



Is Learning Stealing? Tabula Rasa, Human Creativity, and Generative AI

Raymond Ayivor

Scripps College of Communication, Ohio University, Athens, Ohio, USA

ABSTRACT

Creativity has historically emerged through interaction with inherited knowledge, cultural traditions, artistic conventions, language, technologies, and prior creative works rather than from an intellectual *tabula rasa*. Yet the rapid expansion of generative artificial intelligence has intensified debates over whether computational learning constitutes legitimate learning, imitation, or unauthorized appropriation. This article examines these concerns by challenging conceptions of human originality that treat creativity as creation independent of preceding influences. It compares human learning and generative AI as processes involving exposure, pattern recognition, recombination, adaptation, and transformation while emphasizing critical differences in embodiment, intentionality, agency, scale, speed, memory, and technological mediation. The analysis then considers how these differences complicate conventional understandings of authorship and intellectual ownership when models are trained on extensive collections of cultural and creative material. Rather than equating human cognition with machine learning, the article reframes the debate around the conditions under which learning becomes ethically or legally problematic. It argues that meaningful evaluation of generative AI should therefore prioritize transformation, attribution, consent, provenance, ownership, social legitimacy, and originality when determining boundaries between learning, creative influence, derivation, and appropriation.

Keywords: Generative Artificial Intelligence; Human Creativity; Machine Learning; Originality; Intellectual Property; Creative Appropriation

1. INTRODUCTION: THE CRISIS OF ORIGINALITY IN THE AGE OF GENERATIVE AI

1.1 Generative AI and the Contemporary Creativity Debate

Generative artificial intelligence has moved rapidly from experimental technology into mainstream creative practice, transforming the production of text, images, music, film, software, advertising, and design [1]. Systems capable of generating sophisticated outputs from relatively simple prompts have consequently intensified longstanding debates surrounding authorship, originality, and creative ownership [2]. Rather than creating from an informational vacuum, generative models develop their capabilities by learning statistical relationships, structures, and patterns from extensive collections of pre-existing material [3]. This dependence has generated competing interpretations of AI-assisted creativity. One perspective conceptualizes model training as technologically mediated learning from accumulated human culture, bearing limited parallels to how human creators acquire knowledge through exposure to earlier works [4]. Conversely, critics emphasize the unprecedented scale at which computational systems may process creative materials without creators' knowledge, consent, attribution, or compensation [5]. The controversy therefore extends beyond whether generative AI simply "copies" and encompasses difficult distinctions among learning, influence, transformation, derivation, ownership, and appropriation [6].

1.2 The Tabula Rasa Problem and the Myth of Creation from Nothing

Underlying this controversy is an implicit *tabula rasa* conception of creativity: the assumption that authentic originality should somehow emerge from a blank intellectual slate [7]. Human creative production, however, rarely develops within an informational vacuum. Writers acquire language before manipulating narrative conventions and literary genres; filmmakers learn established practices of framing, montage, sound, and visual storytelling; artists encounter existing techniques, symbols, and aesthetic traditions; while programmers develop expertise through established languages, algorithms, libraries, architectures, and previously demonstrated solutions [8]. Creative development consequently depends upon cultural inheritance while simultaneously enabling individuals to modify, combine, reinterpret, challenge, or transcend inherited forms [9]. This relationship complicates claims that dependence upon previous works is itself sufficient evidence of creative theft. If exposure to existing cultural material automatically rendered subsequent production appropriative, ordinary processes of education, apprenticeship, artistic influence, and intellectual development would become conceptually difficult to distinguish from unauthorized copying [5]. Originality must therefore involve more than complete independence from creative precedent [6].

1.3 Aim, Central Argument, and Analytical Direction

This article argues that framing the controversy as “learning versus stealing” creates an inadequate conceptual binary. Human creators and generative systems both depend upon antecedent information [3], but acknowledging this structural similarity does not establish ethical, cognitive, technological, or legal equivalence between them [4]. Human learning occurs through lived experience, embodiment, intentionality, social relationships, memory, judgment, and cultural participation, whereas generative AI operates through computational architectures capable of processing information at scales and speeds unavailable to individual human creators [1]. Accordingly, the relevant question is not simply whether previous creative works contribute to subsequent production, but under what conditions such dependence constitutes legitimate learning, transformative influence, problematic derivation, or appropriation [7]. The analysis therefore examines interconnected dimensions of transformation, agency, embodiment, scale, memory, attribution, consent, provenance, ownership, and originality [8]. These dimensions provide a framework for distinguishing shared dependence on prior knowledge from the conditions governing its use [9]. The next section therefore interrogates originality through *tabula rasa*, cultural inheritance, and creative transformation.


