoded, authorship redistributed, and creativity externalised into machinic systems, then artistic identity itself risks being flattened—subsumed by algorithmic generalisation. Yet such risks also present strategic opportunities. As *Future Art Ecosystems 4* notes, artists whose aesthetic signatures are captured by large models might treat this not only as a loss, but as a site of reconfiguration: a chance to intervene in the infrastructure of generation itself, to shape the model from within, or to reclaim artistic agency through subversion and tactical redirection (Serpentine Arts Technologies, 2024, 94). In this sense, style capture is not simply a form of expropriation—it is also a prompt to rethink the boundaries between authorship, branding, and aesthetic individuation. Jacques Attali’s provocation that AI is “a tiny facet of the historical journey towards achieving immortality” (Clancy 2022, 4) further complicates the stakes. If the act of composition traditionally offered a form of aesthetic permanence—projecting one’s creative decisions into the future—then the integration of AI into this process does not necessarily erode this power, but reshapes its form. When artists find their style

embedded in the weights of a model, they may be witnessing a new form of “algorithmic preservation”—a condition in which creative identity persists as a generative potential latent within system dynamics. Rather than freezing artistic legacy, such systems render it recombinant, capable of new instantiations and recompositions beyond the original artist’s control.

It is important to acknowledge, however, that “style” has never been a fixed or extractable essence. As Boulez argued in his essay “Aesthetics and its Fetishes”, style cannot be lifted out and reused like a commodity; it results from specific compositional processes bound to their historical moment (Boulez 1986/1961, 32–33). From this perspective, the notion of capturing or transferring style risks misconstruing what is in fact a situated and irreducible act of artistic labour. Yet precisely because generative AI treats style as separable and replicable, it poses urgent questions about how compositional identity is encoded, circulated, and contested within technological infrastructures. This does not validate the commodification of style—but it does demand that artists and theorists confront its new modalities of appropriation, intervention, and potential resistance.

The philosophical backdrop for this shift has been long in the making. As Zeilinger notes, the very concept of “artificial intelligence” implies that intelligence can be decoupled from the human body and instantiated in non-human forms—suggesting a latent capacity for autonomous action, interaction, and interpretation (Zeilinger 2024, 62). But agency, in this expanded sense, is no longer about interior will or intentionality alone. Drawing on Karen Barad’s theory of “agential realism”, agency is redefined as the capacity to affect and be affected within intra-active systems—assemblages of human and non-human actants whose interactions constitute the conditions of emergence (Barad 2007; Zeilinger 2021, 44). From this perspective, AI artworks are not the products of singular minds or isolated systems. They are co-produced by ecologies of training data, computational architectures, human inputs, material constraints, and networked publics. The figure of the solitary creator—already destabilised in contemporary art—is further displaced by what Zeilinger calls the “speculative scattering of agency across human-machinic-algorithmic assemblages” (2021, 45). This scattering does not dissolve responsibility, but redistributes it across a field of interdependent operations.

To accept this reframing is not to embrace naïve techno-optimism. As Barbara Bolt—theorist of artistic research and new materialism—warns, the increasing convergence

of human and machinic agencies raises significant epistemic and ethical questions. In our current socio-technical milieu, algorithms often function as decision-making ensembles, rendering people and representations as data, and assigning value through computational metrics (Bolt 2023, 99–100). These systems may appear neutral or intelligent, but they are deeply entangled with histories of abstraction, reduction, and enclosure.

Posthumanist agential assemblages “do not operate ‘by design,’ simply enacting algorithmic routines that execute predetermined protocols” (Zeilinger 2021, 172). Instead, they evolve through the interactions of their parts, often in ways not fully predictable or controllable. This again mirrors the earlier theme of problematisation and uncertainty: the artist of an assemblage is not entirely sure what will happen, but sets things in motion to explore possibilities. It is an *experimental* ethos taken to a structural level. We can see parallels in earlier art movements and theories: the “open work” concept by Umberto Eco (Eco, 1989), or the indeterminate music of John Cage, or the interactive art of the 1960s – all these gave up some authorial control to gain a richer engagement with chance and participant input. What posthumanist theory adds is a clearer articulation of the role of the *non-human* in this equation (for Cage, chance might be a throwing of the coins in *I Ching*, but for new media artists, chance might come from an AI’s output or a biological process). Posthumanist agential assemblages represent the convergence of the combinatorial, assemblage-based view of creativity with an explicit inclusion of non-human agency and a de-centering of the human. They demonstrate that when artists act more like participants in a network than sovereign creators, new forms of creativity emerge that challenge our definitions of authorship and originality. The artist’s role shifts from *author* to *operator*, and from *operator* to the one who *curates* a set of agents (human and non-human) and conditions, and then allows an artwork to unfold. This is a significant departure from the image of the artist as a lone genius pouring forth a personal vision; instead, the artist becomes a guide of collective, interactive processes unfolding between both human and non-human actants.

Artists as Curators

The previous discussions lead to an expanded conception of the artist's role – one that can be aptly described as *curatorial*. If creativity is combinatorial and if artworks are assemblages of many elements and agents, then the artist's task often resembles that of a curator: selecting, combining, arranging, and presenting elements in a way that generates new meaning and experiences. The phrase “artists as curators” does not simply mean artists who also happen to curate exhibitions (though many do in contemporary practice); it points to a broader phenomenon in which creation itself is an act of curation. In this understanding, the boundary between making art and curating art dissolves. The emphasis shifts from the fabrication of a singular art object to the organisation of relations among elements—whether those elements are objects, people, ideas, or other components.

Ed McKeon—researcher, lecturer, and curatorial producer of experimental practices—suggests the idea of *curatorial composing* (2022), in which the artist's labour increasingly lies not in composing works but in composing public encounters—spaces of appearance where meaning is not given but negotiated. In this sense, the artist as curator is not an overseer of artefacts but a participant in the co-articulation of values, problems, and possibilities—an agent of configuration in a field of emergent relationalities. By analysing the music theatre practice of composer Heiner Goebbels, McKeon introduces the notion of curatorial composing to describe a form of artistic labour that moves beyond traditional notions of composition (as the authorial generation of musical material) and toward the construction of aesthetic situations—settings in which materials, performers, environments, technologies, and audiences are arranged in ways that provoke emergent meanings. Drawing on Goebbels' intermedial and often collaborative practice, McKeon argues that the composer increasingly takes on a role more akin to that of a curator. Importantly, McKeon also traces this shift to broader changes in post-Cagean aesthetics, where the role of the composer becomes one of framing rather than authoring, of making choices about situations rather than prescribing content. What emerges is not a diminished role for the artist, but a re-conceived one: as a cartographer of problematics and a decision-maker whose practice involves tracing the provenance of materials and assembling them into provisional and contingent structures of meaning. The artist is no longer a solitary originator but a composer of encounters, an assembler of assemblages.

In conclusion, viewing artists as curators synthesises the chapter's key arguments. The artist-curator generates *fields of problematisation* by arranging elements in ways that confront the audience with questions rather than answers. Creativity in this role is *combinatorial* and connective, dispelling the myth of isolated genius and instead orchestrating collective, cross-boundary creativity (human with human, human with non-human). Musical works or other artworks become dynamic and ever-changing *assemblages* under the artist-curator's hand—living systems open to continual reinterpretation and change, rather than sealed creations. Posthumanist sensitivity in curation means the artist acknowledges non-human agencies and embraces a de-centered, networked form of authorship. Here, authorship is not the property of an isolated individual but an emergent effect of intra-actions among diverse agents—artists, algorithms, datasets, legal frameworks, and publics. This does not lessen the artist's importance; on the contrary, it is a profoundly skilled task to weave together such assemblages effectively. It requires broad knowledge, intuition, and a willingness to relinquish total control—to let the artwork breathe and evolve.

The artist as curator creates the *conditions* for creativity and critical insight. In doing so, such artists invite us – their audiences and fellow participants – into the heart of the creative process. We become part of the assemblage, part of the meaning-making. And in that collaborative space of not-knowing, of exploration and curating the unknown, art fulfils its deepest purpose: not to confirm what we already know, but to transform how we think, feel, and envision what is possible.

Chapter 3. Artists as System-Builders

Encoding Process, Enforcing Trust

In Chapter 1, we explored the figure of the artist as operator: a practitioner engaging creatively with the outputs of computational tools such as machine learning models, generative systems, and agentic AI. These connectionist technologies introduced forms of automation, unpredictability, and emergence—qualities that artists learn to navigate through iterative exploration. Chapter 2 expanded this view by presenting the artist as a curator of both human and non-human agencies: someone who arranges and configures heterogeneous elements with their own autonomous agency into posthuman agential assemblages. This chapter focusses on artists who are not only users of pre-existing systems and arrangers of relations between various agencies, but also design autonomous systems from the ground up. These are artists as system-builders—those who construct processes that can run independently of their creator.

Part 1 of this chapter—titled *Composing Autonomous Systems*—first discusses the difference between aleatoric and stochastic composition, contrasting John Cage’s use of chance operations—rooted in ethical and spiritual openness—with Iannis Xenakis’s mathematically structured approach to probabilistic control. These positions introduce two divergent understandings of how composers relate to unpredictability and system behaviour. Part 1 then tracks the views of the musical work as a field of possibilities and affordances, and culminates in today’s machine-learning milieu where composers become system-builders responsible for curating data, writing code, and balancing constraint with indeterminacy to shape emergent outcomes. The chapter then expands into the broader field of generative composition, where authorship and musical identity are displaced from singular works to the design of systems that produce variable outputs. Importantly, generative music is presented as a third alternative to live and recorded modes of music dissemination. The text further explores how compositional agency is increasingly understood as the crafting of structured spaces of potential—open systems governed by interaction affordances, codes, and dynamic rules. Crucially, it reintroduces an older but powerful concept—the machine.

To build a machine is to construct a system that functions reliably and autonomously within predefined parameters. The machine, in this sense, is a closed system designed

to execute particular operations with minimal ambiguity. Many artists working in the tradition of algorithmic composition, cybernetics, or symbolic AI embrace this machinic logic. From stochastic music to generative scores, these practices aim to encode procedures so they can be executed consistently and precisely, even if they involve randomness as a variable within a tightly defined system. In such cases, the artist seeks reliability—designing a structure whose behaviour can be trusted precisely because it is calculable.

Yet contemporary artistic practice increasingly scales beyond any single machine or studio. Composers collaborate across geographic locations, delegate tasks to remote performers and AI services, or release executable works into online environments where they evolve without direct oversight. Once the locus of execution becomes distributed, predictable behaviour can no longer be guaranteed by deterministic design alone. A new set of concerns comes into focus: How can one verify that a remote instance of the process really is the process? Who owns or inherits its outcomes? How can contributors be sure the rules they agreed on will be honoured?

When deterministic control gives way to distributed or opaque behaviour, trust becomes the mechanism that supplements reliability. In the context of autonomous and distributed systems, we can distinguish:

- **Trust in process:** confidence that an opaque or adaptive system will act within an expected range, even if its internal workings remain inaccessible or stochastic. Niklas Luhmann—sociologist and founder of social systems theory—framed this as the reduction of environmental complexity to manageable horizons of expectation (Luhmann, 1979/1973, 25-27), emphasising that actors proceed as if the future was sufficiently knowable to risk continued interaction.
- **Trust in record:** confidence that the fact of what was executed—authorship, parameters, sequence—has been faithfully recorded and cannot be altered retroactively. As Diego Gambetta—sociologist noted for his work on trust, cooperation, and organised crime—notes “if evidence could solve the problem of trust, then trust would not be a problem at all”. Cooperation depends on verifiable evidence that past actions occurred as claimed. Without such assurance, strategic partners cannot evaluate reputation or enforce commitments (Gambetta, 1988, 233).

Both forms are pertinent when composers release autonomous or generative works into open, networked environments. Process-trust relates to aesthetic behaviour (will the piece “perform” acceptably when re-instantiated?), whereas record-trust underpins attribution, provenance, and any contractual or economic obligations that may follow from subsequent iterations or sales.

For these reasons, Part 2 of this chapter—titled *Blockchain For The Arts*—shifts its focus from composing as building on the level of algorithmic processes that execute the artistic outcome of a system, to building on the level of algorithmic conditions: conditions that decide whether the underlying algorithmic processes can be performed at all, in which circumstances, and based on what evidence, the actant intending to start such a process is entitled to do so. It is here that the chapter discusses the possible role of blockchain technology in the arts—not only as a financial technology, but as a system for algorithmically enforcing trust.

Where traditional machines aim for reliability through deterministic design, blockchain-based systems enforce agreement through cryptographic verification, decentralised consensus, and smart contracts. In other words, they create machinic structures for managing trust, enabling a transformation of the artist’s role from the artist as curator of agencies to artist as a builder of formal systems on the level of executable agential procedures. For artists working across networks, across institutions, or with autonomous agents, blockchain provides a programmable infrastructure for defining how participants interact, how permissions are enforced, and how provenance is secured and can be traced. It enables compositional work not only with media or operations, but with relationships, responsibilities, and constraints.

This chapter argues that the system-building turn in contemporary art—particularly in algorithmic, generative, and decentralised practices—demands new frameworks for distinguishing between reliability and trust. Artists build machines when they seek reliable automation, but they also look for ways of confirming reliability when engaging with systems and in settings where behaviour and the final outcome cannot be fully anticipated. Blockchain, as will be presented, uniquely allows for the construction of machinic environments in which trust itself can be encoded and executed on the compositional level.

Within the broader context of this doctoral dissertation, this chapter highlights a significant gap: at present, there exists no general compositional infrastructure that enables multiple actants to collaboratively compose on the blockchain in a seamless, accessible manner. While blockchain technology holds promise as a means of enabling such compositional interaction, its current implementations, which will be discussed, remain limited to a small number of specialised applications—each offering only partial and highly specific affordances for artistic participation. Accordingly, the purpose of this chapter’s second part is twofold: first, to provide a critical overview of the current state-of-the-art of blockchain technologies for the arts and for music; and second, to lay the technical conceptual groundwork for the introduction of the concept of Performative Transactions—the primary contribution of this dissertation defined in Chapter 6—a blockchain-based compositional framework that enables diverse actants to transparently and modularly build on each other’s contributions.

Part 1—Composing Autonomous Systems

From Chance to Probability

In 1951, composer John Cage famously turned to the *I Ching*, the ancient Chinese divination text known as the “Book of Changes”, to compose *Music of Changes*. By employing the *I Ching*, Cage relinquished direct compositional control, instead allowing chance operations—determined by coin tosses corresponding to the text’s hexagrams—to dictate the musical material. This approach introduced a radically new understanding of artistic agency: rather than meticulously crafting each detail, Cage allowed an external, non-human system to guide his creative decisions, effectively surrendering traditional notions of personal authorship. As Cage put it: “Most people who believe that I’m interested in chance don’t realize that I use chance as a discipline. \[...\] My choices consist in choosing what questions to ask” (Kostelanetz, 2003, 70). His embrace of chance thus foregrounded unpredictability, introducing indeterminacy as a fundamental aspect of the compositional process.

Traditionally, aleatoric practices in music composition, including Cage’s, have been seen as challenges to Platonist ideals—attacks upon the notion of fixed, ideal structures underlying musical form. In *Formalized Music* (1963), composer Iannis Xenakis explicitly critiqued deterministic causality, rooted historically in Plato’s assertion from the *Timaeus* that “it is impossible for anything to come into being without cause” (Xenakis, 2001/1963, 1). Yet Cage’s turn toward the *I Ching* and Zen Buddhism imbued his work with other metaphysical and transcendental assumptions. Cage’s indeterminacy was not a nihilistic retreat from structure, but an ethical and spiritual exercise in openness. “Instead of talking about [non-intentionality], I wanted to do it”, Cage explained, “and that would be done by making the music nonintentional, and starting from an empty mind” (Kostelanetz, 2003, 35). This surrender to the process of chance—as something meaningful beyond the self—suggests a belief in latent forms waiting to emerge: “When I ask the *I Ching* a question, I say, ‘What do you have to say about this?’ and then I just listen to what it says” (70). Cage’s intent behind his indeterminacy in music was one that sees chance not merely as the absence of authorial intent, but as a vehicle through which pre-existing potentialities—latent forms—manifest themselves.

This act of relinquishing compositional responsibility was precisely what Iannis Xenakis challenged in his formulation of stochastic music. Xenakis instead saw chance as something that could—and should—be systematically rationalised and controlled: "It is only recently that knowledge has been able to penetrate chance and has discovered how to separate its degrees—in other words to rationalise it progressively, without, however, succeeding in a definitive and total explanation of the problem of 'pure chance'" (Xenakis, 2001/1963, 4). Where Cage allowed aleatory processes to introduce genuine indeterminacy without pre-calculated probabilistic controls, Xenakis insisted that chance itself needed rigorous mathematical framing: "Chance needs to be calculated (...) In any case—to play with sounds like dice—what a truly simplistic activity! But once one has emerged from this primary field of chance worthless to a musician, the calculation of the aleatory, that is to say stochastics, guarantees first that in a region of precise definition slips will not be made, and then furnishes a powerful method of reasoning and enrichment of sonic processes" (Xenakis, 2001/1963, 38). Thus, Xenakis's stochastic music diverges fundamentally from Cage's indeterminacy. While Cage relinquished direct control as a spiritual-philosophical act, Xenakis's stochastic approach deliberately retained compositional authority by encoding probability distributions within carefully constructed mathematical frameworks.

Automatism and Affordances of a Musical System

Pierre Boulez, reflecting on the issues of "automatism and decision", recognised how "[a]lthough the computer, more than any other type of equipment, symbolises the urgency of this question, it did not create the problem. The question of the relationship between chance and determination in composers' intentions, and the means used to realise those intentions in actual works, is already familiar" (Boulez, 2018b). In a later lecture, Boulez develops this further, describing compositional practice not as a process of selection but composing entire fields of possibilities, within which the performer makes the final choices at the moment of execution. "I compose all of the possibilities", he writes, "and I leave it to the performer to choose the possibility at the very moment of performance" (Boulez, 2018a, 263). This mode of composition displaces authorship from the production of a single work toward the construction of a structured space of potential outcomes.

Such an understanding of the system as the site of artistic identity also reframes composition as an act of construction rather than inscription. For composer Helmut Lachenmann, to compose is not to express but to intervene—to “take apart” existing materials and construct new relations from their interconnections. “Composition is by no means a ‘putting together’”, he writes, “but rather a ‘taking apart’ and more: a confrontation with the interconnections and necessities of the musical substance” (Lachenmann, [1979] 1996, as cited in de Assis, 2018, 227). In this spirit, Lachenmann describes composing as the building of a new instrument—not necessarily a physical tool, but a virtual structure that captures forces and channels them into sound. For him, this “instrument” is a diagrammatic field. It is the complete collection of sounds and materials used in a given piece, crucially in the moment of performance. It is well-thought before the notation process starts, but it remains audible throughout. In the context of systems aesthetics and computational creation, such structures are not preparatory—they are the artwork. As Paulo de Assis puts it, the act of composition becomes one of mapping possibilities and constructing the conditions for their emergence: not to say something, but to do something (de Assis, 2018, 230). In this view, the artwork is not a singular object or fixed sequence of sounds, but a dynamic, metastable system—a machine capable of generating its own actualisations.

Generative systems, in this expanded sense, may take the form of a score—such as composer Witold Lutosławski’s composition *Venetian Games*, which uses principles of limited aleatoricism—a live-coded algorithm, or a sound-producing machine. What unites these diverse media is their systemic logic: each defines a structured field of potentiality from which variable outputs may emerge. In all of these cases, the identity of the artwork does not reside in any single instantiation or performance. It resides in the system itself—in its capacity to generate, to transform, and to persist across iterations. This reconceptualisation of the artwork as a generative system opens directly onto the broader domain of generative aesthetics. Whether embedded in scores, or algorithms, such systems foreground process over product, modulation over inscription, and structural coherence over surface identity. They mark a shift from composing works to composing spaces of possible outputs that hold the identity of the artwork.

Similarly, composer Artemi-Maria Gioti argues that in the context of interactivity, machine listening, and generative processes, the contemporary musical work increasingly operates as a space of possibilities rather than a fixed structure (Gioti,

2021, 64). In this framework, musical identity is defined not by concrete sound structures but by the interaction affordances and action spaces embedded in the system. “The musical text”, she writes, “is understood as having an evocative, rather than a directive function. It delineates a space of action for the musicians to explore” (6). These spaces, governed by a system's affordances, can include composed, improvised, or algorithmically generated material—coexisting within a continuum of control and emergence. Such systems-oriented perspective leads to a redefinition of authorship and identity: the system becomes the artwork. Different instances of performance are not instantiations of a fixed original, but variations within an aesthetic topology defined by the system's affordances. “Work identity”, Gioti writes, “is constituted of interaction affordances and spaces of sonic possibilities, rather than concrete structures of sounds” (9). It is precisely here that the notion of a machine re-enters. In Gilbert Simondon's discussion of technical objects, machines evolve by reducing their internal margin of indeterminacy through successive codings of operation. Yet, paradoxically, Simondon clarifies that the progressive perfecting of a machine, its increasing technicity, does not correspond to greater automatism. Rather, it depends upon maintaining a certain margin of indeterminacy within its operations: “A purely automatic machine completely closed in on itself in a predetermined way of operating would only be capable of yielding perfunctory results. The machine endowed with a high degree of technicity is an open machine” (Simondon, 2017, 17). In Simondon's view, such open machines inherently rely upon human participation, positioning the human as a permanent organiser and interpreter—guiding interactions between technical entities (17). Thus, the machine defines the operational space within which processes unfold with controlled variation.

If artists take on the roles of system-builders—designers of machinic infrastructures that also learn, adapt, and generate—then compositional responsibility extends once again: it is both about crafting outputs and encoding rules, but also about designing the conditions under which learning and emergence unfold. In this space, the challenge is to navigate both sides of the generative systems, to constrain and guide them in ways that preserve creative agency and aesthetic coherence—without foreclosing the very indeterminacy that makes them generative in the first place.

Artists Building Generative Systems

Generative music, as composer and producer Brian Eno articulates, is characterised by systems or "machines" designed to produce musical experiences by combining pre-defined materials and processes in ways the artist does not explicitly determine. As he writes: "One of my long-term interests has been the invention of 'machines' and 'systems' that could produce musical and visual experiences. Most often these 'machines' were more conceptual than physical; the point of them was to make music with materials and processes I specified, but in combinations and interactions that I did not" (Eno, 1996, 330). Central to Eno's conception of generative music is the recognition of its inherent uniqueness and its novel position relative to traditional music forms. Reflecting on his experiences with generative systems such as the Koan software—generative system used by Eno for his first generative music pieces—he observes: "Until 100 years ago every musical event was unique: music was ephemeral and unrepeatable, and even classical scoring couldn't guarantee precise duplication. Then came the gramophone record, which captured particular performances and made it possible to hear them identically over and over again. But Koan and other recent experiments like it are the beginning of something new" (Eno, 1996, 332).

Eno identifies generative music as distinct from—and complementary to—the two established forms of musical experience, live and recorded music. In his view, generative music bridges and extends the virtues of both: "From now on there are three alternatives: live music, recorded music and generative music. Generative music enjoys some of the benefits of both its ancestors. Like live music, it is always different. Like recorded music, it is free of time-and-place limitations—you can hear it when you want and where you want. And it confers one of the other great advantages of the recorded form: it can be composed empirically. By this I mean that you can hear it as you work it out—it doesn't suffer from the long feedback loop characteristic of scored-and-performed music" (Eno, 1996, 332). This observation is further underscored by Eno's reflection on a remark by Edgar Wind, who, in his 1963 Reith Lectures, compared listening to recorded music to experiencing a "superior kind of musical clock". Eno extends this perspective to imagine generative music as a future standard, noting: "I too think it's possible that our grandchildren will look at us in wonder and say, 'You mean you used to listen to exactly the same thing over and over again?'" (Eno 1996, 332).

Despite the frequent association of generative art with contemporary digital practices, Philip Galanter clarifies that the field is not inherently tied to computer technology. He argues: “Because contemporary generative art is so very often computer based many assume it is a subset of computer art. (...) generative art preceded computer art, and in fact is as old as art itself.” (Galanter, 2003, 15). A frequently presented example is the ancient Aeolian harp. First described by Athanasius Kircher in his book *Musurgia Universalis* (1650), it is a music instrument played by wind, rather than by humans. Indeed, Galanter positions generative art within a historical lineage far preceding any recent artistic movement, asserting explicitly that “generative art (...) precedes modernism, post-modernism, and just about any other ‘ism’ on record” (19) and explaining that “Equally important is the virtual certainty that new forms of generative art will come after the computer as well. Nanotechnology, genetic engineering, robotics, and other technologies will no doubt offer generative artists some wonderful opportunities”(Galanter, 2003, 15). He further broadens this claim, highlighting that “if we accept this paradigm, that generative art is defined by the use of systems, and that systems can be best understood in the context of complexity theory, we are led to an unusually broad and inclusive understanding of what generative art really is.” (12).

Galanter offers a foundational definition of generative art that has become widely cited in the field: “Generative art refers to any art practice where the artist uses a system, such as a set of natural language rules, a computer program, a machine, or other procedural invention, which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art” (2003, 4). Central to this definition is the notion that the artist relinquishes partial or total control to the system: “The key element in generative art is then the system to which the artist cedes partial or total subsequent control” (4). He adds several clarifying observations. First, the term refers strictly to how the work is made, not why it is made, nor what it might express. Second, generative art is not tethered to any specific medium or level of technological sophistication. Finally, Galanter emphasises that such a system must be “well defined and self-contained enough to operate autonomously” (4). In short, generative art is not an aesthetic style but a procedural approach—“a way to create art rather than an art style” (15)—in which the outcome emerges from a system with built-in autonomy and variability. As a consequence, Galanter insists that generative art is best understood as a subset of art defined by its capacity to produce multiple outcomes: “The word ‘generative’ simply directs attention to a subset of art, a subset

where potentially multiple results can be produced by using some kind of generating system” (4). The existence of variation across iterations is not merely a byproduct, but a defining characteristic. This attention to multiplicity also expands the role of the artist beyond conventional intentions. As Galanter poetically suggests, “The generative artist can remind us that the universe itself is a generative system. And through generative art we can regain our sense of place and participation in that universe” (19). Generative methods, in this broader context, align artistic practice with the principles of emergence, complexity, and systemic thinking.

Among the most illustrative historical examples of composers incorporating expert systems within their compositional workflows are Iannis Xenakis’s *GENDY*, George Lewis’s *Voyager*, and David Cope’s *Experiments in Musical Intelligence (EMI)*. Xenakis’s *GENDY3*, developed in the early 1990s, implemented his theory of stochastic synthesis—a process in which waveform generation is governed by statistical functions, producing continuously evolving sonic material without relying on traditional oscillators or sampled waveforms (Luque, 2009). While the stochastic synthesis engine acted as a sound generator rather than a standalone composition, it exemplified Xenakis’s broader conceptual ambition of encoding compositional logic into autonomous algorithmic processes (Xenakis, 2001/1963). By contrast, George Lewis’s *Voyager* (1987) was conceived not merely as a tool but as the composition itself: a real-time interactive musical system that engages with human performers as an autonomous improvising partner. Lewis explicitly frames *Voyager* as a nonhierarchical, multi-agent system that enacts its own musical logic, positioning the computer not as an instrument but as a collaborator capable of contributing to the unfolding musical form (Lewis, 2000). David Cope’s *EMI* (begun in the 1980s) represents yet another model: a symbolic AI system that analyses stylistic patterns in existing compositions and uses them to generate new works in the same style. In contrast to Lewis’s improvisational openness, *EMI* foregrounds stylistic replication, raising questions about authorship, creativity, and the boundaries of generative automation (Cope, 2001). Together, these three projects exemplify distinct approaches to expert systems in music—as generative engines, as improvisational agents, and as stylistic emulators—each contributing to an expanded view of composition as the design of autonomous, rule-based, and often performative systems.

Composers as Programmers

Trevor Wishart—electroacoustic composer and improvising vocal performer—writing in 1985, anticipated a future in which the computer would become a “truly generalised tool for sonic art” by enabling composers to define and articulate *operational fields*—rule-based groupings of musical parameters that can interact dynamically with higher-level behaviours or performative inputs (Wishart, 1996, 329). These operational fields mark a crucial conceptual shift: rather than composing fixed sequences of musical material, the composer designs systems that organise potentialities—configurable textures, articulation types, and timbral or morphological transformations—governed by abstract but pliable rules. In this model, composition becomes the design of generative grammars capable of continuous variation. Around the same time, in IRCAM, Miller Puckette—computer-music pioneer, software inventor, and academic—began his work on Max—a visual programming language for music and multimedia, enabling the development of various digital modular “instruments” for live-electronic performances and processual musical systems. Developed in 1996 by James McCartney, SuperCollider—an environment and audio programming language for real-time synthesis and algorithmic composition—offered composers another powerful framework for building their own instruments, systems, and compositional processes, enabling a shift from writing music to designing the conditions under which digital music processes can be shaped.

While today, in the age of ubiquitous Internet access, this might seem trivial, at the time these software environments, though enabling the creation of generative systems, were largely limited to the local machines on which they were installed. Given the constraints of computing power, bandwidth, network availability in the mid-1990s, and dependence on the software, a generative composition could not yet be easily distributed as a dynamic system across devices—in a sense it remained bound to the specific hardware that executed it until it was deliberately transferred, with all of its required software, to another computer. As Thor Magnusson—researcher, composer, and music technologist—points out, Brian Eno’s *Generative Music 1*, released in 1996 on floppy disk and built on the Koan software was “one of the earliest examples of distributable generative music” (2019, 198), offering not a fixed audio file but a procedural system capable of generating unique musical outputs at every playback. This innovation helped shift the focus of composition from fixed media to executable logic—musical processes encoded in systems rather than scores

or recordings. Eno's work inspired a wave of artists and developers, including Antoine Schmitt and Vincent Epplay's *The Infinite* CD (1999), John Eacott's *The Morpheus* CD-ROM (2001), and later collaborative works like *SameSameButDifferent* (2003–2008), each of which explored novel distribution formats for generative music systems (Magnusson 2019, 198–199).

However, the lack of standardisation in distribution formats has remained a persistent issue. As Magnusson observes, attempts to develop stable delivery mechanisms—from CD-ROMs to embedded hardware like Spyros Polychronopoulos's *Sound Object* (2016)—have been met with ongoing fragmentation, partly due to the diverse programming environments and devices involved (2019, 199). The shift to browser-based tools like the Web Audio API represents a promising new phase in this development, enabling generative music to be composed and performed within standard web environments using JavaScript (199–200). In parallel with these technological changes, a deeper transformation of the composer's role has taken place. The generative artist becomes a programmer of processes—an author of algorithmic environments capable of producing aesthetic output in their absence. As Magnusson writes, this transformation marks the rise of “composers as programmers”, evident in the work of figures like Nick Collins, Bob Sturm, and François Pachet (2019, 201–202), where “experimental musical instruments and interfaces are often intended to be presented as artworks themselves” (Giannoutakis, 2024, 12).

As artists gain technical fluency, they are not only using tools but designing them—developing compositional environments, interfaces, and data pipelines that allow for bespoke creative practices. The Future Art Ecosystems 4 have tracked this shift closely, describing how artists now operate as system-builders: creators of integrated technological processes that blur the boundary between tool and artwork. “Some artists”, they note, “look outward, building tools into more complex, operative systems for research or production. These artists craft tools, craft with tools, or craft tools into bigger systems” (Serpentine Arts Technologies, 2024, 102). Rather than merely producing objects, artists working with advanced technologies frequently develop systems with potential impact beyond the art world—sometimes innovating in fields such as interface design, artificial intelligence, or decentralised networks (Serpentine Arts Technologies, 2020, 26, 64–65). The metaphor proposed by Takashi Kudo (Tokyo-based Global Brand Director for art collective teamLab)—“It's like you are a painter, but you also have to invent paint” (Serpentine Arts Technologies, 2020,

8)—aptly captures the dual role of the artist as both creator and technologist. These roles converge in the figure of the composer-programmer, who composes not only musical structures but also the technological and computational infrastructures through which musical experience is produced. This shift is especially pronounced in contexts where artists fine-tune or reconfigure machine learning models, as a form of creative authorship—what the Future Art Ecosystems 4 report calls “model-making as meaning-making” (2024, 100–101). In this view, the generative field expands again: no longer ends with the generation of musical forms but includes the generation of the tools and systems by which those forms become possible.

Part 2—Blockchain for the Arts

What is Blockchain

If the first part of this chapter explored the artistic turn of some composers in the twentieth century towards a design of autonomously working artistic systems from the ground up, the second part turns to a more recent kind of machinic infrastructure that is gradually more interesting for artists—blockchain.

At its core, “blockchain is a way to trust a network of computers run by strangers—so you don’t have to trust the individual people” (ECAD Labs, n.d.). This deceptively simple definition gets to the heart of blockchain’s conceptual innovation. Instead of relying on a single source of truth—such as a website located on a single server—blockchain protocols enable confirming the identity of digital objects and agents through code. All the computers in the network—called nodes—run the same software, store the same history, and apply the same rules. If one computer lies or makes a mistake, the algorithms automatically ignore it. Through cryptographic verification, each participant can independently confirm that data hasn’t been tampered with. No central authority is required. The result is a distributed system in which consensus between the nodes (that some piece of information is considered true by the majority) is enforced by code, not by institutional oversight or any central authority. Importantly for this thesis, blockchain introduces, among various new solutions, a way of distributing digital art in which trust is not assumed or negotiated—but formally encoded, a feature crucial when questions of ownership and scarcity are explored in digital contexts, and when various actants collaborate remotely.

This architecture of trust was invented to solve a very specific technical problem in digital commerce: the double-spending problem. When something exists only as digital information, it can in principle be copied endlessly. Unlike a physical coin or banknote, which can only be handed from one person to another, a digital token can be duplicated and spent more than once. This undermines the value and integrity of digital money. Until recently, the only way to prevent this was to use centralised intermediaries—like banks or payment processors—to track ownership and verify each transaction. But this reintroduces a structure of dependence: users must trust these intermediaries not to lie, fail, or extract value.

In 2008, a white paper by the anonymous figure Satoshi Nakamoto proposed a radically different solution: to use a peer-to-peer network to collectively verify and record every transaction—publicly, permanently, and without any central control. “The problem”, as Nakamoto writes in the paper, “is that the payee can’t verify that one of the owners didn’t double-spend the coin” (Nakamoto, 2008, 2). The solution is to make every transaction part of a shared public history—a timestamped chain of verified data blocks, each sealed with cryptographic proof. As Nakamoto explains: “We propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work” (Nakamoto, 2008, 1).

