Organised Sound (2026), 1–8
doi:10.1017/S1355771826101538

Check for updates

CAMBRIDGE UNIVERSITY PRESS

Article

Beyond automation: Artificial intelligence and embodied craft in artistic practice

João Pedro Oliveira
Music Department, UCSB, USA

Abstract

This article examines artificial intelligence (AI) in artistic practice not as an autonomous creative subject or a mere labour-saving device, but as a field of models, affordances and constraints through which artists may think, test and transform their materials. Against commercial narratives that identify intelligence with efficiency, productivity and the automation of economically valuable work, it argues for the continuing centrality of embodied craft intelligence: the tacit, historically situated, perceptual and ethical knowledge developed through practice. Drawing on Leo Apostel's theory of models, James Gibson's concept of affordances, Michael Polanyi's tacit knowledge, Walter Benjamin's aura, Theodor W. Adorno's critique of fetishism and Isaac Asimov's story Profession, the article proposes a critical, feedback-based understanding of human–AI interaction. It distinguishes serendipitous exploration from template-driven production and suggests that the artist's responsibility lies in judging when machine outputs become meaningful deviations rather than seductive shortcuts. The future of AI in the arts depends less on the technology's generative power than on whether artists can preserve slow perception, historical memory, critique and responsibility within increasingly automated cultural infrastructures. Its central claim is that AI becomes artistically valuable only when kept visible as a contestable tool rather than accepted as a substitute for artistic judgement.

Keywords: Affordances; artificial intelligence; aura; embodied craft; fetishism

(Received 14 May 2026 UTC; revised 16 July 2026 UTC; accepted 23 July 2026 UTC)

1. From embodied craft intelligence to artificial intelligence

Music composition is the product of an embodied, craft-based and historically situated way of knowing. It is built through years of listening, working with instruments, learning repertoires, imagining forms, testing technologies and collaborating with performers. This knowledge is not a fixed stock of rules. It changes through practice, failure, repetition and revision. Much of it is tacit: a form of 'knowing how' rather than 'knowing that', in Michael Polanyi's sense (Polanyi 1966). It is sedimented in the body, in the ear, in the hand and in habits of observation. For an artist, intelligent behaviour is therefore not only the solution of a clearly specified problem. It is also the capacity to navigate open-ended aesthetic situations: to decide which materials deserve attention, to sense when a form requires balance or disruption, to imagine the consequences of a decision before it is written down, and to understand the capacities of performers, audiences and technologies within specific contexts.

What counts as intelligent behaviour has always shifted alongside technological change. The so-called 'AI effect' describes a recurring pattern: while a task remains beyond machines, it is often treated as evidence of genuine intelligence; once machines can perform it, the achievement is reclassified as routine or merely mechanical (McCorduck 2004). In this way, we tend to mythologise whatever machines cannot yet do as 'true' intelligence, while quietly downgrading earlier milestones once they become technically achievable or culturally familiar. Turing's proposal to test machine intelligence through conversational behaviour defined intelligence in terms of observable performance rather than inner essence (Turing 1950). If a machine is able to sustain a conversation in a way that is indistinguishable from a human, then the question of intelligence should be approached through observable performance rather than through speculation about whether the machine possesses a human-like mind, consciousness or interior mental life. Early artificial intelligence research, especially the symbolic programmes of the 1950s to 1970s, conceived intelligence largely as the rule-based manipulation of symbols (Newell and Simon 1972). Haugeland later described this tradition as 'Good Old-Fashioned Artificial Intelligence' (Haugeland 1985). In parallel, cybernetics and systems theory modelled intelligent behaviour through feedback, regulation and adaptation (Wiener 1948; Ashby 1956).

Composers quickly absorbed these ideas. Hiller and Isaacson's Experimental Music: Composition with an Electronic Computer explored rule-based procedures and probabilistic selection (Hiller and Isaacson 1959). Xenakis used stochastic and game-theoretical models to formalise musical decisions (Xenakis 1971). Later generative systems, including grammars, constraints and evolutionary algorithms, translated compositional operations into explicit procedures within a framework whose rules and limits had been established, or at least validated, by a human designer.

Email: jppo@ua.pt

Cite this article: Oliveira JP. (2026). Beyond automation: Artificial intelligence and embodied craft in artistic practice. *Organised Sound*, 1–8, https://doi.org/10.1017/S1355771826101538

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press

Expert systems extended this rule-based paradigm in the 1980s by attempting to reproduce the judgement of specialists through explicitly formulated rules. Their architectures encoded decision trees and related logical structures, formalising judgement as a sequence of discrete, verbalisable procedures (Winston and Prendergast 1984). Within this framework, intelligence was understood as the correct application of rules within a predefined domain. Between the expert-system paradigm of the 1980s and the large neural models of the present, creative computing developed through a heterogeneous field of probabilistic, statistical and adaptive methods. In music, this intermediate period included Markov chains and other stochastic procedures, grammars, constraint-based systems, evolutionary algorithms, case-based reasoning and hybrid approaches that combined symbolic and statistical techniques. David Cope's Experiments in Musical Intelligence, for example, used pattern recognition, recombination and style modelling to generate music that raised questions about authorship, stylistic imitation and the boundary between analysis and composition (Cope 1996). Markov models and probabilistic grammars offered ways of producing musical continuity without requiring every event to be specified in advance. Evolutionary composition systems introduced selection and mutation criteria, allowing musical material to develop through iterative variation and evaluation. Constraint-based and algorithmic approaches similarly treated composition as a field of possibilities shaped by formal limits.

