MEDWIN PUBLISHERS
Committed to Create Value for Researchers

Philosophy International Journal
ISSN: 2641-9130

When Generation Becomes Creation: Distinguishing Human Creativity from AI Derivation

Choi B*

Department of Arts, College of Arts and Design, Wonkwang University, Republic of Korea

***Corresponding author:** Byungkil Choi, Department of Arts, College of Arts and Design, Wonkwang University, Republic of Korea

Conceptual Paper
Volume 9 Issue 3
Received Date: September 07, 2026
Published Date: September 30, 2026
DOI: 10.23880/phij-16000377

Abstract

Generative AI has unsettled a basic distinction in artistic evaluation: the difference between generating an image and creating a work. Existing approaches usually locate that distinction in novelty, intention, transformation, human control, or distributed agency. Each captures a genuine feature of creative practice, but none by itself explains when aesthetic achievement in a human–AI process is attributable to the human participant. This article proposes generative agency as a threshold concept. Generative agency is the capacity to make aesthetically consequential judgments within an indeterminate generative process such that the resulting work becomes attributable to a trajectory of artistic decisions rather than merely to the productive capacities of the system. The article distinguishes generation, derivation, creativity, authorship, plagiarism, and legal infringement; argues that derivation is not the opposite of creativity; and differentiates generative agency from transformational creativity, creative-control theory, and distributed or entangled accounts of authorship. Generation becomes creation when computational contingency is not merely received, selected, or supervised, but becomes answerable to aesthetic judgment through intentional direction, consequential intervention, responsive iteration, selective commitment, and integrative responsibility.

Keywords: Generative AI; AI art; Creativity; Generative agency; Derivation; Authorship; Plagiarism; Aesthetic judgment

Introduction: When Does Generation Become Creation?

The Problem of Creative Attribution

Generative artificial intelligence makes it possible to produce visually complex images from instructions that may specify only a fraction of the resulting form. A prompt can identify subject matter, atmosphere, style, or viewpoint while leaving composition, spatial relations, lighting, textures, gesture, and innumerable local determinations to the model. The user plainly participates in causing the image to appear, but causal participation does not settle whether the aesthetically significant features of the image are the user's creative achievement. This gap between causing an output and creating a work is the central problem of this article.

The problem should not be reduced to the older question of whether machines can be creative. Boden's distinction among combinational, exploratory, and transformational creativity, and Wiggins's attempts to formalize creative systems, widened the conceptual field by showing how novelty can emerge without creation ex nihilo [1,2]. Yet such accounts mainly classify kinds of creative production; they do not determine whose achievement is instantiated when a human and a generative system jointly contribute to an image. Hertzmann's historical comparison with photography and animation likewise usefully resists the assumption that automation eliminates artistry, but the tool analogy becomes inadequate when the tool supplies aesthetically salient formal solutions that the user did not individually specify [3].

CC BY

The distinction becomes sharper if the object of attribution is specified. In ordinary discussion, the statements "the user made the image," "the model generated the image," and "the artist created the work" are often treated as interchangeable. They are not. The first may report participation in a workflow, the second a causal fact about image production, and the third an evaluative claim about artistic achievement. Generative AI separates these claims because the system can contribute composition, lighting, texture, physiognomy, spatial relations, and local detail that the user neither individually represented nor directly executed. The resulting problem is not that human intention has vanished, but that intention and realization are no longer connected in the familiar way. A theory of AI art must therefore explain how a human intention can remain artistically significant when many of the expressive particulars through which the work succeeds arise contingently within the generative process.

Nor can the problem be solved by looking only at the final image. Novelty does not by itself establish creativity, because an unprecedented configuration can arise without an attributable creative achievement. Similarity does not by itself negate creativity, because quotation, parody, repetition, and appropriation can be constitutive of artistic practice. Krauss and Crimp showed, in different ways, that modern and postmodern art repeatedly complicate the opposition between originality and repetition [4,5]. The relevant question is therefore not whether an output is novel or derivative in the abstract, but how antecedent and generated possibilities are taken up within a practice of judgment.

This is why output-centered tests are especially unstable. A visually unprecedented image may be the first acceptable sample from a minimally specified prompt, whereas a highly derivative-looking image may result from sustained critical work directed toward quotation, repetition, or historical citation. The visible object underdetermines the history of judgment that produced it. Conversely, process alone cannot settle the matter either, because a long workflow may consist largely of search, technical troubleshooting, or optimization toward a result whose decisive formal relations are still supplied by the model. The proper unit of analysis is consequently neither the isolated output nor the mere duration of production, but the relation between emerging outputs and the sequence of aesthetic judgments through which some possibilities are made consequential and others are refused.

Legal discussions sharpen the issue without resolving it aesthetically. The U.S. Copyright Office distinguishes between AI used assistively and situations in which the system determines expressive elements, and it has concluded that prompts alone, with currently generally available systems, ordinarily do not provide sufficient human control over the resulting expressive details [6]. That distinction is important because it separates specifying a desired result from determining its expressive realization. But copyrightability is an institutional threshold for assigning rights; it cannot simply be imported as a theory of artistic creation.

The more promising theories focus on process. Liu and Huang's creative-control theory argues that authorship can persist where human artistic imagery thinking guides creation from conception to execution despite generative unpredictability [7]. Bomba and De Angeli, by contrast, reject a single sovereign center of agency and describe generative practice through relational agency and entangled authorship among artists, data, algorithms, and material-discursive conditions [8]. Both positions identify something essential: the first preserves the continuity of human direction; the second recognizes that generative systems are not passive transmitters. Yet neither fully answers the threshold question. Direction may persist while the system supplies the decisive aesthetic success, while relational distribution may explain production without telling us which contribution is artistically constitutive.

Generative Agency: Thesis and Contribution

This article therefore proposes generative agency. By generative agency I mean the capacity to exercise aesthetically consequential judgment within an indeterminate generative process such that the work becomes intelligibly attributable to a trajectory of artistic decisions rather than merely to the system's productive capacities. The qualification aesthetically consequential is decisive. Human activity is not enough; control is not enough; repeated interaction is not enough. The human participant's judgments must make a difference to what the work becomes.

The phrase "aesthetically consequential" carries the argumentative burden. It excludes a merely quantitative account of participation. A user may perform many operations without making a decision that matters to the identity of the work, while another artist may make a small number of decisions that reorganize the significance of everything generated around them. The theory therefore does not reward labor, technical complexity, or frequency of intervention as such. It asks whether human judgments explain why aesthetically important relations in the final work acquired the form and role they possess. In this respect generative agency is narrower than a general theory of human involvement and more demanding than a theory that equates continuing control with authorship.