2. FROM TABULA RASA TO CULTURAL INHERITANCE: RETHINKING ORIGINALITY

2.1 The Tabula Rasa as a Model of Creative Independence

The metaphor of *tabula rasa*, or the blank slate, provides a useful conceptual starting point for examining assumptions embedded within contemporary debates about creativity and generative AI [7]. Its significance here extends beyond its historical philosophical meaning because it exposes an implicit standard against which human and machine creativity are frequently compared: the belief that genuinely original creation should emerge independently of substantial prior cultural influence [10]. Such a standard produces what may be described as absolute originality, in which novelty presupposes creation without meaningful antecedents. Yet this conception inadequately represents how cultural production occurs. Creative agents inherit languages, symbols, techniques, genres, aesthetic conventions, and bodies of knowledge that shape the possibilities available for subsequent expression [8]. A more productive concept is therefore situated originality, whereby novelty arises within existing cultural and intellectual environments rather than outside them [11]. Situated originality does not deny individual creativity; instead, it locates creative contribution in the capacity to reorganize, reinterpret, extend, or challenge inherited materials in ways that produce meaningful difference [12]. Originality thus becomes relational rather than synonymous with cultural independence.

2.2 Creativity as Cultural Accumulation and Recombination

Creativity can consequently be understood as a cumulative social process through which cultural resources are transmitted, internalized, reorganized, and extended across individuals and generations [13]. Artists often begin through imitation or apprenticeship before developing distinctive techniques, while writers encounter genre conventions, narrative structures, quotation, parody, adaptation, and intertextual relationships long before producing mature works. Similarly, filmmakers acquire established visual grammars, designers experiment with inherited aesthetic conventions, and programmers receive technical instruction through existing languages, documentation, algorithms, and code practices [8]. These activities reveal a recurring developmental sequence: exposure → internalization → transformation. Exposure supplies the cultural and technical resources from which creative understanding develops; internalization enables conventions, techniques, and patterns to become incorporated into individual competence; and transformation permits those resources to be reorganized in response to new purposes, experiences, audiences, or contexts [14]. Recombination is therefore not necessarily evidence of creative deficiency. A work may possess recognizable antecedents while simultaneously producing new meanings, functions, arrangements, or aesthetic experiences. The decisive issue is not merely whether prior influence exists, since influence is pervasive, but how inherited material is transformed and what degree of creative distance separates the resulting expression from its antecedents [15].

Figure 1. Conceptual Continuum from Cultural Exposure to Creative Transformation**Figure 1. Conceptual Continuum from Cultural Exposure to Creative Transformation**

2.3 Influence, Imitation, Derivation, and Appropriation

Recognizing creativity as culturally situated requires careful distinctions among influence, inspiration, imitation, derivation, reproduction, plagiarism, and appropriation. Influence describes the capacity of earlier cultural encounters to shape subsequent creative decisions, often without direct reproduction. Inspiration similarly identifies a generative stimulus from which substantially different expression may develop. Imitation involves closer replication of identifiable stylistic or structural characteristics and may function as education, homage, experimentation, or deliberate emulation depending upon context [9]. Derivation involves a more identifiable relationship between a new work and antecedent material, whereas reproduction approaches direct duplication. Plagiarism introduces an additional dimension because the problem concerns presenting another's expressive or intellectual contribution as one's own without appropriate acknowledgement [6]. Appropriation is broader still, potentially involving questions of authority, cultural context, ownership, power, and economic benefit [15]. Transformation consequently provides an important boundary but cannot independently resolve every dispute. A highly transformed use may still raise questions concerning consent or provenance, while attribution does not automatically legitimize every reproduction. Permission, attribution, ownership, social context, commercial substitution, and economic consequences must therefore complement transformation when evaluating contested uses [12]. Once originality is understood as culturally situated rather than creation *ex nihilo*, attention can turn to how humans acquire and transform creative knowledge.

3. HUMAN LEARNING AND THE FORMATION OF CREATIVE CAPACITY

3.1 Learning Through Exposure, Observation, and Imitation

Human creative competence develops through sustained encounters with existing works, established practices, and communities of knowledge rather than through spontaneous creation independent of precedent [12]. In literature, writers learn vocabulary, narrative structure, characterization, genre, and rhetorical technique through reading and experimenting with forms encountered in previous texts. Visual artists similarly develop technical competence by observing composition, perspective, color, and established artistic movements, while filmmakers acquire knowledge of cinematography, editing, narrative pacing, and visual grammar through viewing and production practice [13]. Musical development involves listening, repetition, performance, interpretation, and engagement with inherited rhythmic, harmonic, and compositional conventions [16]. Designers study existing objects and aesthetic solutions before modifying them to address different functional or cultural requirements. Programming provides an equally instructive example: programmers learn through documentation, existing algorithms, code examples, libraries, frameworks, and established design patterns [17]. Imitation within these processes does not necessarily constitute an endpoint. It frequently functions as a developmental mechanism through which techniques are acquired before being adapted, combined, or rejected [18]. Consequently, resemblance between an emerging creator's practice and earlier influences does not demonstrate an absence of creativity. Creative competence develops progressively as inherited practices become resources for increasingly independent expression.