Deep learning emerged from this long trajectory in which composers and researchers had already been testing different ways of modelling musical behaviour, style, memory, variation and decision-making. Contemporary AI systems are built from neural networks trained to perform tasks once associated with human perception, judgement or creativity. They do not usually depend on explicitly programmed instructions for each operation, instead, they extract statistical regularities from data and organise them within what is often called a latent space: a high-dimensional field of learned relationships among timbres, harmonies, gestures, images, words or phrases. To move through such a space is not simply to retrieve examples, but to interpolate, transform and combine patterns according to learned relations. Representations are learned rather than explicitly designed by humans or machines. Instead of a composer-programmer defining features in advance, such as interval classes, rhythmic cells, formal functions or timbral categories, the system develops internal representations that are optimised for a predictive or generative task.

Behaviour therefore emerges through data-driven optimisation. A defined objective function establishes what counts as success, and the network gradually adjusts its parameters across large datasets in order to approach that goal. Unlike rule-based systems, whose operations can often be inspected step by step, deep neural networks usually function as black boxes, making it difficult to determine why a particular result has been produced. Data become the fundamental constructive material and the choice of training corpus – which compositions, recordings, notations and aesthetics are included or excluded – strongly shapes the tendencies, limits and expressive profile of the model being used.

In music, this usually appears in platforms that can fill in harmonies, imitate a genre, continue a fragment in the style of a composer or generate entire songs from a prompt. Audio models can render full productions with melody, harmony, timbre and space. Deep learning thus shifts the meaning of artificial intelligence from explicit symbolic design to pattern recognition and high-dimensional interpolation across large quantities of information.

For a human composer, however, intelligent behaviour includes dimensions that exceed statistical patterning. It includes the capacity to frame problems rather than merely solve them: to decide that a work should question the hierarchy between noise and pitch or subvert a particular formal structure. It includes long-range aesthetic memory and anticipation: the feeling that a texture introduced now will resonate with something heard minutes earlier, or that continuing a gesture will create a more intense formal tension. It includes ethical and contextual judgement: deciding how to use a historical style, whether appropriation is at stake, and how a composition positions itself in relation to history, technology, performance and audience.

Despite the technical shift from rules to deep networks, both expert systems and deep learning tend to treat creativity as something inferred from existing patterns and recombined. In that sense, they remain in tension with embodied craft intelligence, which is not only about recombination but also about deciding which challenges matter. The composer's intelligence is exercised in deciding when to accept, transform or reject a technological output, and in inventing forms of collaboration in which human and machine are treated as fallible, exploratory partners rather than sources of finished solutions.

A useful comparison can be made with other moments in the history of music technology. The arrival of recording did not simply preserve performance; it changed ideas of precision, repeatability and presence. The synthesiser did not merely imitate instruments; it created new timbral categories and new expectations about control. As mentioned above, AI belongs to this longer history, but with one crucial difference: it tends to present itself not simply as an instrument, medium or archive, but as a simulation of judgement itself. This is why the distinction between generated material and aesthetic decision is so important. The system may provide material, but the artistic decision concerns why that material belongs, what it does to the work and what form of responsibility it entails.

The position developed in this article is therefore neither a rejection of AI nor an endorsement of its uncritical adoption. Rather, it argues that AI becomes artistically valuable when approached with sustained critical skepticism: as a tool whose assumptions, outputs and constraints remain open to interrogation rather than accepted as authoritative. Such skepticism is not external to artistic practice, but part of the composer's craft, shaping how feedback, co-creation, judgement and responsibility interact throughout the creative process.

This distinction also affects pedagogy. If students encounter AI mainly as a tool for producing finished-looking results, they may learn to value polished results before understanding structure, technique or listening. The danger is not that they will use technology, since artistic education has always involved tools, but that they may be encouraged or seduced to bypass the slow and continuing formation of criteria.

This risk becomes visible in a small experiment with image generation: the system can produce a polished and culturally recognisable result almost instantly, before the user has had to analyse its structure, question its derivation or develop criteria for judging its artistic necessity. Starting with Leonardo da Vinci's Mona Lisa, I asked an AI system to generate a prompt describing its

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press

visual style. I then used that prompt to create a related image on an AI image-generation platform[1]. The result appeared new, but it was also connected to pre-existing visual forms and to patterns learned from data. While the original *Mona Lisa* is marked by ambiguity, restraint and historical density: its surface, *sfumato* transitions and enigmatic expression that create a sense of distance and aura, the AI-generated image is more immediately legible and polished, with sharper beauty, stronger contrast and a cinematic atmosphere. Yet this polish also makes it less ambiguous: the face, lighting and pose are closer to contemporary fantasy portrait conventions (Figure 1).

The experiment suggests that AI often produces not originality in a strong sense, but successive acts of translation, abstraction and recombination. Its value depends on what the artist does with that condition. And its pedagogical relevance is that surface coherence can appear before the student has learned to judge structure, necessity or historical relation.

2. Exploratory interaction, models and affordances

A fruitful human–AI collaboration should be understood as exploratory and experimental. The artist probes the system's potentials and limitations; the system's behaviour feeds back into the artist's evolving knowledge and aesthetic criteria. Creating with AI is then less like issuing commands to a machine and more like entering a gradually refined dialogue in which both the tool and the artistic project are re-modelled over time. I use co-creation here particularly for such iterative exchanges; a single or even a multiple prompt-response remains a responsive tool use, not a shared artistic agency.