The argument proceeds cumulatively. Section 2 establishes the conceptual distinctions that prevent derivation, plagiarism, generation, and creation from

Copyright© Choi B.

collapsing into one another. Section 3 compares the principal accounts of intention, transformation, creative control, and distributed or entangled agency and isolates their common attributional gap. Section 4 explains why the prompt–image gap and computational contingency make that gap unavoidable. Section 5 develops generative agency in three stages: from causal contribution to aesthetic constitution, through the internal structure of orientation, intervention, response, commitment, and integration, to the attributional threshold itself. Section 6 tests that threshold by distinguishing practices that remain below it, practices that cross it through responsive transformation, and boundary cases in which the level or value of the achievement complicates attribution. Section 7 specifies what the theory does not establish: creative transformation does not erase derivation, creative agency does not require exclusive authorship, and creative attribution does not determine legal permissibility.

Creation, Generation, and Derivation

Creativity Does Not Require Non-Derivation

Human creation has never been wholly non-derivative. Artists work through inherited media, conventions, genres, techniques, iconographies, and prior works. If originality means complete independence from antecedents, almost no artistic practice qualifies; if every recombination of inherited material counts as creativity, the concept loses discriminating force. The more useful account treats originality as the artistically significant reorganization of antecedent possibilities. Irvin's analysis of appropriation is important here because it rejects innovation or perceptual distinctiveness as necessary conditions of authorship and instead emphasizes the artist's responsibility for the objectives pursued through the work [9]. Her account relocates artistic significance from visible novelty toward what the artist is doing through the relation to antecedent material.

That insight must nevertheless be refined for generative AI. Responsibility for choosing an objective is not yet responsibility for every aesthetically significant feature through which a work succeeds. A user can decide to seek a melancholic portrait or an image in a particular style while the model supplies the composition, gesture, chromatic relation, or expressive accident that gives the output its force. The relevant distinction is therefore between setting an objective and becoming responsible for the aesthetic articulation through which the objective is realized.

The point is not to rehabilitate a romantic ideal of originality by moving it from the object to the artist's mind. Human artists themselves often discover rather than foresee decisive features of a work. Painters respond to accidents of material; photographers respond to events they did not cause; improvisers respond to other performers; conceptual artists may establish procedures whose concrete realizations remain open. What matters in these cases is not complete anticipation but the capacity to recognize a contingency as significant and to make subsequent decisions answer to it. Generative AI belongs within this longer history of non-sovereign creation, but it radicalizes the problem because the source of contingency is an artifact capable of supplying highly articulated expressive solutions. The analogy with chance is therefore illuminating only if it does not conceal the increased formal contribution of the system.

Derivation, Appropriation, and Plagiarism

This distinction also clarifies the relation among influence, derivation, appropriation, and plagiarism. Influence describes the formative effect of prior works and traditions, often without conscious quotation. Derivation is broader and normatively neutral: it denotes dependence upon antecedent material, whether a particular image, a style, a genre, a procedure, or a cultural repertoire. Appropriation is more specific because the antecedent relation is deliberately made operative within the new work. Plagiarism, by contrast, is normative; it concerns an improper relation of attribution, representation, or appropriation under the norms of a practice. Sarkar is therefore right to resist the equation of cultural reuse with 'mechanised plagiarism' and to emphasize that reuse is endemic to creativity [10]. But the conclusion that all creativity involves reuse is only a starting point. It does not tell us whether a particular dependence is critically transformed, merely inherited, improperly concealed, or aesthetically insignificant.

Two axes should accordingly remain independent. One concerns derivative dependence: how strongly and in what manner does a work depend upon antecedent material? The other concerns creative agency: how strongly do an agent's aesthetic judgments account for what the work becomes? A work can be highly derivative and highly creative, as appropriation art demonstrates; weakly derivative and weakly creative, as in a generic first-pass output; highly creative yet normatively problematic; or non-plagiaristic while still involving little human creative agency. This is why the article speaks of AI derivation rather than AI plagiarism: derivation identifies dependence without prejudging its normative status.

Keeping the axes independent also prevents a recurrent mistake in debates over AI training. The fact that a model depends upon large bodies of antecedent cultural material raises serious questions about provenance, permission, attribution, and compensation, but those questions do not automatically answer whether a particular user exercised

Copyright© Choi B.

creative agency in a later workflow. Equally, evidence of strong human creative agency does not retroactively legitimate every antecedent relation on which the system or the final work depends. A theory of creation should therefore be capable of saying both that a work is substantially creative and that some of its derivative relations remain normatively objectionable. That possibility is not inconsistency; it follows from treating creative achievement and derivative dependence as different dimensions.

Why Generation is Not Yet Creation

Generation must likewise be distinguished from creation. Generation is a causal-production category: a process brings an output into existence from inputs, parameters, learned regularities, and computational operations. Creation is stronger because it attributes an achievement. Direct manual execution is not required; conceptual art, photography, architecture, collaboration, and chance procedures already show that artists can create works whose features they do not physically manufacture or completely predict. What matters is whether aesthetically relevant determinations can intelligibly be referred back to an artistic practice of decision and judgment.

Conceptual art makes the difference especially clear. An instruction can itself specify the rule through which a work is realized. A generative prompt, by contrast, often specifies a desired semantic or stylistic region while leaving the system to supply the expressive particulars. The common presence of language does not make the two cases equivalent. A prompt can certainly become part of a sophisticated conceptual procedure, but ordinary prompting should not be treated as authorship merely because the instruction initiates production.

This distinction is particularly important for prompts because prompts occupy an ambiguous position between instruction and material. Some prompts function like ordinary requests: they name a subject, mood, style, or compositional preference and leave the system to realize those demands. Others may themselves form part of a rule-governed artistic procedure in which the artist deliberately fixes a linguistic constraint, accepts the first output, or stages the gap between instruction and realization as the work's subject. In the latter case, a sparse prompt can belong to a highly articulated artistic project. The relevant achievement, however, may lie at the level of the procedure or conceptual framing rather than in human authorship of every internal feature of the generated image. This level-specific attribution will recur throughout the argument.

The asymmetry between human and machine derivation should therefore be stated functionally rather than metaphysically. Human artists and models both depend upon antecedents, and both can introduce novelty. The relevant difference is that human artistic derivation can become situated within a practice of selective attention, historical address, normative uptake, revision, and responsibility. A predecessor can be admired, resisted, quoted, misremembered, or deliberately reactivated. A model's statistical relation to training data, by itself, does not establish such an accountable relation. This does not prove that only humans can ever be creative; it identifies the structure of present human–AI practice. Machine derivation becomes part of human creation only when generated possibilities enter a further trajectory of aesthetic judgment.

Existing Accounts and the Attributional Gap

Intention and Transformation: Necessary but Non-Attributive

The first group of accounts identifies conditions without which creativity would be unintelligible, but those conditions do not yet solve attribution in a human–AI process. Gaut treats creativity as an exercise of agency involving originality, value, and flair, while Stokes emphasizes non-accidental agency without requiring the creator to intend the precise result eventually produced [11,12]. Their common achievement is to prevent creativity from collapsing into lucky novelty: an occurrence is not a creative achievement merely because it happens to be new or valuable. Yet generative AI exposes the limit of that insight. A user normally intends to produce an image, and may even pursue a recognizable aesthetic aim, while the model supplies the particular composition or expressive relation that makes the output successful. Intention therefore separates purposive initiation from accident, but it does not by itself establish attribution of the resulting aesthetic achievement.