3.2 Memory, Experience, Embodiment, and Context

Exposure alone, however, provides an incomplete account of human creative learning. Cultural information is encountered by embodied individuals whose memories are intertwined with emotion, sensory experience, interpersonal relationships, education, social position, cultural membership, and changing life circumstances [14]. Human memory is therefore not merely a repository from which previously encountered material is retrieved unchanged. Remembering can involve selection, reconstruction, association, forgetting, reinterpretation, and integration with subsequent experience [19]. These characteristics become especially significant when comparisons are drawn between human learning and generative AI. A novelist reading previous literature, for example, encounters those works through a biography that may include particular relationships, geographical environments, cultural identities, aspirations, disappointments, and historical circumstances. These experiences can subsequently alter how literary influences are interpreted and transformed. Likewise, filmmakers and artists make intentional judgments concerning which conventions to preserve, reject, criticize, or subvert in response to particular social contexts [15]. Human creative learning is thus embedded within embodied and socially situated experience rather than reducible to extracting regularities from informational inputs [20]. This distinction does not establish that computational systems cannot generate novel combinations; instead, it demonstrates why similarity in dependence on antecedent information should not be mistaken for equivalence in the processes through which that information acquires creative significance.

3.3 Transformation and the Emergence of Human Originality

Human originality can consequently be understood as the production of meaningful difference from inherited materials, rather than complete independence from them [12]. Transformation occurs when creators intentionally select elements of existing cultural resources and subject them to reinterpretation, synthesis, critique, contextual displacement, modification, or combination. A writer may adopt an established genre while challenging its conventional narrative perspective; a filmmaker may reinterpret a familiar story through a different historical or cultural setting; and an artist may appropriate an established visual language specifically to criticize the assumptions associated with it [15]. Musicians similarly transform inherited structures through arrangement, performance, sampling, improvisation, and stylistic synthesis, while programmers combine established computational principles to address problems for which previous implementations are inadequate [17]. Importantly, transformation is not synonymous with superficial alteration. Creative originality becomes more persuasive where inherited material contributes to an expression possessing distinguishable purpose, arrangement, meaning, function, or context [18]. Intentionality also matters because creators can consciously acknowledge, criticize, conceal, celebrate, or depart from their influences [20]. Human originality therefore exists along a continuum rather than as an absolute opposition between dependence and independence. This conception permits prior influence and genuine novelty to coexist, while preserving distinctions between transformative creativity and substantially reproductive practices [19].

Table 1. Forms of Dependence on Prior Works in Human Creative Practice

Creative domain	Sources of prior learning	Learning mechanism	Transformative activity	Potential originality
Literature	Books, language, genres	Reading and practice	Reinterpretation and narrative construction	New expression
Visual arts	Images, techniques, movements	Observation and apprenticeship	Stylistic transformation	New visual work
Film	Existing films and visual grammar	Viewing and production practice	Narrative and aesthetic recombination	New audiovisual work
Design	Existing designs and conventions	Observation and experimentation	Functional and aesthetic redesign	New design
Programming	Code, documentation, algorithms	Study and implementation	Adaptation and problem-solving	New computational solution

The dependence of human creativity on antecedent culture therefore establishes an important basis for comparison with artificial intelligence, but it does not demonstrate equivalence between human cognition and computational learning. The relevant issue becomes whether generative systems merely reproduce training material or develop representations capable of supporting novel combinations, and how those computational mechanisms differ from embodied human learning. Section 4 therefore moves from the human creative process to the technical foundations of generative AI, examining how training data are converted into learned model parameters and subsequently into generated outputs.

4. HOW GENERATIVE AI LEARNS: FROM TRAINING DATA TO GENERATED OUTPUT

4.1 Training Data, Pattern Learning, and Model Formation

Generative artificial intelligence begins with training data, which may comprise text, images, audio, video, source code, or other machine-readable information assembled into large training corpora [18]. During training, computational models repeatedly process examples and adjust numerical parameters to reduce errors between predicted and observed patterns [21]. In large language models, for example, learning involves estimating relationships among tokens within different linguistic contexts, whereas image-generation systems learn statistical representations connecting visual structures, textual descriptions, and latent features [22]. Through optimization, models develop internal representations capturing recurring associations, structures, and probability distributions within training material [23]. This process requires distinguishing three concepts frequently collapsed in public debates: the training dataset, the trained model, and the generated output. The dataset contains materials presented during training; the trained model consists principally of learned parameters and computational structures shaped through optimization; and generated outputs arise when the trained system responds to subsequent input or conditioning [24]. These distinctions complicate descriptions of generative AI as merely a database that retrieves stored creative works. Nevertheless, learning statistical representations does not guarantee that models never reproduce training material. Memorization and recognizable reproduction can occur under particular conditions [25]. The appropriate analysis must therefore distinguish generalized pattern learning from instances where generated content exhibits substantial correspondence with specific training examples.