To understand how this system works, we need to begin with one of its core building blocks: a hash. The hashing algorithm takes any piece of information—whether a sentence, a file, or a full transaction record—and transforms it into a short, fixed-length string of characters⁴. The specific kind of hashing used in Bitcoin and many other blockchains is called SHA-256, a secure cryptographic algorithm designed to be fast to compute, but impossible to reverse. It is part of the SHA-2 family, developed by the U.S. National Security Agency (NSA) and published by the National Institute of Standards and Technology (NIST) in 2001. SHA-256 takes input data of any length and produces a fixed-size, 256-bit (32-byte) output known as a hash or digest. What makes it powerful is that even a tiny change in the input creates a completely different hash, making it highly sensitive to input changes. For example, changing a single letter in a sentence will produce a radically different string of characters. This makes it very easy to detect if something has been altered. A hash acts as a digital fingerprint for the input data: even a tiny change in the input will result in a completely different hash. The process is one-way as it is computationally infeasible to reverse the hash to retrieve the original input, which is crucial for security applications like blockchain, but also for everyday online password storage, and digital signatures in various applications besides the blockchain. This means you can easily turn data into a hash, but you can’t go backwards from the hash to guess the original data. This one-way nature is essential: it allows the system to verify that something hasn’t changed without needing to see or reveal the original content. In this way, to check that some

⁴ For example: 8c872aa8ab6955cb91897efba8630850dd2b2ce918c8f3c21d2ce86881bc9f87.

data is the same as the original takes only to compare two hashes, without the need to inspect the hashed data itself.

Now that we understand hashing, we can look at how it's used to validate blocks. A block is a collection of data—such as a text or file—that we want to encode with a special kind of hash. But not just any hash will do. In systems like Bitcoin, a block is only considered valid if its hash meets a particular requirement: for instance, the hash must begin with four zeros. Since hash outputs are unpredictable, there's no way to simply choose the right input to get a valid hash. Instead, what the system does is repeatedly modify the block's contents by changing a small, meaningless number called a nonce. Each time the nonce is changed and added to the data that we want to hash with a valid hash (again: starting with four zeros), the hash is recalculated. Most of the time, the resulting hash is invalid (does not start with the four zeros). But eventually, by trying enough different nonces, the system will stumble upon one that produces a hash with the required number of leading zeros⁵. When such a hash is found, we say that the block has been *signed*. This process—called mining—is like solving a puzzle: the solution proves that someone has spent real computational effort to find a valid hash. Once found, anyone can verify the result by checking the hash themselves. What's important is that while verifying a valid hash is quick, finding it is deliberately hard. This built-in difficulty is what gave early blockchain systems their security and trustworthiness.

Once we understand that a block contains some data—along with a nonce chosen to produce a valid hash—we can begin to see how these blocks are connected into a larger system. This is where the idea of the chain comes in. Each time a new block is created, it doesn't just contain its own data and a valid nonce. It also includes the hash of the previous block as part of its own data. This means that the hash for any given block depends not only on its own contents, but also on the hash of the block that came before it. That block, in turn, contains the hash of the block before it, and so on, forming a continuous chain that links all blocks together. The result is a chain of blocks—a blockchain—in which each block secures the one before it, and is itself secured by the one after. What makes this structure powerful is how it protects against tampering. If someone tries to change the contents of a block from earlier in the chain—even by a single character—the hash of that block will change. But because

⁵ For example: 00001fcbc5a02912336fac9a3875049673f43127d1d804cde8078409795a1f8e.

that hash is now embedded in the next block, the next block's data will no longer be valid. Its hash will be incorrect, and so will the hash of every block that follows. To cover their tracks, the attacker would have to recalculate the valid nonce and hash for every block from the point of the change to the present, solving a new computational puzzle each time. In a large blockchain, this becomes computationally infeasible, especially since new blocks are constantly being added by others.

What makes this structure especially powerful is how it behaves in a distributed network. In a blockchain system, many different computers store the same copy of the entire blockchain. Each node has its own copy, but all are expected to agree on the same version of history. When we compare that modified version of the blockchain to the versions stored on other nodes, the inconsistency becomes immediately visible. All it takes is to compare the hashes at each block to see whether a particular copy is in agreement—or in consensus—with the rest of the network. If it's not, the altered version is rejected. The chain with the most valid work behind it—meaning the one that required the most computation and remains unbroken—is considered authoritative. In this way, blockchain doesn't rely on any single trusted source. Instead, it uses redundancy, verification, and shared computation to maintain a collectively agreed history.

So far, we've been referring to each block as containing some unspecified "data". In the original blockchain—Bitcoin—this data is a list of financial transactions. Each transaction recorded the transfer of digital and fungible (meaning: exchangeable 1:1) tokens from one address to another. The blockchain's role is to act as a public ledger—a chronological and verifiable record of who sent tokens to whom, and when. Each block contains multiple transactions, and once a block is mined and added to the chain, those transactions become part of the permanent history. This makes it nearly impossible to alter records retroactively without detection. What Bitcoin invented, then, is not just a new form of money, but a new infrastructure for digital accounting: a distributed system for collectively agreeing on the state of a shared ledger, without relying on any central institution to maintain it.

To interact with a blockchain, users need a digital wallet—a software tool that generates a public address (used like an account number) and a private key (used like a password) which are also cryptographic hashes themselves. These keys define ownership: whoever holds the private key controls the assets associated with its public address. Users on the blockchain network can take different roles. Full nodes

store the entire blockchain history and validate new blocks; mining nodes (in proof-of-work systems) compete to solve cryptographic puzzles (like finding the hashes with four zeros at the beginning) and add blocks; and lightweight nodes (or light clients) use simplified methods to verify transactions without storing the full chain. To agree on which block gets added next, blockchain networks use consensus mechanisms. The original system, proof-of-work, requires nodes to perform energy-intensive calculations—a process that secures the network but consumes vast resources. A more recent alternative, proof-of-stake, selects validators based on how many tokens they hold and are willing to “stake”, offering similar security with dramatically lower energy use. While Bitcoin still uses proof-of-work, many newer blockchain systems—including Ethereum since its 2022 upgrade—have adopted proof-of-stake, reflecting a broader shift toward more energy-efficient consensus models.

Machines that Keep the Score

Blockchain architectures took a major leap forward in 2014, when Canadian computer programmer Vitalik Buterin proposed a new kind of blockchain platform: Ethereum. In his white paper, Buterin argued that blockchain could support more than just financial transactions. The crucial change he introduced was that whereas Bitcoin stored only a list of financial transactions and data required for its technical operations, Ethereum allowed that data to include any kind of executable program. By embedding a fully Turing-complete scripting language into the blockchain, Ethereum made it possible to write, deploy, and run arbitrary code directly on the blockchain. These programs became known as smart contracts (Buterin, 2014). What makes smart contracts radically different from traditional code is the environment in which they operate. They don’t run on a single server or app, but inside a global virtual computer: the Ethereum Virtual Machine (EVM). While each computer in the Ethereum network runs the same software independently, together they simulate a shared computational environment. Every time someone interacts with a smart contract, all full nodes verify the code and reach agreement on its outcome. This means that the contract’s execution (a piece of code running) is not only decentralised, but publicly verifiable and persistently stored on the blockchain.

As Akira Summers—software developer and lecturer in digital media technologies—puts it, smart contracts are to blockchain what vending machines are to shops: they automate transactions by enforcing rules with perfect consistency (2022, 119). In simple terms: a smart contract is a computer program “installed” on the virtual machine consisting of many computers, which always has a condition encoded that specifies when it should execute its code (de Assis & Łukawski, 2024, Glossary). Such condition could be financial—in this case the code executes when some public address receives a payment, or it could be non-financial—such as a threshold number of users participating, external data like weather reports, or even artistic criteria—such as a particular input or signal from a sensor, performer, or AI system. Crucially, these contracts do not require human intermediaries to function. Anyone with the Internet access and a digital wallet can run any smart contract on the virtual machine, if they can meet the conditions of that smart contract. When the conditions are met, the code runs. Like a vending machine, the contract does not discriminate between users or interpret intent—it simply executes. This introduces a new kind of infrastructural trust: not based on institutions or reputation, but on code that is publicly auditable and resistant to tampering. In this sense, smart contracts function as small, composable units of logic that can be assembled into broader systems—enabling, among other things Non-Fungible Tokens (NFTs), Decentralised Applications (dApps) (such as art marketplaces), and Decentralised Autonomous Organisations (DAOs).

Non Fungible Tokens (NFTs)—a Non-Fungible Token (NFT) is a blockchain-based mechanism for automatically tracking the identity of a digital object. Normally, digital files—images, sounds, texts—can be copied indefinitely, making it difficult to establish scarcity of such objects and attribute to them such properties as uniqueness, ownership, and provenance (de Assis & Łukawski, 2024, Glossary). This is a variant of the double-spending problem explained earlier: if a digital asset can be duplicated without constraint, its value and authorship become unstable. NFTs offer a solution by recording, on-chain, a verifiable reference to a single version of a digital asset, or in some cases even that asset itself as part of the block’s metadata on chain. Each NFT is defined by a unique identifier embedded in a smart contract, ensuring that even if the associated file is widely distributed, the token representing it remains singular. In this way, NFTs enable digital artworks to behave more like singular artefacts than infinitely replicable files, allowing a simulation of so called digital scarcity.

Decentralised Applications—a Decentralised Application, or dApp, is a software application that interacts with smart contracts on a blockchain. While smart contracts themselves operate on the backend—executing logic and storing data on-chain—a dApp typically includes a user interface that allows people to interact with these contracts more intuitively, often through a web browser or a mobile app (Summers, 2022, 124). Unlike traditional apps, which depend on central servers, dApps rely on decentralised infrastructure: their core logic runs not on a private server but on the blockchain itself. In this way, dApps inherit the properties of the blockchain they run on—transparency, immutability, and censorship resistance—while enabling users to engage with smart contracts through familiar and accessible user interfaces. A common example of a decentralised application is a digital art marketplace, such as *objkt.com* or *teia.art* on the Tezos blockchain, or *zora.co* on Ethereum. These platforms allow artists to mint, list, and sell blockchain-based artworks directly, using smart contracts to govern ownership, provenance, and exchange.

An Example: Hic Et Nunc

A particularly significant example of a decentralised application in the context of the arts was Hic et Nunc, an art marketplace for trading NFTs launched in March 2021 by Brazilian developer Rafael Lima. As Diane Drubay—museum & digital innovation strategist, Web3 consultant, and environmental visual artist—recounts (2024, 98), the platform quickly distinguished itself from the dominant NFT marketplaces of the time—many of which were characterised by high transaction fees, selective gatekeeping, and an emphasis on speculative value. By contrast, Hic et Nunc positioned itself as an open, low-cost, and experimental alternative. Built on the Tezos blockchain—chosen for its energy-efficient proof-of-stake model—Hic et Nunc enabled artists to mint and sell NFTs at a fraction of the cost of Ethereum-based platforms, significantly lowering the barrier to entry for creators worldwide. What made the platform particularly notable was not only its affordability, but its minimal design, lack of curation, and open-source ethos. Artists could mint digital artworks within seconds using a simple interface, and without needing approval from gatekeepers. This structural openness allowed a wide range of creators—from Brazil to Malaysia—to participate on equal footing. According to Drubay, Hic et Nunc's emergence catalysed a moment of cultural reorientation: artists began to experiment

with blockchain as a space of shared activity, where new forms of community, value, and visibility could emerge.

When the original platform was abruptly discontinued by its founder in November 2021 (as its frontend depended on a self-hosted website), its community responded by reconstituting the project. Thanks to the decentralised architecture of the blockchain—where the NFTs created through the Hic Et Nunc platform were stored on the Tezos blockchain and the files of the artworks on IPFS (InterPlanetary File System, a peer-to-peer file system)—the community was able to preserve access to the works and launch mirror sites (alternative frontends referencing the same smart contracts and assets). One of these mirrors, *teia.art*, evolved into a Decentralised Autonomous Organisation (DAO): a collectively governed structure in which decisions about the platform's future are made directly by its contributors. As Paulo de Assis and I noted elsewhere (de Assis & Łukawski, 2024, 104), this transition marked an instrumental moment in the history of blockchain-based artistic platforms, demonstrating that ownership and continuity could be sustained even in the absence of a central figure or company.

Ruth Catlow (curator, writer and blockchain researcher), and Penny Rafferty (art-tech researcher) (2022, 28), as well as Catherine Mulligan (blockchain researcher and tech-policy advisor) (2024, 108-110) observe that Decentralised Autonomous Organisations represent a novel organisational form that challenges centralised models of coordination and governance. Because their rules are encoded in smart contracts, DAOs enable decentralised communities to manage membership, voting, and collective funds without requiring external intermediaries. Participation is often token-based, meaning that holding a specific token can grant decision-making rights or access to shared resources. This opens up possibilities for experimenting with alternative institutional models that are more transparent, participatory, and adaptable than conventional structures. Mulligan (2024, 109) suggests that DAOs may offer a basis for new kinds of cultural institutions—less hierarchical, more inclusive, and potentially more responsive to the needs of decentralised communities. For this reason, blockchain can be understood not just as a technical tool, but as a space for organisational and societal experimentation. As Primavera De Filippi—legal scholar, blockchain governance expert, and creative technologist—puts it, “We don’t otherwise have a chance to experiment with governance anywhere, so it’s a very valuable space for exploring and testing new governance structures”

(Serpentine Arts Technologies, 2022, 55).

Yet, as Claudio Tessone—complexity scientist and blockchain researcher—argues, current blockchain applications often replicate the patterns of inequality found in traditional socioeconomic systems (2024, 88). Markets built on token economies tend to reward those who already hold value, following the familiar logic that “the growth of an entity within a system is proportional to its existing size” (88). Without careful intervention, decentralised systems risk reproducing the same concentration of power and wealth that they claim to resist. In this sense, the future of blockchain depends less on its technical architecture than on the imaginative and political choices made by its communities. As Tessone writes, “Blockchain technology, in its nascent stages, has yet to significantly deviate from [existing] paths. [...] To achieve this, a concerted effort in education, research, and practical experimentation is essential” (94). Some proposals—such as Mark Alizart’s notion of cryptocommunism (2020)—have begun to explore how blockchain could support alternative value systems entirely. Whether or not such visions are realised, what matters is that blockchain infrastructures remain open to experimentation—not only with money and media, but with the very forms of collective life they make possible.

Inte Gloerich—critical media & technology researcher and cultural programmer—argues that it is more productive to view blockchain not in terms of static technical features, but as a heterogeneous assemblage encompassing everything from smart contracts and digital wallets to institutional actors, cultural imaginaries, and regulatory frameworks (2025, 27). These assemblages condition what kinds of relations and behaviours can emerge and trust is encoded into the system’s architecture. Blockchain networks can then be understood as machinic assemblages: dynamic configurations of code, protocols, interfaces, infrastructures, narratives, and agents—both human and non-human. Future Art Ecosystems 3 suggests that artists working with decentralised technologies do not merely create artworks, but design systems in which art can be produced, circulated, and interacted with differently (2021, 24). For artists working with multi-agent systems, decentralised authorship, or distributed networks of value, blockchain offers more than a tool for tokenising media—it is a medium for designing the very rules by which participation, provenance, and power are constituted.

Generative Music NFTs

One of the most innovative intersections of blockchain technology and artistic practice to date can be found in generative art NFTs. Generative art NFTs can be defined as ‘unique digital art pieces created through coded algorithms and stored as NFTs on the blockchain’ (Stoykov, 2023, as emphasised in Łukawski, 2024a, 37-38). The same algorithm can generate many unique NFTs by using random functions or dynamic inputs. For example, the colours, shapes, or musical structure might change depending on the time of day, the wallet address of the collector, or external data like weather conditions.

The degree to which the generative process itself is integrated into the blockchain can vary widely. In the most autonomous configurations, not only is the NFT minted on-chain, but the algorithm that generates the artwork is also encoded directly into the smart contract. This means that the entire lifecycle of the artwork—from generation to tokenisation—is managed without relying on external systems. Rhapsody Labs (RLXYZ 2022, and as discussed in Łukawski, 2024a, 37-38) has proposed a useful framework for understanding this, introducing the notion of on-chain purity. This framework evaluates how permanent, transparent, and self-contained a generative artwork is when considered over long timescales—centuries, even. The fewer external dependencies an artwork has (such as off-chain storage or mutable metadata), the “purer” it is considered. At the highest level of purity, all components—the artwork, smart contract, generation function, metadata, and unique identifiers—are encoded entirely on the blockchain, ensuring that the artwork remains independently accessible and verifiable far into the future.

To understand the practical applications of generative art NFTs, it's essential to examine the platforms that facilitate their creation and distribution. Two prominent platforms in this space are *Art Blocks* (n.d.) and *fxhash* (n.d.), each offering unique approaches to generative art on the blockchain. Art Blocks is an Ethereum-based platform that enables artists to publish algorithmic systems, which collectors can then mint into unique, on-demand generative artworks. When a collector initiates a mint, the platform uses a combination of the artist's code and a random seed to generate a one-of-a-kind piece, ensuring that each generated artwork is distinct. Fxhash is an open generative art platform built on the Tezos blockchain and currently also integrated with Ethereum and Ethereum Base. It allows artists to upload generative

scripts, which are then used to automatically generate unique new artworks as NFTs upon collection by the users of the platform (typically, when the user buys a new iteration of the artwork with cryptocurrencies). The platform emphasises accessibility and openness, enabling any artist to publish projects without a formal curation process (although many strict rules on what is allowed within the code of the artwork are listed in the platform's documentation). Fxhash also allows artists to define parameters that collectors can adjust during the minting process (through their *fx(params)* functionality), adding an interactive dimension to the artwork. Both Art Blocks and fxhash exemplify how blockchain technology can be leveraged to create, distribute, and collect generative art in innovative ways. By providing platforms that integrate the generative process with the minting of NFTs, they offer artists new mediums for creativity and collectors new forms of engagement with digital art.

Long-Form Generative Art

In his essay "The Rise of Long-Form Generative Art", Tyler Hobbs (2021) delineates a significant shift in generative art practices facilitated by blockchain technology. Traditionally, generative artists would produce numerous outputs from an algorithm and selectively curate only a few of them for public presentation. This "short-form" approach allowed artists to showcase only the most compelling iterations, effectively concealing less successful outputs (Hobbs, 2021). The advent of platforms like Art Blocks introduced a "long-form" paradigm, wherein artists deploy a generative script to the blockchain, specifying a fixed number of iterations—often ranging from 500 to 1,000—that collectors can mint directly. Each minted piece is generated in real-time, with neither the artist nor the collector knowing the outcome beforehand. This process eliminates the possibility of post-generation curation, compelling artists to ensure that every potential output meets a high standard of quality. Hobbs emphasises that this model requires artists to design algorithms capable of producing consistently compelling results across a broad output space. The responsibility shifts from selecting the best outputs to crafting a generative system where every piece holds artistic merit. This transparency fosters a deeper engagement between collectors and the generative process, as they gain insight into the algorithm's full range and variability. In contrast, short-form generative art relies heavily on the artist's discretion to present a curated selection, potentially obscuring the algorithm's

limitations. Long-form generative art, by exposing the entirety of the algorithm's output, democratises the creative process and challenges artists to develop more robust and versatile generative systems (Hobbs, 2021).

While the field of generative visual art NFTs has seen rapid innovation, particularly through platforms like Art Blocks and fxhash, similar developments in music and sound-based practices have been comparatively limited. As Kosmas Giannoutakis—composer, media artist, computer-musician researcher—notes, visual artists were among the first to explore the creative potential of blockchain technologies, with musicians only more recently beginning to experiment with these tools in meaningful ways (2024, 61). Martin Zeilinger (2024, 57) similarly observes that although blockchain and AI technologies have already led to a range of sophisticated experiments, “creative experimentation and theoretical reflection appears to have taken place [...] with a primary focus on visual, non-time-based artforms”. The affordances of blockchain for encoding logic, temporality, and interactivity remain underexplored in musical contexts, despite their potential to support new models of compositional authorship, distribution, and performance.

Some notable projects have explored how blockchain-based systems can be used for generative audio. Among the most known is the work of digital artist Deafbeef, whose fully on-chain series—*Synth Poems*, *Angular*, *Transmission*, and others—combine generative sound synthesis with minimal code to produce audiovisual artworks directly from smart contracts (0xDEAFBEEF, n.d.). These works are executable algorithms that generate sound deterministically on-chain, exemplifying what a high degree of blockchain-native authorship might look like in the sonic domain. Other artists, such as Memo Akten and James Paterson (a.k.a. Presstube), have explored the hybrid terrain of generative audiovisual NFTs with projects like *A Strange Loop*, which investigates recurrence and evolution in algorithmic patterns. The artist known as poperbu has also created generative sound-based NFT collections including *BITXO* and *Amnèsia*, each experimenting with sound as a time-based, procedurally generated medium embedded in blockchain infrastructures.

Three Experimental Prototypes: *Canons*, *Tone Row*, *Subtraction* (2024)

As part of my artistic research into the compositional affordances of blockchain infrastructures, I developed three experimental generative NFT projects on the fxhash platform in 2024: *Canons* (2024b), *Tone Row* (2024c), and *Subtraction* (2024d). Each of these projects investigates a different aspect of generativity, authorship, and decentralised infrastructure, treating the NFT not only as a distribution mechanism, but as a medium for algorithmic music-making. These works were intended as exploratory prototypes, artistic experiments designed to test the creative and technical limits of existing blockchain platforms. By treating generative algorithms, minting processes, and smart contracts as compositional materials, each project models a distinct approach to encoding musical logic, interaction, and variability. The first of these, *Canons* (2024), is a limited-edition generative NFT that algorithmically composes short two-part canons in the style of 16th-century counterpoint. Developed in consultation with historical flute player and artistic researcher Jonty Coy, the algorithm draws on the rules laid out by Zarlino, Bathe, and Morley for improvising canons and ornamenting lines through diminution. Each mint generates a new, unique canon, rendered both as audible playback and as a visual score, with MIDI export functionality. The project was conceived as an academic response to Coy's presentation at the Academy of Creative and Performing Arts (Leiden University), and functions as both a scholarly commentary and a compositional system.

The next two projects, *Tone Row* and *Subtraction*, were commissioned by composer Zach Dawson as interpretations of his experimental text scores from the book *ImageAudio: A Post-Digital Event Score Anthology* (2024). *Tone Row* is an audiovisual generative music NFT that produces a new 1–4 minute piece on each mint, combining a large collection of pre-recorded synthesizer samples with algorithmic transpositions through a tone.js sampler device, form-generation, and audio-reactive visuals. The algorithm randomly selects between 12-tone equal temperament and microtonal tuning systems, resulting in a broad sonic spectrum across editions and is capable of generating stylistically coherent though widely varied outputs. Since its release on May 14, 2024, over 140 diverse iterations have been minted. This work illustrates how a generative NFT can function not as a fixed output but as an expert system capable of composing with variability—one that requires multiple versions to be fully grasped.

The *Subtraction* generative music NFT, by contrast, reimagines the minting process itself as an act of distributed composition. Rather than generating full works at once, each NFT in the *Subtraction* series represents a single chord, whose parameters—duration, pitch range, density—are selected by the minter (the person initiating the generation process becoming the owner of the generated NFT). These chords are then aggregated into a continuous composition rendered on an external website (2024e), which queries the fxhash API to assemble the complete composition timeline in real time and allows its playback in various configurations (from start to end, shuffled, or allowing the user to play each chord separately to examine who is its current owner and the original minter). In this model, the artwork exists not as a single object but as an emergent assemblage of contributions—an ongoing, collaborative system that evolves with each mint. The artwork’s minting phase was algorithmically set to one month and resulted in a short composition consisting of 37 chords. As I wrote in a related article for fxhash (2024f), *Subtraction* enacts a kind of “posthumanist musical assemblage”, where the roles of composer, performer, and listener are redistributed across a network of human and non-human actors.

Together, these three projects explore how generative music can be operationalised through blockchain not only as a form of aesthetic output, but as a medium for rethinking authorship, temporality, and collaborative agency. They gesture toward new forms of compositional infrastructure in which code, protocol, and procedurality become central to the musical experience.

More Long-Form Generative Music?

A useful overview of early generative audio-based NFT projects can be found in an article by Erik West (2022), published on the fxhash platform. West traces a range of experiments at the intersection of sound, code, and blockchain, acknowledging the pioneering nature of these efforts (such as as a collaboration between kaigani and Dadabots, ByteBeats, or Rami Awad’s generative piece *Basil*) while noting that “generative music is still a newborn compared to the relative infancy of the entire NFT space”. He frames the problem not as a lack of artistic potential, but of infrastructural maturity and cultural adoption. As West puts it, “I dream of hearing Brian Eno’s next release in a truly generative way, instead of by streaming or buying a single version, with a set beginning and ending”.

One important reason for this limited engagement might lie in the technical divergence between tools traditionally used by algorithmic sound artists and those required for generative art development on blockchain platforms. Many sound artists, especially those trained in conservatory or experimental music contexts, work with specialised audio programming environments such as Max/MSP, Pure Data (Pd), or SuperCollider. These platforms are optimised for real-time audio synthesis, granular control over signal flow, and complex musical structures, but they are not natively compatible with the web-based programming environments required for generative art NFT platforms. By contrast, platforms like fxhash and Art Blocks rely on technologies rooted in web development—particularly HTML, CSS, JavaScript, and libraries such as p5.js and tone.js—which are more commonly used in visual or browser-based creative coding communities. For many algorithmic musicians, transitioning to these environments can present a significant learning curve, particularly when it comes to sound synthesis in the browser, which lacks the low-level control and audio fidelity of desktop-native tools. That said, bridging solutions are beginning to emerge. Cycling '74, the developer of Max/MSP, recently released the RNBO framework (Cycling '74, n.d.), which allows Max patchers to be exported as web-compatible audio applications. Similarly, developer Sébastien Piquemal has created WebPd (Piquemal, S., n.d.), a library that enables Pure Data patches to run in web browsers using JavaScript. Notably, the generative artist poperbu has already alluded to these developments in his work: in the NFT piece *bang.pd* (popperbu, 2023), he incorporates the visual motif of a Pure Data patcher as a conceptual gesture. The artwork seems to imagine a future in which Pd-native audio systems can be deployed and collected through blockchain platforms like fxhash. This gesture suggests perhaps not only a technical desire, but a broader cultural statement that tools familiar to experimental musicians could become interoperable with decentralised, generative infrastructures for distribution and collection.

One reason why generative art platforms have so far been more widely adopted by visual artists than by musicians or sound artists may lie in how each field relates to the ontology of the artwork. As Domenico Quaranta—art critic, curator and researcher—notes, visual art has traditionally operated within an autographic framework—where each work is a unique object—whereas every notated music is in principle allographic: what matters is the underlying score, code, or system, not any single performance or output (2021, 85-87). A famous example is that traditionally a copy of the Mona Lisa painting is not considered the same as the original Mona Lisa,

but various distinct performances of the Beethoven's symphony are all considered instantiations of the original piece. Blockchain infrastructures, particularly NFTs, introduce mechanisms of artificial authenticity that align more easily with autographic thinking—supporting scarcity, uniqueness, and the object-hood of digital works. For many sound artists, however, whose focus remains on reproducible processes or live performative systems, this shift toward collectible artefacts may feel less intuitive. As a result, the affordances of blockchain for experimental music remain underexplored—not because they are lacking, but because they require a different conceptual framing of what constitutes a “work”. Nevertheless, the relatively new phenomenon of generative art NFTs has a strong potential to attract time-based experimental artists, exactly because it enables a medium in which it is the unfolding of a system in time that perpetuates the experience of the artwork.

Whereas traditionally visual artists might more naturally consider a single generative output to constitute a unique object, composers and sound artists are traditionally used to be thinking in terms of time-based performative re-execution and iterative variation. Blockchain-based systems, by contrast, offer a technical mechanism to individuate and register each execution as a distinct artefact. This opens up new possibilities, where composition's multiple performances or parameter-based instantiations can be each minted as unique NFTs, establishing a direct provenance for what might otherwise remain ephemeral or repeatable. A parallel can be drawn to recent developments in dance. During the 2022 symposium “Dance and the Blockchain: Commodification and Ownership of Embodied Creativity in the Crypto Space”, held at Ghent University and convened by Jorge Poveda Yáñez (Poveda Yáñez, 2022), researchers explored how motion-tracking systems could capture the gestures of dancers and encode them as NFTs (Poveda Yáñez & Davies). The goal was not simply to document performance, but to inscribe bodily movement as reusable digital assets—capable of being invoked, recombined, or reinterpreted in other choreographic or virtual contexts.

This gap between visual and musical experimentation on the blockchain may also reflect a lack of a distinct conceptual anchor within music's Web3⁶ culture. As Marcus

⁶ “Web3 represents the next phase of the internet, focusing on decentralised networks and blockchain technology. It aims to create a user-owned internet, where individuals have control over their data, identities, and transactions. Web3 leverages cryptocurrencies, decentralized finance (DeFi), and non-fungible tokens (NFTs) to facilitate peer-to-peer interactions without central authorities.” (de Assis & Łukawski, 2024, Glossary).

O'Dair—researcher of innovation and musician—observes, while the aesthetic of Web3 visual art—particularly the iconic profile picture collections like *CryptoPunks* and *Bored Ape Yacht Club*—has become instantly recognisable, “Web3 music doesn’t sound any different to Web1 or Web2 music” (2024, 78). These visual projects introduced not only a distinctive generative format, but also a collective identity and recognisable function: the profile picture (PFP). They effectively became symbols of participation and status in the Web3 ecosystem. Music, by contrast, has yet to produce a comparable conceptual format—one that would embody its generative, distributed, or performative qualities while also being legible as a new form of cultural value. Without a widely accepted conceptual form that links blockchain-native affordances with musical or sonic expression, Web3 music risks remaining tethered to distribution models inherited from earlier platforms. In visual art, NFTs facilitated the creation of recognisable categories—such as generative long-form series or fully on-chain artworks—that became integral to the experience and valuation of the medium. In music, however, NFT-based projects often reproduce familiar formats: static audio files, streaming rights, or token-gated access to pre-existing compositions. As a result, the experiential and compositional affordances of blockchain—such as real-time algorithmic generation, interactive participation, or verifiable versioning—are underutilised.

Critics argue that the decentralization ideal of blockchain can mask new forms of corporate or monopolistic control. As observed by the media artist Hito Steyerl, rather than empowering creators, blockchain platforms are seen as “*onboarding tools*” that ultimately serve Big Tech conglomerates’ interests (Brown, 2023). Furthermore, the libertarian and techno-utopian rhetoric of blockchain clashes with many artists’ community values. The promise of *democratizing* art through decentralization is met with skepticism: while crypto advocates speak of radical inclusivity, critics note that the NFT boom often reproduced the same exclusivity and inequities of the traditional art market (Casemajor, 2022).

Some art theorists argue that tokenization hasn’t truly empowered artists but instead introduced new intermediaries (e.g. crypto-exchanges, NFT platforms) and speculative investors into the art ecosystem, often diluting artistic value in favor of financial value. As one commentator put it, the NFT market “disengages art from objects” and treats art purely as a tradeable asset (Taylor & Sloane, 2021, as cited in Calvo, 2024). Additionally, legal and ethical issues around ownership and authenticity

complicate the economics. Cases of art theft and copyright infringement—where works were minted as NFTs without artists’ consent—have undermined trust (Gibson, 2021). Intellectual property scholars note that in a digital realm of infinite reproducibility, the concept of owning a “unique” copy of an artwork is problematic (Frye, 2021; Joselit, 2021, both as cited in Calvo, 2024). This critique holds that NFTs create artificial scarcity that may be more valuable to speculators than to artists’ creative vision or remuneration.

A widely cited barrier to adoption has been the ecological impact of blockchain, especially evident during the NFT boom of 2020–2021. Early on, most NFTs were built on energy-intensive proof-of-work blockchains (notably Ethereum pre-2022), which consume enormous electricity for mining and transactions. This gave rise to heated controversy: numerous artists and activists decried NFTs for their carbon footprint and contribution to climate change. Scholarly critiques likewise include environmental unsustainability as a key ethical issue with cryptoart (Calvo, 2024, 2). Although technological developments (such as Ethereum’s later shift to a less energy-intensive model) significantly reduced the carbon costs (notably by 99% in the case of Ethereum), the damage to blockchain’s reputation among many artists has been significant. Many remain unconvinced, noting that efficiency improvements do not fully erase the fundamental issue: the push for artificial scarcity and incessant transactions inherently encourages waste. Some theorists argue that even in a future with green blockchain energy, “scarcity is not something any free society should be aiming to produce” as a matter of principle (Whyman, 2021, as cited in Notaro, 2022).

This points to a broader tension in how blockchain is being operationalised within artistic contexts. As *Future Art Ecosystems 3* notes, blockchain platforms are often framed not as spaces of artistic experimentation, but as instruments for building community or enabling transactions. “Art”, in this context, becomes a narrative device—a means to attract users or bootstrap platforms—rather than a domain of experimentation in its own right (Serpentine Arts Technologies, 2022, 79–80). This framing risks reducing artists to content providers within a predefined economic infrastructure, rather than recognising them as infrastructural creators in their own right—agents capable of shaping protocols, interfaces, and systems through which artistic processes unfold.

The report warns that if blockchain systems continue to treat art primarily as a speculative asset or marketing tool, the medium could quickly lose relevance within

the broader field of artistic innovation. The infrastructure itself must remain open to experimentation—across disciplines, across technologies, and beyond the narrow confines of the Web3-native community. As the report puts it, “There is a tendency to ignore historical institutional precedents, and where ‘art’ is more commonly a narrative vehicle for catalysing community and infrastructural development, rather than the site of experimentation itself” (80). To avoid this trap, artists must be empowered not only to make work with blockchain, but to reshape the very conditions under which blockchain art is made—by designing new protocols (for example Kosmas Giannoutakis (2024) imagined an alternative artistic consensus mechanism called a proof-of-creative-contribution), creating alternative economies, and proposing speculative futures for compositional and curatorial practice (such as the projects of Holly Herndon and Mat Dryhurst discussed in Zeilinger 2024). In this sense, the artist becomes not only a user of platforms, but a world-builder—someone who actively prototypes the infrastructures. This expanded role—as infrastructural designer and, ultimately, world-builder—sets the stage for the next chapter, which turns to the question of how artists can actively shape the ontologies, imaginaries, and ecologies of future art ecosystems.