Jonathan Harvey's description of composing with a computer remains instructive. In 'The Mirror of Ambiguity' he writes that, at the terminal, a sound is imagined, tried, adjusted and saved, in a procedure structurally close to traditional composition at the piano:

> There is a feeling that we are contemplating a defiant, yet faithful mirror, which constantly poses relevant questions. We modify a parameter and react emotionally to the result. We modify another parameter and this change produces another emotion, subtly different, perhaps even disconcerting. We repeat hundreds of times this back-and-forth path between the objective and the subjective, until a kind of mutual adaptation is reached (Harvey 1986: 175–6).

Harvey anticipates a central intuition about human–AI interaction: the mirror is not passive but interrogative. It reflects our actions back in transformed form, and we refine both our ideas and our sensibility in response.

Leo Apostel's theory of models is helpful here. A model is a selective representation used for purposes such as explanation, prediction, design or decision (Apostel 1961). It involves a source domain, a target domain and a mapping between them. Every model highlights some aspects and hides others. When artists work with AI, they adopt and revise models of what the system is and how it responds: is it a generative 'co-composer', a sophisticated search engine over musical possibilities, a probabilistic style-copying device, an extended instrument or something else? Each metaphor corresponds to a different model and suggests different strategies of interaction and evaluation.

Moreover, the AI system is also structured by models. A deep-learning architecture models patterns and regularities in its training data. Human and machine therefore meet as layered systems of modelling: the composer's aesthetic, theoretical and embodied models encounter the statistical models embedded in the AI system. The composer revises prompts, parameters, datasets or compositional strategies in response to the system's outputs; the system may be fine-tuned, retrained or redirected through human feedback. Over time, the exchange can produce a form of mutual adaptation that cannot be reduced either to human intention alone or to machine computation alone. This mutual adaptation should not be understood as assigning autonomous artistic agency to the AI system. The system does not form intentions, assume responsibility or evaluate the aesthetic consequences of its outputs. Its agency is operational rather than intentional: it responds, constrains, transforms and proposes according to the statistical structures, parameters and interface conditions through which it has been designed. The difference between human-led iterative adjustment and co-creative interaction therefore lies in the degree to which the system's behaviour becomes materially consequential for the development of the work. In a purely responsive tool, the artist already knows the desired result and uses the system to execute or refine it. In a co-creative process, by contrast, the system's misalignments, resistances and unexpected continuations become part of the compositional inquiry itself. Mutual adaptation, then, is an asymmetrical feedback process in which machine behaviour affects the artistic path while human judgment remains responsible for meaning, selection and form.

James Gibson's concept of affordances complements this perspective. An affordance is not simply a property of an object in isolation; it emerges from the relation between an environment and a perceiving, acting subject with particular capacities, needs and concerns (Gibson 1979). A chair affords sitting to some bodies; a steep staircase affords climbing only to those able to climb and a cliff edge affords caution. Affordances are therefore relational, situated and partly prescriptive. They depend on what the subject can do, what the environment makes possible or difficult, and what the subject recognises as relevant.

A music generator may afford rapid exploration of orchestration, genre pastiche or texture variation. For another user, it may afford the exaggeration of timbral extremes or the construction of hybrid sonorities. Interface design shapes these affordances: a text box invites linguistic prompting; a menu of labelled styles invites selection among predefined genres; a slider allows some continuous control even when the underlying model is opaque. A system may in principle allow complex restructuring of its training corpus, but if the interface hides that possibility, the user cannot act on it. Creativity is then narrowed by design, and outputs tend to become more conventional than the system might otherwise allow.

A rich human–AI environment would make affordances both legible and open-ended, allowing composers to form reasonable expectations about the consequences of their actions, while also inviting uses not fully anticipated by designers: errors, detours, misuses and unstable results. The creative value of such interaction lies not in freedom from constraints, but in making constraints perceptible, negotiable and available for artistic exploitation.

[1](A) I requested ChatGPT to give me a prompt that would render a variation of the Mona Lisa:
a mysterious Renaissance woman seated in a three-quarter pose, inspired by the compositional atmosphere of Leonardo da Vinci's *Mona Lisa*, with a subtle enigmatic smile, soft sfumato lighting, delicate hands folded in the foreground, and a dreamlike landscape behind her. Reimagine the scene as a dark cinematic fantasy portrait: elegant aged textures, muted earth tones, deep shadows, faint golden highlights, misty mountains, symbolic clouds, and a timeless expression. Highly detailed oil painting, Renaissance realism mixed with surreal contemporary atmosphere, refined brushwork, soft transitions, subtle chiaroscuro, museum-quality composition.
(B) I applied that prompt in Tengr.ai to generate the image.

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press

Some constraints are technical: the limits of the interface, the degree of control over parameters and the quality of the training data. Others are epistemic: opacity, uncertainty about the training corpus and the difficulty of explaining outputs in musically meaningful terms. Still others are aesthetic and cultural: the tendency of training data to privilege certain genres, production values or standards. Working creatively with AI therefore means not only using what a system affords, but also learning how to read, resist and transform the limits through which those affordances appear.

Economic and institutional constraints further shape the field. Access to powerful models is often gated by subscription fees, corporate platforms and intellectual-property regimes. Even when a system appears to offer infinite stylistic variety, its models encode decisions about what counts as a valid musical object, who owns which materials and what kinds of reuse are permitted. Current generative AI systems already offer composers some useful affordances: iterative prompting, variation, style transfer, parameter adjustment, audio continuation, stem-like manipulation and the possibility of using generated material as a starting point for further compositional transformation. These possibilities are real, but they remain limited by the opacity of most current systems as their conditions of operation are not sufficiently visible. An ideal system for artistic work would offer greater transparency at several levels. It would provide clearer information about the training process, the kinds of materials included or excluded, the stylistic and cultural biases of the model and the probable sources or families of sources influencing a given output. It would allow the composer to understand more precisely what is being accepted, resisted or transformed.