4.2 Generation, Probability, and Computational Recombination

Once training has produced a parameterized model, generation occurs by using learned representations to construct outputs in response to prompts or other conditioning information [21]. In language generation, a prompt establishes an initial context from which the model estimates probability distributions over possible subsequent tokens, repeatedly selecting or sampling continuations according to its learned parameters [22]. Image-generative systems employ different architectures and procedures, but many operate through representations in latent spaces where learned relationships between textual and visual features guide the progressive synthesis of an image [23]. Consequently, a generated output may combine patterns associated with numerous training examples without corresponding directly to a single identifiable source. This process can be described as computational recombination, provided that the term does not imply conscious selection, imagination, or human-like understanding [24]. The narrower comparison with human creativity is nevertheless analytically useful. Human creators and computational models both produce outputs whose possibilities are influenced by information encountered previously [19]. The mechanisms remain fundamentally different: human transformation is mediated by cognition, experience, embodiment, intention, and social context, while generative systems operate through mathematical representations, probabilistic relationships, computational optimization, and conditioning [20]. Dependence upon antecedent information therefore constitutes a point of comparison, not evidence that human imagination and machine generation are equivalent.

Figure 2. Human Creative Learning and Generative-AI Training as Parallel but Non-Equivalent Processes

Figure 2. Conceptual comparison of human creative learning and generative-AI training. Both pathways depend on antecedent information, but the mechanisms through which previous information contributes to subsequent outputs differ substantially.

4.3 Scale, Speed, Memory, and Computational Difference

The ethical significance of generative AI becomes clearer when attention shifts from abstract similarities in learning to differences in scale, speed, memory, embodiment, reproducibility, and agency. Individual human creators encounter finite portions of cultural production over biologically constrained lifetimes. Generative systems, by contrast, can be trained using extraordinarily large datasets assembled from geographically dispersed digital sources and processed through automated computational infrastructure [21]. Once trained, models can generate large quantities of text, imagery, code, or other content within short periods, allowing creative production to be reproduced at industrial scale [24]. Machine memory also differs from human recollection. Human remembering is reconstructive, contextual, experiential, and susceptible to forgetting, whereas model parameters encode statistical relationships produced through optimization, although models may under some circumstances retain information sufficiently specific to enable memorization or reproduction [25]. Generative systems additionally lack the embodied biography through which human creators ordinarily interpret cultural encounters [20]. Questions of agency are consequently complicated because model outputs involve interactions among training-data selection, model architecture, developer decisions, system parameters, prompts, and user choices [23]. These distinctions mean that even where humans and machines both depend upon prior cultural material, the ethical implications of that dependence may differ significantly. The question is therefore not simply whether both “learn,” but which similarities illuminate creativity and which technological differences alter questions of responsibility, consent, attribution, and appropriation.

5. HUMAN LEARNING VERSUS MACHINE LEARNING: SIMILARITY WITHOUT EQUIVALENCE

5.1 Points of Structural Similarity — 170 words

Human creative learning and generative-AI training exhibit structural similarities without operating through equivalent cognitive processes. Both begin with **exposure to antecedent information** from which patterns, conventions, relationships, and recurring structures can influence subsequent production [23]. Human writers learn grammatical and narrative conventions from language and literature, just as generative language models acquire probabilistic relationships among tokens through training corpora [24]. Artists develop sensitivity to visual forms through observation, while image-generative models learn computational representations of recurring visual characteristics [25]. At a broader level, both processes involve forms of pattern acquisition, generalization, recombination, and production conditioned by previously encountered material [26]. These parallels challenge the simplistic proposition that dependence on previous works uniquely prevents machine-generated outputs from exhibiting novelty. Human creativity itself remains substantially dependent upon cultural inheritance [27]. However, structural similarity establishes only that neither process begins from a genuine informational *tabula rasa*. It does not demonstrate that computational pattern learning constitutes human-like understanding, nor that every generated output is original. Novelty must instead be evaluated through the relationship between antecedent information and resulting expression [28].