Transformational accounts move the argument from intention to change. Boden's combinational, exploratory, and transformational creativity explains how new possibilities can emerge by recombining elements, exploring a conceptual space, or altering its governing constraints [1]. Bomba and De Angeli similarly describe advanced AI practice as transformational insofar as artists work through data curation, generated surprises, and reciprocal interaction that can reshape their creative spaces [8]. These accounts explain why derivation and unpredictability are compatible with creativity. Their difficulty is different from the intentionalists': transformation is not self-attributing. A generative model can produce a radical formal transformation in response to minimal human direction, so the magnitude of change does not tell us whether the human, the model, or their interaction should receive credit for what is aesthetically achieved.

Copyright© Choi B.

Sarkar's discussion of reuse and critical integration occupies an important intermediate position. He correctly resists treating reuse, randomness, or AI assistance as automatic evidence of non-creativity and shifts attention from direct production toward the human work of selecting, contextualizing, and integrating generated material [10]. But integration itself can remain curatorial. A user may competently assemble, select, or contextualize outputs whose decisive internal aesthetic relations remain system-supplied. The unresolved question is therefore not whether intention, transformation, or integration matters; each clearly does. The question is when any of them becomes constitutive of the particular artistic achievement attributed to the human participant.

Creative Control: Direction without Constitution

Liu and Huang's creative-control theory comes closest to this attribution problem because it explicitly asks how human authorship can persist despite the generative black box. They argue that artistic imagery thinking can guide creation across a continuum from conception to execution and that unpredictability does not necessarily defeat human control [7]. This is an important advance over tests that isolate the initial prompt or inspect only the final output. It recognizes that artistic agency can be dynamic, iterative, and compatible with incomplete prediction.

The difficulty is that control can operate at different levels. A user may control the decision to generate, specify subject and atmosphere, revise prompts, reject outputs, choose when to stop, and select a final image. All of these actions are genuine forms of direction. Yet the model may still supply the composition, spatial tension, gesture, chromatic relation, or other formal solution that explains why the selected image is aesthetically compelling. If every sustained capacity to reject or redirect is sufficient for authorship, control risks becoming a theory of supervision: it establishes who manages the process without establishing who is responsible for its decisive aesthetic achievements.

The central distinction is therefore between control as direction and control as constitution. Direction constrains what the system is asked to do; constitution concerns whether human judgments become part of the best explanation of why the work has these aesthetically significant relations rather than others. This refinement does not reject creative control. It identifies the condition that control must satisfy if it is to ground strong creative attribution. A hundred prompt revisions may remain directive search, while one aesthetically decisive intervention may reorganize what the work is doing as a whole.

Creative-control theory thus supplies a necessary insight that generative agency retains: unpredictability does not eliminate human authorship. But it leaves unresolved the threshold at which continued human direction becomes aesthetically constitutive. The present argument begins precisely at that threshold. Control matters not because the human remains in charge in a general sense, but because particular exercises of control become consequential to the identity of the developing work.

Distributed and Entangled Agency: Relation without Attribution

Distributed and relational accounts correct the opposite tendency. Goodfellow places AI-mediated art on a spectrum of distributed authorship and emphasizes the changing role of the artist as translator or curator; Dzhimova and Tigre Moura use an Actor-Network perspective to describe feedback among human and non-human agencies, calculation, randomness, and technical affordances; Bomba and De Angeli formulate entangled authorship through relations among artists, data, algorithms, and other material-discursive conditions [13,14]. These accounts are valuable because they reject both the sovereign-human model and the fiction that the generative system is merely a passive tool.

The three positions should nevertheless be distinguished. Goodfellow primarily pluralizes the location of authorship within AI-mediated production. Dzhimova and Tigre Moura distribute agency ontologically and procedurally across a network of actants and feedback relations. Bomba and De Angeli go further by treating agency and authorship as relationally constituted through intra-action, mutual prompting, data practices, and transformational encounters. What unites them is not a single theory of authorship but a common insistence that the process cannot be adequately described as a one-way transmission of a fully formed human intention.

Their common difficulty appears when causal distribution is converted into creative attribution. A dataset, model architecture, interface, programmer, prompt, user, and generated contingency may all be indispensable to how a work comes about, but indispensability does not mean that each contribution has the same aesthetic status. Co-production therefore does not automatically entail co-creation at every level. The more successfully agency is distributed, the more urgently a theory needs a principle for weighting contributions: which factors merely enable production, which constrain it, and which become constitutive of the work's aesthetic identity?

Generative agency accepts the distributed account of production but rejects an undifferentiated distribution of creative responsibility. Its claim is not that the human

Copyright© Choi B.

secretly remains the sole cause of the work. The claim is that attribution requires a second-order judgment within distributed production: generated and infrastructural contributions must be distinguished from the human judgments through which some possibilities are recognized, redirected, committed to, integrated, and ultimately assumed as part of an artistic achievement. Relational production and differentiated attribution are therefore compatible rather than contradictory.

The Missing Attributional Threshold

The comparison now reveals a common structure. Intentionalist theories explain why creativity cannot be accidental, but intention can remain too general to account for the particular aesthetic achievement. Transformational theories explain how novelty and significant change can emerge, but transformation can be supplied by the model. Critical-integration accounts explain why human creativity may migrate toward selection and incorporation, but integration can remain curatorial. Creative-control theory explains how human direction can survive unpredictability, but direction can remain supervisory. Distributed and entangled theories explain why production is relational, but relationality alone does not weight contributions. These are not interchangeable failures. Each theory identifies a necessary dimension of the problem while stopping one step before differentiated attribution.

This suggests that the missing concept must satisfy two requirements simultaneously. It must be permissive enough to acknowledge that artistic production can be materially and computationally distributed, contingent, and only partially anticipated. Yet it must be discriminating enough to explain why not every causal participant, act of selection, or exercise of control constitutes the same creative achievement. The question is therefore no longer simply whether the human intended, transformed, controlled, or participated in the process. It is whether the human's aesthetic judgments became sufficiently consequential to what the work became.

The decisive opposition is between creative control and entangled agency. Creative control risks moving too quickly from continuing human direction to human authorship; entangled agency risks moving too quickly from distributed causal participation to distributed creative status. Generative agency occupies the missing middle. It preserves the first theory's concern with an attributable human trajectory while accepting the second theory's insistence that the trajectory develops through machine-supplied possibilities that exceed unilateral human determination. Its task is to identify when distributed possibility becomes differentiated artistic achievement.