This is especially important in electroacoustic and audiovisual composition, where the work often emerges from a dialogue between imagined sound or image, technological procedure and perceptual testing. A composer may begin with an abstract desire – a density, a colour, a pressure, a behaviour – and only later discover the technical means by which it can be realised. AI can intensify this process because it returns possibilities that are both connected to and estranged from the initial intention. The result is not simply a faster version of composition. It is a changed process of trials, errors and evaluations, one in which the artist must continually ask whether the system has revealed a genuine possibility or merely offered a persuasive surface. I refer here mainly to text-to-audio, audio-transformation, source-separation, image/video-generation and model-assisted sampling, in opposition to the unedited generation of complete musical works.

In electroacoustic composition before generative AI, evaluation was a slow and recursive process. A sound was recorded, synthesised, transformed, edited, spatialised, tested in different contexts, compared with other materials, and often rejected or reworked many times before it became part of a piece. The composer did not evaluate the sound only as an isolated object, but in relation to gesture, morphology, energy, memory, formal function and listening experience. Evaluation meant asking not only 'does this sound good?', but 'what does this sound do in the work?' Generative AI does not necessarily eliminate this process but changes what must be evaluated: one is not dealing with neutral material, but with material that may already contain histories, references and the labour of others. The question is therefore not only whether the output is attractive, but whether it can be transformed sufficiently, whether its sources or stylistic debts matter, and whether its use remains artistically and ethically responsible.

3. Seduction, fashion and fetish

If we move from the studio to the broader culture, AI appears as a powerful seduction. On platforms such as Facebook, Instagram and TikTok, AI-generated images, reels, synthetic voices and songs circulate as fashion. They are consumed rapidly, liked, shared and forgotten. The criteria of success are speed, novelty and virality rather than depth of inquiry. The viewer is presented with a glamorous surface, not the exploratory process, the failures, the concepts or the ethical and aesthetic questions behind it.

In 'On the Fetish-Character in Music and the Regression of Listening', Adorno argues that musical commodities acquire a fetish character when exchange value, brand and market position overshadow genuine experience (Adorno 1938/2002). In current AI culture, something similar happens when outputs are valued primarily because they are associated with the latest technology. The statement that 'AI made it' becomes a fetishising label. Today's AI songs, designed around style transfer, genre pastiche and instantly gratifying timbres, fit this logic almost perfectly. Viewers and listeners are encouraged to marvel at the technical feat rather than consider the implications of replacing human performers, eroding authorship or displacing responsibility.

However, the label 'AI-made' is also ambivalent. It does not always function as a sign of technological fascination or novelty. In many contexts, it also operates as an accusation: a marker of low effort, ethical suspicion, derivative production or degraded cultural value. The same label may therefore produce both attraction and distrust. This ambivalence complicates, but does not weaken, the Adornian argument. Fetishisation does not necessarily mean admiration alone. It means that attention is displaced from the concrete labour, process and inquiry of artistic production onto the object's social and technological status. The output can be potentially judged through its mode of production before it is understood as a work. It may be fetishised as a technological miracle, dismissed as 'AI slop', consumed as cheap content, or condemned as ethically compromised.

Adorno radicalises the problem at the level of compositional technique (Adorno 1949/2006). In Philosophy of New Music, fetishism appears as the hardening of musical material into cliches, formulas and stylistic masks. If this is transposed into AI-driven production, prompt libraries, style templates and 'make this sound like X' interfaces play a comparable role. They reify musical possibilities into selectable products: lo-fi beats, cinematic trailer music, 1980s synthesis, Ligeti-like choir and so on. AI art can then produce a double fetishism. First, outputs are fetishised as autonomous magical objects consumed without regard to the processes that produced them. Second, the models, apps and interfaces become fetish objects in themselves: mysterious and powerful things into which users project fantasies of creativity, freedom and technological marvel. Desire is translated into sound without the friction of technique, notation, rehearsal, performance, failure, re-evaluation or reconstruction.

Lawrence Kramer's work on musical meaning helps to situate this fetishism within a broader interpretive frame. Kramer argues that the modern musical work often appears as a privileged, autonomous and timeless object, yet it remains bound up with social narratives, ideological formations and forms of subjectivity (Kramer 1990). Listening is not the contemplation of a neutral object but a process through which meanings are projected, negotiated and contested (Kramer 2002). In the AI context, what is fetishised is not only the output but the fantasy it sustains: creativity without constraint, productivity without effort and a

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press

machine-subject that composes, paints or writes on our behalf. The AI artefact appears magical, while condensing a network of labour – coding, data curation, interface design, marketing and platform maintenance – that is largely erased in consumption.

Adorno suggests that commodity fetishism can allow the object to substitute for the effort of thought itself. A similar danger appears when the artist accepts a generated result as if it were already a completed artistic decision. A clever prompt or parameter setting may stand in for the longer process of developing, testing and working through ideas. The consequence is a subtle alienation: the artist recognises less of themself in the work and more in the act of triggering and curating seductive outputs.