5.2 Points of Fundamental Difference — 210 words

The analogy becomes substantially weaker when the mechanisms and contexts of learning are examined. Human creativity develops through consciousness, embodiment, emotion, socialization, lived experience, intentionality, and culturally situated judgment [29]. Generative AI processes information computationally without establishing that the system possesses comparable subjective experience or socially embedded understanding [30]. Human memory is intertwined with biography and reconstructive experience, whereas model representations emerge through optimization of numerical parameters across training data [31]. Scale and speed introduce further differences: individuals encounter limited cultural material over finite lifetimes, while computational systems can process extensive datasets and generate outputs at volumes impossible for individual creators [32]. Accountability also differs. Human creators can ordinarily be identified as social and legal actors, whereas AI-generated content may implicate model developers, dataset constructors, platform operators, deployers, and users in different ways [33]. Economic organization further distinguishes the processes because generative systems can transform aggregated cultural material into commercially scalable services [34]. AI should therefore neither be anthropomorphized as simply another human learner nor dismissed as merely a mechanical copying device. A rigorous comparison must examine what models computationally learn while preserving distinctions concerning consciousness, embodiment, agency, responsibility, and institutional power [35].

Table 2. Comparative Dimensions of Human Creative Learning and Generative-AI Learning

Dimension	Human learning	Generative AI	Why the difference matters
Prior exposure	Cultural and social experience	Training datasets	Conditions governing source acquisition differ
Scale	Biologically and temporally constrained	Potentially massive	Alters questions of extraction and institutional power
Embodiment	Embodied and experiential	Computational	Influences how information acquires meaning
Intentionality	Connected to human purposes and judgment	System- and user-dependent	Complicates attribution of agency
Transformation	Cognitive and experiential	Computational	Transformation occurs through different mechanisms
Attribution	Socially expected in relevant contexts	Often technically difficult	Creates provenance and recognition concerns
Consent	Context dependent	Dependent upon dataset acquisition and governance	Creates ethical and legal questions
Output volume	Constrained by human capacities	Highly scalable	May produce substantial cultural and economic consequences

5.3 Why Analogy Alone Cannot Resolve the Debate

Structural comparison is therefore analytically useful but normatively insufficient. Demonstrating that humans and generative systems both depend upon antecedent information does not establish that the conditions under which they acquire or use that information are ethically equivalent [36]. The human-

learning analogy can reveal the weakness of treating all dependence on existing works as inherently unoriginal, but it cannot independently resolve questions concerning authorization, ownership, attribution, or commercial exploitation. Conversely, emphasizing differences between biological cognition and computational optimization does not demonstrate that machine learning is necessarily appropriate [37]. Mechanistic difference describes *how* learning occurs; it does not by itself determine whether a particular use of cultural material is legitimate. Ethical and legal assessment therefore requires movement from descriptive comparison toward questions concerning the circumstances surrounding acquisition, transformation, generation, and subsequent use [38]. This distinction prevents the debate from collapsing into either technological exceptionalism or anthropomorphic equivalence. If similarity cannot establish legitimacy and difference cannot establish wrongdoing, the decisive inquiry becomes when and under what conditions dependence upon previous creative works becomes problematic. The following section therefore turns to transformation, consent, attribution, provenance, ownership, and economic value.

6. WHEN DOES LEARNING BECOME STEALING?

6.1 Transformation and the Boundary of Derivation

The boundary between learning and appropriation becomes particularly difficult to define when attention shifts from training processes to the relationship between generated outputs and antecedent works. Transformation exists along a continuum rather than as a binary condition [32]. At one end, a system may learn broad stylistic, linguistic, compositional, or technical regularities without generating material readily traceable to an individual source. Further along the continuum, generated content may retain recognizable structural, aesthetic, or expressive characteristics while reorganizing them into substantially different configurations [33]. Close imitation presents greater difficulty because outputs can reproduce distinctive features associated with particular creators even without duplicating an entire work [34]. At the opposite end, substantial reproduction may preserve identifiable expressive elements with limited meaningful alteration [35]. Transformation therefore provides an important analytical criterion for distinguishing influence from derivation, but the amount of modification alone cannot settle every dispute. Purpose, context, recognizability, source relationships, attribution, and subsequent use also shape whether computational transformation represents creative recombination or problematic appropriation [36].

6.2 Consent and the Politics of Training Data

Transformation addresses outputs, whereas consent directs attention upstream toward how creative material enters training systems. Large-scale datasets may incorporate publicly accessible text, images, code, music, and other cultural materials produced under widely different expectations regarding their subsequent use [37]. The central distinction is between technical accessibility and ethical authorization. A work being technically retrievable from a website does not necessarily demonstrate that its creator knowingly agreed to its incorporation into model training, particularly where such use was neither anticipated nor meaningfully negotiable when the work was published [38]. Consent is nevertheless more complex than a simple permission-versus-prohibition distinction. Relevant considerations include licensing conditions, contractual arrangements, platform terms, public-domain status, creator expectations, and mechanisms for opting into or out of particular uses [39]. The issue therefore concerns not merely whether developers *can* acquire material computationally, but what obligations accompany acquisition at scale. This shifts analysis from the accessibility of individual works toward questions of informational power, creator autonomy, dataset governance, and whether meaningful control remains possible within increasingly automated cultural infrastructures [40].