This is why generative agency should not be understood as a seventh theory added to the literature. It is a threshold account generated by the unresolved relation among the existing ones. From Gaut and Stokes it retains purposive, non-accidental agency; from Boden it retains the significance of transformation; from Sarkar it retains critical integration; from Liu and Huang it retains process-wide direction; from distributed and entangled accounts it retains relational production. But each retained insight is made more demanding by an attributional qualification: intention must orient judgment, transformation must be artistically attributable, integration must become formative, control must be aesthetically consequential, and distributed production must permit differentiated responsibility.

Account	Locus of creativity	Insight retained	Attributional gap
Gaut / Stokes	Purposive, non-accidental agency	Creativity cannot be mere accident	Intention may not explain the particular aesthetic achievement
Boden	Transformation of a conceptual space	Novelty can emerge through transformation	Transformation can be machine-supplied
Sarkar	Reuse and critical integration	Reuse/randomness can belong to creativity	Integration can remain curatorial or managerial
Liu & Huang	Process-wide creative control	Unpredictability need not defeat human direction	Direction can remain supervision rather than constitution
Goodfellow / Dzhimova–Tigre Moura	Distributed human-technical agency	Production is not sovereign-human causation	Distribution does not weight contributions
Bomba & De Angeli	Relational / entangled agency and authorship	Human and AI recursively transform the process	Entanglement does not identify constitutive contribution

Table 1: From Existing Accounts to the Attributional Threshold.

Copyright© Choi B.

The theoretical transition can therefore be stated in one sentence: generation becomes human creation not when the human is merely present, intentional, controlling, or situated within an entangled process, but when human aesthetic judgments become constitutive of the developing identity of the work. Section 5 will develop this transition through five dimensions—aesthetic intentionality, consequential control, iterative responsiveness, selective commitment, and integrative responsibility—each of which answers a specific gap identified in the comparison above.

Generative Indeterminacy: From Constraint to Artistic Uptake

Generative-image production differs from conventional execution because an instruction can constrain an output without determining many of its expressive particulars. A prompt such as 'a solitary figure standing beside a ruined building at dusk' establishes a semantic field and perhaps an atmosphere, but it does not ordinarily determine the exact posture of the figure, geometry of the ruin, distribution of visual weight, lighting relation, texture, or chromatic organization. The prompt–image gap is therefore not merely a difference between a few words and many pixels. It is a gap of expressive determination.

This gap persists even when prompting becomes technically sophisticated. Longer prompts, negative prompts, reference images, seeds, guidance parameters, or control modules may narrow the possibility space, but narrowing is not identical with determining the aesthetic relations through which a particular image succeeds. The U.S. Copyright Office's analysis of prompts is useful on this limited point: instructions can communicate what the user wants while leaving the system to determine expressive details [6]. The aesthetic conclusion is not that prompts are non-creative, but that semantic specification should not be confused with expressive authorship.

The prompt–image gap should therefore be understood as a gap of expressive determination rather than a simple mismatch between the number of words and pixels. A short instruction can be expressively decisive if it establishes a rule that organizes the work, while a long instruction can remain expressively weak if it merely accumulates descriptors that the model realizes in ways the user does not meaningfully evaluate. What matters is not how much information the prompt contains but which aesthetically significant relations it determines and how the user responds to those it leaves open. This also explains why advances in prompt adherence will modify but not eliminate the philosophical problem. More faithful execution can increase human determination of some features while leaving other achievements to the model.

Indeterminacy is also structural to probabilistic generation. Diffusion systems generate outputs through learned distributions and iterative denoising procedures rather than by executing a fully articulated human plan [15,16]. The important aesthetic consequence is not that the system is 'random.' Generated outputs are constrained by prompts, model architecture, training, and sampling conditions. The more accurate relation is constraint plus contingency: human input shapes a field of possibility within which particular formal configurations remain underdetermined.

Artistic practice has long incorporated chance, material resistance, improvisation, delegation, and rule-governed indeterminacy, so unpredictability cannot by itself defeat creativity. But neither does the mere use of contingency establish it. The decisive issue is what the artist does with what emerges. An unexpected feature may simply be accepted because it is attractive, or it may be recognized as a problem, tested against alternatives, developed, suppressed, repeated, or made to redirect the work. Unexpectedness can enrich creation without automatically transferring authorship of the unexpected result to the person who encounters it.

The relevant comparison with historical chance practices is thus not that AI is "just another form of chance." The stronger point is that contingency has never been incompatible with authorship when it is embedded in a practice of uptake. John Cage's procedures, photographic timing, material accidents, and improvisational encounters differ radically from one another, yet each demonstrates that artists can create without controlling every causal antecedent. Generative AI adds a new difficulty because contingency arrives already aesthetically articulated. The user does not merely receive a stain, noise, or accidental event but may receive a composition that appears finished. The more complete the supplied solution, the more important it becomes to distinguish recognizing its quality from becoming responsible for its artistic organization.

This distinction also explains why prompt engineering cannot be evaluated by quantity. One carefully structured prompt can belong to a genuinely conceptual procedure; hundreds of prompt revisions can amount to little more than search optimization. The relevant unit is not the isolated prompt but the prompt–generation–judgment trajectory. Once an output becomes a reason for revising the next decision, prompting moves from simple direction toward responsive formation.

The distinction between terminal and formative selection identifies the point at which generative indeterminacy becomes philosophically productive rather than merely problematic. If an image is selected because it is the most

Copyright© Choi B.

attractive among alternatives, the act may display taste while leaving the internal achievement largely where it arose. If selecting one possibility commits the artist to preserve a spatial relation, revise a figure, alter a palette, reject later outputs that weaken an emerging tension, or reorganize a series around the selected feature, the contingency has acquired a normative role in the work's development. The question raised by generative AI is therefore not whether contingency can coexist with creation—it plainly can—but what structure of judgment converts received contingency into attributable artistic formation. That question requires the positive theory developed next.

Generative Agency: From Contingency to Attributable Creation

Generative agency is the capacity to exercise aesthetically consequential judgment within an indeterminate generative process such that the resulting work becomes intelligibly attributable to a trajectory of artistic decisions rather than merely to the productive capacities of the generative system. The concept is designed to answer the attributional gap identified in Section 3. It accepts that generation may be distributed and that important formal possibilities may originate outside the human's prior intention, but it asks when those possibilities become organized through a human trajectory of artistic judgment.

From Causal Contribution to Aesthetic Constitution

The starting distinction is among causal, directive, and constitutive contribution. Causal contribution explains whether the human made a difference to the existence of the output; entering a prompt is already causal in this sense. Directive contribution constrains the possibility space through subject, style, parameters, references, or rejection. Constitutive contribution is more demanding: human judgments help explain why aesthetically significant features or relations have the form and function they do in this particular work. Generative agency is principally concerned with this third level.

The distinction is also object-relative. A human contribution may be constitutive of a procedural project, a series, or an installation without being constitutive of every internal formal relation in an individual generated image. Conversely, a localized intervention can become constitutive of an image even when the larger workflow remains heavily dependent on model-supplied material. The attribution question must therefore specify both the object of the creativity claim and the level at which the human judgment is said to explain that object's identity. This prevents participation in a project from being treated as automatic authorship of all of its components.