This point should be distinguished from earlier moments in the history of music technology. Electronic studios, tape machines, modular synthesisers, samplers, sequencers, digital workstations, VST instruments have all allowed composers to obtain sounds that they could not have produced without those tools. However, in most of these cases, the marvel surrounding the technology remained connected to an acknowledged form of labour: programming a synthesiser, cutting and splicing tape, designing a patch, editing samples, integrating the result into a larger compositional structure. Even when presets or ready-made sounds were used, they normally entered a process in which the composer still had to arrange, transform, contextualise and evaluate the material. Generative AI alters this situation because it can appear to subsume a much larger portion of the creative chain. It may produce not only a sound, a texture or a preset, but a complete stylistic surface. The listener or viewer may no longer be able to identify if the work is the result of sustained compositional inquiry, rapid prompting, automated recombination or the hidden absorption of other artists' materials into training data. This affects not only AI-generated works but also the reception of human-made works, which may now be met with suspicion: was this thoughtfully composed, or fast generated, assembled from a model, or merely selected from automated outputs?

In Adorno's analysis of the culture industry, musical fetishism also contributes to the 'liquidation of the individual' (Adorno 1938/2002). In AI art, a comparable situation may occur: the model appears to occupy the position of the creative subject, possessing agency and authority, while human beings risk being reduced to functional roles within the system that produces, distributes and legitimates the work.

Fetishism wants the object to be opaque and self-sufficient: an idol that simply works. A critical relation to AI should therefore treat opacity as a problem to be investigated rather than a mystery to be accepted. It foregrounds process: trial and error, unexpected failure, ethical friction and conceptual doubt. To de-fetishise AI is to treat models and outputs as historically situated tools whose capacities and limits can be probed, contested and reshaped.

The problem is intensified by the temporal structure of platforms. Their interfaces reward immediate recognition: a style must be legible within seconds, a result must be shareable, a prompt must produce an image or track quickly enough to sustain attention. This temporality is very different from the time of composition, where uncertainty, hesitation and dissatisfaction are often productive. Many works become stronger precisely because they resist premature completion. AI platforms, by contrast, tend to convert uncertainty into options and options into deliverables. The result is not only an aesthetic narrowing but a change in the artist's relationship to difficulty. Difficulty, which has often been the site of invention, is reframed as an obstacle that the system promises to remove.

4. Serendipity, templates and artistic identity

'Serendipity', since Horace Walpole's coinage in 1754, refers to the discovery of something unexpected and valuable while looking for something else: the Three Princes of Serendip were 'always making discoveries, by accidents and sagacity, of things which they were not in quest of' (Walpole 1754; Merton and Barber 2004). In artistic practice, a strange glitch, timbral collision or unforeseen formal pattern may appear, and the artist may recognise it as latent material. Serendipity therefore presupposes a situated artistic identity: someone able to say that an accidental result reveals something they did not yet know they wanted.

Yet serendipity also carries a danger. The accidental may replace conceptual thought rather than enrich it. An artist may arrive with one intention, encounter a seductive output and redirect the work not because the output deepens the original idea, but because it offers an immediate attraction: 'I did not want this, but it fascinates me, so I will use it instead'. This is not necessarily negative; many important works arise from deviations. But the artist must recognise when a substitution is taking place. Otherwise, the accidental ceases to be a productive disturbance within thought and becomes a shortcut that quietly displaces it. The same platforms that encourage fetishistic consumption can, under specific conditions, host serendipity. The artist can cultivate situations in which the unforeseen can emerge, and then assumes responsibility for what that emergence means.

Apostel's account of models is again relevant. To compose with AI is to compose with models: to recognise that one is navigating representations structured by prior data, design choices and implicit world views (Apostel 1961). Surprising outputs are not merely random gifts. They are points at which the system's internal model diverges from the artist's expectations. These divergences can become sites of inquiry. They reveal where the artist's concept of the work, the system's statistical model and the cultural models embedded in the data fail to coincide.

This has consequences for artistic identity. Under the conditions of the culture industry, Adorno recognised the danger that identity might be reduced to a brand: a recognisable signature style or commodity profile. Yet artistic identity has historically been more complex and conflicted. It has involved interiority, authenticity, singularity of voice, experimentation, negation and innovation. Postmodern and critical perspectives, including those of Kramer and Taruskin, further emphasise that artistic identities are constructed through social roles, institutions, genres, technologies and discourses. The composer, performer, producer or DJ is not a timeless essence but a position within a network of practices and expectations (Kramer 1990, 2002; Taruskin 1995, 2009).

In human–AI interaction, artistic identity might then be rethought along three axes. The first is curatorial: the artist as selector of data, prompts, procedures and outputs. Identity appears through choosing, filtering, cutting, juxtaposing and framing. The second is processual: the artist as designer of processes and conversations with systems. The work becomes the trace of experiments, revisions and feedback loops, and the artist's voice is heard in the pattern of questions posed to the machine and in the treatment of its answers. The third is critical: the artist as someone who makes the operations of AI visible, exposing training data, biases, labour conditions and economic structures behind the seductive surface. These identities sit uneasily with the template logic of platform culture, which favours quickly legible styles and branded outputs.

These three identities should not be understood as a new reality, nor as a prediction that all composers or artists will move in this

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press

direction. They are better understood as tendencies already visible in current AI-assisted practice, where artists increasingly work by selecting datasets, designing prompts, curating outputs, building workflows and critically exposing the conditions of production. In this sense, they extend older roles associated with experimental music, sampling, studio composition and media art, rather than replacing the composer as such.

Artists have always worked with inherited forms, templates, genres, schemas and conventions. A fugue, a sonata, a raga, a blues progression or a cinematic genre can all function as templates in a broad sense. The difference is that historical forms are transformed through use, resistance, misuse and interpretation. Platform templates, by contrast, often conceal the conditions of their formation and restrict the user's capacity to intervene in them. A critical AI practice must therefore recover the historical thickness of the materials that templates tend to flatten.