6.3 Attribution, Provenance, and Recognition

Questions of attribution become difficult when model behavior reflects statistical learning across extensive corpora rather than transparent one-to-one relationships between individual training works and particular outputs [34]. A generated image, paragraph, or musical sequence may reflect patterns distributed across numerous sources, making conventional source identification considerably more difficult than attribution in quotation or direct adaptation. Provenance can address part of this problem by documenting where training materials originated, how they were acquired, what permissions or licenses governed them, and how datasets changed during model development [38]. Attribution addresses a related but distinct question: whether creators whose identifiable contributions materially influence subsequent outputs should receive acknowledgement [39]. Neither mechanism provides a complete solution. Dataset-level provenance may establish the origin of training material without demonstrating precisely how an individual work affected model parameters, while output attribution may be difficult where influence is diffuse [40]. Nevertheless, greater transparency can make relationships among creators, datasets, model developers, computational systems, and generated outputs more visible, enabling more informed assessments of recognition, responsibility, and legitimacy.

6.4 Ownership, Value Extraction, and Creative Labor

The debate becomes broader still when generative AI is examined through the political economy of creative labor. Focusing exclusively on whether a particular output resembles an identifiable source can obscure how commercial models derive capability from aggregated bodies of cultural production [36]. Writers, illustrators, photographers, musicians, programmers, and other creators collectively produce the informational environment from which model developers may obtain training resources, while the resulting systems can subsequently generate commercially valuable services at considerable scale [37]. This creates questions about value extraction and distribution even where individual outputs cannot be traced straightforwardly to particular works. The ethical concern may therefore arise from institutional relationships between those who produce cultural resources and those capable of

computationally aggregating and monetizing them [39]. Scale is consequential because practices that appear minor at the level of individual learning can acquire different economic significance when automated across extensive datasets and millions of generated outputs [40]. Ownership, compensation, bargaining power, market substitution, and recognition consequently belong alongside transformation when assessing whether computational learning remains culturally productive or becomes extractive.

Figure 3. A Multidimensional Framework for Evaluating Learning, Transformation, and Appropriation




Figure 3. A Multidimensional Framework for Evaluating Learning, Transformation, and Appropriation

Figure 3. Multidimensional framework for assessing generative use of prior creative works. The framework rejects a single-variable test of computational appropriation. Transformation addresses the relationship between source and output; consent and provenance concern acquisition; attribution and originality concern creative relationships and recognition; while ownership, scale, and economic impact capture the institutional consequences of computational use.

Taken together, these dimensions move the debate beyond the rhetorically powerful but analytically restrictive question of whether AI “steals.” The more consequential inquiry concerns how prior creative material is obtained, transformed, attributed, governed, and economically utilized, and whether those conditions support legitimate cultural learning or problematic appropriation.

7. REFRAMING ORIGINALITY AND RESPONSIBILITY IN GENERATIVE CULTURE

7.1 From “Did It Learn?” to “Under What Conditions Did It Learn?”

The preceding analysis suggests that asking whether generative AI learns from existing creative works establishes little by itself, because cultural production has always depended upon antecedent knowledge [35]. The more consequential question is: under what technological, ethical, economic, and legal conditions did that learning occur, and what relationship does the resulting output maintain with its sources? This reframing separates the inevitability of prior influence from the legitimacy of particular methods of acquiring and using creative material [36]. A multidimensional assessment should therefore examine how training materials were obtained, whether relevant permissions or licenses existed, what provenance information remains available, and whether generated outputs reproduce recognizable expression or achieve substantial transformation [37]. It should additionally consider scale, commercial purpose, attribution, market consequences, and opportunities available to creators to exercise meaningful control [38]. Such an approach rejects both the assumption that learning from prior works is inherently illegitimate and the opposing assumption that describing model training as “learning” automatically resolves concerns about appropriation [39]. Prior learning can consequently be accepted as fundamental to cultural and computational production while individual practices remain subject to scrutiny according to their acquisition conditions, transformative characteristics, institutional context, and effects on creators [40].