This three-level distinction condenses the comparative result of Section 3 more accurately than a simple opposition between human and machine causation. Intentionalist accounts distinguish purposive activity from mere accident; creative-control theory shows how human direction can persist across an unpredictable process; distributed and entangled accounts show why production cannot be reduced to a single sovereign agent. Yet none of these propositions by itself identifies the constitutive level. The threshold of creation therefore cannot be inferred from how many human acts occur, how technically elaborate the workflow is, or how indispensable the human is to initiating production. It turns on whether human aesthetic judgment becomes part of the best explanation of the work's significant organization.

Constitution should not be misunderstood as exclusive determination. In collaborative and generative practices, more than one source can be constitutive of different aspects of the same work. The claim of generative agency is correspondingly modest but discriminating: the human need not originate every relevant possibility, but the human must do more than cause, supervise, or approve its appearance. The following subsection therefore asks how aesthetic constitution is built across time—through an orientation that generates reasons, interventions that make those reasons effective, responses that revise the problem, commitments that constrain later choices, and integration that makes the resulting configuration answerable as a work.

The Structure of Generative Agency

Generative agency should not be understood as five independent criteria added together after the fact. Its dimensions describe a temporal-logical movement through which a merely causal relation to generation can become aesthetically constitutive: orientation gives the process reasons; consequential intervention makes some of those reasons effective in the developing form; responsiveness allows unforeseen results to revise the problem; selective commitment converts chosen contingencies into constraints on later decisions; and integrative responsibility makes the resulting configuration answerable as an artistic achievement. The sequence is therefore orientation → intervention → response → commitment → integration. Aesthetic intentionality is its first moment. It is weaker than a complete mental blueprint and stronger than a generic intention to make an attractive image: an artist may begin with a problem, tension, atmosphere, relation, or procedural commitment without knowing the final solution, provided that this orientation can generate reasons within the process. This preserves Stokes's refusal to make creativity depend on intending the precise result in advance, while requiring

Copyright© Choi B.

a stronger account of how unforeseen results become attributable within an ongoing practice [12].

The difference can be seen when an unexpected relation appears. A user who merely likes the result may have a preference, but not yet an aesthetic orientation capable of explaining the work. An artist who recognizes that the relation bears on an unresolved formal problem and allows that recognition to govern subsequent choices exhibits a stronger form of agency. Aesthetic intentionality does not predetermine contingency; it makes contingency intelligible as material for judgment. In this way the theory avoids both the excessive demand of prior specification and the excessive permissiveness of mere purposive initiation.

Consequential control answers the limitation of creative-control theory. Liu and Huang are right that human control may persist across an unpredictable process, but control becomes evidence of creation only when it makes a non-trivial aesthetic difference. The relevant counterfactual is not whether the process would have occurred without the human, but whether the work's significant organization would plausibly have remained the same without the human decision at issue. Fifty prompt changes may amount to search; one intervention may become constitutive if it reorganizes the work's spatial, chromatic, expressive, or conceptual relations [7].

This requirement changes the object of evaluation from operational command to aesthetic consequence. A user can determine when generation begins and ends, which outputs are discarded, and what broad result is desired while still relying on the model for the solution that gives the image its force. Consequential control therefore does not ask whether the human remained 'in control' in a managerial sense. It asks what aesthetically important difference the human's control made. This is the precise point at which generative agency is more demanding than creative control while retaining its insight that unpredictability need not destroy authorship. Yet consequence alone is still insufficient if the artist's criteria remain fixed throughout the process. Generative creation becomes more distinctive when what emerges can alter the reasons governing what happens next. Consequential control must therefore open into responsive judgment rather than terminate in successful execution.

Responsive judgment is the next moment of the same structure. Bomba and De Angeli are right that advanced generative practice often involves mutual prompting: outputs alter human intentions, which in turn reshape later inputs [8]. Yet interaction becomes artistically significant only when the sequence acquires normative continuity. The relevant structure is not generation–generation–generation–selection, but generation–judgment–redirection–generation–further judgment, in which earlier decisions create reasons that constrain what counts as an adequate later response.

Responsiveness therefore distinguishes genuine development from repeated search. A hundred outputs can be generated under one unchanged criterion, whereas three iterations can embody substantial creative development if each changes the artist's understanding of the work. Generated contingency becomes more than an external stimulus when it modifies the criteria governing subsequent decisions. This is how a distributed interaction can acquire an attributable human trajectory without denying that the system contributed possibilities that the human did not originate.

Responsiveness becomes formative through selective commitment. Sarkar is right that AI can shift human creativity toward evaluating, curating, and integrating generated material, but selection and integration do not automatically transfer the internal achievement of a selected output to the selector [10]. A collector selects, a curator selects, and an editor integrates material produced by others. What matters is whether a selected possibility ceases to be merely preferred and becomes a reason that reorganizes subsequent decisions.

Commitment is therefore stronger than preference. If an artist preserves an unexpected spatial relation, rejects later outputs that weaken it, adjusts other elements around it, or reorganizes a series because of it, the selected feature has become a normative constraint within the work. Terminal selection ends a search; formative commitment establishes a reason for future decisions. This distinction also explains why a single selection can sometimes be more artistically consequential than choosing one image from thousands: the issue is not the quantity of options but the role the choice acquires in the developing structure.

The Attributional Threshold

The sequence culminates in integrative responsibility, and it is here that generative agency becomes an attributional threshold rather than merely a description of process. Distributed production remains real: the model, dataset, interface, and other conditions participate in what comes about. The question is which generated and inherited possibilities the human ultimately assumes as part of an organized artistic achievement. Irvin's emphasis on responsibility for artistic objectives points toward this dimension, but generative AI requires a stronger formulation because significant formal relations may arise outside the initial objective [9]. The artist's defensible claim is therefore not 'I originated every feature here' but 'I am responsible for

Copyright© Choi B.

how these possibilities became this work.'

This formulation joins credit to answerability. The same integration that supports attribution also makes retained relations part of what the artist must be prepared to defend: why a resemblance was preserved, why an unexpected feature was made central, why one transformation rather than another was adopted. Responsibility is thus differentiated within distributed production rather than assigned exclusively by causal origin. The model can remain a genuine source of formal possibility while the human becomes responsible for the role those possibilities acquire in the final configuration.

The five dimensions are not an additive checklist. Aesthetic intentionality without consequence remains aspiration; control without consequence remains operation; iteration without responsiveness remains search; selection without commitment remains preference; and retrospective responsibility without formative involvement remains appropriation after the fact. Generative agency arises when these dimensions converge sufficiently for human judgment to become constitutive of the developing work.

The Threshold Principle can therefore be stated as follows: generation becomes human artistic creation when the human participant's aesthetically oriented judgments become sufficiently consequential to the developing identity of the work that the resulting configuration is better explained as the outcome of an attributable trajectory of artistic agency than as a merely selected product of the generative system. The threshold is crossed not when computational contingency disappears, but when contingency becomes answerable to artistic judgment.