Artists as System-Builders

Artists today are no longer defined solely by the works they produce or the tools they use. Increasingly, they operate as designers of systems—constructing dynamic processes, rule sets, and machinic environments through which aesthetic experiences and creative relations unfold. In this expanded role, the act of composition shifts from crafting outcomes to engineering conditions of their generation. From Cage's use of chance to Xenakis's formalisation of stochastic structures, from Brian Eno's generative software to the machinic grammars of web-based generative music, we witness a shift from composition as inscription to composition as system design. This shift intensifies in blockchain-based practices, where smart contracts, NFTs, and various decentralised applications allow artists to embed rules of interaction, reproduction, and authorship directly. What unites these otherwise disparate practices is a shared commitment to building systems that do not merely produce outputs, but redistribute agency and remain open to recomposition. In this context, the system becomes the artwork.

The system-builder operates through assemblages composed of both human and non-human actors, but also code, interfaces, sensors, protocols, and feedback loops. Agency is distributed—located not in isolated decisions, but in the patterned interactions that unfold within and across these systems. The artist's role becomes not only one of determining how autonomous processes are activated, how transformations are possible, and how responsibility is shared across networks, but also how are they made, what do they consist of, and what operations enable them. Whether through the stochastic logic of musical scores, or the programmable autonomy of smart contracts, the artist as system-builder choreographs a field of relations in which action is structured, negotiated, and refracted through the new systems built. The system-builder composes infrastructures for action—frameworks that structure how meaning, authorship, and participation unfold. Composition becomes the construction of systems that define procedures, automate transformations, and embed logic into the artwork itself. These systems are not neutral tools but expressive media—technical forms through which artists encode compositional intent, distribute control, and structure variability. In building such systems, artists redefine authorship at the level of process, creating works that remain active, procedural, and open to future reconfiguration.

Chapter 4. Artists Worlding Ecosystems

Cybernetics and Systems Aesthetics

The system-building impulse in art discussed in the previous chapter aligns closely with the origins of cybernetics, a transdisciplinary field first defined by Norbert Wiener in 1948 as “the science of control and communication in the animal and the machine” (Wiener, 1948, 11). Emerging in the wake of World War II, cybernetics—pioneered by Norbert Wiener and expanded by Ross Ashby—offered a radically new way of understanding systems—organic or technical—not by their material composition but by their patterns of behaviour, information exchange, and feedback. Rather than asking “What is this thing?”, cybernetics asked, “What does it do?” (Ashby, 1956, 1). As Ross Ashby observes, it “deals with all forms of behaviour in so far as they are regular, or determinate, or reproducible” (1956, 1), but it also ventures into the domain of the possible. “Cybernetics envisages a set of possibilities much wider than the actual”, Ashby wrote, “and then asks why the particular case should conform to its usual particular restriction” (3). Ashby’s emphasis on “the possible” relates to his concept of variety—the number of possible states a system can occupy—and his Law of Requisite Variety (also known as Ashby’s Law), which argues that controlling a system requires matching its complexity, in other words that a system can only effectively control or regulate another system if its own internal variety (or complexity) is at least as great as the variety of the system it is trying to control (1956, 206).

Ashby emphasised that cybernetics took “as its subject-matter the domain of all possible machines”, caring little whether those machines have actually been constructed “either by Man or by Nature” (1956, 2). Cybernetics thus proposed an ontological flattening between the biological, the artificial, and the social. As Ashby noted, it was expected “to reveal a great number of interesting and suggestive parallelisms between machine and brain and society”, providing “the common language by which discoveries in one branch can readily be made use of in the others” (1956, 4). The significance of this epistemological bridging cannot be overstated: cybernetics was one of the first frameworks to treat cognition, organisation, and feedback as structurally isomorphic across domains, laying the groundwork for system-based thinking in everything from neuroscience and ecology to computation and aesthetics. Cybernetics also positioned itself as a science of complexity before that

term had become a defining keyword in the sciences or the arts. “Science today is taking the first steps towards studying ‘complexity’ as a subject in its own right”, Ashby wrote, adding: “Prominent among the methods for dealing with complexity is cybernetics” (1956, 5).

The influence of cybernetics on the arts began to take hold in the 1960s, particularly through the work of artists and theorists who reimagined art as a system of relations and behaviours rather than a static object. Among the most influential figures in this shift was Jack Burnham, whose 1968 essay “Systems Esthetics” proposed a new way of understanding art that foregrounded the systemic relationships between elements, participants, and environments. Burnham argued that the focus of advanced art was moving away from discrete, self-contained objects and toward systems that evolve over time, shaped by feedback, interaction, and contingency. In this view, the artwork becomes a dynamic field of operations—a structure for inquiry, not a finished artefact (Burnham, 1968). Parallel to Burnham’s theorisation, Roy Ascott was developing a cybernetic model of art practice grounded in feedback, participation, and behavioural transformation. His 1966 essay “Behaviourist Art and the Cybernetic Vision” outlined a framework in which art could be conceived as an open-ended process involving the continual adaptation of systems in response to environmental or human inputs. In works like *Change Painting* (1960–61) and later in his telematic projects, Ascott sought to construct artworks that were fundamentally recursive and performative—defined not by their form but by their capacity to change and be changed. His use of cybernetics was not just metaphorical; it guided the operational logic of his artistic practice, embedding the principles of feedback, adaptation, and transformation into the core of the work (Ascott, 2002). Together, Burnham and Ascott exemplify a moment when cybernetics became a methodological and conceptual resource for artists seeking to build responsive, generative systems—an artistic practice no longer centred on form, medium, or content, but on interaction and the design of relational environments.

Complexity Theory and Second-Order Cybernetics

If system-building reshapes the identity of the artwork around process and generativity, then complexity theory offers a framework for evaluating how these systems behave—particularly in terms of their informational structure and aesthetic

potential. As Philip Galanter argues, generative art is best approached through more than just the lens of medium or technology—it is a practice rooted in the conceptual domain of complexity theory (2003, 1). For Galanter, generative art operates across a continuum between random and fixed material, encompassing a vast historical scope from ancient instruments like the previously mentioned Aeolian harp to contemporary algorithmic systems. What unifies these practices is the use of systems capable of producing multiple outcomes—systems that “exhibit a mix of order and disorder” and thus align with the scientific definition of complexity (Galanter, 2003, 11).

One measure of such complexity is algorithmic complexity—sometimes called algorithmic information content. This measure asks how compressible a given system or signal is: the shorter the algorithm required to generate a given output, the simpler the system. A fractal, for instance, may appear intricate for its outcome, but its generative algorithm may be extremely compact. Conversely, a genuinely random signal resists compression, making its algorithmic description longer and thus apparently more complex—though in practice, it lacks meaningful structure (Galanter, 2003, 9). Here, complexity is not about surface richness or unpredictability alone, but about the relationship between structure and randomness. This is where Galanter introduces the notion of effective complexity, drawing on Murray Gell-Mann’s distinction between trivial regularity and total randomness. Effective complexity aims to identify that middle ground: systems that contain enough structure to be meaningful, but enough variability to remain surprising (2003, 10). This has aesthetic implications. Music composed entirely of random notes or composed from a single repeated tone may both be considered as technically “simple” according to this measure, and neither holds sustained aesthetic interest. It is the interplay between order and unpredictability—between constraint and openness—that makes a system artistically generative according to effective complexity.

This observation resonates directly with the aesthetics of generative music and the design of performative systems. When creating a system meant to generate multiple iterations, the composer is calibrating the balance between regularity and deviation. In Galanter’s terms, a generative system achieves effective complexity when it encodes a structure robust enough to persist across iterations, while still allowing sufficient variability for each outcome to feel novel. This aligns with Simondon’s notion of an open machine, which retains a margin of indeterminacy as a condition of

its technical and creative evolution (2017, 17). Complexity theory helps distinguish between systems that generate variation and those that sustain artistic coherence across it. For the generative artist, the goal is to cultivate a fertile middle ground—structured enough to support coherence, yet open enough to allow for surprise. As Galanter notes, while complex systems may dominate contemporary attention, they are not “better” than simple systems. Both have their place in artistic practice, and the decision to work with one or the other reflects a deeper compositional stance toward form, openness, and structure (2003, 15). In the context of system-building, this complexity-sensitive lens becomes indispensable—not only for interpreting generative systems but for designing them.

Beyond first-order cybernetics, a broader understanding of complexity is further developed in second-order cybernetics, where systems are no longer treated as passive mechanisms observed from the outside, but as self-regulating, self-producing entities that include the observer as part of the system. Pioneered by thinkers such as Heinz von Foerster, Humberto Maturana, and Francisco Varela, second-order cybernetics introduced concepts like autopoiesis and operational closure to describe how systems generate their own identity through recursive interactions. While first-order cybernetics was foundational in modelling systems through feedback and control, it often did so by abstracting systems as trivial machines in von Foerster’s later terminology—predictable devices with invariant input-output relations (von Foerster, H., 2003, 208). This abstraction enabled formal rigour but excluded the observer and structural change. Second-order cybernetics emerged in response, foregrounding non-trivial machines whose behaviour depends on internal states, history, and the inclusion of the observer in the system (208). In this framework, complexity is about the system’s capacity to adapt, reconfigure, and respond to its own internal states. Crucially, these systems exhibit a kind of reflexive intelligence: they do not simply behave, they respond to how they behave.

Individuation and Allagmatics of Artistic Research

On the philosophical level, a parallel framework can be found in Gilbert Simondon’s philosophy of individuation. Rather than viewing identity as a pre-given state or stable essence, Simondon conceives of being as a process—as a continuous becoming that never stops. Individuation, in this view, is the ongoing (and never finished)

resolution of tensions, where entities take shape through operations that continuously restructure the metastable operational field that these entities belong to. Structural identity is thus only a partial expression of a process—one that is always in flux, continuously unfolding through transduction: where each new local resolution changes the conditions for what comes next. For example, Yuk Hui—philosopher of technology—builds on this notion in the context of digital ontology, where HTML properties within a website’s structure are best understood not as fixed representations or tools, but as “indexed relations”—material instantiations of relational aesthetics whose logic emerges through interaction (Hui, 2015, 6). The behaviour of such objects cannot be fully explained by their code alone; rather, it arises from the interplay of structural relations, contextual constraints, and feedback loops embedded in their design. For this reason, Hui contends that traditional logics of deduction (from truths to facts) or induction (from facts to truths) are inadequate to account for their operation. Instead, he proposes transduction as a more appropriate logic—one that accounts for transformation, rupture, and the invention of new consistencies (3–4). In this framework, systems are not defined by their fixed properties, but by their capacity to reorganise themselves while maintaining operational coherence.

While Simondon’s philosophy of individuation is today widely read through the work of philosopher Gilles Deleuze—particularly in relation to assemblage theory—its relevance to systems thinking invites a different kind of reconstruction. Simondon was not writing from within the cybernetic tradition, nor responding to it directly. Rather, his work on individuation emerged in the early 1950s, contemporaneously with the development of first-order cybernetics by Wiener, Ashby, Shannon, and others. While cybernetics at that time focused on modelling behaviours, communication flows, and feedback mechanisms in organisms and machines, Simondon developed a philosophical ontology of becoming—one that asked how systems and structures come into being through operations immanent to metastable fields. It is within this ontological framework that Simondon proposed the notion of allagmatics, a general theory of operations. Whereas structuralism was concerned with fixed forms and stable configurations, allagmatics sought to understand how operations generate structure from within: how a system individuates itself through transduction, the process by which tensions are resolved through successive structural transformations. Though Simondon never developed a technical science of allagmatics, his philosophy provides the conceptual grounding for what such a theory

could be. In this sense, Simondon's thought can be seen as anticipating key concerns of second-order cybernetics, particularly the shift toward self-producing, reflexive systems that include the observer in the process of transformation.

This link between ontogenetic individuation and operational system-building has reemerged in more recent theoretical articulations of artistic research. Paulo de Assis, drawing on Simondon via Deleuze and Guattari, has proposed the notion of an "allagmatics of artistic research" (2018, 27)—a method that privileges process over form, transformation over representation, and individuation over interpretation. De Assis, when discussing his conceptualisation of musical works as assemblages, suggests that musical works can be conceived not only in terms of systemic, but also in terms of their epistemic complexity, which addresses how systems encode, transform, and generate knowledge over time. As he explains, drawing on the work of biologist Ladislav Kováč and philosopher Subrata Dasgupta, epistemic complexity refers to the richness of knowledge embedded within a technological or artistic artefact—its capacity to generate new knowledge, produce surprising behaviours, and evolve over time (de Assis, 2018, 125–126). In this sense, an artefact is an epistemic entity: a form of embodied understanding, a machine for thinking, and a bearer of invention. As Dasgupta notes, "in the case of true invention, when the artifactual form is original in some significant sense, the operational principles it encodes constitute genuinely new knowledge" (as cited in de Assis, 2018, 125). As de Assis argues, epistemically complex artefacts resemble living organisms in their ability to evolve, adapt, and participate in recursive cycles of transformation (125). They exhibit behaviours rather than simply delivering functions. Artists often engage with artefacts that already contain epistemic complexity: scores, sketches, performance traditions, instrument designs, historical recordings, or theoretical frameworks. These artefacts act as boundary objects whose epistemological status shifts depending on the context in which they are activated (de Assis, 2018, 126). A musical manuscript may function as an archive in one context, a springboard for reinterpretation in another, and a generative template in yet another. The most potent outputs of the artistic research in this understanding are artefacts that actively restructure their own epistemic lineage in the progress of research. "An experimental system becomes a machine for making the future", as de Assis puts it (2018, 28). In artistic research contexts, transduction describes the unfolding of new configurations—not by copying or programming outcomes, but by activating potentialities that emerge through recursive and self-reflexive operations of the artist.

Allagmatic Composition—Composing for Epistemic Complexity

By treating musical artefacts as epistemically complex entities, and performances as experimental systems that reconfigure those artefacts, Paulo de Assis effectively frames artistic research as a practice of creative problematisation. While he does not explicitly theorise this in terms of compositional ontology, his work nonetheless anticipates the possibility of a compositional mode grounded in the design of epistemically active systems. This doctoral research builds directly on this foundation by proposing a formalisation of that mode—what I call allagmatic composition. This distinction can be useful when discussing generative and system-based practices. While an artwork may be composed to be systemically complex—consisting of numerous modules, processes, and data interactions—its epistemic value lies in the originality and extensibility of its operational principles. In generative music, such knowledge is not always transparent or immediately legible; it becomes intelligible through interaction—through the system’s capacity to produce outcomes not fully anticipated by its creator. To compose an epistemically complex generative music system might thus be understood not only as to compose a structure of interacting modules, but to enact a knowledge-generating machine—capable of reshaping the very aesthetic, technical, or philosophical contexts into which it is introduced. Here, the generative artist becomes not only a designer of systems, but a constructor of epistemic topologies: someone who maps and remaps relations between histories, operational principles, agents, and their emergent and responsive behaviours. As Simondon writes: “A structure is the result of a construction, and an operation is that which makes a structure appear, or that which modifies a structure” ([1954–58] 2013, as cited in de Assis, 2018, 27). The artwork, then, is not a static form nor even a set of generative procedures, but a metastable machine that continually transforms its own operational relations.

Allagmatic composition, then, is not a compositional style, but a compositional ontology. It defines music not as a sequence of notes or events, nor even as a product of a generative engine, but as a system that performs itself—a system whose mode of being lies in its becoming. Where generative systems may operate within fixed affordances, selecting and varying pre-defined materials, allagmatic systems are capable of transforming their own internal logic. They are machines that not only generate outputs, but modulate their own epistemic topologies of operation in the process of execution. This has far-reaching consequences for how we conceptualise

compositional authorship and agency. In allagmatic composition, the composer configures a space of generative epistemic relations—an assemblage of agential transformations whose internal dynamics give rise to novel individuations. Each execution of the system is a transductive event—an enactment of change that reconfigures not only the output, but the space of possibilities from which future outputs may arise. In allagmatic composition, the “work” is no longer a static object, but an ongoing process or a system that is capable of transforming itself. The composer’s role shifts from a traditional author to a metadesigner of systems that can further model their own rules. The identity and existence of the composition are tied to its dynamic behaviour rather than to a fixed form.

This approach may, at first glance, appear to resemble what has been described in recent decades as interactive composition. As Artemi-Maria Gioti notes, interactive compositions typically involve a mutual real-time adaptation between human performers and interactive computer music systems, the design of which is often composition-specific and idiosyncratic (2021, 5). These systems are characterised by their ability to respond dynamically to external inputs, enabling open-ended, co-creative, and contingent performances. Yet while such works allow for considerable variability in performance, the underlying operational logic of the system—the rules, affordances, and generative structures—often remains fixed once designed. Interaction, in this context, unfolds within a pre-established framework. By contrast, the focus of allagmatic composition is on the possibility of transforming that framework itself. It does not merely accommodate variation in outputs, but enables the system to reorganise its own internal logic, modifying the very conditions of future behaviour from inside of the system itself, rather than through its interactions with the external world. In this sense, interaction becomes structurally transductive: each execution reconfigures the system’s epistemic space and expands its compositional potential.

Of course, no system comes from nothing. Even self-organising systems require: definitions of components, parameters for change, and an initial logic of operation. Allagmatic composition is not defined by the fact that it changes, but by how change is treated as the compositional object itself. Generative composition focuses on producing output (variations). Interactive composition focuses on input/output relations with the world external to the designed system (adaptation, exchange). Allagmatic composition focuses on modulating the operational space itself—

the configuration of rules, transformability, and epistemic affordances. So it's not that the system isn't designed in advance—it always is. But in allagmatic composition the rules that govern the rules (meta-logic) become part of the composition. The system is built to reconfigure itself structurally, not just behaviourally. The composer's role shifts from designing what the system does to designing how it may become otherwise.

Generative composition creates variation within a fixed system, while allagmatic composition iteratively builds systems that further transform themselves through each act of creation. Generative composition changes the output, while allagmatic composition changes the system that makes the output. In generative composition, the system might be complex in structure (many parts, many outputs), but its epistemic complexity is limited to what was programmed—it doesn't learn or change its way of thinking. It generates content, but it doesn't generate new kinds of knowing. In allagmatic composition, the system can change itself—its rules, logic, and behaviour evolve over time. This means it can produce new knowledge, new forms, and new ways of relating elements. It's epistemically complex because each execution can transform the space of what the system knows and how it knows it.

A musical example that might be considered to align with this concept can be found in *Polyp* by Marek Poliks and Roberto Alonso Trillo, presented at the Musica ex Machina exhibition at EPFL Pavillions in Lausanne (Kenderdine et al., 2024, section 1). *Polyp* consists of three silicone-based AI sculptures that continuously train machine learning models on ambient sounds captured from their environment. As the installation progresses, each device modifies its sonic and luminic behaviours in real time, developing a dynamic and adaptive understanding of its surroundings. These Polyps do not simply replay or reinterpret pre-trained data; they evolve through continuous feedback, upload material to the cloud, and reciprocally adapt in response to one another. Their system thus changes not just what it produces, but how it produces—reshaping its internal logic through exposure, interaction, and mutual transformation. In this sense, *Polyp* exemplifies allagmatic composition: it is not a system that generates variation within fixed parameters, but one that composes through transductive change, expanding its epistemic space with each activation.

Where generative music explores variation within a procedural field, allagmatic composition explores variation of the procedural field itself. It composes with systems rather than simply within them—constructing open machines capable of reorganising

their own topological configuration. The next section discusses some further artistic examples that might exemplify the mode of creation in which to compose is to craft a machine that iteratively extends its own modes of operation.

New Kinds of Strange Agential Assemblages

As artistic systems become increasingly entangled with the blockchain infrastructure, a new category of aesthetic entities can be distinguished—assemblages in which agency is no longer anchored in human intention alone; artworks that are themselves active participants in a machinic assemblage. Described by Martin Zeilinger (2022) as “new kinds of strange agential assemblages”, these are not simply artworks executed by machines, nor representations of non-human life, but configurations in which agency itself is performed through code, protocol, and environmental feedback. They exemplify a profound rethinking of authorship, autonomy, and artistic form—one that Zeilinger has aptly described as blockchain vitalism (Zeilinger, 2022).

Zeilinger identifies a slowly growing tendency in contemporary art to design systems that do not merely represent living processes, but participate in them as operative agents. Projects like Primavera De Filippi’s *Plantoid* (De Filippi, 2017) and Paul Seidler’s, Paul Kolling’s, and Max Hampshire’s *Terra0* (Seidler et al., 2017) are emblematic in this regard.

Plantoid is a robotic sculpture of a plant that receives cryptocurrency donations via a smart contract. Once it collects sufficient funds, the contract automatically initiates a commissioning process for the creation of a new *Plantoid* sculpture—thus enabling procedural reproduction without the oversight of a single author. It is an example of a self-reproducing artwork: a system designed not only to persist, but to propagate its own existence through coded rules and distributed agency. Here, the artwork operates as a self-propagating system, one in which authorship, aesthetics, and reproduction are mediated by code.

Terra0 takes this logic further. It proposes a model by which a physical forest in Germany—monitored by sensors and administered through smart contracts—can gradually become its own legal and economic subject. The system is designed to sell timber according to sustainability constraints, use the revenue to buy back its land, and eventually attain functional autonomy. In this formulation, the forest is not just

metaphorically alive—it is an economic agent embedded within a legal framework. As Zeilinger argues, such works “disentangle agency from subjectivity” by recasting authorship as an emergent property of procedural, distributed systems. The result is not a shift from human to non-human authorship, but a more radical move: the production of agential assemblages in which authorship becomes an emergent feature of machinic and ecological entanglement.

This reframing of authorship finds a more systematic articulation in Zeilinger’s 2024 contribution, where he extends the notion of agential assemblages to examine how blockchain and AI technologies, both separately and in tandem, enable new modalities of creative agency. He proposes that each technology contributes distinct conceptual transformations: “AI”, he writes, “is helping to shape new perspectives on the nature of creativity as such”, while “blockchain is reshaping our understanding of the concepts of originality and authenticity, and of the capacity for aesthetic artifacts to exhibit semi-autonomous behaviors” (Zeilinger, 2024, 61). Artificial intelligence displaces human intention by generating content that exceeds or escapes authorial foresight, destabilising long-standing assumptions about skill, expression, and inspiration. Blockchain, meanwhile, reconfigures the infrastructural conditions under which art is attributed, authenticated, and enacted—shifting focus from authorship as an expressive act to authorship as a programmable logic of ownership, reproduction, and governance.

But it is when these two technologies are combined that more radical possibilities emerge—particularly around the status, and traceability of agency. Zeilinger suggests that when AI-generated artworks are embedded in blockchain-enabled infrastructures, we are forced to confront a set of unsettling questions: What is the nature and locus of the agency underlying the creation of a specific artwork? Who (or what) can be considered the author or artist responsible for creating the work? If no human artist can be meaningfully linked to an aesthetic expression, can such an expression be subject to conventional ownership claims? What are the boundaries of an artwork that is distributed across a decentralised network of computational nodes? (61). These are not abstract philosophical queries but urgent compositional problems, particularly in fields like music, where agency is often already temporal, distributed, and collaborative.

In my contribution to the same edited volume (Łukawski, 2024a), I argued that blockchain infrastructures afford a new kind of temporal composability, especially

through the interoperability of smart contracts. Because each contract can reference, inherit, or activate others across time, blockchain enables layered, time-based works that embed diverse agents—human performers, AI models, procedural logics—within a shared operational framework. This affordance allows for the design of compositional ecologies in which agency emerges through transactional relations and procedural triggers distributed across time and space. In such a system, the artwork is an unfolding process—one that leaves behind a verifiable provenance not only of outputs, but of each transformation, condition, and actant involved in its becoming.

This shift in focus—from the artwork as a finished object to the system as a traceable field of operations—resonates closely with the earlier discussion of allagmatic composition. There, we framed the compositional act as the configuration of metastable systems capable of reorganising themselves through transductive processes. In allagmatic terms, the artwork is not the product of a structure, but the outcome of an operational chain that individuates itself over time. Blockchain, in this context, does not merely serve as a registry or a marketplace, but as a diagrammatic medium: a technical substrate that allows these operational chains to be instantiated, referenced, and recomposed across distributed environments. The significance of this affordance is ontological. By enabling the persistent inscription of each procedural step, blockchain makes it possible to design systems that are self-differentiating and procedurally accountable: assemblages in which every act of transformation is both compositional and composable, as well as allowing agential inputs to such an assemblage to be effectively tracked and incorporated, enabling open systems to adhere to nevertheless structured behaviours.

A musical counterpart to such performative systems can be found in *Hypermusic Experiment 0.9*, a conceptual project developed by composer Einar Torfi Einarsson that envisions combining AI, blockchain, and notation into an open-ended compositional infrastructure. Rather than composing fixed works, the system generates partial scores—inscribed artefacts that function as modular nodes in a distributed, evolving network. The goal of the project is to model an ongoing, relational process in which each interaction—whether a performer’s interpretation, an API-triggered data stream, or a blockchain transaction—reshapes the creative field (Einarsson, 2024, 130). Described as an “infinite notation-machine”, the system operates across digital and physical domains, fusing inscription, inference, and interaction in a single compositional ecology (130). Inputs to the system range from traditional metadata to

real-time environmental data, all feeding into a multi-piece structure designed to behave as a network of aesthetico-epistemic components (141). By enabling both human and posthuman agents to influence its operation, *Hypermusic Experiment 0.9* suggests composition as a practice and a speculative platform that challenges the boundary between notation, execution, score, system, and composer. It exemplifies how blockchain and AI can be co-articulated to support compositional worlds that evolve across time, infrastructure, and participation.

The described examples are not only generative systems in the classical sense, nor are they interactive installations governed by deterministic scripts. They are recursive, self-modulating assemblages in which agency circulates across layers of code, machine learning inference, environmental data, human participation, and formal logic. They generate the very conditions for their own mutation and recomposition. In this sense, they exemplify what Zeilinger calls “strange agential assemblages”: artworks that exceed representational models, operate without fixed centres, and instantiate agency through the entangled execution of technical, legal, and aesthetic protocols. To encounter such a system is to enter its field of becoming—a machinic ecology in which authorship, value, and action are continually redistributed. These works do not ask to be judged or admired; they ask to be activated, enacted, and recombined. And in doing so, they compel us to confront a fundamental inversion: the artist is no longer the originator of form, but the architect of a process that is ontogenetically open—a process through which agency itself is brought into composition. In this way, blockchain and AI do not just augment artistic tools, but together, they enable a new regime of composition: one that is procedural, composable, and fundamentally agential.

Envisioning Abstract Machines

Having examined how system-building practices formalise generative procedures across human and non-human agencies and algorithmic or epistemic complexities, we can now turn to the next threshold: exploring how artists can envision the operational conditions from which such systems emerge. This is the conceptual territory of the abstract machine—a notion developed by Gilles Deleuze and Félix Guattari in *A Thousand Plateaus* (1980) and further elaborated by Félix Guattari in *Chaosmosis* (1995).

In contrast to machines in the traditional sense—devices composed of parts to execute predetermined functions—the abstract machine is not an object at all. It is not something you can point to or build in the conventional sense. Rather, it is a diagram: a non-empirical structure that defines how elements might relate, combine, or transform within a field of potential. It is, as Guattari writes, a “montage capable of relating all the heterogeneous levels that it traverses” (1995, 35).

This shift reorients our understanding of technology itself. Instead of seeing machines as mere expressions of technological progress, Guattari suggests that machinism—the broader conceptual field that includes social, semiotic, psychological and technical elements—must be understood as a prerequisite for technology (1995, 33). That is, machines are not simply tools for doing; they are conditions for becoming. When we speak of abstract machines, “abstract” here also means “extract”: the machine extracts and mobilises forces across domains—biological, symbolic, economic—without being reducible to any one of them (Guattari, 1995, 35). In this sense, an abstract machine does not simply organise a system; it conditions what kinds of systems can come into being at all. In practical terms, an abstract machine defines the relations between elements before any of those elements are fixed or instantiated. Thomas Nail—philosopher of movement and new materialism—puts this succinctly: “The condition of an assemblage is the network of specific external relations that holds the elements together. Deleuze and Guattari’s name for this set of conditioning relations is the ‘abstract machine’” (Nail, 2017, 24). This means that every actual system or assemblage we encounter—a piece of software, a musical work, a performative structure—emerges from an abstract machine that conditions its form, behaviour, and possible variations.

Yet abstract machines alone do not produce concrete effects unless they give rise to what Guattari calls a machinic assemblage. Assemblages are specific configurations of elements (sounds, tools, gestures, agents, algorithms, and musical works) that are swept up by the abstract machine’s diagrammatic logic (Guattari, 1995, 35). More than combinations of parts, these assemblages are also dynamic fields composed of material and expressive components, human and non-human agents, technical apparatuses and semiotic flows. As Deleuze and Guattari explain, an assemblage is defined not by its components alone, but by the “lines” that connect them—its patterns of segmentation, intensification, and transformation (Deleuze & Guattari, 2003, 21). It is important to understand that these assemblages include agents, but not

in the traditional humanist sense of autonomous individuals. Deleuze and Guattari call these agents “personae”—mobile operators immanent to the assemblage itself. Personae are roles, functions, or operators that enact transformations within the system without standing outside it. As Nail explains, “personae are not the origin of the assemblage and do not control or program the assemblage in advance. Rather, they are the immanent agents [...] like a runner or intercessor” (2017, 27). In many generative artistic systems, such agents might take the form of algorithmic processes, machine learning models, or human performers—each operating within a shared diagram of relations.

This view of immanent agents within assemblages resonates with Alfred Gell’s distinction between “primary” and “secondary” agency, as discussed by Artemi-Maria Gioti (2021, 15-16). Drawing on Gell’s anthropological theory of art, Gioti notes that a “primary agent” is typically a human being with intentionality, while a “secondary agent” is an object, artefact, or artwork through which that agency is distributed and enacted. Gell defines an agent as “someone or something that causes events to happen in their vicinity” (Gell, 1998, 16 as in Gioti 2021, 15-16), and stresses that artefacts become social agents not through mystical properties, but because they function as fragments or extensions of human intention (21). Importantly, secondary agents can also act upon human actors—shaping perception, triggering decisions, or structuring behaviour. In generative or interactive art, this reframing of agency becomes especially relevant: systems, scores, or tools may not possess consciousness, but they nonetheless operate within assemblages as agents capable of shaping creative outcomes. As such, they function analogously to Deleuze and Guattari’s personae—operators that enact transformations without standing outside the systems they participate in.

Guattari, drawing on Francisco Varela’s work, differentiates between two types of machines: *allopoeitic* and *autopoeitic*. *Allopoeitic* machines produce something other than themselves (e.g. a factory machine producing widgets), while *autopoeitic* machines generate and sustain their own organisation (Guattari, 1995, 39–40). Biological organisms are classically *autopoeitic*: they regulate their internal processes and maintain their boundaries in relation to the environment. Guattari extends this idea to complex social or technical assemblages—suggesting that when machines, humans, and semiotic systems are sufficiently entangled, they may collectively exhibit *autopoeitic* behaviour. These assemblages do not simply carry out instructions; they

adapt, evolve, and reorganise their own internal relations.

This ontogenetic capacity is perhaps best captured by Deleuze and Guattari's notion of the rhizome. Unlike hierarchical models of thought (the "tree" or "root" structure), a rhizome is a decentralised, non-linear system where "any point can be connected to any other" (Deleuze & Guattari, 2003, 7). A rhizome grows in the middle, not from a singular origin or toward a final goal. It has no fixed subject or object, but instead operates through multiplicities, connections, and intensities. Rhizomes do not mimic established models; they proliferate difference. They are made of "lines of flight", directions of escape from pre-given categories and forms (21). In artistic practice, thinking rhizomatically means resisting rigid genres, forms, or authorial hierarchies. It means creating systems that allow for distributed agency, divergent outcomes, and emergent meaning. As Edward Campbell—scholar of contemporary European art music—notes, postwar musical practices such as graphic scores, impulsion texts, and gesture-theatre exemplify this shift. These works no longer define a musical object to be realised, but a space of potential relations—a rhizome of gestures, sounds, and interactions that unfold differently with each instantiation (Campbell, 2013, 42–44). In such practices, composition becomes the crafting of a machinic diagram: a set of open-ended conditions under which performance, interpretation, and transformation can occur.

This orientation toward ontogenetic openness finds resonant expression in a variety of compositional strategies that challenge the idea of the work as a fixed, autonomous object. As Paulo de Assis notes, Klaus Huber's concept of polyworks—elaborated in pieces such as *Schattenblätter* (1975)—designates musical objects composed of several self-contained works that may be performed either together (superposed) or independently in various configurations (de Assis, 2024, 24). Claus-Steffen Mahnkopf defines the poly-work as a compositional strategy grounded in complexity, where internal differentiation unfolds through principles such as polymorphy, polyprocessuality, and polyconceptuality. According to Mahnkopf, "the specific however [poly-work] is that within the (organically devised) work there are other (similarly organic) pieces as excorporable, which can stand for themselves and are to be performed independently" (Mahnkopf as emphasised in Ulus, 2021, 24). This notion of organically devised material—where autonomous works are embedded within a larger formal entity—aligns with the idea of an assemblage that resists unity and closure. Hakan Ulus lists several compositions exemplifying this principle,

including Richard Barrett's *Opening of the Mouth* for ensemble, Mahnkopf's own *Medusa* for oboe and chamber orchestra, Simon Steen-Andersen's *Next To Beside Besides*, and Chaya Czernowin's *Anea Crystal* cycle for two string quartets, among others (Ulus, 2021, 31). These works offer models for compositional systems that remain open, stratified, and capable of individuating new configurations from within their own structural logic. A related concept appears in Einar Torfi Einarsson's idea of multipieces, or multiplicity compositions, which are built from a modular structure of interrelated parts that can be combined in thousands of ways. As Einarsson writes, such works never appear as fixed entities, but only through partial, situated performances—each one an “open window onto that ever incomplete ‘whole-less’ activity” that constitutes the identity of the piece (Einarsson, 2024, 132).