In this context, template logic is meant to be interpreted as the reduction of artistic decision-making to the selection of pre-formatted styles. It is not caused by AI alone; it is intensified by commercial interfaces that make certain actions frictionless while making other forms of exploration technically or economically inaccessible. In this context, the value of a system lies less in the quality of its first response than in the resistance and misalignment that response introduces into a larger compositional process.

Isaac Asimov's short story Profession offers a useful allegory (Asimov 1957). In its future society, education has been reduced to direct mental uploading through a machine-brain interface. Reading is transferred instantaneously at the age of eight; a decade later, at eighteen, citizens also receive, instantaneously, their full professional formation. One's profession is determined by tests. The protagonist, George Platen, expects to become a computer programmer but is told that his brain is unfit for any profession. Sent to what seems to be an institution for the feeble-minded, he later discovers that it is actually an Institute of Higher Studies. Those who cannot be fitted into direct upload are the minority capable of original thought: the people who will create new knowledge and design the future machines that others will use to receive pre-defined knowledge.

Asimov's parable is striking in the age of generative AI. It warns against a culture in which knowledge becomes a template delivered from outside, and in which genuine learning is replaced by instant functionality. It reminds us that the true profession may lie in questioning the system rather than serving it; Adorno and Kramer remind us that fetishism is a social and perceptual danger; Apostel reminds us that no model is final. To insist on exploratory collaboration with AI is therefore an act of de-fetishisation: a way of reclaiming artistic thought within an environment increasingly populated by seductive images and sounds.

5. Aesthetic perception and aura

At the centre of embodied craft intelligence – Curtis Roads has recently proposed the term 'artisanal intelligence' – lies aesthetic perception (Roads 2025). Aesthetic perception is not purely rational. It is entangled with intuition, unconscious yet informed processes and ways of sensing form before they can be fully articulated. Over years of practice, artists learn to hear minute differences of articulation, feel when an image is overworked, recognise when a structure is decorative rather than necessary and sense when a material has become exhausted.

This perception is not simply taste. It is embodied memory, formed through countless encounters with works, tools, performers, audiences, successes and failures. It functions as an internal compass, helping artists decide which problems are worth pursuing, which materials resonate with their world-view and which risks are necessary. Models may be trained to predict human aesthetic judgements, classify sounds or images as appealing, optimise for engagement or mimic fashionable styles. Such systems can simulate a form of aesthetic consensus by interpolating between many responses. But they remain second-order approximations: statistical reflections of aggregated reactions rather than lived aesthetic experience.

Embodied craft intelligence is shaped by vulnerability, biography, historical memory and risk. It is anchored in a singular trajectory: an artist's history of observing, listening, criticising and connecting works to other works. It depends on an embodied sense of duration, attention and expectation, and on the composer's willingness to allow a work to find its necessity gradually. The danger arises when we confuse these two modes, treating the outputs of a model as if they carried the same risk-exposed judgement as artisanal work. At that point, there is not only a technical but an ethical displacement: decisions that should be grounded in responsibility and vulnerability are projected onto a system that cannot itself be responsible or vulnerable.

Walter Benjamin's concept of the artwork's aura marks a crucial tension. He argues that mechanical reproduction devalues the unique 'here and now' of the artwork, its singular presence in time and space (Benjamin 1935/1969). Aura is the experience of distance in proximity: the sense that the work comes from another time, another world, another horizon of meaning, even as we stand directly in front of it. Mechanical reproduction detaches the work from its ritual or cultic context and multiplies it as copies. The work becomes available everywhere, but at the cost of that singular, unrepeatable presence. What once demanded a pilgrimage to the Louvre or to Bayreuth can now be printed on a postcard or streamed in a playlist.

Template-based, AI-driven production radicalises this process in at least two ways. First, instead of reproducing a pre-existing work, it generates endless new variants from the same model, compressing differences into a narrow corridor of stylistic and procedural presets. The logic shifts from copying to sampling a distribution: we do not have replicas of the same painting or performance, but statistically plausible surrogates for a style or genre. Second, the locus of uniqueness is displaced. Aura no longer clings to the singular artwork, since there may be thousands of equally 'authentic' outputs generated in seconds; instead, it migrates to the model, the style template or the dataset. From Benjamin's perspective, this displacement is ambivalent. On one hand, it continues the de-auratic trend of mechanical reproduction: the individual work becomes a transient item in a feed. On the other hand, AI systems can acquire a pseudo-aura: the glamour of opacity and the fascination of what cannot be fully understood. Their outputs are consumed as though they carried a trace of this mystery, even when they are structurally interchangeable. Aura is not simply missing in such cases; its absence becomes perceptible as a kind of flattening, a lack of depth.

Conversely, careful human–AI collaborations can re-stage aura by embedding AI processes in situations that restore singularity: live decisions that cannot be replayed with exactitude. Rather than accepting the model's best guess as a finished product, the artist treats each output as a question. Is this mine? Does it belong to the work? What does this accidental gesture reveal?

This feedback process must remain embedded in human perception and emotion. The machine can propose, but it cannot care.