Three diagnostic questions make the principle operational without turning it into a score. Counterfactually, would removing the human judgments at issue leave the work's significant aesthetic identity largely unchanged? Explanatorily, do those judgments account for why important relations were preserved, rejected, transformed, or integrated? Integratively, did generated contingencies become commitments governing later choices, or were they merely accepted because they happened to be attractive? Strong answers to these questions establish why generative agency is neither another name for human involvement nor a denial of distributed production. It is the attributional threshold within it.

Testing the Threshold: From Generation to Creation

The threshold is qualitative, but AI practice can still be described through a graduated model. The categories are ideal types rather than technological stages, and they do not measure labor, prompt count, or aesthetic quality. They track the changing role of human judgment in explaining what the work becomes.

Below the Threshold: Prompting, Direction, and Selection

Mechanical generation involves initiation and acceptance with minimal aesthetically constitutive agency. A user may enter a generic prompt and accept the first successful result. The user causes the image to appear and may recognize its quality, but recognition of an achievement is not identical with producing that achievement. Directed generation adds intentional constraint. Detailed prompting, reference images, or parameter control can establish a more determinate orientation, yet a system may still supply the decisive formal articulation. Determining the task is artistically relevant without necessarily determining its realization.

Directed generation occupies a distinct but still sub-threshold position. Here the user meaningfully constrains subject matter, style, atmosphere, or formal parameters, perhaps through sophisticated prompting or reference inputs. The process is no longer mechanically initiated, yet the distinction between determining the task and determining its realization remains. A director can specify what is wanted without being the author of every solution by which another agent realizes it. Directed generation becomes more strongly creative only when the user begins to treat emerging realizations not simply as successes or failures against a fixed request but as material that can reshape the work's developing identity.

Difficult cases are useful because the same technical workflow can fall on different sides of the threshold. A one-prompt image ordinarily provides weak evidence of human creative attribution when the prompt describes a desired result and the system supplies the aesthetically decisive relations. Yet one prompt can belong to a strong conceptual procedure if the artist has structured in advance the rule, conditions, and deliberate acceptance of contingency. The relevant distinction is not one prompt versus many prompts, but contingently received success versus artistically structured contingency.

Consider first the one-prompt case. Ordinarily, a single descriptive prompt followed by acceptance of the first impressive output supplies weak evidence of human creation at the level of the image, because the system realizes most of the expressive particulars. But the number one is not disqualifying. An artist may construct a conceptual procedure in which a fixed prompt, a rule of accepting the first output, and a historically or institutionally meaningful presentation

Copyright© Choi B.

constitute the artistic act. The creative achievement would then lie principally in the procedure and framing rather than in human determination of every visual feature. This example shows why the theory must always specify what object the claim of creativity concerns.

Extensive prompt engineering is similarly ambiguous. A user may spend hours refining instructions toward a fixed target, exercising sophisticated directive control while the system continues to supply the particular aesthetic solution. Another user may revise prompts because generated results have changed the artistic problem itself. The first case can remain advanced search; the second can become iterative responsiveness. Post-generation manual editing is therefore neither necessary nor sufficient for creation.

The same distinction applies to elaborate prompting. Fixed-target refinement can remain a sophisticated form of search: the user repeatedly adjusts language until the model approximates an already desired type of image. Output-responsive refinement is different. An unexpected result can reveal a problem or possibility the artist had not formulated, after which later prompts are designed not merely to get "closer" to an initial target but to develop the emerging relation. The second practice provides stronger evidence of generative agency because the process has acquired normative continuity. The work begins to generate reasons for its own development.

Repeated generation and curatorial selection expose the same issue. Selecting the best-looking image from thousands can demonstrate taste without transferring authorship of its internal formal achievements. Selection becomes formative when a retained feature guides subsequent decisions, later outputs are judged against it, and the choice becomes part of the developing identity of the work. Curatorial creativity may exist at the level of a sequence or installation even where the individual generated images remain weakly attributable to the curator as their creator.

Crossing the Threshold: Responsive and Transformative Practice

Crossing the threshold begins when outputs cease to function merely as candidates and become reasons within the process. Iterative generation is therefore not defined by repetition as such. Its defining structure is responsive revision: generated results modify the artist's understanding of the problem, which in turn changes what later generation is asked to do. This is stronger than repeated search because earlier judgments acquire normative force. A later output is not evaluated only against an unchanged target; it is evaluated against commitments that arose through the developing work.

The distinction can be stated counterfactually. In fixed-target search, replacing one generation with another that satisfies the same antecedent specification ordinarily leaves the human criterion unchanged. In responsive practice, by contrast, a particular emergent relation can alter what counts as success: an initially unwanted distortion may become necessary to the work's tension, or an apparently successful composition may be rejected because it undermines a newly clarified conceptual relation. The work begins to acquire a history in which later judgments are intelligible only through earlier encounters with contingency.

Image-to-image transformation makes attribution especially difficult because visual change can be extensive while human agency remains thin. A model can transform a photograph radically in response to a minimal instruction. Stronger human agency appears when the artist determines what relations should be preserved, displaced, intensified, or suppressed and makes the source relation part of a developing artistic problem. This also shows why visual transformation and human creative attribution must be kept separate.

For this reason, the magnitude of visible transformation is a poor proxy for the strength of human agency. A single mask, crop, or compositional decision can be aesthetically constitutive if it reorganizes the relations through which the image functions, whereas extensive automated variation can remain procedurally dramatic but artistically thin. The relevant question is whether transformation is merely something the system performs on material or something the artist makes answer to a developing aesthetic problem. In the latter case, responsiveness and transformation converge: generated difference becomes part of an attributable trajectory rather than merely a sequence of altered outputs.

The threshold is therefore crossed neither at the first moment of iteration nor at the first substantial alteration. It is crossed when a sufficiently continuous relation emerges among judgment, revision, and commitment that the work's organization can no longer be explained adequately as the fortunate survival of model-supplied solutions. This is why a short but normatively connected process can exhibit stronger generative agency than a much longer workflow dominated by trial, rejection, and technical optimization.

At the Threshold: Level-Specific Attribution and Aesthetic Success

Generative creation names the level at which aesthetic intentionality, consequential control, iterative responsiveness, selective commitment, and integrative responsibility converge sufficiently for human creative attribution. The machine contribution need not disappear.

Copyright© Choi B.

Indeed, the strongest case may involve possibilities the human did not initially foresee. Human agency arises through their artistic incorporation. The progression can therefore be described in terms of the status of contingency: received, constrained, interrogated and redirected, reorganized, and finally integrated and assumed.