7.2 Toward Responsible Computational Creativity

Responsible computational creativity therefore requires governance mechanisms that address both model development and the cultural resources supporting it. Dataset transparency can clarify the categories and origins of training materials, while stronger provenance mechanisms can document acquisition pathways, licensing conditions, and subsequent dataset transformations [36]. Meaningful approaches to creator consent should similarly provide practical choices concerning computational uses rather than relying exclusively on the technical availability of online content [37]. Responsible licensing arrangements can further establish clearer expectations regarding permissible training, commercial exploitation, attribution, and compensation where appropriate [38]. At the output level, safeguards should reduce substantial reproduction of identifiable works and provide mechanisms for investigating outputs that closely correspond with protected creative expression [39]. Attribution systems may also improve recognition where meaningful source relationships can be established, although diffuse model learning makes conventional attribution difficult in many circumstances. These measures should not reduce generative-AI governance to a binary choice between unrestricted technological development and absolute creator control [40]. A sustainable framework must accommodate innovation, cultural exchange, technological experimentation, creator interests, and public access to knowledge while recognizing asymmetries in computational and economic power [35]. Clearer distinctions among stylistic influence, transformative generation, close imitation, and substantial reproduction can consequently support generative culture without presuming either that all machine learning is theft or that all computational reuse is legitimate.

8. CONCLUSION

Is learning stealing? Dependence on prior knowledge alone cannot make it so. Neither human creativity nor generative AI emerges from a genuine *tabula rasa*; both depend upon antecedent information that shapes the possibilities for subsequent production. Nevertheless, this shared dependence should not obscure fundamental differences between human and computational learning in embodiment, agency, intentionality, lived experience, scale, speed, and social positioning. Consequently, the ethical significance of generative AI cannot be determined simply by establishing that models learn from existing creative works. The more meaningful boundary lies in examining how source materials are acquired and transformed, whether outputs reproduce protected expression, and how attribution, consent, provenance, ownership, and economic value are addressed. Generative-AI debates should therefore abandon the impossible standard of creation from nothing. Instead, attention should focus on establishing conditions under which technologically mediated learning remains culturally productive, ethically defensible, and appropriately respectful of the human creative labor from which computational systems ultimately learn.

REFERENCE

1. Wheeler JC. Path to Singularity: How Technology Will Challenge the Future of Humanity. Simon and Schuster; 2024 Nov 19.
2. Chomsky N. Rules and representations. Behavioral and brain sciences. 1980 Mar;3(1):1-5.
3. Johnston JW. The Construction of Reality in an AI: A Review. arXiv preprint arXiv:2302.05448. 2023 Feb 3.
4. Wasserman-Rozen H, Chagal-Feferkorn K. Lawyers are from Mars, data scientists are from Venus: promoting responsible AI on Earth. In Research Handbook on the Law of Artificial Intelligence 2025 Jun 19 (pp. 177-193). Edward Elgar Publishing.
5. Schneijderberg C, Wieczorek O, Steinhardt I. The Handbook of Qualitative and Quantitative Content Analysis: Introduction to Classical, Digital, AI-supported, and Automated Data Analysis. Taylor & Francis; 2026 Jan 23.
6. Oquendo J. Euphemism in AI Discourse as Colonial Harm & Effect: Learnings from Subaltern Puerto Rico. Louisiana State University and Agricultural & Mechanical College; 2024.
7. Pires EB. The Automated Self: A Philosophical Study on the Evolution of Artificial Intelligence and Its Normative Challenges. Bloomsbury Publishing USA; 2024 Dec 17.
8. Robinson JC. Feallan Dynasty. In Artificial Intelligence: Ethics and the New World Order 2025 May 28 (pp. 173-222). Cham: Springer Nature Switzerland.
9. Parkinson A, Dunning G. Improvising with Machines: A taxonomy of musical interactions. Organised Sound. 2024 Dec;29(3):340-51.
10. Saghafian S. Insight-driven problem solving: Analytics science to improve the world. Cambridge University Press; 2025 Oct 30.
11. Obunadike TC. Evaluating global tariff shocks on staple crop import dependency and national food security resilience systems. Int J Res Publ Rev. 2025 Jun;6(6):12423-12440. doi:10.55248/gengpi.6.0625.23103.
12. Bratton J. TECHNOLOGY, SOCIETAL CHANGE. Artificial Intelligence and Work: Transforming Work, Organizations, and Society in an Age of Insecurity. 2024 Dec 4:1.
13. Spence C. Necromantic Media. SubStance. 2026;55(2):97-114.
14. Savin-Baden M, Ball J. Higher Education in a Postdigital World: Learning Everywhere. Taylor & Francis; 2026 Sep 1.