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press

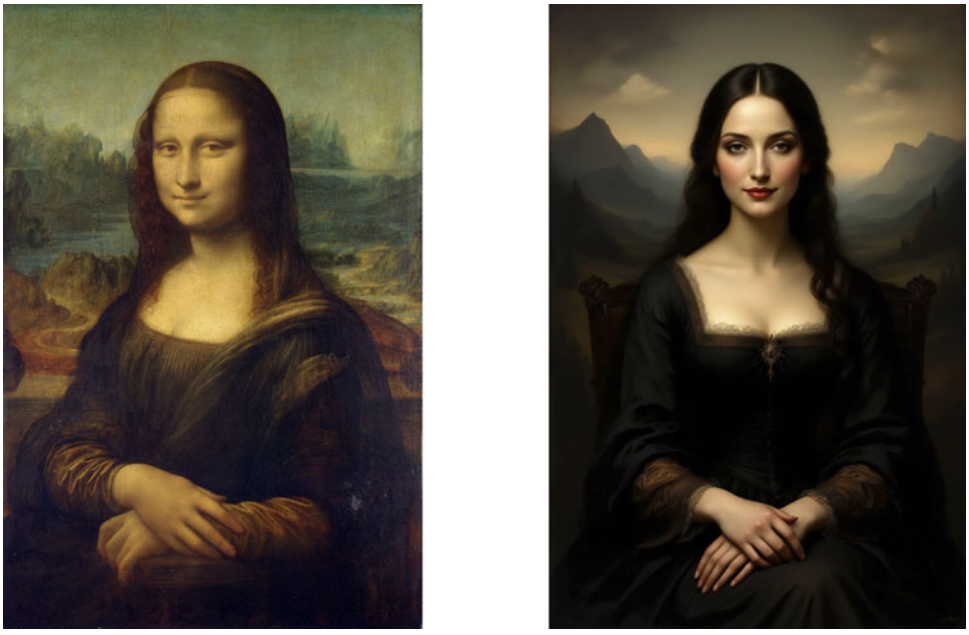


Figure 1. Mona Lisa and My Own Lisa. Comparative illustration of Leonardo da Vinci's Mona Lisa and an AI-generated image produced from a prompt derived from an AI-generated description of the original painting's visual style. The AI-generated image is used here as an example of translation, abstraction and recombination rather than as a claim of autonomous originality.

It can generate many variants, but it cannot feel that one of them carries a specific tension, memory or risk. The opposition between embodied craft intelligence and artificial intelligence is therefore also a conflict between temporalities and economies of attention. Craft intelligence grows slowly, through repetition, critical awareness and revision; commercial AI thrives on speed and instant feedback.

In the age of AI, aura may not appear as the untouchable uniqueness of a painting or score. It may reappear as the singularity of a specific human–machine relation: a historically situated way of working with models, grounded in embodied craft. The future of aura depends less on the technology itself than on the persistence of cultivated, vulnerable, risk-exposed aesthetic perception.

6. AGI and the politics of intelligence

OpenAI's charter defines artificial general intelligence (AGI) as 'highly autonomous systems that outperform humans at most economically valuable work' (OpenAI 2018). This is only one definition among many; others emphasise breadth of domains, reasoning capacity or human-level flexibility. Yet when AGI is framed explicitly in economic terms, intelligence risks being reduced to productivity, efficiency and revenue.

This reduction has concrete consequences for the arts. If intelligence is defined primarily as the capacity to outperform humans at economically valuable work, artistic labour risks being reclassified not as a domain of irreducible human expression but as a cost to be minimised. Corporate strategies that replace musicians with royalty-free AI tracks, virtual performers, synthetic voices or digital avatars are not accidents. They follow logically from a definition of intelligence that equates value with efficiency and cheap labour. Within such a framework, the ideal music system generates infinite, stylistically compliant content with minimal legal friction and no labour claims: background music without royalties, voices that never age, performers who never refuse to play and bands that never go on strike. This argument applies most clearly to music treated as a functional commodity: background tracks, library music, stock audio or generic mood-based content designed for speed, low cost and licensing simplicity. Film music produced with the help of AI, however, often operates slightly differently. It can support narrative development, shape emotional expectation, create cultural or historical associations and adapt to recognisable stylistic presets that help audiences interpret a scene. In such cases, music is not only a replaceable sonic product, but part of a larger dramaturgical and cultural system.

There is also a labour dimension that should not be separated from aesthetics. Training data are rarely neutral. It may contain the traces of countless artists, performers, recording engineers, designers and writers whose work has been absorbed into models without meaningful consent, attribution or compensation. To speak of AI as if it were an isolated creative agent is therefore to erase the social and material conditions that made its outputs possible.

If AGI is presented as the future standard of intelligence, human creativity runs the risk of being tacitly redefined as a temporary workaround: a charming but obsolete phase before automated content. Even when AI is marketed as democratising creativity, this often means democratising access to production tools without democratising control over platforms, data or revenue.

Against this background, treating human–AI interaction as a feedback-based mode of inquiry is not only a technical design principle but an ethical and political stance. It assumes that the value of AI lies not in replacing human effort, but in opening spaces of experimentation, perception and understanding. Such a stance asserts that human intelligence – embodied, craft-based, historically situated and aesthetically informed – cannot be reduced to economically valuable output. It foregrounds dimensions of artistic work that resist quantification: the slow training of perception, the cultivation of judgement, the responsibility for what one puts into the world and the commitment to results that may never become profitable. To embed that conception in our tools and institutions is to resist the reduction of intelligence to profit and to insist that working with AI remains a space for thought rather than another channel for extraction. AI did not create economic pressure in the arts; it intensifies and reconfigures older systems of commissions, markets, institutions. The danger is when speed, scale and low cost become the primary measures of creative value.

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press

7. Conclusion

AI can expand sonic and visual possibilities, introduce unexpected connections and accelerate certain kinds of experimentation. It can reveal patterns difficult to detect by ear or eye alone and act as a vast, noisy memory of stylistic tendencies. It can also affect literary and text-based practice, narrative structure, metaphor, dialogue, genre and conceptual framing, especially if results are treated as experiments to be questioned, revised, abandoned or accepted.