Curatorial selection is especially revealing because curation can itself be creative without making the curator the author of each selected object's internal form. The same possibility exists in AI workflows. A user who assembles a sequence of generated images may exercise strong creative agency at the level of sequencing, juxtaposition, installation, or conceptual framing even when agency over each image remains weak. Generative agency therefore supports level-specific attribution rather than an all-or-nothing verdict. The question should be: creative with respect to what—the internal image, the selection, the series, the installation, or the procedural project?

At this level the model's contribution does not disappear. Indeed, a work can count as generative creation precisely because important possibilities were not fully anticipated by the artist. What changes is their status within the practice. They are no longer merely supplied possibilities that happened to survive selection; they have become elements whose preservation, alteration, relation, and significance are explained through the artist's developing judgments. The graduated model therefore measures neither "how much AI" nor "how much human labor" is present. It measures the degree to which generated possibility becomes organized within an attributable trajectory of aesthetic agency.

Derivation does not literally end at the final stage. Human art remains dependent upon antecedents, and AI-assisted work remains computationally and culturally mediated. What changes is the relation to those conditions. Derivation ceases to be the dominant explanation of the human participant's relation to the work when generated and inherited possibilities are integrated into an attributable structure of aesthetic judgment. The distinction between derivation and creation is therefore not a distinction between two kinds of material but between two relations to material.

Sustained human–AI iterative practice provides the strongest positive case. An artist begins with a provisional orientation; an unexpected generated relation redirects the project; later results are rejected or retained according to newly emerging demands; prompts, references, masks, edits, and selections progressively integrate the unforeseen possibility. The final work is neither the execution of a fully formed human plan nor merely a successful discovery supplied by the model. It is better explained as an integrated trajectory of judgment in which machine-generated contingency has become artistically answerable.

The strongest cases are those in which the criteria of judgment themselves develop through interaction. The artist no longer asks only "Which output is best?" but increasingly "What does this work now require?" That change marks a transition from search to formation. An emergent feature becomes a commitment; later generations are judged in relation to it; contradictions are deliberately introduced or removed; and the final configuration records a history of decisions that cannot be explained merely by the model's capacity to generate attractive alternatives. Here distributed production and differentiated creative responsibility coexist rather than compete.

Level-specific attribution leads to the most important limiting case. An aesthetically successful image can fail to constitute a substantial human creative achievement. A first output can be beautiful, original, and culturally important while the user has done little more than request and recognize it. Conversely, a process can involve substantial human generative agency and produce a mediocre work. Aesthetic merit and creative attribution are therefore independent. This is not an exception to the threshold theory but one of its strongest tests, because it prevents the quality of the output from being used as retrospective evidence of the user's agency.

This conclusion may seem severe, but relaxing it would make aesthetic success do the work of a theory of authorship. If a model improves, the same minimal human input may suddenly produce much better images. It would be strange to conclude that the user's creativity increased simply because the system became more capable. The theory therefore treats quality as evidence about the work, not about the location of the achievement. An excellent generated image may be aesthetically valuable while the user's creative contribution remains slight; a difficult, imperfect work may embody substantial human agency. Creativity and aesthetic value intersect, but neither can serve as a proxy for the other.

Normative Limits of Generative Agency

Creative Transformation Does Not Erase Derivation

Once generative agency has been established, a different question begins: what follows from creative attribution for the work's derivative and normative relations? The answer must be deliberately limited. Strong human creative agency does not make a work non-derivative, and visible transformation does not by itself legitimate the relation to an antecedent. Similarity can identify a dependence requiring explanation, but the explanation may involve quotation, appropriation,

Copyright© Choi B.

convention, influence, coincidence, or machine-mediated statistical dependence. Generative agency therefore evaluates what the human made of a derivative relation; it does not erase that relation.

The source material also matters to the derivative relation without automatically determining creative agency. If a user supplies a personally authored photograph, the input contains human-authored expression, but subsequent model-generated transformations may still introduce aesthetically decisive features not attributable to the user. Conversely, working from another artist's image may involve strong human generative agency while simultaneously raising serious questions of appropriation or infringement. The axes remain independent: source legitimacy, degree of derivation, and human creative achievement must each be assessed on their own terms before being brought into a final evaluation.

Substantial resemblance to an identifiable artist or work adds a normative dimension but does not alter the threshold. A user can deliberately request stylistic imitation while exercising little creative agency, or can make resemblance itself the object of a critical practice by analyzing, exaggerating, or reorganizing the model's stereotyped reconstruction of a style. Similarity therefore does not settle creativity. Nor does creativity settle plagiarism or infringement: a work may involve strong generative agency and still raise serious questions about appropriation. This separation between creative achievement and the normative status of appropriation is consistent with the broader literature on reuse and appropriation [10].

Transformation is equally insufficient. Three senses must be distinguished: output transformation, procedural transformation, and artistically attributable transformation. The first describes visible difference; the second describes what the generative process does; the third asks whether human aesthetic judgment explains why the antecedent was transformed in these particular ways and what those changes are doing within the work. Only the third directly supports strong human creative attribution.

Creative Agency Does Not Require Exclusive Authorship

Creative agency also need not coincide with conventional or exclusive authorship. Distributed production can contain differentiated creative responsibility, and accounts of distributed or entangled authorship are right to resist the assumption that one sovereign human source must explain every aesthetically salient feature [13,14]. A human can exercise strong generative agency while acknowledging that the model supplied important formal possibilities. Conversely, institutional attribution of a work to a named artist does not prove that every aesthetically significant feature is that artist's achievement. The object of the creativity claim must therefore be specified: an individual image, a procedure, a sequence, an installation, a generative system, or a larger conceptual project.

This distinction requires separating at least three senses that ordinary discussion often compresses into the word authorship. Creative attribution concerns whose judgment explains an artistic achievement; conventional authorship concerns whom a practice recognizes or names as the maker of a work; legal authorship concerns a rights-bearing status defined by positive law. These categories can overlap, but they need not coincide. Generative agency addresses the first. It can therefore support a strong aesthetic attribution at one level of a project without claiming that every causal contribution has become human-authored, and without predetermining the legal question considered in the next subsection.

This is where the present theory differs from both exclusive humanism and undifferentiated co-authorship. A work may be co-produced by human and model while creative responsibility remains asymmetrical and level-specific. The model may supply formal possibilities the human did not invent; the human may organize those possibilities into a project whose aesthetic identity depends upon judgment, commitment, and integration. Calling the result simply "human-authored" may conceal the machine contribution, while calling every contributor a co-author may conceal the differentiated structure of achievement. Generative agency offers a middle vocabulary: distributed production with attributable, and potentially nonexclusive, creative responsibility.

Nor does nonexclusive creative agency imply symmetrical co-authorship. The fact that two sources are constitutive does not show that they are constitutive in the same manner or at the same level. A model may supply a formal possibility that the human could not have predicted, while the human may determine how that possibility is related to other elements, whether it is retained, and what role it comes to play within a larger project. Differentiated attribution is therefore preferable to both extremes: it avoids claiming exclusive human origination where none exists, but it also avoids treating every enabling or generative contribution as equivalent to an attributable artistic achievement.