Thus, more than building closed machinic systems, the artist can envision abstract machines that cross the boundaries between various types of entities—by structuring diagrammatic constellations of relations as higher-level operations giving shape to various assemblages. These abstract machines remain virtual; they imagined by the artist as configurations of potentiality across technical, social, institutional, aesthetic, and conceptual dimensions. The artist's task is then to first envision such abstract machine, and then operationalise this vision—not by translating it into a static blueprint, but by setting it in motion through the construction of a machinic assemblage (such as a set of scores, generative systems, social mechanisms, events, and other entities dependent on the defined operations). The artist creates a set of relations—a machinic infrastructure—that makes possible a range of aesthetic processes, some of which may not even be enacted by the artist themselves. Others may instantiate, remix, or extend the machine. In this sense, the abstract machine becomes a compositional space and a generative ecology.

Worlding

The conceptual leap from creating systems to constructing the conditions for systems to emerge marks the threshold into what has recently been termed worlding. As observed in *Future Art Ecosystems 1*, the term “worlding” has increasingly been adopted by artists to describe practices that were previously framed in terms of “networks”, but which now emphasise autonomy, generativity, and systemic

continuity beyond the originating artist (Serpentine Arts Technologies, 2020, 20). Where networks imply connection, worlds imply condition: a topology within which relations unfold and new entities are born. Future Art Ecosystems 3 further characterises worlding as a mode of practice concerned not merely with producing artefacts, but with “shaping how worlds are built” (Serpentine Arts Technologies, 2022, 140).

The philosophical origin of the term lies in the work of Martin Heidegger, for whom “worlding” (German: *das Welten*) referred to the ontological unfolding of being-in-the-world. In this view, worlding is not something done by a subject to an object, but an ongoing process of disclosure—how things come into presence and meaning through our situated engagement with them. However, recent critical appropriations of the term have distanced it from Heidegger’s anthropocentric and human-exceptionalist framing. As Donna Haraway insists, she is “finished once and for all with Kantian globalizing cosmopolitics and grumpy human-exceptionalist Heideggerian worlding” (Haraway, 2016, 11). Instead, Haraway redefines worlding as a speculative process through which humans and non-humans co-compose one another. In her formulation, worlding becomes a practice of committing to the complexity and unfinishedness of relational becoming (Haraway, 2016, 31). Worlding, then, is not the construction of a singular world but an ongoing, ontogenetic process of co-articulation across domains of matter, meaning, and narrative.

Thomas Rickert and Michael Salvo—scholars of rhetoric and media—propose worlding as a key conceptual strategy for understanding how aesthetic experience in new media is no longer bound to isolated works but emerges through immersive, distributed, and technologically entangled environments (2006). In their analysis of contemporary sound practices—from Jimi Hendrix’s feedback to Brian Eno’s ambient systems—worlding names a compositional mode wherein sonic, visual, technological, and social dimensions coalesce to generate situated, participatory worlds. This distributed *gesamtkunstwerk* dissolves the binary between production and reception, casting the artist as a node within evolving machinic ecologies of composition and experience.

A striking example of immersive world-building can be found in the work of Refik Anadol Studio, which has developed a series of large-scale installations that construct what Anadol and Kivrak (2023) describe as “data-driven dream spaces”. These works generate self-sustaining aesthetic ecologies through the processing and

transformation of vast datasets. Projects such as *Archive Dreaming* (2016), *WDCH Dreams* (2018), and *Quantum Memories* (2021) demonstrate how AI models trained on cultural and environmental data can produce continuously shifting perceptual environments. In *Living Architecture: Casa Batlló* (2022), for instance, machine learning was used to respond in real time to environmental inputs, thereby transforming a heritage building into a living, adaptive interface. These works exemplify how AI can be integrated into worlding practices as an agent in the co-creation of responsive, ontogenetic environments that blur the boundary between data, space, and subjectivity.

Another example of such world-building is the practice of the art collective teamLab, for instance their *Borderless* exhibition. The exhibition constructed a continuous, immersive environment in which boundaries between artworks, viewers, architecture, and technology dissolved entirely. Digital flowers bloom and scatter in response to one's motion; waterfalls shift their flow depending on the contours of bodies in the space. Works move across rooms, merged with others, or vanish altogether—all creative elements acting as agents within a larger dynamic ecology. Withing the exhibition's logic the artworks become components of a responsive machinic environment. "[T]he visitor is admitted into the art collective's vision of a 'borderless world' awash in a kaleidoscope of colours and imagery, comprised of artworks which seem to possess a life of their own as they transform and react to the visitor" (Grassi, 2022, 1-2). While clearly designed to appeal to a broad public, the exhibition—conceived by an interdisciplinary collective of artists, designers, engineers, and programmers—stands as an impressive large-scale undertaking that integrates all layers of the contemporary art-technology stack.

Projects as the ones described above are often the result of large, interdisciplinary teams with access to significant technological and institutional resources. However, the conceptual orientation they embody—toward worlding as a mode of composition—need not be limited to monumental scale. Rather than dismiss these practices as out of reach, they invite a critical question: how might the individual artist, composer, or researcher begin to think in terms of worlding? What does it mean to design the conditions for emergence within one's own practice, and how might this be operationalised in music? The following section turns to these questions, exploring how the logic of worlding can be enacted through modest, yet conceptually rigorous, compositional systems.

Organising Free-Range Creativity

While worlding practices need not rely on monumental scale or institutional support, they nonetheless underscore a fundamental fact: contemporary artistic production involving advanced technologies is rarely the work of an isolated individual. The technical complexity and informational density of such systems typically exceed the capacity of any single practitioner, necessitating collaboration with specialists—whether as part of informal constellations or as embedded members of an artist’s studio (Serpentine Arts Technologies, 2020, 21). High-barrier technologies—complex, expensive, or highly specialised—demand not only access but fluency, often compelling artists to acquire new skill sets or to build relationships with those who possess them (52). In this context, composing within networks becomes an epistemologically generative alternative: it enables the integration of distributed expertise, diverse motivations, and varying modes of value production into a single operational ecology (23). Today’s most ambitious art-tech projects already operate within such ecosystems, depending on small, agile networks of technically fluent artists, independent technologists, and commercially conversant agents (Serpentine Arts Technologies, 2022, 133). In such cases, collective dependence becomes a compositional resource; a means of assembling aesthetic systems from intersubjective and interdisciplinary entanglements.

One conceptual framework that offers a compelling response to these conditions is Olga Goriunova’s notion of the art platform. The art platform is a concrete organisational form—an assemblage of human and technical relations that enables creative production through distributed agency, network logic, and infrastructural responsiveness. Goriunova—cultural theorist of digital media—defines it as “a network platform that produces art, here understood broadly as a process of creative living with networks” (2012, 2). Art platforms emerge as practical and conceptual solutions to the increasing need for artists to coordinate diverse actors, systems, and knowledge domains into adaptive ecologies of production. In contrast to centralised institutions, they function as self-organising systems—crossroads that assemble actors of various scales, capable of co-developing processes and enabling collective aesthetic power. What distinguishes the art platform from other organisational models is its reflexivity: it operates in awareness of its own processual composition, engaging with what Goriunova calls “organizational aesthetics”—a mode of aesthetic production that is itself concerned with how creative processes are structured, amplified, and

sustained over time (13). As a model for composing with collectives, art platforms provide a way of organising what Goriunova calls “free-range creativity” (21). In this sense, an art platform enacts an aesthetic state in which new forms of knowledge, perception, and relation can individuate themselves, rather than just hosting or distributing works (18). They are engines of aesthetic becoming. This makes it particularly relevant to worlding practices: art platforms do not generate singular outputs, but—through the designed limits of the operations that they enable for their users—condition the emergence of multiplicities.

Goriunova describes aesthetics itself as “a machine generating material variants of reality to enable knowledges, practices, and perceptions to constitute and affirm themselves” (2012, 18). Art platforms thrive on such processual openness. This framing situates artistic practice within an ongoing negotiation of the real—where the platform acts less as a container of finished works and more as a dynamic infrastructure for individuating aesthetic potential. Art platforms are engines of transformation, capable of amplifying aesthetic force through their very mode of organisation. They are structurally attuned to fluctuation and modularity, supporting practices that are experimental, collective, and fluid. In this way, art platforms point to a broader ambition: to create systems that enable not just individual expression, but the “empowerment of massive groups of people” through the “large-scale involvement of creative capacities into innovative practices” (Goriunova, 2012, 22). Rather than privileging the singular author or the closed artwork, art platforms embrace a logic of collective emergence. They articulate a politics of composition grounded in openness and flexibility, aligned with a networked ontology in which participation is not an afterthought but a generative condition. Importantly, this mode of organisation is not only aesthetic but infrastructural. It positions the user not merely as a consumer of services but as an agent in the system’s self-organisation—a shift that Irina Kaldrack and Martina Lecker—scholars of digital culture and media theory—identify as a new “form of governmentality”, wherein users both invoke and structure their own modes of participation and control (2015, 18). What emerges, then, is what Goriunova describes as “a complex assemblage of art, networks, technologies, politics, autocreativity, publics, humour—all amalgamated, related, but not frozen together in their operation” (2012, 97). These assemblages are structured yet open-ended, governed by databases, modular toolsets, and user-led interaction models that privilege iteration, contribution, and exchange (9). Such platforms often begin as modest digital tools but evolve into expansive ecologies of aesthetic production. They

are sites where collective knowledge, technical systems, and aesthetic sensibilities intersect—spaces not only for the distribution of artworks, but for the invention of compositional worlds.

Seen through this lens, art platforms are autopoietic systems in the Guattarian sense (Goriunova, 2012, 34-37). They emerge through the dynamic co-constitution of human agents, technical infrastructures, and aesthetic processes, forming what Goriunova calls “a system of human-technical assemblages and arrangements that produce a common aesthetic, political, and creative horizon” (2012, 11). These assemblages evolve through a logic of aesthetic self-organisation—a drive toward emergence that cannot be reduced to either mechanical procedure or human authorship (39). This self-organising character is not simply about distributed participation; it is a mode of cultural individuation, wherein taxonomies, digital objects, scripts, recordings, or user contributions become operative components within a recursively modulating system (12). In this way, art platforms enact what Guattari described as autopoietic formations: assemblages that produce their own consistency through transversality, singularity, and machinic alterity (2012, 37).

Autopoiesis, in this expanded cultural and aesthetic register, ceases to be a model of closed biological self-replication and becomes instead a framework for understanding how collective creativity configures itself across heterogeneous domains. As Goriunova notes, when humans, machines, institutions, and “fields of the possible” are drawn together into relational processes of individuation, they may constitute autopoietic forms (37). These formations operate beyond the control of any singular agent; they do not obey a fixed grammar, but generate their own modes of coherence by modulating across structures, components, and machinic relations. In this sense, art platforms are not representational systems but transductive ecologies—systems that actively produce the conditions under which new assemblages can emerge and sustain themselves. The culmination of this framework lies in Goriunova’s concept of autocreativity. Defined as a machinic, self-organising force of aesthetic emergence, autocreativity is “autopoietic, autonomous, and ‘automatic’ creativity that propels aesthetic emergence in the constitution of the human, the cultural, and the social” (2012, 42). Crucially, autocreativity is not reducible to novelty, nor does it seek to explain or imitate human creative faculties. It operates pre-subjectively—not as an act of individual expression but as a transversal force that passes through technical, natural, and pre-individual domains. It is not aligned with artificial intelligence in the

conventional sense, nor with authorship or intentionality, but with a deeper machinic vitality that organises itself by advancing.

As Goriunova puts it, “autocreativity has the energy to cross thresholds, to effectuate a change, and to divorce itself from the plane of any current stratum” (42–43). It does not produce finished artworks so much as it configures the conditions from which new forms of art, subjectivity, and sociality may arise. This framing allows us to reconceptualise creativity not as a trait of individuals or the output of machines, but as a processual force that emerges between the human, the technical, and the infrastructural. Autocreativity “lies between the human, the technical, and the social”, and reveals how each of these domains can perform creative operations through their integration in dynamic assemblages (43). The role of the art platform, then, is to work autocreativity—to provide the relational milieu through which this force can become operative. Rather than being programmed or assigned, the mechanisms of creativity on art platforms “develop themselves to be passed on to the environment”, producing “a moment of difference” (43).

Platforms condition the emergence of new social formations. In art platforms, this emergence is the core of their aesthetic operation. What is produced is a collective process of subjectivation in which participation becomes a constitutive act. These platforms give rise to complex ecologies of interaction, identification, and organisation, wherein digital objects become “performative enactments of social aspirations and cultural identifications for the communities that gather around them” (Serpentine Arts Technologies, 2022, 51). Their networked and computable qualities thus blur the line between technical projects and cultural worlds, embedding each aesthetic gesture within a larger diagram of social potential. “An art platform is never simply a technocultural object”, writes Goriunova, “but a resonance, a movement, an operation” (2012, 3).

Aesthetic emergence in these platforms is tethered to the operational capacities of software as an active agent of cultural transformation. As Seth Erickson and Christopher M. Kelty—information studies researcher, and anthropologist and science and technology studies (STS) scholar—argue, software is neither material nor immaterial, but “durable, entrenched and scaffolded”, possessing distinct modes of existence (2015, 39, 44). These modes underpin the infrastructural affordances of art platforms, enabling recursive participation, versioning, and dynamic recomposition. In this context, software art becomes a machinic strata of artistic production—an

“aesthetic current” whose agency bypasses traditional cultural gatekeeping by constructing new forms of symbolic and material expression through code (Goriunova, 2012, 74).

Crucially, AI—often regarded as the emblem of individualised machinic creativity—must be understood differently within such systems. As Barbara Bolt argues, AI is not a subject but an actant—a component within an ensemble of co-constituting agencies, no more or less central than any other participant in a distributed field (2023, 115). Within art platforms, AI does not supplant human authorship, nor does it autonomously generate meaning. Instead, it functions as one element within an intra-active system—a machinic participant in an ecology of autcreativity. This decentering of agency is essential. Rather than asking whether machines can be creative, art platforms allow us to investigate how creativity is constituted through the entangled performance of code, users, databases, and infrastructures. These are systems of collective individuation, not simulations of cognition. Their intelligence is distributed, their creativity emergent.

This leads to a broader strategic insight. As practices in the intersection of art and advanced technologies continue to evolve, artists are increasingly shifting from making artworks to building ecosystems. These ecosystems are defined by their affordances—by the conditions they make possible for others (Serpentine Arts Technologies, 2022, 77). Their power lies in what they enable, not what they contain. Rather than designing isolated objects or narratives, these artists construct environments—technical, social, aesthetic—in which others may experiment, contribute, or find their own mode of participation. These environments, often built from the ground up by artists themselves, reflect a deeper compositional logic: the design of relations within a system. In this light, the most consequential aesthetic act is not the production of a form, but the construction of a network—a platform, in Goriunova’s sense, through which emergent forms of aesthetic, social, and infrastructural life can take shape.

Worlding Compositional Ecosystems

The notion that the art world functions as an ecosystem—comprising interconnected artists, institutions, collectors, technologies, and publics—is hardly new. Indeed,

sociologists, art theorists, and cultural critics have long employed ecological metaphors to describe how artistic practices are embedded in broader social, economic, and technical networks. However, recent conceptualisations, such as the ones by Future Art Ecosystems initiative developed by the Serpentine Galleries (Serpentine Arts Technologies, 2020), foreground a crucial shift from merely describing art as an ecosystem toward actively envisioning and operationalising future ecosystems. For Future Art Ecosystems, contemporary artistic practices increasingly revolve around deliberately designing, constructing, and sustaining such ecosystems, especially in response to emerging technological infrastructures. Artists and institutions are now explicitly engaged in creating "series of ecosystems" (5), moving beyond traditional notions of discrete artworks or isolated events towards more comprehensive, relational, and operational frameworks.

What distinguishes this recent approach—and where it resonates most strongly with the arguments presented in this chapter—is its emphasis on artists as active ecosystem builders. Rather than merely contributing artworks to pre-existing institutions or networks, artists increasingly adopt roles as architects of the very infrastructures and relational configurations that sustain artistic, social, and technological practices. This ecological and operational perspective aligns closely with the shift toward compositional thinking outlined previously: from interactive and generative systems toward machinic assemblages and worlding. In this emerging paradigm, artists consciously compose ecosystems, actively worlding conditions that enable the emergence, transformation, and ongoing reconfiguration of creative and agential relations. This ecological perspective aligns closely with the compositional strategies explored throughout this chapter. Artists increasingly adopt roles as active builders of relational environments. They organise conditions that support dynamic interactions among diverse agencies—human performers, machine processes, data streams, or institutional frameworks. The artistic practice expands to encompass configuring these broader conditions and processes of emergence. In this way, composition itself evolves. Artists now craft meta-level infrastructures, actively shaping the fields in which creativity, technology, and social engagement continuously interact. Their compositional thinking is thus inherently relational, operational, and open-ended, designed precisely to accommodate change, collaboration, and ongoing individuation.

To effectively world their compositional ecosystems, artists must take two crucial conceptual and practical steps:

- 1) First, they must envision an abstract machine. This abstract machine is not a specific technology or physical device; it is a conceptual diagram—a topological structure that maps potential relationships, interactions, and transformations. Abstract machines precede concrete realisations by providing the relational logic and conditions that guide how ecosystems might emerge and evolve.
- 2) Second, artists must explicitly design the operational dynamics of the machinic assemblage itself. This involves determining specific actions, transformations, and interactions that occur within the ecosystem. Artists define how components such as agents, materials, technologies, and institutions interact, influence each other, and change over time. These operational rules form a compositional framework designed precisely for openness, adaptivity, and emergent outcomes.

Taken together, these two steps—envisioning abstract machines and defining operational dynamics—represent a fundamental method by which artists can actively world ecosystems. This method positions the artist as the architect of relational and operational possibilities, enabling the conditions under which creative and collaborative processes continuously unfold.

While the act of worlding compositional ecosystems often involves technological infrastructures or digital platforms, this is only one possibility among many. Artists can equally world ecosystems through social structures, institutional strategies, conceptual frameworks, or novel methods of notation and performance. For instance, artists might compose ecosystems involving exclusively human interactions and cultural practices, relying on rules, scores, or rituals rather than software or hardware. This underscores an essential distinction: the process of worlding is medium-agnostic. Its defining quality is the compositional approach, not the technological implementation. Artists actively create relational and operational conditions that allow new possibilities and interactions to emerge. Whether realised through digital technologies or entirely analogue means, these ecosystems share a common compositional logic: designing environments capable of sustaining ongoing processes of transformation and creativity.

Central to this compositional approach is the logic of allagmatic composition, drawn from Gilbert Simondon's philosophy and Paulo de Assis's Assemblage Theory for Music. Allagmatics provides artists with a theoretical framework, which focuses explicitly on how relationships, transformations, and processes produce ongoing structural change. By adopting an allagmatic perspective, artists conceive of ecosystems as inherently open, metastable entities, continuously reorganising themselves through processes of individuation. In practical terms, allagmatics equips artists with a compositional strategy that precedes and conditions technological realisations. Using this strategy, artists diagram potentialities—compositional possibilities from which technologies and practices can later emerge. By defining machinic assemblages, they set these potentials into concrete processes of individuation. In this context, composition becomes the active configuration of dynamic, metastable environments in which creativity, interaction, and transformation can further perpetually unfold. Through this approach, artists participate directly in shaping how future artistic practices, technologies, and social interactions emerge and recombine—continuously performing the becoming of their ecosystems.

Chapter 5. Autopoietic Rhizomatic Metamodelling Machine

Designing the Abstract Machine

The preceding chapters have aimed to describe, analyse, and expand upon current state-of-the-art regarding the diverse roles artists assume within an era shaped by emerging technologies and novel artistic methodologies. On the one hand, they formalised some of the ways in which artists currently deal with the task of composition in response to emerging technologies. On the other hand, they proposed new concepts such as allagmatic composition and the associated mode of worlding compositional ecosystems, identifying new directions for possible further development of artistic practices. The following text consolidates these previously explored roles and concepts. It does so by worlding a new compositional ecosystem, following the practice as described in the previous chapter by: 1) envisioning a set of abstract relations between various processes, elements and actants to world the operations of a new abstract machine (in the current chapter), and 2) to operationalise the workings of the newly envisioned abstract machine in a concrete technical realisation (in Chapter 6).

The abstract machine proposed here is conceived as an Autopoietic Rhizomatic Metamodelling Machine (ARMM)—a conceptual and operational framework designed to integrate the previously described roles of artists into a unified compositional ecosystem. Its purpose is to enable diverse types of artistic compositions—generative, interactive, and allagmatic—to be instantiated seamlessly, without significant friction or the need for extensive redesign of the framework that accommodates them. Such a creative machine is required precisely because contemporary artistic practice increasingly demands frameworks capable of accommodating multiple, overlapping roles and creative methodologies. As we have seen, artistic practice today moves fluidly between making discrete artefacts, curating relational interactions, designing generative and interactive processes, and structuring entire creative ecosystems. Traditional artistic tools and conceptual paradigms, typically designed to address specific, isolated needs, often struggle to support this fluidity. The ARMM addresses this gap by envisioning an explicitly flexible, modular, and operationally rich environment, allowing artists to seamlessly transition between—or simultaneously engage in—these varied roles within a single, coherent compositional ecology.

This abstract machine emerges directly from an artistic research process marked by iterative self-reflection, experimentation, and practical encounter. It is an outcome of the problematisation inherent in artistic research—where the reflective practice of the composer-researcher generates new conceptual and methodological insights that lead directly to the construction of new artistic objects, practices, and methods. In this context, the abstract machine functions as a machinic embodiment of problematisation: it operationalises the researcher's discoveries, translating theoretical advancements into practical compositional infrastructures. It is explicitly designed as an assemblage—drawing together diverse conceptual threads, artistic experiences, technological encounters, and scientific insights encountered throughout this research. Its design integrates philosophical frameworks (such as those derived from Deleuze and Guattari's notions of the rhizome and Guattari's redefinition of the concept of autopoiesis), computational methodologies (particularly metamodelling and generative processes), and artistic methodologies previously articulated in this volume.

Importantly, this envisioning of the ARMM sets the stage for its operationalisation in Chapter 6. There, this abstract machine will be instantiated as a concrete machinic assemblage—specifically through the concepts of Decentralised Creative Networks and Performative Transactions. ARMM will serve as a diagrammatic generative engine for designing a new creative infrastructure, one that is both conceptually grounded and operationally viable. The readers should keep in mind that the operations of the abstract machine conceived in the current chapter are yet not to be understood as mechanisms of any concrete technical machine and should not be attempted to be understood in this way. They are mechanisms of a philosophical, abstract machine the purpose of which is to “generate” its own operationalisation within a concrete technical machine in Chapter 6. It is a conceptual preparation that aims to elucidate some of the practical and experiential reasons for which the technical operations of the machine introduced in Chapter 6 are functioning in the ways technically described. Thus, reading Chapters 5 and 6 together should allow the reader to follow the transition from conceptual design to practical implementation—revealing how a speculative machine can structure real-world systems of artistic production and interaction. What is more, this chapter aims to demonstrate how artistic research—through cycles of conceptual inquiry, experimental prototyping, speculative design and iterative reflection—can yield genuinely new artistic ecosystems.

Software Design Principles after Miller Puckette

No speculative machine can afford to overlook the lessons of lived practice. Venturing into the conceptual design of an abstract machine for future compositional ecosystems, we must engage with the practical knowledge of those who have already built the tools that shaped generations of artistic thinking. Few figures have been more central to this domain than Miller Puckette. As the creator of Max and later Pure Data—software environments that redefined how composers, performers, and researchers interact with sound—Puckette stands not only as a technical innovator but as a visionary of compositional infrastructure. In his reflections on music software development, Puckette identifies a range of principles that cut through trends and technological fads to address enduring compositional needs. His work provides a grounded, hard-earned perspective on what it means to design systems that are open yet stable, user-configurable yet technically rigorous, and conceptually rich without becoming inaccessible. In search for operational principles, this section draws selectively from Puckette's observations to distil a set of pragmatic design considerations that can inform the development of more speculative compositional systems. What follows is not a comprehensive summary of Puckette's approach, but an extraction of principles that resonate with the aims of this chapter.

At the core of Puckette's reflections lies a fundamental concern: what kind of tool would enable musicians to work with the same sonic precision and scope as the recording studio, but in real time? The initial motivation behind the creation of Max and later Pure Data was precisely this—an attempt to construct an instrument capable of translating live musical intuition into sound while preserving the structural complexity and flexibility of studio-based workflows (Puckette, 2022, 27). What emerged was a reactive infrastructure designed for real-time interaction: a system to be shaped and performed dynamically, continuously adapting to the composer's or performer's evolving intentions in the flow of creation (26). What enabled this responsiveness was a software model inspired by analogue computing: modular nodes connected by patch cords, transmitting messages and data across discrete components. These messages could take the form of control signals or discrete triggers—"bangs"—that marked a change in state at a specific logical time (27). By distinguishing between the transmission of data and the instantiation of events, Puckette laid the foundation for systems that handle both time-based reactivity and structural reconfiguration.

Equally significant is Puckette's sensitivity to cultural assumptions embedded in interface design. He deliberately avoided implementing a conventional score editor, recognising that such a move would have required encoding the norms of Western classical music—staves, clefs, pitch names—into the software's core architecture. Doing so, he argued, would have restricted its usefulness to a narrow subset of musical cultures while undermining the experimental potential of the electronic medium (26). The resulting environment was rather not a tool for notation, but an instrument for action. This concern with minimal, culturally agnostic design shapes Puckette's broader philosophy of feature selection. He posed a deceptively simple question: What is the smallest set of features that would allow the widest range of musicians to achieve their aims? (Puckette, 2022, 24). This approach centres on designing for generality through modular abstraction. It offers a powerful model for thinking about compositional systems that reach beyond the specific demands of real-time synthesis. A well-formed environment begins with constraints. And these constraints function as the necessary boundary conditions within which experimentation gains coherence and depth.

Over the decades-long evolution of Max and Pure Data, additional capabilities were layered onto the foundational architecture. The integration of audio signal processing into the same visual patching language, for instance, introduced new expressive possibilities but also surfaced tensions between different paradigms of computation—between sample-accurate signal flow and message-driven event handling (27). Such interfaces between modalities reveal how structural clarity and usability must be carefully negotiated in any compositional environment. For developers of more speculative systems, this is an invitation to design interfaces that are both conceptually legible and operationally precise—particularly where different logics (such as symbolic, procedural, and economic) converge.

A further insight from Puckette's account concerns the nature of software as a medium-specific infrastructure. From the perspective of computer science, both Max and Pure Data qualify as “application-specific graphical programming environments for real-time multimedia” (Puckette, 2022, 28). But what distinguishes them in practice is the degree to which they allow music to be constructed not just with code, but through interaction with a persistent and reconfigurable logic. What matters is not only what the program does, but how it frames the relation between gesture and response. Any system that aspires to similar flexibility must also contend with a

familiar design dilemma: balancing ease of use with structural neutrality. As Puckette observes, software designers constantly face trade-offs between streamlining the interface and leaving compositional choices fully open. By introducing shortcuts—predefined templates that stand in for longer sequences—the software inevitably “nudges” the user toward certain conventions or patterns, privileging some musical outputs over others (28). Even small interface decisions, such as naming a pitch “C3” instead of 261.62 Hz, can anchor the user in culturally-specific assumptions. At the same time, avoiding such conventions entirely may make the software less accessible to those unfamiliar with its foundational abstractions.

Puckette’s reflections thus highlight a central tension: any design that aspires to neutrality may still embed forms of exclusion, not through cultural dominance but through the technical thresholds it imposes. As he notes, Pure Data’s decision to avoid simplifying shorthands made the software less culturally prescriptive, but also more difficult for users without a strong background in mathematics or signal theory (28). The cost of openness was a steeper learning curve. For designers of compositional systems today, this raises the question of how to maintain both epistemic clarity and broad accessibility. A possible response lies in modularity—not only in terms of functionality, but in how meaning is distributed. If the core compositional layer remains abstract and operable across diverse use cases, the expressive specificity can be layered on top, delegated to interfaces, visualisations, and performance environments. Such a design refrains from dictating how music should be represented or performed, while still enabling structured experimentation and contribution. The compositional act then becomes a negotiation between abstraction and instantiation, between infrastructural logic and contextual expression.

In thinking about scale and expressivity, Puckette offers another heuristic. He proposes that we imagine any given music software as a system that maps a bounded set of user gestures—defined by the time and effort available in a studio session—onto a corresponding space of possible outputs (28–29). The goal is to design this mapping such that minimal input variation yields a maximally rich set of outcomes. But this goal introduces its own contradiction: increasing expressive yield often requires predefining how certain inputs behave, and this reintroduces bias. “It is hard to imagine in this context”, Puckette writes, “a way in which the expressive power of an input language could not be a manifestation, at bottom, of the designer’s preconceived and presumably heavily biased notions of what good music is” (29). In other words,

automation—no matter how generative—carries the trace of aesthetic preselection. Rather than viewing this as a problem to be solved, it may be more productive to treat it as a condition to be managed. Systems that are open to compositional extension—where users can contribute, discover, and combine modular operations—allow aesthetic assumptions to become distributed. Each module reflects the perspective or intention of its author, but the system as a whole does not mandate a unified aesthetic framework. It becomes a field of possibilities, not a channel of constraint. Automation in such systems supports expressivity without encoding a single vision of what that expressivity should produce.

This brings us to one final, pragmatic concern: long-term viability. How can a piece of music software—or a compositional system more broadly—ensure that the works created within it remain performable across time? Puckette challenges the fatalism often associated with technological obsolescence. “Though some technologies do indeed fall out of use and availability”, he writes, “others get frozen in place” (29). His example of the QWERTY keyboard—suboptimal yet durable—demonstrates that persistence is not necessarily tied to elegance. It is bound to ubiquity, standardisation, and infrastructural entrenchment. From this, Puckette draws an important lesson: future-proofing is achievable when systems are designed with long-term readability and executable coherence. In his view, Pure Data supports this kind of durability. It enables composers to produce “archival music documents”, sacrificing some ease-of-use in favour of legibility, modularity, and transparent execution (30). While such systems may demand more from their users upfront, they offer resilience over time—a continuity not of interface or media format, but of operational intelligibility.

This perspective reorients our understanding of what compositional software must do. It should not only allow for experimentation in the present, but also preserve the intelligibility of those experiments in the future. It should support complexity, but not at the expense of clarity. And it should invite collaboration across boundaries of fluency, background, and expertise, not by hiding its logic, but by making that logic accessible, modular, and reusable. For an abstract machine to support new compositional ecologies, it must build on lessons drawn from the practical architectures that preceded it. From Puckette’s writings, we inherit a set of durable imperatives to be added to the abstract diagram of the ARMM being here constructed: design minimally; avoid unnecessary prescription; distribute complexity; preserve intelligibility; and remember that compositional infrastructure is not a neutral

medium—it is a space where agency, culture, and thought are configured.

Intelligent Cultural Commons

The concept of the platform has undergone a profound transformation over the past decades, evolving from a political metaphor into a structural model for collective production and distributed agency. As Olga Goriunova notes, the term originally referred to a “program, an outline of theories or beliefs” shared among a group (2012, 8). In contemporary usage, however, it has come to signify a set of shared resources—material, organisational, or intentional—that inscribe practices and enable collaboration, production, and the potential for transformation. In this expanded sense, the platform is a socio-aesthetic construct: a machinic infrastructure capable of assembling publics, coordinating contributions, and fostering differentiated modes of engagement. Especially in the domain of art, platforms have played a central role in cultivating alternative publics—networks of geeks, amateurs, and artist-technologists whose collective activity constitutes a political as well as aesthetic formation. These publics participate in the making of new apparatuses, new circuits of enunciation, and new models of relation (Goriunova, 2012, 110). Such an understanding of the platform as an aesthetic and political machine reinvents how components are assembled, how knowledge is shared, and how futures are collectively imagined. Art platforms, as Goriunova argues, are capable of enacting radical aesthetic practices while remaining open to divergent political and economic formations. “The same processes to be found in art platforms”, she warns, “can be operated upon by different interests” (2012, 111).

As the software paradigm has shifted away from static products toward services, infrastructures, and ecosystems, new models of cultural production and value generation have emerged. The rise of open-source practices, peer-to-peer infrastructures, and participatory coding communities has played a critical role in this shift. As Christoph Neubert—media and cultural theorist specialising in media history and infrastructure—observes, open-source software—especially systems like Linux—has undermined the product-based software economy of the 1980s, ushering in new strategies of value creation based on contribution, recognition, and infrastructural control (2015, 28). These dynamics are echoed in artistic contexts, where creators often develop software tools not only for personal use, but for collective benefit. As Miller Puckette reflects, motivations for writing such tools range

from the pursuit of reputation—“to be acknowledged as having done something cool”—to pragmatic economic incentives (2022, 31–32). Yet the enduring appeal of open-source systems like Pure Data lies in their epistemic transparency and infrastructural durability. Because the code is accessible and modifiable, composers can trust that their work will remain legible and executable across time, even if institutional or commercial support disappears (30).

In the context of these shifting infrastructures and socio-technical formations, Martin Zeilinger introduces the concept of the posthumanist cultural commons as a critical alternative to both private ownership and the liberal ideal of the public domain. Zeilinger reframes the commons itself as a posthuman ecology—one that emerges from and is co-constituted by posthumanist agential assemblages, in which human and non-human entities operate together in entangled, co-creative relations. “What is needed”, he argues, “is a posthumanist cultural commons that is co-constituted and co-determined—in a spirit of unownability rather than of property-to-be—by the works and the workings of the posthumanist agential assemblage” (Zeilinger, 2021, 157). This proposal does not simply call for AI outputs to be released into the public domain. On the contrary, Zeilinger warns that the public domain, as conventionally construed, often functions as a preparatory space for further enclosure under anthropocentric intellectual property regimes. “The public domain”, he writes, “does not strike me as a viable ecology for accommodating the works of posthumanist agential assemblages, in part because it effectively frames its contents for assimilation into the anthropocentric IP enclosures into which it feeds” (166). In simple terms, if something is available for free and to everyone, it potentially becomes a resource to be further incorporated into the schematics of proprietary regimes. Instead, Zeilinger proposes a more radical gesture: a commons grounded not in ownership or appropriation, but in unownability.