At the same time, it can seduce artists into fetishising the tool, normalising template-based creativity, eroding aura and replacing aesthetic perception with algorithms optimised for engagement and profit.

A model treated as a magical oracle encourages fetishism and a regression of listening. A model treated as a fallible interlocutor can encourage critical engagement. If AI is used within practices that foreground embodied craft, slow perception and historical memory, it can become a catalyst for reflection and invention. Apostel reminds us that every model is selective and serves particular aims. Gibson reminds us that the meaningful properties of an environment are its affordances for action, to be discovered, critiqued and reshaped.

In that environment, the artist's knowledge and intuition remain central. They generate the questions worth asking of the system; they judge when serendipity reveals something truly new and when it is merely noise; they filter accidents; they decide which glitches to elevate into form and which to discard; they protect aesthetic perception against the dulling effects of automation.

This is why refusal is an essential part of human–AI practice. To refuse an output is not to reject technology; it is to affirm that the work has criteria that exceed the system's offer. AI may multiply possibilities, but art still requires the more difficult question of necessity: why this sound, this image, this structure, here and now.

There is a paradox here: the more an artist uses powerful technology, the more skeptical they must become about it. This means keeping the tool visible as a tool, interrogating training data, slowing down operations that platforms are designed to accelerate and sometimes working deliberately against the grain of the model. Such skepticism is not hostile to creativity; it is part of the craft.

The debate around AI in the arts should therefore shift. Instead of asking whether machines will surpass us, we should ask what kinds of relations we want to build with them. Instead of measuring success by speed and volume, we should attend to the quality of attention, the depth of listening and the specificity of the worlds that works open up. Technology always arrives inside a field of choices about perception, value and power. The task is not to decide once and for all whether AI belongs in the studio, but to negotiate continuously how it belongs there, so that it can participate in the fragile construction of artistic freedom.

Acknowledgement of AI-assisted editorial preparation

The author used OpenAI's ChatGPT (GPT-5.5 Thinking, accessed 12 July 2026) for English-language editing, condensation and formatting assistance in preparing this manuscript. All intellectual claims, citations, examples and final editorial decisions remain the author's responsibility.

References

Adorno, T. W. 1938/2002. On the Fetish-Character in Music and the Regression of Listening. In R. Leppert (ed.) *Essays on Music*, translated by S. H. Gillespie. Berkeley: University of California Press, 288–317.

Adorno, T. W. 1949/2006. *Philosophy of New Music*, translated by R. Hullot-Kentor. Minneapolis: University of Minnesota Press.

Apostel, L. 1961. Towards the Formal Study of Models in the Non-Formalized Sciences. In H. Freudenthal (ed.) *The Concept and the Role of the Model in Mathematics and Natural and Social Sciences*. Dordrecht: Reidel, 1–37.

Ashby, W. R. 1956. *An Introduction to Cybernetics*. London: Chapman & Hall.

Asimov, I. 1957. Profession. Astounding Science Fiction, July. Reprinted in *Nine Tomorrows: Tales of the Near Future*. New York: Doubleday, 1959.

Benjamin, W. 1935/1969. The Work of Art in the Age of Mechanical Reproduction. In H. Arendt (ed.) *Illuminations*, translated by H. Zohn. New York: Schocken Books, 217–51.

Cope, D. 1996. *Experiments in Musical Intelligence*. Madison, WI: A-R Editions.

Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Harvey, J. 1986. The Mirror of Ambiguity. In S. Emmerson (ed.) *The Language of Electroacoustic Music*. London: Macmillan, 175–90.

Haugeland, J. 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.

Hiller, L. A. and Isaacson, L. M. 1959. *Experimental Music: Composition with an Electronic Computer*. New York: McGraw-Hill.

Kramer, L. 1990. *Music as Cultural Practice, 1800-1900*. Berkeley: University of California Press.

Kramer, L. 2002. *Musical Meaning: Toward a Critical History*. Berkeley: University of California Press.

McCorduck, P. 2004. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*, 2nd edn. Natick, MA: A K Peters.

Merton, R. K. and Barber, E. 2004. *The Travels and Adventures of Serendipity: A Study in Sociological Semantics and the Sociology of Science*. Princeton, NJ: Princeton University Press.

Newell, A. and Simon, H. A. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

OpenAI. 2018. OpenAI Charter. https://openai.com/charter/ (accessed 12 May 2026).

Polanyi, M. 1966. *The Tacit Dimension*. Garden City, NY: Doubleday.

Roads, C. 2025. Music and AI: A Personal Journey. Paper presented at Music and Machines, Accademia Musicale Chigiana, Siena.

Taruskin, R. 1995. *Text and Act: Essays on Music and Performance*. Oxford: Oxford University Press.

Taruskin, R. 2009. *Music in the Late Twentieth Century*. Oxford: Oxford University Press.

Turing, A. M. 1950. Computing Machinery and Intelligence. *Mind* **59**(236): 433–60. https://doi.org/10.1093/mind/LIX.236.433.

Walpole, H. 1754. Letter to Horace Mann, 28 January.

Wiener, N. 1948. *Cybernetics, or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press.

Winston, P. H. and Prendergast, K. A. (eds.) 1984. *The AI Business: The Commercial Uses of Artificial Intelligence*. Cambridge, MA: MIT Press.

Xenakis, I. 1971. *Formalized Music: Thought and Mathematics in Composition*. Bloomington: Indiana University Press.

https://doi.org/10.1017/S1355771826101538 Published online by Cambridge University Press