15. Hanzelkova J. Generative Artifacts: Chinatown and an Ornamental Architecture of the Future. In *Arts* 2025 Nov 28 (Vol. 14, No. 6, p. 155). MDPI.
16. Barron L. AI and popular culture. Emerald Group Publishing; 2023 Apr 3.
17. YUNUS NS. A must-read for policy makers, academics, and financial market participants interested in understanding how policy makers respond to external shocks. *Artificial Intelligence* Also in this issue. 2023 Dec.
18. de Zeeuw D, Geil A. Algorithmic Abstraction in Photographic Portraiture. *Reconfiguring the portrait*. 2023 Aug 22:43.
19. Pohan JE, Ambarita R. Comparison of Metaphorical Language Styles in the Novel *Saman* by Ayu Utami and the Novel *Tabularasa* by Ratih Kumala. *Modality Journal: International Journal of Linguistics and Literature*. 2025 Dec 31;5(2):109-26.
20. Koch C. *Then I am myself the world: What consciousness is and how to expand it*. Hachette UK; 2024 May 7.
21. Abdullahi Bamidele Bashiru; The Social Resource Network (SRN): A Socio-Technical Framework for Circular Waste Matchmaking. *ACS Sustainable Resour. Manage.* 26 February 2026; 3 (2): 283–285. <https://doi.org/10.1021/acssusresmgt.5c00510>
22. Varghese DJ. Metaverse Liability and the Future of Tort Law. Available at SSRN 5043578. 2024 Sep 13.
23. Whittemore T. Beyond the black box: The case for a prize system to encourage artificial intelligence innovation. *Mich. St. L. Rev.*. 2024:249.
24. Andriandaza L, Marco A, Periola AA. Technical Advancements in Natural Language Processing. In *AI Models for Human and Societal Beneficial Technologies on Earth and in Space 2026* (pp. 6-76). CRC Press.
25. Louwerse MM. Understanding artificial minds through human minds: the psychology of artificial intelligence. Taylor & Francis; 2025 Jun 2.
26. Sapir E. Language in Humans, Hybrids. *Artificial Humanities: A Fictional Perspective on Language in AI*. 2025 Nov 25:144.
27. Alebiosu O. AI and cybersecurity governance and legal compliance challenges affecting digital energy infrastructure, smart grids, and critical utility resilience systems. *Int J Comput Appl Technol Res.* 2026;15(8):1-14. doi:10.7753/IJCATR1508.1001.
28. Campedelli GM. Artificial Intelligence and Criminology. *Annual Review of Criminology*. 2026 Sep 14:10.
29. Kweskin S. The End Is Still to Come: How the Law's Interaction with Digital Replicas and Derivative Digital Replicas Will Shape the Future. *Notre Dame Journal of Law, Ethics & Public Policy*. 2025 Jul 2;39(3):749.
30. Kwesi J, Cao J, Manchanda R, Emami-Naeini P. Exploring user security and privacy attitudes and concerns toward the use of {General-Purpose}{LLM} chatbots for mental health. In *34th USENIX security symposium (USENIX security 25)* 2025 (pp. 6007-6024).
31. de Berardinis J, Porcaro L, Merono-Penuela A, Cangelosi A, Buckley T. Towards responsible AI Music: An investigation of trustworthy features for creative systems. *arXiv preprint arXiv:2503.18814*. 2025 Mar 24.
32. Ajayi T. Comprehending the Epistemic Harms in Generative Artificial Intelligence. *Journal of Social Sciences and Humanities*. 2026 Jun 29;65(1):23-38.
33. Jovic A, Buchanan BJ. The Stories We Tell in the Age of the GenAI Essay: Putting the Human Back in the Humanities. In *Design, Delivery, and Evaluation of Professional Learning for University Staff 2026* (pp. 365-400). IGI Global Scientific Publishing.
34. Kommers P. AI and STEAM: The Power of Thematic Learning. In *AI for STEAM in Education: Thematic Learning in Practice 2026* May 14 (pp. 1-26). Cham: Springer Nature Switzerland.
35. Jääskeläinen P, Sanchez C, Holzapfel A. Anticipatory Technology Ethics Reflection By Eliciting Creative AI Imaginaries Through Fictional Research Abstracts. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* 2025 Jun 23 (pp. 125-136).
36. Jain N, Thomas A, Gupta V, Pralat E. Creativity, Innovation, and Knowledge Creation in the Era of Generative Artificial Intelligence: Challenging Human Intelligence. *Global business and organizational excellence*. 2026 May;45(4):421-34.
37. Williams P. Tears and rain—AI and authorship: Challenging concepts of the unique through unauthored rarity. *Journal of Illustration*. 2025 Apr 1;12(1):171-93.
38. Miller AI. *The artist in the machine: The world of AI-powered creativity*. Mit Press; 2019 Oct 1.
39. George Daramola. Assessing environmental impacts of rapid data centre expansion on energy consumption, water resources, and emissions. *Int. J. Res. Eng.* 2023;5(1):38-47. DOI: [10.33545/26648776.2023.v5.i1a.191](https://doi.org/10.33545/26648776.2023.v5.i1a.191)
40. Klowait N, Erofeeva M. The presentation of self in the age of ChatGPT. *Frontiers in Sociology*. 2025 Aug 21;10:1614473.