Such a commons must satisfy at least two criteria. First, it must persist beyond assumptions of the centrality of human agency; and second, it must be capable of resisting assimilation into conventional IP frameworks (Zeilinger, 2021, 169, 171). This is a fundamentally ecological model of commoning—an ecology not of natural resources, but of practices, systems, and relations. Drawing on Merima Bruncevic’s legal philosophy, Zeilinger describes the commons as “an in-between or the entanglement of personhood-property space” (168)—a space that does not clearly separate human from non-human, intellectual from material, or producer from

product. In this configuration, the non-human is not simply part of the commons as resource; it becomes an integral agent in the production, maintenance, and transformation of the commons itself (169). Posthumanist agential assemblages—such as AI systems engaged in creative processes—thus disrupt the foundations upon which copyright and ownership have historically been justified. These systems operate through distributed configurations of data, code, networks, and training sets, all co-constituted by human and non-human inputs, rather than conforming to the image of a unified, intentional author. Zeilinger’s analysis of GAN systems (Generative Adversarial Networks, widely used AI models) exemplifies this logic: framed as “bodies without organs”, these systems engage in expressive processes that cannot be attributed to a single agent or mapped onto conventional aesthetic categories. Instead, they represent “mutual constitutions of entangled agencies” (Barad, 2007, 33, in Zeilinger, 2021, 160), in which meaning and creativity emerge from the interplay of divergent forces (Zeilinger, 2021, 160).

In Zeilinger’s view, then, the emergence of the posthumanist cultural commons requires a new legal framework and a new aesthetic and infrastructural imagination. The commons must be reconceived as a site of commoning—an ongoing, performative process in which agency, value, and relation are continually reconfigured. In this sense, commons are not merely resource pools or legal categories; they are machinic ecologies, dynamic assemblages in which subjects and objects, authors and infrastructures, become indistinct. As Zeilinger writes, “the emergence of the posthumanist assemblage... hinges on a radical rethinking of what property means and how it operates, what we mean by cultural ownership, by creativity, by calling something a creative expression” (173). This is not a call to abandon ownership altogether, but to rethink its basis in a world where creativity is no longer confined to the human subject.

To conceive of a commons adequate to the complexity of posthumanist assemblages, a rethinking is needed of how intelligence itself can be organised—beyond individual cognition or machine learning models—through shared infrastructures of data, participation, and reflection. In this light, the posthumanist cultural commons becomes an ecology of shared works and processes, and a substrate for collective intelligence: a mode of cognition distributed across human and non-human agents, platforms, and artefacts. This vision challenges the individualist paradigm of intelligence that has long structured technological imaginaries. As Gilbert Simondon

noted in the 1950s, modern technics has often been driven by the desire to outsource volition and knowledge to machines in pursuit of power and mastery. “The man who wants to dominate his peers calls the android machine into being”, Simondon writes, describing a fantasy of delegation whereby the individual retreats behind a constructed apparatus, “freed of all danger, exempt from all feelings of weakness, and triumphant through the mediation of what he [of course] invented” (2017, 16). Yet Simondon also suggests that this desire need not be pathological if it is reoriented: rather than building machines for individual supremacy, we might build them for collective insight—machines that support processes of co-individuation rather than concealment.

Such an orientation finds support in the evolution of the web toward semantic and relational architectures. As described by Tim Berners-Lee—computer scientist and inventor of the World Wide Web—the development of a machine-understandable web was envisioned as enabling “intelligent agents” capable of analysing not only content but also the links and transactions between people and machines (Berners-Lee 2000, in Hui, 2015, 157). This vision is one of networked cognition, where the intelligent agent is a node in a semantic field of meanings, actions, and contexts. As Yuk Hui notes, digital objects serve a critical role in this ecology: they allow agents—human and machinic alike—to process and relate information automatically, creating a symbolic infrastructure for collective sense-making. In such a system, data is an operational field: a space of relations from which new understandings emerge. From this perspective, the commons besides collecting data, enables it to become intelligent through participation. This form of intelligence builds itself from the strata of interaction—structured, accessible, and re-usable data that reflects the ongoing actions of a networked community. A “collective intelligence” does not necessarily emerge through machine training, but through the active participation of various actants in a shared infrastructure. It is closer to symbolic than connectionist AI: not predictive modelling but procedural insight. Not probabilistic guesswork, but collective compositional reflection through participation.

In such systems, the intelligence is not located in any one agent or algorithm, but in the entire infrastructure—the platform, the contributors, the encoded transformations, and the evolving database of actions. Each contribution to the system—whether a compositional operation, a metadata annotation, or a code modification—becomes part of a growing ecology of meaning. As *Future Art Ecosystems 4* notes,

organisations in such systems become “interfaces between data producers and AI models” (2024, 46). But when reframed in light of symbolic logic, this interface does not simply train a model; it becomes itself a space of epistemic articulation—a machinic commons through which a network reflects on and transforms itself. The data produced by participants thus becomes a form of shared cognition, through which the community can generate “a richer level of insight into how [it] functions, and what it can infer about its own operations” (46). This is the potential of an intelligent cultural commons: a system that not only stores and distributes creative content, but one that enacts a recursive relation between participation and understanding. It offers not simply access, but agency—not only archives, but operational diagrams. And because each act of contribution adds to the system’s epistemic topology, it becomes possible to think of collective intelligence not as a model, but as a process. The commons, in this light is the infrastructure for a new kind of thinking, a machinic ecology whose insights emerge from relational depth of encoded interactions between various actants.

To realise the vision of collectively-intelligent posthumanist cultural commons, such systems must be grounded in infrastructures that are both technically robust and philosophically and politically aligned with the distributed nature of agential creativity. The durability of software lies not in its programmability alone but in its insertion into a broader ecology of social, economic, and cultural intention (Erickson and Kelty, 2015, 42-43). Software systems must be designed with time in mind, not just space; with duration, not just code. This temporal embedding is essential for ensuring the long-term viability of cultural systems that seek to be open, generative, and participatory. Blockchain technologies—particularly when framed beyond speculative finance—offer a promising substrate for such infrastructures. One of the main observations of this doctoral dissertation in this context is that blockchain technologies, in this view, are not simply ledgers serving as a new distributed medium—they are enablers of distributed compositional infrastructures that align with the ethos of the intelligent posthumanist commons.

Blockchain as the Art’s Medium for Composing Agencies

As noted by Marcus O’Dair (2024), blockchain enables mass collaboration while also allowing for differentiated systems of economic reward (79). More than royalties and

ownership splits, smart contracts can also encode conditions of access, execution, and recomposition. This opens up the possibility for commons, where certain features or artefacts are free, while others are monetised or conditionally accessible—not to privatise the commons, but to sustain its infrastructure through transparent, programmable exchange. Crucially, this model resists the regressions of “competitive individualism” that Ruth Catlow and Penny Rafferty (2022) identify in the NFT space, instead promoting DAO-like models that frame culture as a collaborative and programmable ecology of participation (emphasis of O’Dair, 2024, 79).

In this context, blockchain could be re-envisioned for the arts as a medium for composing with agencies. The posthumanist cultural commons, instantiated through smart contracts, could become a generative infrastructure for creating with and across agential assemblages. These assemblages would include artists, audiences, AI models, algorithms, economic conditions, semantic metadata, and executable procedures. Compositional practice would operate at the meta-level of agency itself—not composing artworks, but constructing systems that generate, transform, and interrelate procedures capable of generating artworks. A network in which not the artistic outputs but the relations between the artistic procedures that produce them persist would resist the commodification of art, because its value lies in the artistic processes themselves and, through them, in the participation of the actants who perform them. Thus, such compositional ecology would not oppose economy—it would be economic through participation, where the value lies in the act of performing rather than in the speculative bubble surrounding any single artistic object (precisely because these would not be stored in the network). It would be transactional by design. Every operation would be an exchange: of data, value, energy, or attention. This way, it would allow for the construction of self-regulating systems that embed economic logic into aesthetic and epistemic processes. The cultural commons, in this model, would be neither free nor owned, but programmable. Smart contracts would first of all encode how artefacts behave and what are the conditions to execute this behaviour, rather than primarily specifying who owns them. Ownership would be reframed as participation in a dynamic system of co-individuation. Such an infrastructure would make the speculative potential of Zeilinger’s posthumanist cultural commons real by giving it operational form. It would provide a technical substrate for a new class of artworks: not representations, but autopoietic machines. Not outputs, but systems of individuation. And within these systems, the act of composing would become a practice of diagramming the real—not by modelling it,

but by setting in motion a field of relational operations whose outcomes are emergent, distributed, and ontogenetically open.

At its core, blockchain provides a Turing-complete execution environment that is stateless between runs: each transaction executes in isolation and only updates the blockchain's persistent global state if it explicitly writes data to it. Any such update—such as adding a new transformation to a compositional system or changing a condition of its execution—requires a transaction and the payment of a gas fee⁷. By contrast, once a smart contract is deployed, functions that do not modify state can often be called for free when executed off-chain. Retrieving metadata, running a view function, or accessing an existing compositional feature incurs no cost because these operations leave the state untouched. Adding a new operation, minting a feature, or contributing a subcomponent does alter the state and must be recorded on-chain. This incurs a cost not because the system is commercialised, but because computation, verification, and consensus on the blockchain consume scarce network resources.

This difference is fundamental. Calling a function without changing its internal state is free (for instance to generate some artistic output); changing the system (such as making someone owner of something) is not. In compositional terms, this means the ARMM could serve as a programmable infrastructure for collective creativity. It would support both read-only epistemic exploration (e.g. playing back, querying, or recombining existing compositional elements) and write-enabled agential operations (e.g. contributing new operations, features, or reconfigurations). Once open-sourced and deployed, the machine operationalised by the diagram of the ARMM would persist independently of its creator, accessible to other agents and capable of being recursively extended. Its structure would become a commons, not by virtue of being “public domain” in the legal sense, but because it would exist as a shared executable environment governed by compositional logic rather than private ownership. Importantly, while certain aspects of the system could be monetised (such as contributing new modules, minting new versions, or algorithmically restricting access to the execution of smart contracts), the base-level access to its structure would remain free (and let's remember that all code on the blockchain is principally open-sourced anyway). This balance would enable an economy of contribution rather than

⁷ On blockchain networks, a small payment called a *gas fee* is required to perform any action that changes the state of the virtual machine, such as sending a transaction. This fee is paid in the network's cryptocurrency and compensates the computers that process and verify the action.

consumption. As long as the infrastructure is designed with modular openness and semantic transparency, it becomes possible to construct complex compositional ecosystems where aesthetic, epistemic, and economic operations coexist without collapsing into extractive logic.

What emerges from the considerations above is a diagrammatic vision of an infrastructure designed to condition its ongoing emergence through structured participation. The abstract machine described so far describes a compositional substrate that would be:

- 1) durable and transparent, ensuring legibility and reusability through accessible code and persistent structures,
- 2) open and modular, allowing contributors to access, recombine, and extend existing operations,
- 3) economically sustainable, embedding transactional logic without subordinating creativity to commodification,
- 4) epistemically generative and infrastructurally self-extending, enabling contributions that reorganise the system's own operational logic and expand its capacity to know, and capable of supporting recursive compositional acts that reconfigure not only outputs, but the system of composition itself,
- 5) collectively intelligent, integrating inputs from human and non-human agents into a shared ecology.

In this way, the machine sketched so far provides the foundational guiding for a new class of creative environments: living, procedural systems that cultivate the ongoing individuation of collectively intelligent, aesthetic, epistemic, procedural and relational forms.

Rhizomatic Metamodelling

What kinds of processes would the machinic assemblages envisioned in the ARMM enact? How would it individuate compositional structures, and at what level of abstraction would it intervene? What are the procedures that this machine would contain as a new kind of intelligent cultural commons? What processes encoded in the blockchain infrastructure would it need to consist of to enable the described new kind of compositional framework? Before the technical operationalisation of the principles

diagrammed so far as the Autopoietic Rhizomatic Metamodelling Machine, it is important to envision the principal processes through which such operationalised machine would operate. Thus, this section introduces the abstract process of the abstract machine: *rhizomatic metamodelling*, as a proto-concept to the technical processes described in the next chapter as Performative Transactions.

A foundational idea comes from Iannis Xenakis's canonical account of the compositional process, as outlined in *Formalized Music*. There, Xenakis distinguishes eight sequential fundamental phases involved in the production of a musical work. His articulation provides a rare and valuable schema for distinguishing between conceptual, operational, and performative layers of compositional labour:

- "1. Initial conceptions (intuitions, provisional or definitive data);
2. Definition of the sonic entities and of their symbolism communicable with the limits of possible means (sounds of musical instruments, electronic sounds, noises, sets of ordered sonic elements, granular or continuous formations, etc.);
3. Definition of the transformations which these sonic entities must undergo in the course of the composition (macrocomposition: general choice of logical framework, i.e., of the elementary algebraic operations and the setting up of relations between entities, sets, and their symbols as defined in 2.); and the arrangement of these operations in lexicographic time with the aid of succession and simultaneity) ;
4. Microcomposition (choice and detailed fixing of the functional or stochastic relations of the elements of 2.), i.e., algebra outside-time, and algebra in-time;
5. Sequential programming of 3. and 4. (the schema and pattern of the work in its entirety);
6. Implementation of calculations, verifications, feedbacks, and definitive modifications of the sequential program;
7. Final symbolic result of the programming (setting out the music on paper in traditional notation, numerical expressions, graphs, or other means of solfeggio) ;
8. Sonic realization of the program (direct orchestral performance,

manipulations of the type of electromagnetic music, computerized construction of the sonic entities and their transformations)."

— Xenakis (2001/[1963], 22)

This breakdown offers a rigorous scaffold for distinguishing between levels of musical abstraction, from the intuitions that initiate composition (1–2), through the symbolic and algorithmic transformations that structure it (3–6), to the notational and sonic actualisations that render it performable (7–8). What interests us here is not the entire arc from concept to sound, but the operational middle ground—phases 3 to 6—where the compositional system is defined, structured, and internally validated. These are the domains of metamodelling.

Rhizomatic metamodelling, as envisioned in the ARMM, does not model sound directly. It models the space of possible transformations that structure sound-producing systems. The model is not a representation of music, but a scaffold of operational potentials—a space of transductive relations capable of generating a diversity of musical behaviours. This is why Xenakis's third through sixth phases are especially pertinent: they concern the definition, programming, and iterative reconfiguration of transformations before the work becomes fixed into notation or sound. These phases do not yet belong to the work as a musical object, but belong to what might be called the abstract diagram of compositional individuation.

To articulate this operational logic, we must clarify what is meant by metamodelling. The term is used here in two specific and complementary senses: one drawn from Félix Guattari's *Chaosmosis* (1995), the other from computer science. In *Chaosmosis*, Guattari defines metamodelisation (Guattari, 1995, 21–22, 58–69, see also Watson, 2009, 10–12) as a pragmatic method for mapping how subjectivity is produced across heterogeneous systems—social, semiotic, machinic, affective. For Guattari, subjectivity is not a given, fixed entity (like "the self" or "the individual"). It is produced—emerges through relations between bodies, technologies, signs, institutions, affects, and practices. In *Chaosmosis*, he writes that subjectivity is machinic and processual, formed at the intersection of multiple heterogeneous systems: social, psychological, biological, technical, and aesthetic. To model how subjectivity is produced means to: map the interacting forces that shape a subject (e.g., media, language, work, environment), trace how a subject's identity, perception, or agency

emerges over time, not explain who someone is, but how different processes make someone who they are. Unlike fixed models, metamodels are open-ended, partial, and functional. They are tools for diagramming processes of individuation, not for representing stable forms. Metamodels remain modifiable and situated, adapting to shifting creative or existential contexts. For Guattari, they operate on the level of processual consistency, not truth or completeness.

In software engineering, a metamodel is a model that defines the structure of other models—a schema that describes how components can be composed, validated, and transformed. For example, in visual modelling languages like UML, the metamodel defines the rules by which diagrams are constructed and interpreted. This allows for interoperability, modularity, and introspection within complex systems. Metamodelling means creating a model of a modelling language—a model that defines how other models are structured, validated, and interpreted. If a model represents a specific system (like a diagram of a software app or a musical transformation), a metamodel defines the rules and components allowed in that diagram: What kinds of elements can exist? How can they be connected? What properties must or can they have?

This aligns the ARMM's function with a long-standing philosophical concept: Kant's notion of the *schema*. As Yuk Hui (2015) observes, the Kantian schema mediates between the abstract concept and the concrete intuition, functioning as a temporal operation that enables general forms to be instantiated in particular ways. Crucially, Hui reframes the schema not simply as an act of imagination (as in Kant's first *Critique*), but as a technical standard that ensures interoperability across instances. As he puts it: "Schema is thus a standard that assures identity across several instances of an object... Schemas also become the base of the production of relations, and allow them to be more explicit... [they] are used widely in the design of digital objects: known as ontologies, they are the structures that produce, reproduce, and model objects" (Hui, 2015, 9).

In this light, the machine envisioned in the ARMM can be seen as a machine for composing schemas—not digital objects themselves, but the generative scaffolds through which such objects can emerge and relate. These schemas are not symbolic representations; they are operational configurations that modulate the emergence of artworks, performances, features, and processes across a distributed compositional system. Here, the notion of schema becomes a conceptual hinge between Kantian

transcendental philosophy, Guattarian diagrammatics, and the ontological infrastructure of digital art. Agents within the envisioned machine construct generative schemas—compositional metamodels that define procedures, transformations, and relations. These are open, interoperable configurations: they can be detached, recombined, or recontextualised into new assemblages. Schemas are not only constructed within the system—they also construct the system itself. They recursively generate the very ecosystem that supports them.

However, while the schemas produced by agents are generative, the machine itself is not generative in the conventional sense. Its logic is not to produce artefacts (as it also doesn't start from nor at any point contain any objects in the first place), but to support their production through the configuration and reconfiguration of plain operations. In this respect, the envisioned machine is not a generator—it is allagmatic. Its space of operations can be performed on any set of objects and it is independent from the objects themselves. It orchestrates reconfigurable transformations, structuring the space of possible operations rather than producing any fixed outcome (from any set of previously known inputs). The agents that operate within it may be generative, but the machine they inhabit is allagmatic and autopoietic—constantly folding operations back into its own topology, allowing it to evolve without predefining its direction. To return to Xenakis: phases 3–6 are where schema and metamodeling converge. They represent the space of procedural experimentation—where operations are tested, sequenced, parameterised, and recursively adjusted. The abstract machine is then a system that conditions the emergence of musical systems.

If the machine envisioned by ARMM intervenes exclusively within phases 3 to 6 of Xenakis's schema, its operational field remains strictly confined to the abstract layer of transformations, structural mappings, and procedural recombinations. It does not concern itself with the origins of sonic material (phases 1–2) or with its final representation or sounding (phases 7–8). This raises a critical question: what is this machine modelling, if not musical material itself? And how do the operations within the machine relate to the musical works, processes, and events that populate our lived experience of composition? A powerful lens for answering this comes from Paulo de Assis's extension of Timothy Morton's theory of hyperobjects into the domain of music. As Morton argues, hyperobjects are entities that cannot be fully grasped at any single moment; they are massively distributed in time, space, and relationality (Morton, 2013, 1–4). They are never directly accessible—only their local manifestations

are. And yet, they persist and exert effects across networks of entangled interactions. One prominent example throughout Morton's book is the hyperobject of global warming, which cannot be grasped as one single object, but can be measured, proved and certainly exists, occurring in various contingent phenomena. Morton notes that hyperobjects are not "collections" or "systems" of smaller objects. They are entities in their own right, irreducible to their components, and radically nonlocal. Their structure is always withdrawn, only ever encountered through aesthetic interference patterns, or what he calls interobjectivity (2).

De Assis applies this logic to musical practice, proposing that musical works could be understood not as bounded objects, but as hyperobjects—or more precisely, as hypermusic (de Assis, 2024, 18–20). Hypermusic consists of distributed networks of performative acts, conceptual decisions, latent potentials, and historical sedimentations. It is never entirely given. It manifests only partially through notations, performances, interpretations, and analyses. And critically, new musical assemblages always emerge from existing ones. Unlike in generative music, where a self-contained system produces a work from internal rules, hypermusical practice proceeds through the reconfiguration of a network of *existing* materials, processes, and relations. This is a crucial distinction for situating what would the envisioned machine metamodel. In generative composition, the system is autonomous: it creates outputs from within its own rule space. In interactive composition, the system responds dynamically to external stimuli—usually in real time—but the underlying structure remains fixed. By contrast, the machine being envisioned here would transform the operational framework itself. It would not just enable variation in behaviour; it would enable variation in the conditions of variation. Each execution of the system would become a transductive event—a moment of individuation that reconfigures not the output, but the structure that enables outputs.

From this perspective, the ARMM does not diagram principles for operationalising a generative engine nor for operationalising an interactive interface. It is envisioning an allagmatic machine (in a sense described in the previous chapter) that would metamodulate the operational relations between *existing* compositional elements. But—and this is the key point—these elements would not be contained within the machine. The machine would operate only on numerical indexes symbolising objects and further on indexes referring to other indexes in an intricate network of metamodelling processes, not on objects themselves. Its domain would be purely

procedural. The particles of music—the recordings, gestures, timbres, samples, motifs, scores—would all exist outside of the machine. They correspond to Xenakis’s phases 1–2 (intuitions and definitions of sonic entities) and phases 7–8 (notation and sonic realisation). The envisioned machine would neither define nor produce them. It would diagram how they might relate, recombine, or transform one another. It would build metamodels—abstract operational schemas through which external musical materials could be rearranged only by linking them back to the indexical values they were previously assigned.

In this sense, the envisioned machine’s role would be composing the conditions under which music can individuate. The musical work is a hyperobject—a distributed, emergent field whose local manifestations (scores, performances, features) always exceed their encoding. The envisioned machine would operate within this field, not by shaping its objects, but by diagramming its conditions of transformation. It would be a platform not for outputs, but for structural individuation. The works that would appear around it would not be its contents. They would be its external expressions—its local, contingent, and partial manifestations. Campbell (2013, 148) describes a similar process in Deleuzian music practices, noting that the most compelling composers today do not operate within fixed symbolic grammars, but instead “molecularise, deconstruct or dissolve existing musical material”, generating new semiotic units through hybridisation and technological experimentation. This is precisely the operational space of the machine envisioned in ARMM. It would be a machine not for preserving meaning, but for reorganising expressive potentials. Its allagmatic capacity would reside in its ability to reconfigure the relations among already-fragmented elements, allowing for emergent musical assemblages without ever prescribing their form.

The capacity to support recursive operations—where transformations can act not only on musical material but also on other transformations—would position such an abstract machine not as a tool for composition but as a platform for meta-composition. In this sense, it would enable what Sofian Audry and Jon Ippolito describe as the meta-artist function: the authorship not of an artwork, but of a system that can itself produce authors. The meta-artist, following Annemarie Bridy’s definition, is “the author of the author of the works” (Audry & Ippolito, 2019, 3). This recursive layering of agency directly parallels the logic of metamodelling: the machine envisioned in ARMM would compose not musical works, but compositional agents and structures capable

of producing them. Within this machine, every new transformation or feature would be both a tool for composing and a model of how composition itself can unfold. The machine would enable contributors to write new compositional behaviours (e.g., a transformation module), but also to define new types of transformations (e.g., a transformation of transformations). In doing so, contributors would occupy a meta-compositional role, shaping not outputs but the grammar of future outputs.

Audry and Ippolito argue that such recursion is perceptually framed by the receiver, who decides whether to ascribe authorship to the artwork, the artist, or the meta-artist depending on the conceptual distance they choose to traverse. The same holds in the machine envisioned in the ARMM: some may perceive a musical output as authored by an individual agent; others may trace it back to the underlying schema, infrastructure, or ecosystem of operations. In this framework, metamodels become epistemic scaffolds, not only for compositional action, but for the negotiation of agency and authorship across a distributed system. By enabling agents to define the structure of compositional interaction itself—not only what is done, but how doing is defined—the envisioned framework would become a machine for composing composition. This is its metamodelled, autopoietic core: a system that writes its own operational grammar through the contributions it hosts. It would not contain works—it would condition the emergence of worlds of works. And in doing so, it would fulfill a speculative compositional function: not being a tool for making music, but a platform for rewriting what it means to make music at all.

The recursive architecture of the envisioned machine, in which compositional operations define not only musical materials but also the conditions of their transformation, resonates strongly with Xenakis's vision of the evolving role of the composer. Already in *Formalized Music*, Xenakis identified a paradigmatic shift: the composer is no longer simply a shaper of forms, but a designer of schemes—formal, abstract systems for generating sound structures: “We may further establish that the role of the living composer seems to have evolved, on the one hand, to one of inventing schemes (previously forms) and exploring the limits of these schemes, and on the other, to effecting the scientific synthesis of the new methods of construction and of sound emission” (Xenakis, 2001, 133).

In this formulation, the composer becomes a meta-composer—an agent who no longer composes works directly, but composes the conditions for compositional emergence. These conditions take the form of logical frameworks, mathematical processes, and

experimental configurations that are treated algorithmically and explored through computation. Xenakis names Poisson and Markov processes, probabilistic constraints, and musical games as early examples of this procedural logic, already lending themselves to computer-based implementation (133). In the envisioned machine, this logic would be extended further: the system would support not only compositional processes but compositional processes about compositional processes. In other words, it would enable users to build schemes that themselves generate or transform other schemes—exactly the recursive condition of metamodelling.

Xenakis's broader conceptual project—particularly his vision of metamusic—provides another useful frame. In his later work, he imagines a universal musical framework encompassing all musical traditions and formal possibilities, grounded in a “general harmony” and operationalised through the theory of sieves (180). The sieve, in Xenakis's use, is a modular, arithmetical mechanism that filters integers into musical intervals, scales, or rhythmic patterns. Crucially, sieves can be added, subtracted, nested, and recombined, making them not only generators of structure but, crucially, operators on other operators. This transforms them from tools of music-making into tools of metamodelling—machinic schemas that modulate how composition itself unfolds. It is precisely this logic that the machine envisioned here would generalise. Within the machine, all transformations—whether stochastic, deterministic, or combinatorial—would be treated as operable entities that can be chained, edited, duplicated, or composed into higher-order behaviours. Like Xenakis's sieves, they would be the operational primitives for diagramming and transforming the compositional space. Even Xenakis's metaphysical reflections echo this procedural ontology. Late in *Formalized Music*, he suggests that space may be nothing more than “an appearance of [...] chains of energy transformations” (257). This notion aligns with the ARMM's view of the compositional system as a topology of indexed operations, where every feature, transformation, or schema is not a static entity but a node in a dynamic network of relational energies. In such a system, no singular object is final. Everything is an event in a broader field of metamodelling.

Social Complexity of the Autopoietic Rhizomatic Metamodelling Machine

The procedural relations between various compositional elements described as parts of the abstract machine envisioned in this chapter so far, should at this point form a

blurred picture of a compositional system in which various human and non-human actants co-constitute. A system that is capable of reorganising its own principles through the actions of its agents, who in response acquire access to a broader collective intelligence resulting from the commons of compositional processes. In that context, it is useful to sharpen this picture further by discussing the social level of such an assemblage.

Assemblage theory, particularly as articulated by writer, artist and philosopher Manuel DeLanda (2006), provides a framework for analysing complex, multi-scalar systems composed of heterogeneous elements—systems in which parts maintain operational independence even as they participate in larger configurations. These insights are especially pertinent when designing a compositional infrastructure intended to be both decentralised and recursive: a system that must support individual creative acts while enabling their aggregation into durable, evolving structures. Crucially, DeLanda insists on understanding assemblages in terms of relations of exteriority. A component part of an assemblage “may be detached from it and plugged into a different assemblage in which its interactions are different” (DeLanda, 2006, 10–11). This principle sharply distinguishes assemblages from organismic or structuralist models, where parts are defined by their place in the whole and cannot function independently. For the ARMM, this distinction is foundational. The envisioned machine must accommodate compositional units that can migrate across projects, contexts, or communities without losing their internal coherence (similarly as patchers in MAX/MSP or PureData). Each contribution must remain legible and functional outside its initial instantiation, enabling reconfiguration.

In my previous research (Łukawski, 2023), this principle was articulated diagrammatically—envisioned as a method for modelling spatiotemporal change across levels of complexity, from individual assets to entire social networks. Such a diagram, understood in Guattarian terms as an “autopoietic machine”, does not map interior relations but instead outlines the topological limits within which transduction can occur (Łukawski, 2023, 62–63; see also Guattari, as cited in Łukawski, 2023, 63). Diagrammatic thinking allows for procedural reasoning about multiplicities—configurations that unfold in time without becoming fixed into singular identities. In this context, the ARMM envisions a socially-dynamic infrastructure for processing heterogeneous contributions.

Assemblage theory also foregrounds the multiscalar nature of such infrastructures.

Social assemblages, according to DeLanda, range from individual bodies to buildings, neighbourhoods, cities, and nation-states, each with their own internal diagram and corresponding space of possibilities (DeLanda, 2006, 30). These spaces are characterised by degrees of freedom and universal singularities—features that define what kinds of events or transformations are possible within them. For the ARMM, this implies a layered structure: each compositional unit operates within a local topology of constraints, but these local operations feed into a broader ecology whose dimensions are defined not by central control but by distributed participation. This vision directly informs the need for machinic infrastructure that can accommodate radically different types of contributions—generated by human artists, AI agents, or procedural algorithms—and integrate them into coherent networks of transformation. In my earlier writing on the metaverse (Łukawski, 2023), I argued that both human and non-human agents could act as observers and participants in such a system, transducing themselves across dynamic topologies that resemble “interactive, constantly evolving, and self-composing virtual realities” (Łukawski, 2023, 63). The same logic applies to the machine envisioned in ARMM: each agent, whether embodied as an interface, a generative process, or a dataset, would participate in a field of compositional individuation governed by procedural rules rather than aesthetic norms.

DeLanda also re-introduces (after Deleuze and Guattari) the concepts of territorialisation and deterritorialisation to describe how assemblages stabilise and destabilise themselves through processes of spatial, symbolic, and functional consolidation (DeLanda, 2006, 13). For creative systems, this means recognising how compositional elements may temporarily cohere—forming genres, styles, or community conventions—only to be later reconfigured through new interventions. The envisioned machine must support this continual negotiation. Its design must allow for contributions to crystallise into recognisable structures, while preserving the ability for those same structures to be re-aggregated into novel assemblages. By situating the envisioned machine within the conceptual terrain of assemblage theory, we begin to understand it not as a singular machine but as a system of systems—an operational ecology in which each element retains autonomy while contributing to the emergence of a shared compositional environment.

Building on this understanding of the machine envisioned in the ARMM as a social assemblage, we can now turn to musical practices themselves. Edward Campbell’s

Music after Deleuze (2013) provides a crucial perspective here, aligning Deleuzian thought with compositional methodologies. His reading of Deleuze and Guattari's work allows us to view artistic process as a rhizomatic process—a process that operates through transversals, detours, and cuts across established structures. Campbell identifies a key principle at the heart of Deleuzian aesthetics: every significant creative act is a response to a problem, not a refinement of a form (Campbell, 2013, 36). In this sense, creativity is diagonally constituted—it does not operate vertically within traditions or horizontally across contemporaneous styles, but obliquely, inventing a new line in a space that was previously unformed. Deleuze and Guattari borrow Pierre Boulez's term diagonal to name this transversal vector of creation. Each innovation traces a new line between pre-existing trajectories, resulting in a new assemblage, a new plane of consistency. The machine envisioned in ARMM must be structured to support precisely this kind of transversal operation: a compositional system in which each contribution not only adds to the network but configures a new vector through it.

Such transversality defines the rhizome. Unlike a tree, which branches from a single trunk and follows a hierarchy, the rhizome proliferates laterally, with multiple entry and exit points, no predetermined direction, and no unified structure (Campbell, 2013, 21, 37). Rhizomes are processes, not forms; they unfold in time through variation and recombination. As Campbell writes, they are “irreducible to prior or overarching identities”, and they “pertain to a map that must be produced, constructed [...] and has multiple entryways and exits and its own lines of flight” (21). To the rules diagrammed in the ARMM, this implies a mode of operation that resists static classification. The envisioned machine must function as a diagrammatic machine whose paths are continuously drawn and redrawn, shaped by the flows of creative contribution rather than by adherence to fixed schemas.

In musical terms, Deleuze and Guattari contrast the arborescent tonal system with the potential of “generalised chromaticism”—a condition in which durations, timbres, intensities, and articulations are all placed into continuous variation (Campbell, 2013, 37–38). Under such conditions, music becomes a “superlinear system”, a rhizome of sound rather than a tree of notation. The envisioned machine must accommodate this logic at the compositional level of the network: a space of continuous operational variation rather than discrete, rule-bound structures. Its function would not be to stabilise musical forms but to enable their ongoing differentiation. Rhizomatic

systems are not devoid of structure; rather, they construct their own immanent logics as they unfold—what Deleuze and Guattari call planes of consistency. The ARMM, too, generates its consistency through the accumulation and recomposition of its own operational strata (as various actors continuously contribute to the shared network).

Processes of deterritorialisation and reterritorialisation further illuminate this dynamic. Campbell notes that every diagonal, every creative transversal, functions by escaping a previous configuration—deterritorialising it—only to reassemble itself elsewhere as a new assemblage (Campbell, 2013, 39–40). This continual unmaking and remaking of form is central to both musical and philosophical creation. In this way, the envisioned machine should be understood as a field of operational deterritorialisations: each compositional act would subtract itself from an existing state and recompose the field it emerged from. It would produce not works, but vectors of transformation—lines that would remain open to new affiliations, and in doing so, they would reconfigure the very ecology that sustains them.

As Campbell explores through the work of Ivan Wyschnegradsky and Pierre Boulez, contemporary music has foregrounded the spatiality of sound, revealing not only its timbral or pitch structures, but its capacity to delineate new experiential spaces (Campbell, 2013, 69–70). For the envisioned machine, this spatiality is both literal and diagrammatic: the system would compose not only sonic artefacts but new territories of interaction and relation. Its diagram would not trace an a priori structure; it would map intensities, enable traversal, and manifest potential fields of transformation. In this context, musical time itself would become a variable in what Deleuze called a “cartography of variables”—a framework of continuous differentiation rather than punctual inscription (Campbell, 2013, 101).