Creative Attribution does not Determine Legal Permissibility

The boundary with copyright law is especially important. The U.S. Copyright Office asks whether sufficient human-

Copyright© Choi B.

authored expression can be identified for protection, while fair-use and derivative-work doctrines pursue different legal purposes [6]. Aesthetic transformation cannot simply be converted into a conclusion about legal permissibility. Nor does legal non-infringement establish creativity. Generative agency may help describe what the human actually contributed, but it is not a copyrightability test, a fair-use defense, or a plagiarism exemption.

Legal categories should therefore be used as points of comparison rather than imported as aesthetic criteria. Copyright systems require administrable thresholds tied to protectable expression, ownership, and statutory doctrine. Aesthetic attribution asks a different question about achievement. A human contribution can be aesthetically constitutive without mapping neatly onto separately protectable expression, and a legally protectable arrangement or modification can exist without supporting a strong claim that the human created the work as a whole. Likewise, a use may display substantial creative agency and still infringe another's rights. Keeping these domains distinct protects both the aesthetic theory and the legal analysis from conceptual overreach.

Integrative responsibility nevertheless has a normative implication. A human who claims generated contingencies as part of an artistic achievement cannot coherently claim credit for unexpected machine-generated successes while treating knowingly retained derivative relations as wholly external whenever attribution becomes inconvenient. Generative agency is therefore not only a theory of credit; it is also a theory of answerability. The artist's strongest claim is not 'I caused every feature in this image,' but 'I stand behind what these generated possibilities have become within my work.'

Conclusion: Creation as Generative Agency

Generative AI makes aesthetic production easier while making creative attribution harder. Existing theories illuminate important parts of the problem: Boden explains modes of transformation; Gaut and Stokes show why creativity requires non-accidental agency; Sarkar demonstrates why reuse and critical integration complicate simple plagiarism narratives; Liu and Huang preserve human control under generative unpredictability; Goodfellow, Dzhimova and Tigre Moura, and Bomba and De Angeli show why AI practice is distributed and relational. The remaining question is more specific: when does human judgment become constitutive of what the generated work becomes?

The broader contribution of the argument is therefore not a new vocabulary for praising AI-assisted art, but a more discriminating account of attribution. Intentionalist theories prevent creativity from collapsing into accident; transformational theories explain how novelty can arise from antecedent structures; creative-control theory preserves human direction under unpredictability; and distributed or entangled theories correctly reject the fiction of sovereign human production. Yet none of these insights alone identifies when the particular aesthetic achievement becomes attributable to the human participant. Generative agency names that missing relation: the point at which aesthetic judgment becomes constitutive of what generated possibilities are made to become.

The threshold is therefore not crossed by prompting alone, by a large number of iterations, by selection alone, by manual editing, or by transformation in the abstract. It is crossed when aesthetic intentionality, consequential control, iterative responsiveness, selective commitment, and integrative responsibility converge sufficiently that the work is better explained through an attributable trajectory of human artistic decisions than as a merely selected product of the generative system. Computational contingency may remain essential; what changes is that contingency becomes answerable to artistic judgment.

This account has broader implications for aesthetics. AI art reveals that creativity has never been identical with direct execution, complete prior intention, novelty, or labor. Artists have always worked through materials, conventions, collaborators, accidents, and procedures that exceed full control. What matters is not freedom from antecedence or contingency but the formation of a trajectory in which possibilities become reasons for further artistic action. AI makes that structure unusually visible because it introduces a powerful source of externally generated possibility directly into the creative process.

Seen from this perspective, AI art is not an isolated anomaly but an unusually clear test case for a longstanding aesthetic problem. Artists have always worked through conditions they did not create and with materials that resist or exceed intention. What generative AI changes is the degree to which the external condition can itself propose highly articulated formal solutions. The resulting pressure is philosophically useful: it forces aesthetics to distinguish execution from judgment, novelty from achievement, and causal contribution from creative attribution with greater precision than many pre-AI practices required.

The theory also avoids defining creativity biologically. If an agent were capable of forming aesthetic orientations, responding to emergent results, making consequential selections and transformations, and integrating those decisions into a sustained practice of responsibility, the structure described here would be relevant regardless of substrate. The immediate argument is narrower: in present

Copyright© Choi B.

human–AI practice, the human remains the participant who can make generated possibilities answerable to reasons, commitments, revisions, and aesthetic responsibility.

The final claim can therefore be stated simply. Creation does not begin where derivation ends. It begins where derivation becomes answerable to generative agency. Generative AI has not made the distinction between generation and creation obsolete; it has made the need to draw that distinction more precise.

References

1. Boden MA (2004) The Creative Mind: Myths and Mechanisms. 2nd(Edn.), Routledge, London.
2. Wiggins GA (2006) A Preliminary Framework for Description, Analysis and Comparison of Creative Systems. Knowledge-Based Systems 19(7): 449-458.
3. Hertzmann A (2018) Can Computers Create Art? Arts 7(2): 18.
4. Krauss RE (1981) The Originality of the Avant-Garde: A Postmodernist Repetition. October 18: 47-66.
5. Crimp D (1980) The Photographic Activity of Postmodernism. Vol 15: 91-101.
6. U.S. Copyright Office (2025) Copyright and Artificial Intelligence, Part 2: Copyrightability. U.S. Copyright Office, Washington, DC.
7. Liu W, Huang W (2025) Authorship in Human-AI Collaborative Creation: A Creative Control Theory Perspective. Computer Law & Security Review 57: 106139.
8. Bomba F, De Angeli A (2025) Agency and Authorship in AI Art: Transformational Practices for Epistemic Troubles. International Journal of Human-Computer Studies 205: 103652.
9. Irvin S (2005) Appropriation and Authorship in Contemporary Art. British Journal of Aesthetics 45(2): 123-137.
10. Sarkar A (2023) Exploring Perspectives on the Impact of Artificial Intelligence on the Creativity of Knowledge Work: Beyond Mechanised Plagiarism and Stochastic Parrots. Proceedings of CHIWORK 2023, article 13, pp. 5–7, §2.4 (reuse, randomness, and the limits of the plagiarism analogy).
11. Gaut B (2010) The Philosophy of Creativity. Philosophy Compass 5(12): 1034-1046.
12. Stokes D (2011) Minimally Creative Thought. Metaphilosophy 42(5): 658-681.
13. Goodfellow P (2024) The Distributed Authorship of Art in the Age of AI. Arts 13(5): 149.
14. Dzhimova M, Tigre Moura F (2024) Calculated Randomness, Control and Creation: Artistic Agency in the Age of Artificial Intelligence. Arts 13(5): 152.
15. Ho J, Jain A, Abbeel P (2020) Denoising Diffusion Probabilistic Models. Advances in Neural Information Processing Systems 33: 6840-6851.
16. Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B (2022) High-Resolution Image Synthesis with Latent Diffusion Models. Proceedings of CVPR 2022: 10684-10695.

Copyright© Choi B.