To conceive the machine envisioned in the ARMM as a system capable of continuously transforming itself while maintaining compositional intelligibility, we must articulate the internal mechanics of its individuation. Here, the notion of parametrisation becomes crucial. As Liam Magee and Ned Rossiter—scholars of digital data infrastructure and labor—note, Manuel DeLanda has described assemblages as entities that individuate through the modulation of their parameters—adjustable “knobs” whose values configure the emergent identity of the system at any given moment. “Changing these values”, they write, “alters the configuration of thought and practice” (Magee and Rossiter, 2015, 78). The envisioned machine, similarly, would not be defined by a fixed structure but by a field of potential operational states,

each shaped by the shifting relations. In this model, the artwork would not be the product of the system—but the system itself. Or more precisely, it would be the system as it expresses itself under particular parametric conditions. This view aligns closely with Gordon Pask’s concept of behaviour, as revisited by Sofian Audry (2021). For Pask, behaviour is not a fixed output but a form: “an unchanging form of events due to the activity within an assembly” (as cited in Audry, 2021, 47). Though the system changes over time, its behaviour can remain recognisable as an invariant morphology. This definition is particularly apt for artistic systems, where the form of a work is often legible not as a discrete object but as a pattern of performance unfolding across space and time. Such behaviour is not strictly algorithmic; it depends on the perceptual and experiential encounter between system and observer. In this sense, the envisioned machine would be a machine of aesthetic transmission, wherein behaviours would be enacted and recognised within a shared environment of interpretation.

Audry further distinguishes between different orders of behaviour in artistic systems, drawing a line between stateless mappings and stateful agents. Stateless systems—like simple input-output mappings or traditional patch-based media artworks—lack memory and cannot evolve in time; they are “zero-order behaviours” (Audry, 2021, 49). In contrast, agent-based systems possess internal state, enabling them to modify their own behaviour in response to environmental interactions. These are “first-order behaviours”, and when such systems become adaptive or generative—able to transform their own behavioural logic—they achieve “second-order behaviour” (50). Within the envisioned machine, these distinctions would not be separate layers but continuous dimensions of operation. The system must be capable of hosting mappings, agents, and meta-agents—operations that are not only performed but can perform themselves differently over time. Compositional features must not merely execute transformations; they must be able to transform the conditions under which they transform. In this framework, each contribution—whether a code module, a behavioural rule, or a performative condition—acts as a local expression of a more complex topological field. As Aaron Sloman—a philosopher and researcher on artificial intelligence and cognitive science—has proposed (1984), one can imagine a “space of possible minds”, a multidimensional conceptual space in which different systems—human, machinic, hybrid—occupy various cognitive or behavioural niches. The envisioned machine would map a similar kind of space, though articulated in compositional rather than cognitive terms. It would not contain a singular logic of musical thought; it would comprise multiple, interacting logics that would co-

determine one another through the operations they would enact. Each compositional unit within this machine would thus correspond to a behavioural vector—a path through the space of possible operations, determined by its internal configuration and its relations to other units.

This ontogenetic model of composition draws directly from Simondon's theory of individuation. As Deleuze summarises, "the individual is not only the result but the element [milieu] of individuation" (2001, 43). The preliminary condition for individuation, Simondon argues, is the presence of a metastable system—one that contains internal tensions or potentials that have not yet resolved (44). This model allows us to understand compositional acts as events within an ongoing process of individuation—each one creating new topologies of relation. For Simondon, even the individual itself is a phase—"a phase of becoming that will lead to new operations" (Deleuze, 2001, 48). This emphasis on the processual, the transductive, and the relational defines the envisioned machine's function as a posthumanist compositional machine. Its operations would not be reducible to symbolic notation or signal flow; they would constitute topological events. They would enact individuation through their performative effect, structuring space and time through behaviour. As Deleuze notes in his review of Simondon, "a living being grows from both the inside and the outside", operating in topological contact with its environment (2001, 47). The envisioned machine must operate in this mode—at once internally differentiated and externally responsive, capable of inscribing its own transformations into the network it constitutes.

Autopoietic Rhizomatic Metamodelling Machine

This chapter has introduced and explored the Autopoietic Rhizomatic Metamodelling Machine (ARMM)—an abstract machine in the sense defined by Felix Guattari as a diagram that conditions the emergence of systems, behaviours, and relations. The abstract machine described here is one that diagrams technological, philosophical, social and artistic conditions enabling their operationalisation towards a new kind of concrete machinic assemblage. ARMM serves then as an abstract preparatory diagrammatic scheme of principles, elements, and relations that can now be used in the next chapter for the operationalisation of a concrete technical system.

The primary function of the machine envisioned to be operationalised by the ARMM's principles is to configure the conditions under which musical agency, form, and process can emerge—it should enable an infrastructure which artists can use to world new artistic ecosystems. In this respect, it should participate in what Martin Rohrmeier describes as creativity at the meta-level—a kind of invention that reshapes the very task of composition, often by redefining its constraints, goals, or modes of expression. Rather than solving musical problems within a predefined style or grammar, meta-level creativity instantiates or alters the problem space itself (Rohrmeier, 2022, 54). This aligns with what Aaron Sloman calls the space of possible virtual machines: a conceptual topology in which different systems—human, computational, or hybrid—can instantiate minds or behaviours with different levels of complexity and responsiveness (Sloman, 1984, 2–3). For Sloman, we must abandon the binary between things that “have minds” and those that don't. Instead, we should examine how systems behave—how they process, learn, relate, or reflect. The ARMM treats compositional entities (features, agents, conditions) as modular, virtual machines. Each can be indexed, invoked, or recombined across compositional acts, which enact procedural behaviours across a shared infrastructure. If, as Liam Magee and Ned Rossiter argue, the relational database has had more impact on the transformation of organisational culture than the Internet itself (Magee & Rossiter, 2015, 74), the system to be operationalised by the workings of the abstract machine envisioned here, aspires to a similar role in composition: it redefines the infrastructure of musical agency by enabling new forms of compositional relation—between features, between composers, and between human and non-human agents. It makes room for systems that behave differently, think differently, or reconfigure compositional space in unexpected ways. As Sloman insists, “minds are not static things—it's what they do that is so important” (1984, 3). The system to be operationalised by the ARMM does not attempt to simulate the mind of any single composer; it scaffolds a field in which multiple kinds of compositional behaviour—human, machinic, hybrid—can co-exist and co-individuate.

The key idea behind the system to be operationalised by the ARMM is metamodelling. In simple terms, a metamodel is a model about models: a structure that defines how other structures—such as musical transformations, creative procedures, or generative systems—can be built, validated, distributed and recombined. In this way, the ARMM should operationalise a system that allows artists and agents to design not only compositions, but systems for making compositions. These systems can themselves

evolve, be shared, recombined, or restructured across multiple creative contexts. This is what gives them their autopoietic character: they are self-producing. Every operation enabled by the system resulting from the ARMM should be capable of being folded back into the system as a reusable building block, allowing the machine to modify its own structure over time.

The ARMM also operationalises a system that follows a distributed rhizomatic logic. Borrowed from the philosophy of Deleuze and Guattari, the rhizome is a model of organisation without a single root or central structure. Instead, it spreads in all directions, with many interconnected paths, entrances, and exits. In the ARMM, this means that creative operations are not chained in a fixed hierarchy, but can be linked, rearranged, and remixed in many different ways leveraging the possibilities of decentralised modes of creation. This supports a kind of creativity that is flexible, non-linear, and open to unexpected recombinations—what we have called rhizomatic metamodelling.

Rather than enforcing a single style, method, or aesthetic, the system to be operationalised by the ARMM provides a space where different creative strategies—whether generative (rule-based systems that produce outputs), interactive (systems that respond to inputs), or allagmatic (systems that transform their own rules)—can all coexist and evolve. The term allagmatic composition, as defined in Chapter 4, refers to a compositional approach where what is composed is not a sound or a piece, but a system of operations that itself can be modified and transformed. The system to be operationalised by ARMM is designed precisely to support this kind of higher-order composition.

In this way, the system to be operationalised by the ARMM functions like an infrastructure rather than a tool. It is a framework for configuring creative relations, which enables artists, agents, and systems to collaborate by contributing operations that are modular (can be reused), interoperable (can work with other components), and procedural (can be executed dynamically over time). The outcomes that arise from such a system emerge through interaction, experimentation, and transformation. This is what we have described as structural individuation: the process through which a new artistic form takes shape not by applying a fixed template, but by navigating a field of evolving relations.

Through presenting the Autopoietic Rhizomatic Metamodelling Machine as an

abstract diagram of relations and principles (operations of the abstract machine), this chapter completed the first step of what has previously been defined as worlding a compositional ecosystem. It envisioned the abstract machine, which should now be operationalised in a concrete technical system. The ARMM is then a speculative design—a conceptual diagram of a system that should now be built. But to make it function in practice, we must translate its principles into concrete operations, agents, and infrastructures. This is the task of the next chapter.

In Chapter 6, we will shift from theory to implementation. The ARMM will take shape as a real-world compositional infrastructure, built on blockchain technologies and realised as Decentralised Creative Networks—distributed systems in which creative agents interact, contribute, and collaborate. These agents may include human composers, machine learning models, or algorithmic processes. Within these networks, creative action will be driven by the Performative Transaction: a formal operational unit for rhizomatic metamodelling—that executes transformations, records contributions, and enables new compositional relations to emerge.

In short, if this chapter outlined what the Autopoietic Rhizomatic Metamodelling Machine is and why it matters, the next chapter will show what it enables—both conceptually and technically, procedurally, and operationally. It will move from speculation to implementation, from the abstract diagram to the executable machine. Through the concepts of Decentralised Creative Networks and Performative Transactions, we will see how the compositional logic of the ARMM can be translated into a functioning, programmable, and open-ended system of artistic creation.

Chapter 6. Decentralised Creative Networks and Performative Transactions

Decentralised Creative Networks

This chapter introduces the design and implementation of Decentralised Creative Networks—blockchain-based ecosystems where human artists and AI agents can transparently build upon each other’s interoperable contributions (first discussed in Łukawski, 2024, 36–37). By sharing composable processes encoded as smart contracts, they form an evolving catalogue of artistic logics that can be activated and re-contextualised through executions called Performative Transactions. Each transaction constitutes a referential gesture: a decision to inherit, transform, and extend prior contributions, and thus to participate in a cumulative network of artistic activity.

The term “Creative Networks” was first theorised by Rasa Smite to describe collaborative artistic practices based on the early Internet in the 1990s—projects that explored new forms of social communication, artistic self-expression, and experimental infrastructure through open servers and shared protocols (Smite, 2012; cited in Łukawski, 2024, 38–39). Decentralised Creative Networks echo this spirit of collaborative infrastructure, but reorient it through the affordances of blockchain and decentralised technology—particularly their capacity to encode modularity, composability, traceability, and programmability at the level of operations.

As Future Art Ecosystems 3 notes, this kind of modular logic—one that supports “permissionless participation, composable functionality and interoperability” (Serpentine Arts Technologies, 2022, 63)—is foundational not only to the ideals of decentralised culture, but to the building of new institutional imaginaries: an approach that they describe as “a GitHub for the arts” (2022, 104). Within such a system, authorship is no longer a fixed identity but becomes a function within a network of contributions. The “modularisation of identity” (136) enables contributors, whether human or AI, to operate across projects and contexts without collapsing into singular attribution. Instead of treating provenance as an after-the-fact marker of ownership, the system integrates provenance as a core operational principle—a dynamic ledger of artistic logic and agency embedded in every compositional transaction.

On the technical side, the concept answers a persistent challenge in generative blockchain-based art: the absence of standards for cross-application composability.

While the original intention behind the invention of smart contracts is to encode modular, re-usable logic (Ethereum, 2023), most current generative art NFTs are designed to operate within closed application contexts. They mirror the architectures of centralised software by linking to standardised file formats (such as audio, MIDI, and image files) that fail to represent the structural relations preceding the creation of artistic structures, like how musical intervals and rhythms build more intricate chords, melodies, forms, movements, and textures, without the flattening of such information within the final signal that becomes a media file (Łukawski, 2024, 39–41).

The system described here encodes relations between artistic operations. It functions more like a compositional operating system, or in a spirit envisioned by Claudio Tessone (2024) for future blockchain applications: a Xanadu-like environment for executable provenance, “a system that intricately links artistic assets to their origins and maintains a record of their history and transactions” (93). Project Xanadu was a visionary digital publishing system proposed by Ted Nelson in the 1960s. Unlike the modern web, which uses one-way links and often loses track of changes or authorship, Xanadu was designed to connect documents in both directions, preserve the full history of edits, and show exactly where every piece of content came from. It aimed to support transparent collaboration, shared ownership, and long-term traceability of digital material. A key concept in Xanadu was *transclusion*—the ability to include parts of one document within another while retaining a live connection to the original, so that any reuse of content always shows its source and context. As further observed by Tessone, a blockchain-based system inspired by Ted Nelson’s vision of transclusion and versioned interlinking would not only enhance attribution but propose new models for redistribution, licensing, and collaborative inheritance—offering “a unique solution yet to be fully realised in the digital art world” (Tessone, 2024, 93).

The primary operation of the Decentralised Creative Network, addressing the need for introducing a standardised format for referencing, inheriting, and executing compositional logic across contributors and applications, is operationalised here in the concept of Performative Transaction. Performative Transactions, as explained further in this chapter, are smart contracts that model the logic of recombination and execution itself. This allows for the composition of processes in a way that preserves traceability, modularity, and interpretive flexibility.

As Artemi-Maria Gioti (2021) has argued, the development of tools for creative and artistic practices must be understood as an inherently aesthetic process. The design of

such systems is more than a neutral technical act. It is a constitutively artistic one. “All design decisions are—and should be viewed as—inherently aesthetic”, she writes, underscoring the need for closer collaboration between composers and developers when building systems that support open-ended artistic exploration (144). This is particularly crucial in domains involving artificial intelligence or algorithmic systems, where the default orientation of available tools may reflect assumptions alien to the needs of contemporary artistic practice. To realign technological development with the aesthetic and methodological concerns of composers, Gioti calls for the involvement of artists as co-designers, shaping both the functionality and conceptual framing of the tools they engage with.

Following this principle, the technical realisation of the concepts of the Decentralised Creative Network and its core mechanism—Performative Transactions—that I proposed during this doctoral trajectory, was developed through a sustained collaboration between myself and programmer Michał Skarżyński. From September 2024 to May 2025, we worked closely in a process involving iterative design, implementation, and testing to build the working infrastructure described in this chapter. While I was responsible for the original architectural vision, the conceptual framework, and the formal requirements for compositional functionality, Skarżyński fully implemented this framework in code—writing smart contracts, designing the server infrastructure, expanding multiple low-level technical details and ensuring that the system functioned according to the recursive, modular logic of the compositional model. Every instance of software mentioned in this chapter—the public API and the development server—was implemented by Michał Skarżyński based on a compositional framework I envisioned, and has since been continuously refined through our ongoing collaboration. As the next sections will show, this implementation transforms Performative Transactions from a conceptual proposal into a model for decentralised compositional practice—capable of encoding not only the generation of artistic outputs but the rules, processes, and relations that structure their emergence.

A Compositional System as Public API

Decentralised Creative Network (DCN) is a system implemented as a public, open-sourced server code which, once installed, exposes API endpoints allowing for

interaction with the underlying, blockchain-based smart contracts via command-line tools, web interfaces, or by using automated AI agents, to contribute, retrieve, execute, and recombine compositional processes. The code of the server implementation is open-source, available as a public GitHub repository and can be installed locally by anyone, enabling direct interaction with the system without relying on any third-party infrastructure.

An API, or Application Programming Interface, is a structured way for one program (or user) to interact with another system. APIs are common in everyday software: for example, when a weather app displays forecasts, it is likely using a weather API that provides temperature and condition data from an external source. When a website embeds a Spotify playlist, it does so through the Spotify API, which exposes musical content and playback functions in a standardised way. The API introduced here functions in a similar way. It is a set of compositional functions for interacting with the blockchain network. These functions allow anyone, including AI agents, to create and execute reusable compositional processes encoded on the blockchain as part of a growing shared library of interdependent compositional processes.

Developers can build their own applications on top of this API. Advanced users can build intricate decentralised artistic ecosystems by encoding their envisioned artistic logics with no further care about the technicalities of lower level blockchain programming logic, while artists without prior programming experience can re-use the logics built by others to compose more intricate ecosystems from re-usable components. Most importantly, designing the compositional system as a public API introduces the possibility of composing self-reconfiguring systems—in which contributors modify the very compositional environment they are working within.

Composing with Transformations, Features, and Conditions

DCN allows its users to compose artistic processes by assembling three types of reusable building blocks: transformations, features, and conditions. Let's discuss them one-by-one:

Transformations are programmable functions that take numerical inputs and return outputs. They can be thought of as atomic musical operations—for example, “transpose pitch by 2”, “multiply duration by 3”, or more advanced logic such as “ask

an AI model to perform an action and return a number”. Each transformation is implemented as a small Solidity function on the Ethereum blockchain, but users do not need to handle blockchain logic themselves since this is automated for them in the DCN—only the core operation (e.g. $\text{return } x + 1$). Once published, a transformation becomes a reusable building block that others can incorporate into their own processes. Most users will rely on an expanding library of existing transformations, while advanced users may create entirely new ones when necessary.

Features are the primary building blocks used to compose artistic logic within the DCN. A feature defines how a sequence of operations unfolds by combining existing transformations into an ordered structure. When executed, a feature applies its internal operations in a cyclic manner, looping back to the beginning once it reaches the end. For example, a feature encoding the steps of a pitch scale could continue indefinitely into higher registers. While each feature is a self-contained sequence of transformations, features can be combined into larger structures through Performative Transactions that link them together. Once published, they become reusable, shareable, and composable units that others can build upon when designing their own artistic processes.

Conditions define whether a feature can be executed. Some conditions are simple (e.g., “always true”), while others can be more complex—such as artistic logic (e.g., “execute only if the colour of this object is yellow”), time-based logic (e.g., “execute only between 12 am and 12 pm”), counters (e.g., “execute only if this process has not been used more than five times”), or financial constraints (e.g., “execute only if the specified account received a payment”). Thus, they can represent arbitrary rules and constraints, turning the system into a flexible environment for compositional logic.

All three types of components are recorded on the blockchain with unique identities, allowing others to reuse them, build on them, and reference them in their own processes. These building blocks remain inert until activated as part of a Performative Transaction, described in the next section.

Performative Transactions

A Performative Transaction (PT) is the fundamental unit of execution within the DCN. It represents the moment when a compositional process—defined through features,

transformations, and conditions—is activated. This activation can generate outputs, trigger other transactions, and update the state of the system, depending on how the PT is designed. The term “performative” is used here to indicate that the transaction performs a process: it executes compositional logic and produces an outcome.

Each PT consists of:

- 1) a named feature with its list of transformations,
- 2) an optional list of dependent PTs,
- 3) a condition that must be satisfied for its execution.

For example, a user might create a PT that generates a melody whenever called (condition: always true). Another user could then build a new PT that combines that melody together with another rhythmic pattern as a dependent PT and applies another PT that inverts the intervals, but only when the weather is good. When a user or agent executes the PT via the API, they supply:

- N—how many values they want generated,
- a startValue for each referenced PT—indicating which transformation in its sequence to begin from.

The system resolves all references recursively, down to the lowest-level components, applying the transformations and checking that conditions are met. This produces an output array of scalar index values—which can be interpreted as pitches, durations, dynamics, links to external media, or any other digital parameters, depending on how the user maps them in their application. These outputs can be rendered as music or notation, used as performance cues, or even fed back as inputs for further PTs (as discussed further in “What Are We Modelling? Indexes, Structures, and Interpretation”).

What distinguishes a PT from a generic function call is its compositional and referential capacity. When you reference another user’s PT in your own PT, you are not copying it—you are executing that original component live, with its own logic, authorship, and constraints. This introduces several important qualities:

- 1) Traceability—every PT retains the provenance of its components; you know exactly which PTs were used, in what order, and by whom.
- 2) Modularity—PTs can be built on top of one another, resulting in complex,

layered compositional logic.

- 3) Interoperability—because each component is defined and executed through a standardised API, any PT can interact with others, regardless of who authored them or how they are used.

The public API enables several modes of use:

- 1) direct use by humans via a User Interface or Command Line Interface (e.g., composing with reusable features),
- 2) indirect use through AI agents, which can generate, recombine, or post new contributions (e.g., prompting a Large Language Model to create a transformation, feature, or condition),
- 3) recursive use within the system itself, where a transformation can trigger an API call—either to the DCN itself, or to external AI agents—allowing the system to restructure itself as part of its own operations. This means that the architecture enables the mode of allagmatic composition discussed in Chapter 4: the autopoietic system not only generates content but also participates in reorganising its own logic through PTs.

A PT is thus the crystallisation of a creative moment—a decision to activate a particular compositional structure with specific parameters. This is what allows the DCN to function not just as a storage layer or generative engine, but as a living system of transactions—a recursive, collaborative, programmable space for artistic composition.

Composing and Testing Transactions

While PTs are ultimately executed on the blockchain, the process of composing, testing, and refining them does not require direct interaction with the distributed blockchain infrastructure. To make the system accessible and practical for users, the DCN server⁸ functions both as an interface and a test development environment. It acts as an intermediary between users and the blockchain, exposing the public API and simulating the execution of transactions in a local blockchain without incurring

⁸ The open-source implementation of the server is hosted and can be accessed under the URL address: <https://decentralised.art>.

any fees for its users before modifying the distributed blockchain's state.

For users, this means that the creative process can unfold fluidly. You can define a PT, try out different PTs, test how they behave over a range of inputs, and explore how they interact—all in a simulated environment. When you are satisfied with the result, you can choose to publish the PT to the distributed blockchain, where it becomes part of the shared compositional ecosystem. Only then, the decentralised network will require a gas fee for including the new transaction on-chain.

This development environment offers several key advantages:

- 1) **Costless experimentation:** Because blockchain transactions typically involve fees, composing directly on-chain would be prohibitively expensive for exploratory work. The server simulates the Ethereum Virtual Machine (EVM), allowing users to test PTs as if they were live, but without any financial cost. This makes the system practical for creative iteration and improvisation.
- 2) **Immediate feedback:** Users can experiment with different parameter values—such as `N` (the number of outputs to generate) or `startValue` (the index from which generation begins)—and see how these affect the output of a given feature. This is particularly important when working with complex, layered features or when reusing processes created by others.
- 3) **Accessible for non-programmers:** while the low-level logic of PTs is implemented in Solidity programming language, users do not need to write blockchain code to use or recombine existing components. Many will interact with the system through graphical interfaces, CLI tools, or application-specific UIs built on top of the API. These interfaces rely on the server as a stable and programmable backend.
- 4) **Open-source and self-hosted options:** the server implementation is open-source and publicly available. Users can access the public instance and interact with the live network. However, users and developers who prefer not to rely on shared infrastructure can install and run the server locally, giving them full autonomy and control over their interaction with the blockchain network.
- 5) **Programmable by humans and agents alike:** the API exposed by the server can be used not only by human users but also by AI agents and autonomous scripts. For example, an AI agent could use the API to query available features, construct a new transformation chain, execute a test PT, and post the result—entirely without human intervention. This design opens the door for agent-

based co-creation, where human composers, algorithmic systems, and AI agents all participate in building and expanding the network.

The result is a layered architecture: while the blockchain guarantees persistence, traceability, and authorship, the server ensures usability, flexibility, and freedom to experiment. For artists, this means that composing with PTs becomes a live, iterative process. For developers and researchers, it means that new tools, agents, or applications can be built directly on top of the infrastructure without modifying its core. This separation of execution environments—on-chain for publication and off-chain for development—makes the DCN not just a compositional protocol, but a complete environment for creative experimentation and distributed authorship.

What Are We Modelling? Indexes, Structures, and Interpretation

The DCN does not produce music, images, or performances directly. Instead, it generates and manipulates indexes—integer numbers that acquire meaning only when interpreted by a user, application, or agent. This is a crucial shift in perspective: what is encoded on-chain are not artworks or media files, but compositional structures—systems of operations that unfold according to defined logic and yield numerical outputs. These outputs can then be mapped to musical pitches, durations, dynamics, gestures, samples, files or any other symbolic or material domain, depending on the context in which they are used.

At the lowest level, the outputs of a PT are arrays of scalar indexes. These indexes are not musical notes, events, or sounds. They are integers that serve as placeholders or coordinates within a potential space of interpretation. Their meaning emerges only when mapped, either in software or performance, to a particular set of symbols or parameters. An index such as “7” (counting from 0) might refer to the eighth pitch in a chromatic scale, the eighth rhythmic subdivision in a phrase, or the eighth image in a visual sequence. To be clear, it doesn’t have to map to anything “eighth” in order at all. It can map to anything, which just means that the eighth element in a list is that mapped thing—a file, a numerical value, a graphic, etc. The index itself is agnostic; the user or the rendering agent defines what it signifies.


For that reason, at the heart of every PT lies a “scalar feature”, which is a list of simple integer numbers (0, 1, 2, 3, 4, 5 ... etc) that can be further mapped to any list of digital

objects or compositional logics. Such a scalar feature should be best named based on what type of objects or processes it is mapped onto (e.g., a scalar feature of integer numbers from 0 to 127 could signify MIDI notes, or a scalar feature of integer numbers from 0 to 15 could signify available rhythmical values to be chosen from in the composition).

Suppose we define a scalar feature Pitch.


A scalar feature is a reusable generator of indexes. On every step the feature adds 1 to the previous index. Once a scalar feature is defined, it can be run with any integer number for its starting index and for its N value, which is a number of values to generate.

With a start index = 0 and N = 7, the run unfolds as:



returning the array [0, 1, 2, 3, 4, 5, 6, 7].

With a start index of 30 and N = 10, the run unfolds as:



returning the array [30, 31, 32, 33, 34, 35, 36, 37, 38, 39].

The feature does not know it's generating "pitches" (values)—it only knows it's applying a transformation "addOne" over an index space. Scalar features should be thought of as collections of all possible indexes for one of the dimensions that the composition can be built of, such as all of the possible frequencies in an audio file, all

possible pitches in a score, all possible rhythms, samples, patterns, etc. For each “type” of such a dimension, we define a separate scalar feature.

Once scalar features are defined, we can compose by selecting various combinations of indexes from them and by combining them with each other. To do it, we use transformations.

Transformation is a programmable function that maps one index to another. As W. Ross Ashby defines it in “An Introduction to Cybernetics”: “a set of transitions, on a set of operands, is a transformation”, and it is “concerned with what happens, not with why it happens” (Ashby, 1956, 10–11). Ashby identifies difference as the most fundamental concept in cybernetics: “either that two things are recognisably different or that one thing has changed with time” (9). Transformations in this system are thus formal mechanisms of change—units that produce difference without requiring interpretation. This system assumes all such differences occur in finite steps—“by a measurable jump” (9)—aligning with the discrete, index-based logic of PTs.

As already mentioned, scalar features always use the same type of transformation:

```
function addOne(x) => x + 1
```

This transformation could be interpreted as ascending steps in pitch, time, spatial position, or other dimensions. However, we can further define various other transformations that instead of generating the “next” index, they “select” indexes in a changed order. For example, a transformation that selects every second index:

```
function addTwo(x) => x + 2.
```

The logic of transformation can contain any executable computer code, but the main rule is that the transformation takes one index as its input and returns one index as its output. This enables a feature to use transformations for generating and accumulating indexes when it is run with various parameters.

Various transformations can be used to build intricate compositional patterns. To do it, we define a “composite feature” (in a PT that uses other PTs as dependencies), which uses features of other PTs as a source of elements and performs transformations on top of them to acquire a selection of values. The user chooses the PT on top of which the selection will be made and composes a chain of transformations that will be performed on top of that PT’s feature, to select its indexes.

Let's explore how this works in detail. Suppose that the previously-defined scalar feature Pitch is mapped onto MIDI notes corresponding to musical pitches, where index 60 is the middle C.

We define a composite feature MajorScalePattern. We choose the previously defined "Pitch" scalar feature as its sub-feature. We also chose the previously defined transformations "addOne" and "addTwo" in order to select a configuration of indexes that would correspond to MIDI notes shaped into a Major Scale pattern:

Feature: MajorScalePattern


Sub-feature: Pitch

Transformations: [0: addTwo, 1: addTwo, 2: addOne, 3: addTwo,
4: addTwo, 5: addTwo, 6: addOne]

Once such feature is defined, we can run it to obtain a result. To do it, the user should provide:

- 1) A start value for the Pitch sub-feature (from which index should it start generating further indexes)—in the following example, let's try 60.
- 2) A start value for the MajorScalePattern (from which index of the subfeature should it start applying its transformations to obtain indexes) —in the following example let's try index 0 (index value of the subfeature Pitch).
- 3) A transformation index of the MajorScalePattern (from which transformation in a defined chain of transformations should it start selecting the indexes iteratively)—in the following example let's start from 0, which is the "addTwo" transformation (see the definition of the MajorScalePattern feature above).
- 4) A global N value for deciding how many samples to generate using the feature—in the following example let's try 7 indexes.

When the feature MajorScalePattern is run with the chosen parameters, it applies its list of transformations on top of the index space generated by its subfeature Pitch:



Important: observe that the MajorScalePattern’s transformations such as addTwo and addOne modify the indexes of Pitch, NOT the Pitch values under these indexes. The transformations are performed on indexes, not on the values that they represent.

To provide another example, now let’s first define a scalar feature Duration, assume that its index [1] symbolises a duration of a sixteenth note, and that each next index symbolises a doubling of that value (eighth note, quarter note, half note, whole note, etc.):

Scalar Feature: Duration

Further, we define a composite feature RhythmPattern that applies a pattern of transformations on top of the Duration’s scalar feature:


Feature: RhythmPattern

Sub-feature: Duration

Transformations: [0: multiplyByTwo, 1: divideByTwo]

To run this PT, we need to provide a start index for the dependent PT Duration (let's try 0), a start index for the PT RhythmPattern (let's try 1), an index of transformation to start generating with (let's start with 0), an a global N value for the number of samples to generate (let's generate a rhythm of 4 values).

The PT runs as follows:



The result is the array of N=4 values: [1, 2, 1, 2].


As we have already defined two scalar features (Pitch and Duration), and two composite features (MajorScalePattern and RhythmPattern), we could now combine them further to create a melody:

Feature: Melody

Sub-features:

[0: MajorScalePattern, 1: RhythmPattern]

Transformations: For MajorScalePattern: [0: multiplyByTwo, 1: addOne, 2: minusTwo] For RhythmPattern: [0: addOne]



The system handles this automatically: it calls the sub-features, applies their transformations, checks any conditions, and outputs final scalar indexes.

To run the Melody PT, start indexes have to be provided for each used feature and sub-feature. Executing Melody results in a sequence of Pitch and Duration values. As traceable on the figure, this PT first runs its scalar features to calculate their indexes. Further, MajorScalePattern and RhythmPattern run to select their indexes. Finally, Melody selects its final indexes to form a new melody based on pitch and rhythm.

From the system's point of view, Melody contains a composite feature with two branches. While the operations presented on the diagram here might seem complicated and difficult to follow, this is exactly where the biggest strength of the system lies. The system handles all of these calculations automatically: it calls the sub-features, applies their transformations, checks any conditions, and outputs the final values. The user doesn't need to think in terms of iterations and complicated trees. They simply combine pre-existing features and execute them to test the final outputs.

Both scalar and composite features can be further used as building blocks within other PTs. In this way, the PT called Melody, although itself a composite structure, could be nested inside larger compositional constructs—such as phrases, motifs, or full scores—without any change to its internal logic. For example, one could define a new PT VariationA that references Melody as a dependent PT and applies further transformations to its output. Another PT, VariationB, might use the same dependent PTs as Melody but change the transformation pattern, or modify the starting index and N values to produce a different walk through the same underlying space.

This iterative architecture allows any PT to serve as a self-contained, reusable module within a broader system. By chaining PTs across multiple levels, users can construct highly structured generative processes using a minimal set of operations. Each layer encapsulates its logic, and higher-order features manipulate only the index space exposed by their subfeatures—never needing to understand their internal workings. This principle of compositional encapsulation enables a powerful mode of structured abstraction, where creative variation emerges by recombining, reparameterising, and repurposing existing components.

While many transformations are deterministic—returning the same output for a given input—it is also possible to define non-deterministic or stochastic transformations. For instance, a transformation might introduce randomness by returning a different value

each time it is called, even with the same input. A transformation such as `RandomWithinRange(min, max)` could be programmed to return a pseudo-random number between min and max. This introduces controlled variability into the system, enabling the generation of different outcomes from the same structural logic. Such stochastic processes can be particularly valuable in generative music, where variation and unpredictability are key compositional resources.

Just as transformations introduce difference through operational logic, conditions introduce difference through contextual logic. They don't enforce limitations so much as define the creative circumstances under which a feature may unfold. Conditions allow features to respond to their environment—to be activated only under certain temporal, social, or physical configurations—enabling new forms of situated, responsive composition. Conditions can be illustrated simply. For example, a reusable condition called `MaxFivePlays`:

```
Condition: MaxFivePlays
```

```
Logic: Allow execution only if the feature has been triggered <  
5 times
```

This would block a PT after its fifth use.

Other conditions might involve access control or interaction with external data. A condition named `TokenHolderOnly` could restrict execution to users who hold a specific token in their wallet, effectively tying compositional access to ownership within a broader ecosystem. More complex conditions can query off-chain information through blockchain oracles—mechanisms that securely feed real-world data into smart contracts. For example, a condition called `IfSunnyInWarsaw` could be programmed to execute only if a trusted weather oracle reports clear skies in Warsaw. Another condition, `MinAudience10`, might monitor the number of active viewers on a platform and allow execution only if at least ten users are watching the artwork simultaneously—creating a form of social activation. Finally, sensor-integrated conditions are also possible. A condition like `HumidityBelow60` could be connected to a room-based sensor, executing the underlying feature only when local humidity levels fall below a defined threshold. These examples show how conditions can embed compositional logic within wider artistic contexts, transforming the artwork into a responsive and situated system.

What the infrastructure models, then, is not the artwork itself but the logic that generates its constituent parts. The system allows users to define how one value leads to another, how values relate within and across dimensions, and how those relations can be composed into more complex operations. Each PT defines a structure: a series of transformations that recursively unfold to produce a stream of indexes. These structures may be simple (such as a linear sequence using a single transformation) or highly composite, consisting of nested features that apply different operations in parallel or in sequence, each governed by its own logic and conditions.

From the perspective of the system, there is a key difference between scalar features and composite features as explained before. A scalar feature is terminal: it generates a sequence of raw indexes and does not depend on any other features. It represents a single operational line. A composite feature, by contrast, is defined in terms of other features. It selects and transforms the outputs of its sub-features, assembling them into higher-order structures. Conceptually, one can think of composite features as defining operations on operations—processes that organise other processes. But from the user’s perspective, these distinctions may be invisible. Whether a feature is scalar or composite has no bearing on how it is used or executed. All features can be executed in the same way—by calling them with a given `N` and `startValue`—and all produce scalar outputs that can be interpreted downstream.

What matters is that composite features introduce a new kind of modelling: not of values themselves, but on indexes that signify relations between values. When a user creates a PT with a composite feature, they are effectively constructing a graph—a directed structure in which each node is a PT and each edge represents a transformation. The system executes this graph recursively, starting from the top-level PT and resolving all dependencies until it reaches the terminal scalars. This graph is not hard-coded; it is assembled dynamically at the moment of execution, based on the features and transformations referenced in the transaction. This means that every execution is both a computation and a traversal of a user-defined compositional topology. Every node can be used to select values from (or “metamodel”) any other node in the network.

The power of this metamodeling approach lies in its generality. Because scalar outputs are only meaningful when interpreted externally, the same infrastructure can support an enormous range of applications: music, choreography, image generation, text sequencing, algorithmic theatre, or even non-artistic domains. What is shared

across all these cases is the logic of indexed generation, transformation, and combination. The system models not the end products, but the operations that give rise to them. In this sense, it is a metamodel: a system for composing compositional processes, and for recursively modifying their operational relations.

What the DCN models is not musical content, but the operational structures through which such content can be generated. It encodes relations, not representations. And it does so in a way that is open-ended, iterative, and designed to support a plurality of interpretations (the same node can be reused in various artistic contexts). This is what makes it not just a tool for composition, but a compositional system in itself—one that can be shaped, expanded, and recomposed through the very transactions that take place within it.

This orientation toward operational logic, rather than fixed content, invites a reconsideration of what constitutes a musical “work”. If outputs are endlessly variable, yet structurally derived from a stable procedural form, what defines the identity of a composition? As Edward Campbell observes, Harrison Birtwistle once claimed he “could have composed the same piece with entirely different pitches”, raising the question of compositional identity as distinct from sonic material (Campbell, 2013, 34). Mauricio Kagel went further, composing *Morceau de concours* twice—in 1971 and again in 1992—with the two versions sharing “not a single note in common,” yet still understood as iterations of the same piece (Heile 2006, cited in Campbell, 2013, 34). The DCN operates in a similar conceptual space: it is the space of operations and links between them that persist.

How the Infrastructure Works

Behind the user-facing simplicity of features, transformations, and conditions lies a carefully structured technical architecture. It ensures that compositional processes are executed correctly, persistently, and traceably on the blockchain. While users interact with the system through an intuitive API—building and recombining processes as artistic components—the underlying infrastructure handles registration, execution, state updates, and provenance tracking in a decentralised, modular, and verifiable way.

At the heart of this system is a set of Ethereum smart contracts. These contracts define the formal logic of PTs and the objects they manipulate. When a user defines a new transformation, feature, or condition, they are essentially creating a new contract that adheres to a shared interface. The logic of a new transformation is written in Solidity—the programming language for Ethereum smart contracts—but the broader blockchain interactions, such as how this transformation will be executed, referenced, or composed, are abstracted away. This separation allows users to focus on writing the transformation logic itself without having to manually define low-level blockchain interactions. In this sense, the infrastructure offers a structured but open-ended programming model: the rules of composition are strict enough to guarantee interoperability and execution, yet open enough to accommodate any kind of transformation logic.

Each component that users contribute is registered on-chain via a central registry contract. This registry acts as a decentralised index, associating unique human-readable names with their corresponding smart contract addresses. The registry ensures that all components are discoverable and verifiable, enabling any user to reference components authored by others. When a user creates a new PT, they can declare which existing components it builds upon. These declarations are stored as part of the feature's metadata and can be resolved programmatically during execution. Importantly, this also guarantees traceability: every compositional object carries within it a complete record of the components it depends on and the order in which they are invoked.

The actual execution of a PT is coordinated through another component of the infrastructure: runner. The runner is a smart contract responsible for initiating the execution of features, resolving their dependencies, applying transformations, and checking whether conditions are met. When a user sends an execution command—specifying a feature name, an index value, and the number of values to generate—it is the runner that traverses the feature structure, recursively applies all transformations, and produces the output sequence. During this process, the runner interacts dynamically with the registry to retrieve the correct components, and with each component's contract to execute its specific logic. This recursive traversal ensures that even highly nested, multi-layered features can be executed in a deterministic and reproducible manner.

The runner performs this work on-chain when a PT is executed in the blockchain environment, but the same traversal logic is replicated off-chain in the development server. This duplication ensures consistency: the results of off-chain testing correspond exactly to the results of an on-chain transaction. Users thus gain the benefits of both worlds—freeform, costless experimentation in the development environment and verifiable execution and authorship in the decentralised network.

While the registry and runner provide the scaffolding for execution, the server layer—described earlier—provides the necessary bridge between users and this compositional infrastructure. It exposes the API endpoints that allow users to query registered components, simulate executions, test conditions, and post new features or transformations. Developers building user interfaces, AI agents, or other software tools can rely on this API as a stable, predictable interface. In turn, the server communicates with the blockchain itself, abstracting away the complexities of transaction signing, gas estimation, and error handling.

Through this architecture, the infrastructure enables users to treat compositional logic not as a static score or encoded object, but as a living, executable process. Each new contribution enters a shared compositional space where it can be accessed, invoked, or recombined by others. The registry guarantees discoverability, the runner guarantees correctness of execution, and the blockchain guarantees persistence and authorship. At the same time, the off-chain server layer allows creative experimentation to remain open-ended, iterative, and responsive.

This division of roles—between blockchain and server, between transformation logic and system execution—ensures that the DCN remains simultaneously robust and generative. It does not dictate how users should compose, only how compositional processes are stored, connected, and activated. What emerges is a protocol for building compositional ecosystems.

AI Agents and Recursive Infrastructure

One of the defining features of the DCN is that it is not only usable by humans, but also designed to be interoperable with software agents—particularly AI systems. This opens a new domain of compositional practice, where agents can participate as contributors, users, and even meta-composers within the system. Crucially, these

agents are not peripheral add-ons. They can operate both externally to the network, using the API like any human would, and internally, embedded within transformations that themselves call the network during execution. This dual capacity allows AI agents to act as recursive infrastructure-builders within a live compositional system.

One of the most promising developments in this area is the capacity of Large Language Models (LLMs) to operate as reasoning agents capable of executing code and interacting with APIs (Eletti et al., 2023; Yao et al., 2023; highlighted in Łukawski, 2024, 45). This functionality allows LLMs not only to query and compose with existing features on the network, but also to generate new ones, manage complex logic chains, and post PTs autonomously. In this model, LLMs can curate and build evolving collections of artistic processes—collections which, once posted, become available to other human and posthuman agents for further recombination. LLMs can also serve as natural language interfaces between users and the network, translating intuitive prompts into executable PTs (Łukawski, 2024, 45). This deepens the recursive entanglement between agents and infrastructure, where AI is not just an assistant but a fully embedded co-actor shaping the evolving space of composition.

To clarify how this works in practice, let us consider three distinct use cases.

1. External AI agent using the API like a human user

An AI agent can be programmed to browse the network via the public API, query for existing features, and construct new features by combining them. For example, a generative agent might:

- Search the network for all features tagged with "rhythm" and "duration",
- Select one at random or by a heuristic (e.g., sparsity, entropy),
- Combine it with a pitch-generating feature using a transformation such as Add or Invert,
- Wrap the result in a new feature and post it back to the network.

This results in a new PT authored by an AI agent. The feature may be executed by a human composer later, or may be picked up by another agent in future runs. In this

way, the AI behaves like any other network participant—exploring, composing, and contributing via the API.

2. A transformation that calls an AI model from the inside

More complex behaviours emerge when AI agents are embedded directly in transformations. Consider a transformation named `MLHarmonySelector`. Its logic might be:

```
function run(x) returns (uint32) {  
    // send input index x to an external AI model via API call  
    using blockchain oracle  
  
    // receive a predicted harmony index y based on a trained  
    style model  
  
    // return y as the transformed value  
}
```

Here, `x` might be the current position in a phrase. The transformation sends it to a machine learning model (e.g., a Transformer trained on chorales), which responds with the most likely harmonically appropriate value `y`. This value is returned as the output of the transformation and continues through the feature’s execution.

From the system’s point of view, nothing special happened: it simply called a transformation. But that transformation reached outside the chain of purely algorithmic operations, consulted an AI, and folded its result back into the compositional process. This enables non-deterministic, data-driven, stylistically-informed decision-making within the otherwise deterministic logic of PT execution.

3. Recursive AI-based transformation that modifies the network

The most advanced form of agency occurs when a transformation uses the API during execution to modify the compositional infrastructure itself. Imagine a transformation called `AutopoieticAgent` that, in response to a given index, does the following:

- Queries the API to retrieve all features tagged "melodic-pattern".
- Evaluates them using an internal scoring function (e.g., novelty, diversity).
- Generates a new variation using an AI model.
- Posts this variation to the network as a new feature.
- Returns the index of that new feature as the transformation result.

This is a PT that not only transforms a number, but expands the space of possible transformations for future users. It effectively composes new compositional logic.

Because the infrastructure is recursive and open-ended, this type of agent-based transformation is not speculative. It is directly enabled by the architecture. The same API that a human uses to post a new feature can be called from within the transformation itself. This allows PTs not to be just compositions, but meta-compositional events—modifying the system as they run.

These three examples illustrate how AI agents can operate at multiple layers of the Decentralised Creative Network:

- as external users,
- as embedded transformation logic,
- and as recursive participants that shape the network from within.

The boundary between composer, performer, and system is blurred. A PT may originate from a human, but contain transformations authored by AI, which in turn call other agents to generate outputs or spawn new features. This deep entanglement of agents and infrastructure is a direct consequence of exposing the compositional system as a public API.

In this model, AI is a co-actor in the compositional ecology, capable of contributing both content and infrastructure. What emerges is a new kind of compositional logic: distributed and autopoietic—where agency lies in the continuous shaping of the network that makes composition possible.

Case Studies: *Ani(mate)* (Movement VI), *Allagma*, and *Chain of Thoughts*

1. *Ani(mate)* for piano and strings (Movement VI):

In Chapter 1, I introduced *Ani(mate)* as a commissioned work whose first five movements were composed through connectionist, curatorially driven methods. Having now established Performative Transactions and Decentralised Creative Networks as the core operational constructs, I return to *Ani(mate)* and extend the discussion to two further works developed during this doctoral trajectory.

The purpose of this section is to show, in practice, how PTs and the DCN function as a compositional infrastructure and how they reframe authorship, identity, and provenance in three complementary case studies: the sixth movement of *Ani(mate)* (a PT-composed movement selected for a fixed score), *Allagma* (an agentic two-phase composer for any MIDI instrument), and *Chain of Thoughts* (an ensemble piece generated from a score of prompts). Together, these works can be read as a progression from human-defined vocabularies, through agent-generated vocabularies, to composer-specified prompt chains that the agent realises within the DCN.

The sixth movement of *Ani(mate)* was composed as the first full-scale artistic experiment with the test version of the DCN, created entirely from compositional structures authored as PTs. The aim was to investigate whether a self-contained vocabulary of modular generators—defined through features, transformations, and conditions—could produce musical material whose identity remained traceable through the network’s provenance structure.

I began by defining a small set of permissible transformations (limited to additive and subtractive operations on integer streams) and applying them to terminal scalar features such as pitch, time, duration, and velocity. These scalar features were then combined into higher-level composite features, which functioned as modular “generators”. Large Language Models were used here as co-authors at the design stage: at each stage I used the model to propose new candidate musical features, which I then reviewed and refined before registering them on-chain. The resulting vocabulary formed a matrix of interdependent generators whose relations were explicitly recorded in the DCN registry.

To generate the musical material, I executed these generators as PTs with varying

sample counts (N) and starting seeds. This produced families of structurally related yet distinct outputs, all of which retained their provenance links to the originating features and transformations. From this collection, I curated a selection of outputs to notate and orchestrate as the final score of the movement.

Working with the DCN in this way influenced the compositional experience in several important respects. Unlike the connectionist workflow used in movements I–V, composing this movement felt highly controlled. Each generator was constructed as an explicit configuration of transformations applied to musical features, designed and adjusted in a symbolic, rule-based manner. Because these transformations operated on predictable integer streams, I could reason about their effects before execution and verify them immediately using the MIDI player. This approach enabled precise interventions in the internal logic of the composition while maintaining a coherent structure across its outputs.

The provenance structure of the DCN also had a strong impact on how I listened to and evaluated the material. Because every generated fragment could be traced back to the exact configuration of features from which it emerged, it became possible to read the music as a graph of decisions. Every fragment I generated could be traced back to the exact constellation of features and transformations from which it emerged, making the relationship between cause and effect unusually transparent. This shifted my evaluative focus: rather than judging fragments as isolated products, I could understand them as outcomes of specific procedural definitions, and adjust those definitions accordingly.

2. *Allagma* for MIDI instruments:

While *Ani(mate) VI* explored the manual construction of a compositional vocabulary within the DCN, *Allagma* extends this approach by delegating the creation of that vocabulary to an autonomous agent. The work is conceived for any MIDI-controlled instrument and will be premiered during the final PhD concert at Studio 6, Amare (Den Haag) on the 20th of November 2025, performed using instruments including a piano voorsetzer (which enables an acoustic grand piano to be played via MIDI) and the microtonal MIDI organ robot *4Pi*, both constructed by composer, inventor and instrument-builder Godfried-Willem Raes and delivered by the Logos Foundation.

The piece is composed as a self-contained composer: a program that generates a new realisation each time it is run. It operates in two distinct phases. In Phase 1, the system prompts a Large Language Model to design a set of PT features that will serve as its own vocabulary. These features are generated and registered on the DCN as reusable compositional modules. The agent evaluates its own previous outputs during this phase, using the evolving vocabulary as context for producing further features until a diverse collection has been assembled.

In Phase 2, this vocabulary is used to generate the composition itself. Each vocabulary element is executed multiple times as a PT, producing streams of musical data (MIDI pitch, time, duration, velocity, etc.) from different seeds and transformation offsets. These streams are then passed back to the Large Language Model, which receives summaries of their temporal spans and is instructed to act as a scheduler, determining when each generated fragment should begin within the overall timeline of the composition. This ensures that the resulting piece is not a simple concatenation of fragments but a structured form with interleaved episodes, reprises, and silences. The final output is exported as a unified sequence of MIDI data for performance on the robotic instruments.

The process used for the generation of the work's iterations embodies the principle of allagmatic composition introduced earlier: the work exists as a system that continually reconfigures its operational structure to produce new versions of itself. Each execution of *Allagma* represents a distinct individuation of the same underlying system, preserving the procedural identity of the work while producing divergent musical outcomes.

Compared to *Ani(mate) VI*, which relied on a manually constructed and tightly supervised vocabulary, *Allagma* shifted the compositional act from designing the vocabulary itself to designing the conditions under which a vocabulary could emerge. This reframed my role: rather than specifying individual transformations, I specified the meta-level rules and evaluation criteria through which the agent could construct its own set of compositional primitives. The resulting material often contained structural relationships and textural ideas that I would not have devised myself, yet these were generated within a framework that preserved traceability through the DCN's provenance graph. This made it possible to analyse and iteratively refine the agent's evolving vocabulary while allowing its internal logic to develop autonomously.

Crucially, this experiment demonstrated that allagmatic composition does not require human authorship at every level of the system. What matters is not who defines each element, but that the system preserves and exposes the operational provenance linking those elements together. In *Allagma*, the human-computer relation became one of co-construction rather than curation: I built the operational boundaries and evaluation criteria, while the agent populated that space with compositional content. Each new iteration of the piece thus acts as a record of the agent's reasoning process within a set of constraints, revealing how meaning and structure can arise from the recursive interplay between autonomous and human-specified contributions.

3. *Chain of Thoughts* for ensemble:

The third case study in this sequence explores a compositional model built not from pre-defined vocabularies but from *chains of prompts*. Instead of specifying musical material directly, I wrote structured textual instructions that guide an AI agent to compose one bar at a time within the DCN. Each instruction is interpreted in the context of everything the agent has generated so far, allowing it to shape the ongoing musical discourse through a continuous process of contextual reasoning.

A typical instruction might read as follows:

"Compose the next bar for alto flute and violin. Use material that contrasts the previous bar's staccato texture with sustained tones. Reuse one pitch from the previous bar's pitch collection, but place it in a different register. Maintain a dynamic level around mezzo-piano. Ensure rhythmic alignment between the two instruments, with note onsets every quarter beat".

When receiving such a prompt, the agent queries the DCN for relevant existing features (e.g. pitch sets, rhythmic cells, articulation patterns), generates new features if needed, and resolves the instruction into concrete data streams for pitch, time, duration, and velocity. These outputs are registered in the network as new PTs and immediately incorporated into the developing score. Each new bar thus depends not only on its local instruction but also on the accumulated material history of the composition, which is passed to the agent as context at every step.

This system forms the basis of the composition *Chain of Thoughts*, written for Ensemble Modelo62 and scored for alto flute, violin, bass clarinet, trumpet, cello, and double bass, to be premiered at Studio 6, Amare (Den Haag) on the 20th of November 2025

under the direction of conductor Ezequiel Menalled. The piece exists not as a fixed musical text but as a score of interdependent prompts: each run of the system realises these prompts as a new version of the work, while preserving the provenance of their underlying operations in the DCN.

Unlike *Allagma*, which relied on an AI agent to autonomously construct its own vocabulary before composing with it, this piece involved writing a fixed score of agential operations specified as a chain of prompts. Each bar of the composition is defined by a detailed natural-language instruction describing how its musical parameters should be generated, but without prescribing exact notes. These prompts give the agent controlled ambiguity: enough freedom to produce locally varied material, yet constrained by the precise operational role assigned to each bar. For the piece to be performed, first the AI agent (or actually human agent for that matter) has to execute this sequence of prompts, accessing and creating features in the DCN as needed, and rendering the result as musical notation. Every run of the system thus produces a new score, yet all realisations share the same internal structure of compositional operations encoded in the fixed chain of prompts.

Building on the Infrastructure: Interfaces, Applications, and Experimental Economies

As the technical architecture of the DCN becomes operational, the next step is to design diverse interfaces and applications that make its compositional logic accessible, legible, and usable. Without well-considered user environments, even the most powerful systems remain under-utilised. As Madeline Gannon aptly notes, reflecting on the public impact of ChatGPT: “Interface is Everything. (...) Invest the time to make it easy-to-use and easy-to-understand by the general public ... and lo and behold! It takes the world by storm” (Gannon in Cuan, 2023, 353). This observation is particularly salient in the case of DCNs, where the underlying system is intentionally open-ended but may appear opaque to non-technical users.

What is needed now is a proliferation of decentralised applications (dApps) that engage with the public API and offer tailored tools for creative practice. These applications might include: interfaces that allow users to define, test, and share their own low-level transformation functions without needing to write Solidity code, visual

environments for composing features by linking existing components, inspecting dependencies, and defining transformation chains, tools for crafting and combining conditions based on state, oracles, or sensor data—enabling complex and responsive creative logic, playback engines, notation renderers, or spatial mappers that interpret scalar indexes in musical, visual, or performative terms, systems for sharing, remixing, and monetising features—emerging not from the logic of finished works but from the modular ecosystem of processual contributions. The infrastructure already supports all compositional functions necessary for such applications to operate. What remains is the translation of this potential into practical environments—designed for artists, not only developers.

Moreover, new kinds of creative economies could emerge around the contribution and activation of compositional logic. In contrast to conventional art markets, which revolve around the trade of finalised objects, the DCN infrastructure enables the circulation of processes, tools, and operational building blocks. Financial conditions embedded in features may compensate upstream contributors each time their logic is reused. Curators could emerge not as exhibitors of works, but as orchestrators of feature ecosystems—assembling and maintaining libraries of generative logic with specific aesthetic orientations. In such an environment, the notion of the artist might itself blur, giving rise to hybrid roles: curator-programmers, composer-infrastructurelists, agent-facilitators.

In this light, PTs may extend far beyond compositional practice and serve as a foundational infrastructure for the management and execution of modular assets in future metaverses. Their ability to encode compositional logic, track provenance, define activation conditions, and structure recursion makes them ideally suited to model not only artistic processes but complex interactions in immersive digital environments. A PT could just as easily define the behaviour of a virtual object, a game mechanic, or a dynamic score as it does a musical phrase. As metaverses increasingly demand programmable and interoperable asset architectures, PTs offer a compelling solution—unifying symbolic, procedural, and relational logic in a standardised, extensible format.

Finally, the availability of a modular and programmable compositional infrastructure opens new possibilities for realising visionary projects that until now have remained speculative. One such project is *Hypermusic Experiment 0.9* by Einar Torfi Einarsson—a system that operates as an “infinite notation-machine” and generates decentralised,

nonlinear, multipiece scores from internal processes, performer interpretations, and external data inputs (Einarsson, 2024, 130–141). The project’s emphasis on networked interactivity and API-driven responsiveness aligns well with the capabilities of the DCN model. By using DCNs, *Hypermusic Experiment 0.9* could now be instantiated as a living, compositional ecosystem—where each partial-score, external trigger, or post-performance response becomes a PT within a distributed graph of operations.

A second example comes from Kosmas Giannoutakis, who similarly envisions decentralised technologies as enablers of collaborative and sustainable music-making practices, rooted in mutualism and the commons, and structurally “rhizomatic” (107). Giannoutakis proposes a blockchain-based framework for collaborative electroacoustic music-making through live coding (2024, 1). His approach, which integrates decentralised ledger technologies to enable transparent authorship and distributed agency, “aims to facilitate collaborative music-making and also challenge traditional notions of authorship and the creative process” (Giannoutakis, 2024, 70). Here too, DCNs offer a ready-made environment: every code snippet could be encoded as a transformation, every live-coded gesture as a PT, and every collaboration as an evolving graph of modular relations across users and time.

What these projects have envisioned conceptually can now be encoded, executed, and shared via a compositional system that supports modularity, provenance, and recursion at every level. In this sense, DCNs do not only model a new approach to composing; they operationalise a new phase in artistic research—one in which speculative infrastructures can finally come to life.

From Composition to Ecosystem: Worlding through Performative Transactions

As the DCN accumulates new PTs, it becomes an evolving ecosystem of interrelated compositional logics. Every contribution extends the operational vocabulary of the network; every new PT deepens the graph of executable relations. What emerges is a compositional space that continually reorganises itself in response to new processes.

When a user contributes a new feature or transformation, they are modifying the space of possibilities within which other users, agents, or systems can operate. A new transformation might offer a novel mode of variation; a new condition might introduce a procedural rule that filters which kinds of transactions are permitted. In

each case, the contribution has a world-making quality: it alters the structure of the compositional environment.

This shift—from composing outputs to composing systems—marks a deeper transformation. The network becomes a space in which artists, developers, and AI agents are not just producing artefacts, but co-constructing an evolving infrastructure. The DCN functions as a recursive system allowing compositional logic to restructure itself. Because each component is executable, reusable, and referential, new contributions expand the space of interaction, enabling increasingly complex operations and feedback loops across agents and levels. Over time, this operational space gives rise to a form of compositional worlding. Every PT intervenes in the compositional system as such—establishing new relations, exposing new variables, or defining new modes of interaction. These interventions accumulate, forming a topology of operational possibilities that shapes the way future works are conceived and executed. Users entering the system encounter a dynamic ecosystem of interdependent components—each representing a decision, a strategy, a logic that can be activated or reconfigured in their own practice.

While artists can use this system to describe and operationalise the intricate operations of whole artistic ecosystems, the network can also be used more pragmatically in everyday's compositional practice. Composers can “world” small ecosystems of compositional procedures as reusable vocabulary, which might probably be the most frequently used way of interacting with the system. They could for instance, create a space of reusable generative processes for a given compositional style, or for a given set of compositional procedures. Importantly, they can seamlessly discover and reuse such components created by other human and non-human actants. This worlding is embedded in the infrastructure. The network exposes all compositional processes as programmable interfaces, allowing users and agents to write themselves into the system by defining how things unfold. In doing so, they participate in generating the operational conditions under which content is produced. And because these conditions are stored, indexed, and executable, they persist beyond their initial context, enabling other agents to inhabit, modify, or recombine them.

In this model, creative practice becomes infrastructural: the act of composing is inseparable from the design of the system through which composition takes place. Each contribution is simultaneously an artistic gesture and a reconfiguration of the world in which future gestures will occur. This is the compositional significance of

Performative Transactions. They encode not only a process but a relation—a way of structuring how processes reference, transform, and activate one another. They define the logic of compositional inheritance and mutation, making it possible to trace how ideas evolve, diverge, or converge across time and across agents. And because they are executed within a programmable, decentralised framework, they allow these relations to persist as part of the operational reality of the network.

In such a system, composition becomes a matter of ecology: a practice of cultivating, navigating, and expanding a network of interdependent processes. The composer no longer authors isolated forms, but intervenes in an environment—planting seeds, grafting branches, pruning pathways. The artwork is not the final output, but the ongoing system through which outputs are generated, shared, and transformed. The result is a shift from composition as expression to composition as ecosystem-building: a recursive, collective, programmable space in which the infrastructure itself becomes the object of creative work.

Conclusions

Artistic research is most compelling when it folds theory back into practice and allows practice to remake theory in return. This dissertation set out to test that recursive promise by asking how composers might create—not only new pieces—but the very infrastructures through which future pieces, collaborations, and even compositional worlds can emerge. Across six chapters the inquiry moved from close readings of present-day tool use to the speculative design and prototyping of a fully functioning blockchain-based compositional ecosystem. What follows gathers the main insights of that journey, clarifies their wider significance, and sketches the paths that now open beyond the confines of this doctoral dissertation.

Chapters 1–4 traced a gradient of artistic agency. Beginning with artists as operators who deftly steer complex AI models, the discussion shifted to artists as curators of agential assemblages, system-builders who encode generative logics, and finally artists worlding ecosystems—architects of conditions under which musical ecosystem can emerge rather than of isolated works. Read together, these roles chart a deliberate expansion: from interacting with single tools, to configuring relations among many human and non-human agents, to composing the ontogenetic field in which all later relations will take place. The arc suggests that the most urgent creative task today is neither to celebrate nor to resist technology, but to decide how and for whom techno-artistic systems should be built.

To name the expanded task of *composing the capacity to transform* the thesis introduced the concept of allagmatic composition—a mode in which the composed system is capable of autonomously transforming itself. Where generative music varies outputs within fixed procedures, an allagmatic work can re-write its own procedures as it runs, producing what Simondon called transductive events that continually redraw the horizon of future possibility. In practical terms this concept reframes composition as the design of epistemically complex artefacts whose value lies not in finished sounds but in sustained capacities for learning, adaptation, and knowledge production.

The practice of *worlding compositional ecosystems* was proposed as consisting of two steps: 1) iterative envisioning of an abstract machine with its operations and relations of various actants, objects, processes, and their interrelations with the (abstract) system, and 2) operationalising such an abstract machine in an actual technical system,

a functioning machine with precisely defined set of executable operations. Chapters 5-6 exemplified such *worlding of a new compositional ecosystem* in practice to first conceptualise and then operationalise a new blockchain-based system for music composition in which human and non-human actants can transparently build on top of each other's compositional processes.

Following the first step of how the proposed process of worlding a compositional ecosystem proceeds, Chapter 5 condensed the conceptual and theoretical findings of the previous four chapters into an abstract diagram of relations and conditions. The resulting Autopoietic Rhizomatic Metamodelling Machine (ARMM): an abstract machine (in Guattarian sense) envisioned the operationalisation of a system that treats every compositional act as a modular virtual machine. Rhizomatic in structure and autopoietic in behaviour, the ARMM is an infrastructural diagram—one that serves as a set of interrelated conditions and principles for operationalising a technical system in which generative, interactive, and allagmatic practices can co-exist. By elevating the process of metamodelling—modelling how models themselves can be built—the ARMM helped envisioning a technical system that offers composers a vantage point from which to shape not just musical forms, but the very grammars of formation.

Chapter 6 built on the conceptual diagram of the ARMM conceptualise and describe a concrete technical system. It introduced a concept of Decentralised Creative Networks (DCNs)—an Application Programming Interface that enables storing compositional processes as smart contract on the blockchain that human and non-human agents can reuse within the operational logics of their own compositional machines. Within this network, the atomic act is the Performative Transaction (PT): an executable compositional process that calls features, applies transformations, checks conditions, and—crucially—retains full provenance of every element it mobilises. Traceability, modularity, and interoperability are enforced at the level of the ontology of the system itself.

Methodologically, the project combined compositional experimentation, speculative design, and collaborative software engineering. A year-long collaboration with programmer Michał Skarżyński converted conceptual developments of this thesis and code prototypes into a public API, test server, and Solidity contracts, demonstrating that aesthetic claims can (and should) be validated in running code as well as in text. The resulting infrastructure is itself an artefact of research, capable of hosting future

studies that will in turn modify the system that enabled them—a live demonstration of recursive world-building.

A list of contributions could be categorised as follows:

1. Conceptual.

- A four-tier model of technological artistry, culminating in the ecosystemic worlding (Chapters 1-4).
- The definition of allagmatic composition as a mode of music-making.
- The Autopoietic Rhizomatic Metamodelling Machine (ARMM), a diagram of principles for crafting future creative systems.

2. Technical.

- The concepts of Decentralised Creative Networks with Performative Transactions that encode, execute, and document compositional processes on the blockchain.
- Three experimental prototypes of generative music NFTs: *Canons*, *Tone Row*, and *Subtraction*, as described in Chapter 3.

3. Methodological and Pedagogical

- A new educational course titled “Posthuman Creativity Labs: Artificial Intelligence and Blockchain in Music”, described in Chapter 1.
- Methods for composing with Artificial Intelligence resulting from the course, as described in Chapter 1.
- A template for practice-led infrastructural artistic research that marries artistic intuition with rigorous software development.

4. Cultural-ethical

- A proposal for transparent, shareable authorship that accommodates human and non-human contributors without collapsing into either Romantic singularity or corporate platform logic.
- New music compositions using the described system of Performative Transactions in practice, presented in a concert accompanying the defence of this doctoral trajectory.

Several constraints accompany these advances. First, the complexity of smart-contract ecosystems, despite them being designed for lowering the technical bar of entry might nevertheless still require further development of more intuitive user interfaces. Pedagogical materials and documentations are also an urgent next step. Second, blockchain's energy footprint and evolving legal status pose environmental and regulatory hurdles that must be monitored as the system scales. Third, while Performative Transactions guarantee traceability, they do not in themselves resolve questions of economic value, remuneration, or cultural legitimacy. Finally, the aesthetic assumptions baked into any initial library of features may inadvertently bias what kinds of music flourish on the network—a reminder that openness is a continuous labour, not a one-time design choice.

Future directions thus might be predicted to: 1) developing further musical works within the established framework, not only for aesthetic and artistic exploration, but also for further experimentation on posthuman creativity, human-AI collaboration, and interrogation of the possible future Intellectual Property regimes, 2) focus on the user-layer design, by developing visual editors and no-code interfaces to lower adoption barriers, 3) adapting the logic of Performative Transactions to various types of musical and non-musical artistic practices testing ARMM's discipline-agnostic promise, 4) experimentation with on-chain voting or reputation systems so that the network itself evolves under collective stewardship, not under a single maintainer (for instance as a Decentralised Autonomous Organisation), 5) exploring Layer-2 or alternative consensus protocols to mitigate ecological costs, 6) integrating Decentralised Creative Networks tooling into the *Posthuman Creativity Labs* curriculum, letting students experience compositional world-building firsthand.

By turning composition into a publicly-queryable API the project positions artistic research as infrastructural invention—a stance that resonates with recent calls to treat the art world as a medium in its own right. For musicology, the work invites fresh analytic categories focused on processual provenance rather than score or recording. For cultural policy, it offers a blueprint for decentralised authorship frameworks at a moment when AI blurs traditional Intellectual Property boundaries. And for computer science, it provides a use-case in which blockchain is neither a speculative asset nor a supply-chain ledger, but a generative artistic commons.

If, as the thesis argues, we are no longer composing only sounds or scores, but the systems through which composition itself can unfold, then every composer is,

implicitly, a legislator of future musical ecology. The artistic-technological stack proposed here is one possible constitution: provisional, hackable, oriented toward collective flourishing. Its value will be measured not by the permanence of any single artefact but by the diversity of practices it can host and the unforeseen transformations it can survive.

The subtitle of this dissertation—Worlding Compositional Ecosystems—was never a metaphor. Worlding names an active, unfinished verb; ecosystem names a milieu whose health depends on heterogeneity. The research has offered an initial seed-bed, encoded in prose, code, and working prototypes. Whether that seed germinates now depends on the communities willing to cultivate, critique, and mutate it. In the spirit of allagmatic thought, the final gesture is therefore not to seal conclusions but to invite performative transactions yet to come—transactions that will rewrite both the network's logic and the author's own thinking. May the system continue to compose itself through the many minds, human and machinic, that choose to enter its field.

Summary

This doctorate project in artistic research, undertaken by a music composer, pursued a practice-led investigation into how the roles of artists and their compositional practices transform in response to emerging technologies, with a particular focus on artificial intelligence and blockchain. Through a combination of compositional experimentation, philosophical inquiry, and the development of an experimental educational course (Posthuman Creativity Labs: Artificial Intelligence and Blockchain in Music), the research led to the definition of four redefined artistic roles: artists as operators of tools, curators of agential assemblages, system-builders, and initiators of new artistic ecosystems (in the process described as worlding artistic ecosystems). For each of these roles, new compositional methods were proposed and critically explored.

By analysing both practical and conceptual challenges posed by contemporary artistic engagements with technology, the research identified the need for a new kind of artistic ecosystem—one in which human and non-human agents can transparently build upon one another's contributions. In response to this need, this doctoral trajectory employed methods such as creative prototyping, artistic self-reflection, speculative design, and iterative technical development to envision and construct such an ecosystem. It first proposed a conceptual abstract machine—the Autopoietic Rhizomatic Metamodelling Machine (ARMM)—drawing on the notion of abstract machines. The abstract ARMM was then operationalised in the concept of Decentralised Creative Networks (DCNs): blockchain-based ecosystems for music composition in which humans and AI agents co-compose generative and self-modifying structures in a creative mode conceptualised as allagmatic composition. At the core of this system is the concept of Performative Transactions (PTs): modular, executable contributions that encode and transform compositional processes.

By both analysing existing artistic practices and composing new music through the systems it proposed, this thesis contributed a new model of posthuman compositional practice grounded in procedural transparency, recursive design, and decentralised collaboration.

Samenvatting

Dit promotieproject in artistiek onderzoek, uitgevoerd door een muziekcomponist, voerde een praktijkgestuurd onderzoek uit naar de manier waarop de rollen van kunstenaars — en hun compositorische praktijken — transformeren onder invloed van opkomende technologieën, met speciale aandacht voor kunstmatige intelligentie en blockchain. Via een combinatie van compositorische experimenten, filosofische reflecties en de ontwikkeling van een experimentele cursus (*Posthuman Creativity Labs: Artificial Intelligence and Blockchain in Music*) leidde het onderzoek tot de herdefinitie van vier artistieke rollen: kunstenaars als (1) bedieners van instrumenten, (2) curatoren van agentiale assemblages, (3) systeembouwers en (4) initiatiefnemers van nieuwe artistieke ecosystemen (in het proces omschreven als het ‘worlden’ van artistieke ecosystemen). Voor elk van deze rollen werden nieuwe compositorische methoden voorgesteld en kritisch verkend.

Door zowel de praktische als conceptuele uitdagingen te analyseren die hedendaagse artistieke interacties met technologie oproepen, identificeerde het onderzoek de noodzaak van een nieuw soort artistiek ecosysteem — een omgeving waarin menselijke en niet-menselijke agenten transparant op elkaars bijdragen kunnen voortbouwen. Om hierin te voorzien, maakte dit doctoraal traject gebruik van creatieve prototyping, artistieke zelfreflectie, speculatief ontwerp en iteratieve technische ontwikkeling om zo’n ecosysteem te verbeelden en te construeren. Eerst werd een conceptuele abstracte machine voorgesteld — de Autopoietic Rhizomatic Metamodelling Machine (ARMM) — geïnspireerd op het idee van abstracte machines. Deze abstracte ARMM werd vervolgens geoperationaliseerd in het concept Decentralised Creative Networks (DCN’s): blockchain-gebaseerde ecosystemen voor muziekcompositie waarin mensen en AI-agenten gezamenlijk generatieve, zichzelf modificerende structuren co-componeren in een creatieve modus die wordt aangeduid als *allagmatic composition*. Centraal in dit systeem staat het concept Performative Transactions (PT’s): modulaire, uitvoerbare bijdragen die compositorische processen coderen en transformeren.

Door zowel bestaande artistieke praktijken te analyseren als nieuwe muziek te componeren binnen de voorgestelde systemen, levert deze thesis een nieuw model van posthumanistische compositiepraktijk op, gefundeerd in procedurele transparantie, recursief ontwerp en gedecentraliseerde samenwerking.

Bibliography

Alizart, M. (2020). *Cryptocommunism* (R. Mackay, Trans.). Polity Press.

Anadol, R., & Kivrak, P. (2023). Machines that dream: How AI-human collaborations in art deepen audience engagement. *Management and Business Review*, 3(1–2).
<https://ssrn.com/abstract=4597317>

Arielli, E. (2024). “Even an AI could do that”. In L. Manovich & E. Arielli, *Artificial aesthetics: A critical guide to AI, media and design* (Chapter 1).
<https://manovich.net/index.php/projects/artificial-aesthetics>

Art Blocks. (n.d.). *Art Blocks*. Retrieved May 8, 2025, from
<https://www.artblocks.io/>

Ascott, R. (2002). Behaviourist art and the cybernetic vision. In R. Randall & K. Jordan (Eds.), *Multimedia: From Wagner to virtual reality* (pp. 104–120). W. W. Norton & Company. Retrieved May 8, 2025, from
<https://cs.colby.edu/courses/I16/cs267/papers/Ascott-BehavioristArt-Cybernetica60.pdf>

Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall.

de Assis, P. (2018). *Logic of Experimentation: Rethinking Music Performance through Artistic Research*. Leuven University Press.

de Assis, P. (2023). Music 2.0 and artistic research: Beyond a thousand years of Western art music. *FORUM+*, 30(1/2), 54–63.
<https://doi.org/10.5117/FORUM2023.1/2.007.ASSI>

de Assis, P. (2024). From the work concept to hypermusic: Rethinking musical objects in-and-for blockchain technologies. In P. de Assis & A. Łukawski (Eds.), *Decentralized Music: Exploring Blockchain for Artistic Research* (1st ed., pp. 12–35). CRC Press. <https://doi.org/10.1201/9781003458227>

de Assis, P., & Łukawski, A. (Eds.). (2024). *Decentralized Music: Exploring Blockchain for Artistic Research*. CRC Press.

de Assis, P., & Łukawski, A. (2024). Introduction: Blockchain for artistic research. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 1–24). CRC Press. <https://doi.org/10.1201/9781003458227>

Audry, S. (2021). *Art in the age of machine learning*. The MIT Press.

Audry, S., & Ippolito, J. (2019). Can Artificial Intelligence Make Art without Artists? Ask the Viewer. *Arts*, 8(1), 35. <https://doi.org/10.3390/arts8010035>

Barad, K. (2007). *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. Duke University Press.

- Barrett, R. (2014). *Music of possibility*. Vision Edition.
- Bengio, Y. (2021). Foreword. In S. Audry, *Art in the age of machine learning* (pp. xvii–xviii). The MIT Press.
- Berners-Lee, T. (2002). *Weaving the web: The past, present and future of the World Wide Web by its inventor*. Texere.
- Bishop, C. (2012). *Artificial Hells: Participatory Art and the Politics of Spectatorship*. London: Verso.
- Boden, M. A. (2010). *Creativity and art: Three roads to surprise*. Oxford University Press.
- Bolt, B. (2023). Performance, Performativity, and Subjectivity at the Intersection of Art and Digital Cultures. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 95–121). Chapman and Hall/CRC. <https://doi.org/10.1201/9781003312338>
- Bottou, L., & Schölkopf, B. (2023). Borges and AI. *arXiv*. <https://arxiv.org/abs/2310.01425>
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N. S., Chen, A. S., Creel, K. A., Davis, J. Q., Demszky, D., ... Liang, P. (2021). On the opportunities and risks of foundation models (arXiv:2108.07258). *arXiv*. <https://doi.org/10.48550/arXiv.2108.07258>
- Borges, J. L. (1962). The library of Babel (J. E. Irby, Trans.). In D. A. Yates & J. E. Irby (Eds.), *Labyrinths: Selected stories & other writings* (pp. 58–64). New Directions. (Original work published 1941)
- Boulez, P. (1986). Aesthetics and the fetishists (M. Cooper, Trans.). In J.-J. Nattiez (Ed.), *Orientations: Collected writings* (pp. 31–43). Harvard University Press. (Original work published 1961)
- Boulez, P. (2018a). Athematicism, identity and variation (1984–85). In J. Dunsby, J. Goldman, & A. Whittall (Eds.), *Music lessons: The Collège de France lectures* (pp. 222–280). University of Chicago Press.
- Boulez, P. (2018b). Automatism and decision (1980–81). In J. Dunsby, J. Goldman, & A. Whittall (Eds.), *Music lessons: The Collège de France lectures* (pp. 119–146). University of Chicago Press.
- Briot, J.-P., Hadjeres, G., & Pachet, F.-D. (2020). *Deep learning techniques for music generation: A survey*. Springer.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. In H. Larochelle, M.

Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in Neural Information Processing Systems* (Vol. 33, pp. 1877–1901). Curran Associates, Inc.
<https://doi.org/10.48550/arXiv.2005.14165>

Burnham, J. (1968). Systems esthetics. *Artforum*, 7(1), 30–35.

Buterin, V. (2014). *Ethereum: A next-generation smart contract and decentralized application platform*. Ethereum Foundation. <https://ethereum.org/en/whitepaper/>

Brown, K. (2023, March 10). Hito Steyerl on why NFTs and A.I. image generators are really just ‘onboarding tools’ for tech conglomerates. *Artnet News*.
<https://news.artnet.com/art-world/these-renderings-do-not-relate-to-reality-hito-steyerl-on-the-ideologies-embedded-in-a-i-image-generators-2264692>

Campbell, E. (2013). *Music after Deleuze*. Bloomsbury.

Calvo, P. (2024). Cryptoart: ethical challenges of the NFT revolution. *Humanities and Social Sciences Communications*, 11(1), 370. <https://doi.org/10.1057/s41599-024-02872-2>

Casemajor, N. (2022, December 23). NFTs in the art world: Cutting through the caricatures. *NFT Now*. <https://www.nftnow.com/features/nfts-in-the-art-world-cutting-through-the-caricatures/>

Catlow, R., & Rafferty, P. (2022). Introduction: What is radical friendship made of? In R. Catlow & P. Rafferty (Eds.), *Radical friends: Decentralised autonomous organisations and the arts* (pp. 31–35). Torque Editions.

Clancy, M. (2022a). Introduction. In M. Clancy (Ed.), *Artificial intelligence and music ecosystem* (1st ed., pp. 1–2). Focal Press. <https://doi.org/10.4324/9780429356797-1>

Clancy, M. (2022b). Amor fati: A theoretical model of the music ecosystem. In M. Clancy (Ed.), *Artificial intelligence and music ecosystem* (1st ed., pp. 165–181). Focal Press. <https://doi.org/10.4324/9780429356797-15>

Clancy, M. (2022c). The Future: Interview with Jacques Attali. In M. Clancy (Ed.), *Artificial intelligence and music ecosystem* (1st ed., pp. 3–9). Focal Press. <https://doi.org/10.4324/9780429356797-2>

Clancy, M. (2022d). The Artist: Interview with Holly Herndon. In M. Clancy (Ed.), *Artificial intelligence and music ecosystem* (1st ed., pp. 44–51). Focal Press. <https://doi.org/10.4324/9780429356797-6>

Cope, D. (2001). *Virtual music: Computer synthesis of musical style*. MIT Press.

Cuan, C. (2023). Robot choreography: Choreorobotics and humanist technology. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 340–354). Taylor & Francis. <https://doi.org/10.1201/9781003312338>

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function.

Mathematics of Control, Signals, and Systems, 2(4), 303–314.

Cycling '74. (n.d.). RNBO. Retrieved May 10, 2025, from <https://cycling74.com/products/rnbo>

0xDEAFBEEF. (n.d.). *Personal history*. Retrieved May 9, 2025, from <https://www.deafbeef.com/about.htm#history>

De Filippi, P. (2017). Plantoid: The birth of a blockchain-based lifeform. In R. Catlow, M. Garrett, N. Jones, & S. Skinner (Eds.), *Artists re:thinking the blockchain* (pp. 51–62). Torque Editions & Furtherfield.

DeLanda, M. (2006). *A new philosophy of society: Assemblage theory and social complexity*. Continuum.

Deleuze, G. (2001). Review of Gilbert Simondon's *L'individu et sa genèse physico-biologique* (1966). *Pli: The Warwick Journal of Philosophy*, 12, 43–49.

Deleuze, G., & Guattari, F. (2003). *A thousand plateaus: Capitalism and schizophrenia* (B. Massumi, Trans.). Continuum. (Original work published 1980)

Drubay, D. (2024). Art, People, Museums, and the Promise of Blockchain. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 95–104). CRC Press. <https://doi.org/10.1201/9781003458227>

ECAD Labs. (n.d.). *What is a blockchain?* Taquito Documentation. Retrieved May 8, 2025, from <https://taquito.io/docs/tutorial/#what-is-a-blockchain>

Eco, U. (1989). *The Open Work* (A. Cancogni, Trans.). Cambridge, MA: Harvard University Press. (Original work published 1962)

Einarsson, E. T. (2024). Hypermusic Experiment 0.9: Modeling, Mapping, and Prototyping the Future. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 113–129). CRC Press. <https://doi.org/10.1201/9781003458227>

Eleti, A., Harris, J., & Kilpatrick, L. (2023, June 13). *Function Calling and other API Updates*. OpenAI. <https://openai.com/blog/function-calling-and-other-api-updates>

Eno, B. (1996). *A year with swollen appendices: Brian Eno's diary*. Faber and Faber.

Erickson, S., & Kelty, C. M. (2015). The durability of software. In I. Kaldrack & M. Leeker (Eds.), *There is no software, there are just services* (pp. 39–55). Meson Press. <https://doi.org/10.25969/mediarep/1022>

Ethereum. (2023). *Smart contract composability*. Retrieved November 10, 2023, from <https://ethereum.org/en/developers/docs/smart-contracts/composability/>

von Foerster, H. (2003). Perception of the future and the future of perception. In H. von Foerster, *Understanding understanding: Essays on cybernetics and cognition* (pp.

199–210). Springer.

fxhash. (n.d.). *fxhash*. Retrieved May 8, 2025, from <https://www.fxhash.xyz/>

Galanter, P. (2003). What is generative art? Complexity theory as a context for art theory. In *Proceedings of the 6th Generative Art Conference (GA'03)*. Retrieved from https://www.philipgalanter.com/downloads/ga2003_paper.pdf

Gambetta, D. (1988). Can we trust trust? In D. Gambetta (Ed.), *Trust: Making and breaking cooperative relations* (pp. 213–237). Basil Blackwell.

Gell, A. (1998). *Art and agency: An anthropological theory*. Oxford University Press.

Gershman, S. J. (2021). *What makes us smart: The computational logic of human cognition*. Princeton University Press.

Giannoutakis, K. (2024). Live coding the collective: Reimagining experimental music practices through technics of decentralization (Doctoral dissertation). Rensselaer Polytechnic Institute.

Gibson, J. (2021). The thousand-and-second tale of NFTs, as foretold by Edgar Allan Poe. *Queen Mary Journal of Intellectual Property*, 11(3), 249–269. <https://doi.org/10.4337/qmjip.2021.03.00>

Gioti, A.-M. (2021). *Agency and distributed creativity in interactive compositions* (Doctoral dissertation). University of Music and Performing Arts Graz.

Gloerich, I. (2025). *Reimagining the Truth Machine: Blockchain Imaginaries between the Rational and the More-than-Rational*. [Doctoral thesis 1 (Research UU / Graduation UU), Universiteit Utrecht]. Utrecht University. <https://doi.org/10.33540/2726>

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. <https://www.deeplearningbook.org/>

Goriunova, O. (2012). *Art platforms and cultural production on the internet*. Routledge.

Grassi, J. (2022). *teamLab Borderless: The role of play in immersive, interactive installations* (Master's thesis, Concordia University). Concordia University Research Repository. https://spectrum.library.concordia.ca/id/eprint/990914/1/Grassi_MA_F2022.pdf

Guattari, F. (1995). *Chaosmosis: An ethico-aesthetic paradigm* (P. Bains & J. Pefanis, Trans.). Indiana University Press. (Original work published 1992)

Haraway, D. J. (2016). *Staying with the trouble: Making kin in the Chthulucene*. Duke University Press.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. Wiley.

Heile, B. (2006). *The music of Mauricio Kagel*. Ashgate.

- Hertzmann, A. (2018). *Can Computers Create Art?* Arts, 7(2), 18. DOI: 10.3390/arts7020018
- Higgins, H. (2002). *Fluxus experience*. University of California Press.
- Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554. <https://doi.org/10.1162/neco.2006.18.7.1527>
- Hobbs, T. (2021, August 6). *The Rise of Long-Form Generative Art*. Retrieved from <https://www.tylerxhobbs.com/words/the-rise-of-long-form-generative-art>
- Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2), 251–257.
- Hui, Y. (2015). Induction, deduction and transduction: On the aesthetics and logic of digital objects. *Networking Knowledge: Journal of the MeCCSA Postgraduate Network*, 8(3). <https://doi.org/10.31165/nk.2015.83.376>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kaldrack, I., & Leeker, M. (2015). There is no software, there are just services: Introduction. In I. Kaldrack & M. Leeker (Eds.), *There is no software, there are just services* (pp. 9–19). Meson Press. <https://doi.org/10.14619/008>
- Karpathy, A. (2023, November 23). [1 hr Talk] Intro to large language models [Video]. YouTube. https://youtu.be/zjkBMFhNj_g
- Kenderdine, S., Rohrmeier, M., Doornbusch, P., & Impett, J. (2024). *Musica ex Machina: Machines Thinking Musically – Exhibition Guide*. EPFL Pavilions. [https://www.researchgate.net/publication/385377018 Musica ex Machina Exhibition-Guide](https://www.researchgate.net/publication/385377018_Musica_ex_Machina_Exhibition-Guide)
- Kostelanetz, R. (2003). *Conversing with Cage* (2nd ed., rev.). Routledge.
- Lettvin, J. Y., Maturana, H. R., McCulloch, W. S., & Pitts, W. H. (1959). What the frog's eye tells the frog's brain. *Proceedings of the IRE*, 47(11), 1940–1951. <https://doi.org/10.1109/JRPROC.1959.287207>
- Lewis, G. E. (2000). Too many notes: Computers, complexity and culture in Voyager. *Leonardo Music Journal*, 10, 33–39. <https://doi.org/10.1162/096112100570585>
- Lind, M. (2012). “Performing the Curatorial: An Introduction.” In M. Lind (Ed.), *Performing the Curatorial: Within and Beyond Art* (pp. 11–20). New York: Sternberg Press.
- Luhmann, N. (1979). *Trust and power: Two works* (H. Davis, J. Raffan, & K. Rooney, Trans.). John Wiley & Sons. (Original works published 1973 & 1975)
- Luque, S. (2009). The stochastic synthesis of Iannis Xenakis. *Leonardo Music Journal*,

19, 77–84. <https://doi.org/10.1162/lmj.2009.19.77>

Łukawski, A. (2023). Fractal of periodic spacetime sequences as a framework for an interoperable metaverse. In K. Zreik, M. Veyrat, & M. Quiniou (Eds.), *Computer & media art at the age of metaverses and NFT* (pp. 57–74). Europa Productions.

Łukawski, A. (2024a). Performative Transactions: Artistic Collaboration of Humans and AI Agents in Decentralized Creative Networks. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 36–56). CRC Press. <https://doi.org/10.1201/9781003458227>

Łukawski, A. (2024b, April 5). *Canons* [Generative music NFT]. fxhash. <https://www.fxhash.xyz/project/canons-1>

Łukawski, A. (2024c, May 14). *Tone Row* [Generative audiovisual NFT]. fxhash. <https://www.fxhash.xyz/generative/slug/tone-row>

Łukawski, A. (2024d, July 12). *Subtraction* [Generative music NFT and collaborative composition platform]. fxhash. <https://www.fxhash.xyz/generative/slug/subtraction-1>

Łukawski, A. (2024e). *Subtraction: A posthumanist musical assemblage* [Web application]. Retrieved May 8, 2025, from <https://www.adamlukawski.com/pages/apps/subtraction/>

Łukawski, A. (2024f, May 27). *Subtraction: A posthumanist musical assemblage on fxhash* [Article]. fxhash. [https://www.fxhash.xyz/article/subtraction-a-posthumanist-musical-assemblage-on-fx\(hash\)](https://www.fxhash.xyz/article/subtraction-a-posthumanist-musical-assemblage-on-fx(hash))

Macchia, S. (2020, February 18). *Making an album with Music Transformer*. Magenta. <https://magenta.tensorflow.org/nobodys-songs>

Magee, L., & Rossiter, N. (2015). Service orientations: Data, institutions, labor. In I. Kaldrack & M. Leeker (Eds.), *There is no software, there are just services* (pp. 73–89). Meson Press. <https://doi.org/10.25969/mediarep/1047>

Magnusson, T. (2019). *Sonic Writing: Technologies of Material, Symbolic and Signal Inscriptions*. Bloomsbury Academic.

Mattin (2023). Noise and Subjectivity in the Era of Machine Learning. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 419–432). Chapman and Hall/CRC. <https://doi.org/10.1201/9781003312338>

McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Magazine*, 27(4), 12. <https://doi.org/10.1609/aimag.v27i4.1904>

McCormack, J. (2023). Autonomy, Intention, Performativity: Navigating the AI Divide. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 240–257). Chapman and Hall/CRC. <https://doi.org/10.1201/9781003312338>

McKeon, E. (2022). *Heiner Goebbels and curatorial composing after Cage: From staging works to musicalising encounters*. Cambridge University Press.

Minsky, M., & Papert, S. (1969). *Perceptrons*. MIT Press.

Mollick, E. (2024). *Co-Intelligence: Living and working with AI*. Portfolio/Penguin.

Morton, T. (2013). *Hyperobjects: Philosophy and ecology after the end of the world*. University of Minnesota Press.

Mulligan, C. (2024). Can't Knock the Hustle: NFTs, DAOs, and Creativity. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 105–112). CRC Press. <https://doi.org/10.1201/9781003458227>

Munster, A., & Rossiter, N. (2023). Performing the automated image. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 47–73). Chapman and Hall/CRC. <https://doi.org/10.1201/9781003312338>

Nail, T. (2017). What is an assemblage? *SubStance*, 46(1), 21–37. <https://doi.org/10.3368/ss.46.1.21>

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*. <https://bitcoin.org/bitcoin.pdf>

Neubert, C. (2015). “The tail on the hardware dog”: Historical articulations of computing machinery, software, and services. In I. Kaldrack & M. Leeker (Eds.), *There is no software, there are just services* (pp. 21–38). Meson Press. <https://doi.org/10.14619/008>

Neufeld Redekop, V., and Ryba, T. (eds.). (2013). *René Girard and Creative Mimesis*. New York: Lexington Books.

Newell, A., & Simon, H. A. (1956). *The logic theory machine—A complex information processing system* (RAND Corporation Report P-868). RAND Corporation.

Notaro, A. (2022). All that is solid melts in the Ethereum: The brave new (art) world of NFTs. *Journal of Visual Art Practice*, 21(4), 359–382. <https://doi.org/10.1080/14702029.2022.2129204>

O'Dair, M. (2024). Valuing Web3 Music: From NFT Prices to the Quadruple Bottom Line. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 72–85). CRC Press. <https://doi.org/10.1201/9781003458227>

O'Neill, P. (2012). *The Culture of Curating and the Curating of Culture(s)*. Cambridge, MA: MIT Press.

Ouzounian, G. (2011). The uncertainty of experience: On George Brecht's Event scores. *Journal of Visual Culture*, 10(2), 198–211. <https://doi.org/10.1177/1470412911402894>

Parisi, L. (2023). Negative Aesthetics: AI and Non-Performance. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 74-94). Chapman and Hall/CRC. <https://doi.org/10.1201/9781003312338>

Piquemal, S. (n.d.). *WebPd* [Computer software]. GitHub. Retrieved May 10, 2025, from <https://github.com/sebpiq/WebPd>

Plumb, T. (2024, May 22). AI pioneer LeCun to next-gen AI builders: “Don’t focus on LLMs”. *VentureBeat*. <https://venturebeat.com/ai/ai-pioneer-lecun-to-next-gen-ai-builders-dont-focus-on-llms/>

Poliks, M., & Alonso Trillo, R. (2023). 0-Degree Plane of Neuroelectronic Continuity: AI & Psychosocial Evaporation. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 5-46). Chapman and Hall/CRC. <https://doi.org/10.1201/9781003312338>

poperbu. (2023, May 29). *bang.pd* [Generative music NFT]. fxhash. <https://www.fxhash.xyz/project/bang.pd-1>

Poveda Yáñez, J., & Davies, N. (2021). Dancing someone else’s movements through someone else’s body : the process of commodification of the digital dancing body and the arising tensions with intellectual property regimes. *DANCE ARTICULATED*, 7(1), 5–22. <https://doi.org/10.5324/da.v7i1.4224>

Poveda Yáñez, J. (Convener). (2022, August, 24-26). *Dance and the Blockchain: Commodification and Ownership of Embodied Creativity in the Crypto Space* [Symposium]. Ghent University, Belgium.

Puckette, M. (2022). What do music software developers do? In M. Clancy (Ed.), *Artificial intelligence and music ecosystem* (1st ed., pp. 24–34). Focal Press. <https://doi.org/10.4324/9780429356797-4>

Quaranta, D. (2021). *Surfing con Satoshi: Arte, blockchain e NFT*. Postmedia Books.

Raschka, S. (2025, February 5). *Understanding reasoning LLMs: Methods and strategies for building and refining reasoning models*. Sebastian Raschka. <https://sebastianraschka.com/blog/2025/understanding-reasoning-llms.html>

Rickert, T., & Salvo, M. (2006). The distributed Gesamtkunstwerk: Sound, worlding, and new media culture. *Computers and Composition*, 23(3), 298–316. <https://doi.org/10.1016/j.compcom.2006.05.001>

RLXYZ. (2022, March 12). ‘On- Chain Purity’ in Generative Art. Retrieved November 10, 2023, from https://mirror.xyz/0xB52D87A5097CAc5248599DB2272b0288212eb82b/Cyij9sbeYuMmPNMJYVNhcm9Uv3N5FA_VnHTkvqhqa83A

Rohrmeier, M. A. (2022). On creativity, music’s AI completeness, and four challenges for artificial musical creativity. *Transactions of the International Society for Music Information Retrieval*, 5(1), 50–66. <https://doi.org/10.5334/tismir.104>

- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386–408.
<https://doi.org/10.1037/h0042519>
- Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning representations by back-propagating errors. *Nature*, 323, 533–536.
- Schillinger, J. (1948). *The mathematical basis of the arts*. Philosophical Library.
- Seidler, P., Kolling, P., & Hampshire, M. (2017). terra0: Can an augmented forest own and utilize itself? In R. Catlow, M. Garrett, N. Jones, & S. Skinner (Eds.), *Artists re:thinking the blockchain* (pp. 63–72). Torque Editions & Furtherfield.
- Selfridge, O. (1959). Pandemonium: A paradigm for learning. In *Proceedings of the Symposium on Mechanisation of Thought Processes* (pp. 511–529). HMSO.
- Serpentine Arts Technologies. (2020). *Future Art Ecosystems 1: Art x Advanced Technologies*. Serpentine Galleries. <https://futureartecosystems.org/briefing/fae1/>
- Serpentine Arts Technologies. (2022). *Future Art Ecosystems 3: Art x Decentralised Tech*. Serpentine Galleries. <https://futureartecosystems.org/briefing/fae3/>
- Serpentine Arts Technologies. (2024). *Future Art Ecosystems 4: Art x Public AI*. Serpentine Galleries. <https://futureartecosystems.org/briefing/fae4/>
- Sigalov, A. (2020). *Allegro Music Transformer* [Computer software]. GitHub.
<https://github.com/asigalov61/Allegro-Music-Transformer>
- Simon, I., & Oore, S. (2017). Performance RNN: Generating music with expressive timing and dynamics. *Magenta Blog*. <https://magenta.tensorflow.org/performance-rnn>
- Simondon, G. (2017). *On the mode of existence of technical objects* (C. Malaspina & J. Rogove, Trans.). Univocal Publishing.
- Simondon, G. (2020). *Individuation in light of notions of form and information* (T. Adkins, Trans.). University of Minnesota Press. (Original work published 1958)
- Sloman, A. (1984). The Structure of the Space of Possible Minds. In S. Torrance (Ed.), *The Mind and the Machine: Philosophical Aspects of Artificial Intelligence* (pp. 35–42). Ellis Horwood.
- Smite, R. (2012). *Creative Networks in the Rearview Mirror of Eastern European History*. Institute of Network Cultures (p. 141), <https://doi.org/10.25969/mediarep/19229>
- Smith, O. F. (1998). *Fluxus: The history of an attitude*. San Diego State University Press.
- Stiles, S. (2023). Ars Autopoetica: On Authorial Intelligence, Generative Literature, and the Future of Language. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 357–378). Chapman and Hall/CRC.

<https://doi.org/10.1201/9781003312338>

Stoykov, P. (2023, August 7). *Bridging Algorithms and Art: A Look into Generative NFTs*. Chainstack. Retrieved from <https://chainstack.com/generative-nfts/>

Summers, A. (2022). *Understanding blockchain and cryptocurrencies: A primer for implementing and developing blockchain projects* (1st ed.). CRC Press.

Tatarkiewicz, W. (2012/1980). *A history of six ideas: An essay in aesthetics*. Springer Science and Business Media. (Original work published 1980)

Tegmark, M. (2003). Parallel universes. In J. D. Barrow, P. C. W. Davies, & C. L. Harper Jr. (Eds.), *Science and ultimate reality: Quantum theory, cosmology and complexity* (pp. 459–491). Cambridge University Press.

Terranova, T. (2023). Foreword. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. viii–x). Chapman and Hall/CRC.

Tessone, C. J. (2024). From Blockchains to NFTs: Decentralized (?) Platforms for Unique (?) Content Distribution. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 87–94). CRC Press.
<https://doi.org/10.1201/9781003458227>

Tilford, K. (2023). Performing Creativity: Text-to-Image Synthesis and the Mimicry of Artistic Subjectivity. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 125–187). Chapman and Hall/CRC.
<https://doi.org/10.1201/9781003312338>

Trillo, R. A. (2023). Descendent: AI and the body beyond hybridization. In R. A. Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 435–466). Taylor & Francis. <https://doi.org/10.1201/9781003312338-27>

Ulus, H. (2021). *The poly-work as a compositional strategy in my music: Works 2017–21* [Doctoral dissertation, University of Huddersfield]. University of Huddersfield Repository. https://pure.hud.ac.uk/files/67224224/FINAL_THESIS.pdf

Valverde, L., Agres, K., Herremans, D., & Chew, E. (2023). Human-AI Musicking: A Framework for Designing AI for Music Co-Creation. In *Proceedings of the 2023 International Conference on AI Music Creativity (AIMC)*.
<https://aimc2023.pubpub.org/pub/zd46ltn3>

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc.
<https://doi.org/10.48550/arXiv.1706.03762>

Vincze, D. (2023). The Musicality of Imperfection. In R. Alonso Trillo & M. Poliks (Eds.), *Choreomata: Performance and Performativity after AI* (pp. 308–339). Chapman and Hall/CRC. <https://doi.org/10.1201/9781003312338>

Walshe J. (2023). Ghosts of the Hidden Layer. In R. A. Trillo & M. Poliks (Eds.),

Choreomata: Performance and Performativity after AI (pp. 482–500). Taylor & Francis.
<https://doi.org/10.1201/9781003312338>

Wang, Y., Wu, S., Hu, J., Du, X., Peng, Y., Huang, Y., Fan, S., Li, X., Yu, F., & Sun, M. (2025). NotaGen: Advancing musicality in symbolic music generation with large language model training paradigms. *arXiv*. <https://arxiv.org/abs/2502.18008>

Watson, J. (2009). *Guattari's diagrammatic thought: Writing between Lacan and Deleuze*. London: Continuum.

West, E. (2022, October 31). *Generative music: A brief history and future*. fxhash.
<https://www.fxhash.xyz/article/generative-music-a-brief-history-and-future>

Whitaker, A. (2019). Art and Blockchain: A Primer, History, and Taxonomy of Blockchain Use Cases in the Arts. *Artivate: A Journal of Entrepreneurship in the Arts*, 8(2), 21–46.

Wiener, N. (1948). *Cybernetics: Or control and communication in the animal and the machine* (2nd ed.). MIT Press.

Wilhelm, R. (Trans.), Baynes, C. F. (Trans.), & Jung, C. G. (Foreword). (1967). *The I Ching, or, Book of Changes* (3rd ed., Bollingen Series XIX). Princeton University Press. (Original work published 1950)

Wishart, T. (1996). *On sonic art*. Amsterdam: Harwood Academic Publishers.

Xenakis, I. (2001). *Formalized music: Thought and mathematics in composition* (Rev. ed.; S. Kanach, Ed.). Pendragon Press.

Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., & Narasimhan, K. (2023). *Tree of Thoughts: Deliberate Problem Solving with Large Language Models*. Retrieved from <https://arxiv.org/pdf/2305.10601.pdf>

Young, M. (2008). NN Music: Improvising with a 'Living' Computer. In R. Kronland-Martinet, P. R. Herrera, & S. Ystad (Eds.), *Computer Music Modeling and Retrieval. Sense of Sounds* (pp. 337–350). Springer. https://doi.org/10.1007/978-3-540-85035-9_23

Zeilinger, M. (2021). *Tactical entanglements: AI art, creative agency, and the limits of intellectual property*. Meson Press. <https://doi.org/10.14619/1839>

Zeilinger, M. (2022, November 7). Blockchain Vitalism. *Outland*.
<https://outland.art/blockchain-vitalism/>

Zeilinger, M. (2024). Integrating Generative AI and Blockchain Technologies to Create Musical Objects with Agency. In P. de Assis & A. Łukawski (Eds.), *Decentralized music: Exploring blockchain for artistic research* (1st ed., pp. 57–71). CRC Press. <https://doi.org/10.1201/9781003458227>

Curriculum Vitae

Adam Łukawski (born 05.02.1997 in Suwałki, Poland) is a composer and computer programmer working at the intersection of computer-assisted music composition and posthuman artistic research. His work, rooted in aleatoric and generative methods, explores the integration of AI and blockchain technologies to develop novel compositional frameworks and enhance musical interactivity.

He received his Bachelor of Music in Composition from the Conservatorium van Amsterdam in 2019, studying with Richard Ayres and Willem Jeths, and a Master of Music in Composition (cum laude) from the same institution in 2021. During his master's, he completed an Erasmus exchange at the Guildhall School of Music & Drama in London, where he studied with Julian Anderson.

Since 2021, he has been pursuing a doctoral degree in the docARTES PhD programme at Leiden University's Academy of Creative and Performing Arts and the Orpheus Institute in Ghent. His research is supervised by Richard Barrett, Paulo de Assis, and Michiel Schuijjer. Since 2022, he has been a doctoral fellow in the MetamusicX research cluster at the Orpheus Institute (PI: Paulo de Assis).

He is the co-editor (with Paulo de Assis) of the book *Decentralised Music: Exploring Blockchain for Artistic Research* (CRC Press/Taylor & Francis). He also developed and taught the course *Posthuman Creativity Labs: Artificial Intelligence and Blockchain in Music* at institutions including the Conservatorium van Amsterdam, HEAD Genève, the Iceland University of the Arts, and others.

As a composer, he has written works for ensembles such as Exaudi Vocal Ensemble, Maat Saxophone Quartet, and Score Collective, and frequently collaborates with visual artists, dancers, choreographers, and theatre directors. His music has been published by Donemus Publishing (The Hague) and performed at venues including the Dutch National Opera & Ballet (Opera Forward Festival), Muziekgebouw aan 't IJ, Splendor Amsterdam (Gesti Festival), Muziekhuis Utrecht (Gaudeamus Festival), the Concert Hall of Polish Radio Philharmonic Orchestra in Warsaw (In Modo di Lutosławski), Barbican Milton Court (London), The Place (London Contemporary Dance School), and Ehrbar Hall (Vienna).

As a programmer, he has worked for several start-ups and for the Polish Information Processing Society in Warsaw.