

# Reconstructing Originality in Human–AI Collaboration under US Copyright

## Law

Terry Taorui Guan

### Abstract

US copyright law protects only ‘original works of authorship’ created by human beings. That premise is doctrinally settled and constitutionally grounded. But as generative AI becomes part of everyday creative workflows, the standard inquiry, specifically who exercised ‘ultimate creative control’ over the expressive elements, has proven increasingly under-inclusive. The US Copyright Office has rightly denied protection to fully autonomous machine outputs and have registered human contributions in hybrid projects. Yet in the common middle ground, where authors guide, edit, or curate AI-generated content, there is no administrable framework for finding and determining the extent of human authorship within complex, tool-mediated processes.

This Article proposes a structured, precedent-faithful solution: a Fourfold Framework of Originality—*Conception, Orchestration, Curation, and Inscription*. These four dimensions describe the roles that human authorship plays in hybrid works. Each is traceable in the work and supported by evidentiary records. The framework clarifies what applicants should register, what they must disclaim, and how to ensure that copyright protects what it has always protected: perceptible, original, human

expression, and nothing more. Without altering the law's threshold or expanding its subject matter, this approach provides an implementable rubric for registration, adjudication, and platform design in the age of AI-assisted creativity.

### **Keywords**

copyright; originality; authorship; human–AI collaboration; generative AI; copyright registration; US Copyright Office

## 1. Introduction

American copyright law is built on an established but under-specified premise: Only human beings are authors, and only ‘original works of authorship’ merit protection.<sup>1</sup> Recently, the US Copyright Office has recognised this using a crude proxy, ultimate creative control, that asks who determined the work’s expressive particulars.<sup>2</sup> The proxy worked tolerably when creative tools were transparent and deterministic, but it misfires across today’s iterative, stochastic, and collaborative workflows, especially where humans create with generative systems.

Recent examples illustrate the scope of the problem. The DC Circuit’s decision in *Thaler v. Perlmutter*<sup>3</sup> denies protection to fully autonomous machine outputs.

Similarly, the US Copyright Office’s guidance requires applicants to disclose, and disclaim, AI-generated material.<sup>4</sup> Those guardrails are sound at the poles. But in the hybrid middle, such as *Zarya of the Dawn*,<sup>5</sup> *Théâtre d’Opéra Spatial*,<sup>6</sup> *SURYAST*,<sup>7</sup> and countless everyday projects, the control proxy obscures rather than illuminates

---

<sup>1</sup> 17 USC § 102 (a).

<sup>2</sup> US Copyright Office, *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence* 88 Fed Reg 16190, 16192 (16 March 2023).

<sup>3</sup> *Thaler v Perlmutter* 130 F4th 1039, 1041 (DC Cir 2025).

<sup>4</sup> US Copyright Office (n 2) 16193.

<sup>5</sup> US Copyright Office, ‘Re: Zarya of the Dawn’ (21 February 2023) 1

<<https://www.copyright.gov/docs/zarya-of-the-dawn.pdf>> accessed 1 January 2026.

<sup>6</sup> US Copyright Review Board, ‘Re: Second Request for Reconsideration for Refusal to Register Théâtre D’opéra Spatial’ (5 September 2023) 1.

<sup>7</sup> US Copyright Review Board, ‘Re: Second Request for Reconsideration for Refusal to Register SURYAST’ (11 December 2023) 3, 8.

where human originality resides. It risks under-crediting the human conception, procedural steering, curation, and inscription that shape the finished work.

Scholars have tried to fix this, but each move leaves important gaps. Vision centred accounts, especially Jane Ginsburg and Luke Ali Budiardjo's conception and execution model, sensibly value planning, but risk crossing the idea/expression line when 'vision' is asserted without a demonstrable link to the expression fixed in the work.<sup>8</sup> Other commentators, including Robert Denicola, would recognise the initiating or directing user as author; but 'master' or art direction analogies fail when the human does not determine the work's expressive particulars, and cannot supply authorship where the system supplies the traditional elements of expression.<sup>9</sup>

Causation focused frameworks, such as Shyamkrishna Balganesh's account of 'authorial causation'<sup>10</sup> and Dan Burk's analysis of causation, volition and intention,<sup>11</sup> clarify why some inputs feel authorial but, without a workable taxonomy of which choices count and where to look in the work, risk becoming case by case intuition.

Finally, disclosure based regulatory proposals (eg Jacob Noti-Victor) and calls for *sui generis* protection for outputs generated by AI (eg Firas Massadeh and coauthors)

---

<sup>8</sup> Jane C Ginsburg and Luke Ali Budiardjo, 'Authors and Machines' (2019) 34 Berkeley Technology Law Journal 343, 346–47, 352; Pamela Samuelson, 'Allocating Ownership Rights in Computer Generated Works' (1986) 47 University of Pittsburgh Law Review 1185, 1201.

<sup>9</sup> Robert C Denicola, 'Ex Machina: Copyright Protection for Computer Generated Works' (2016) 69 Rutgers University Law Review 251, 251.

<sup>10</sup> Shyamkrishna Balganesh, 'Causing Copyright' (2017) 117 Columbia Law Review 1, 1, 53.

<sup>11</sup> Dan L Burk, 'Thirty-Six Views of Copyright Authorship, by Jackson Pollock' (2020) 58 Houston Law Review 263, 263, 275.

either sidestep authorship or risk overinclusive protection for machine determined details to which copyright has never applied.<sup>12</sup>

Unless we clarify what is protectable in AI-assisted works, the costs of uncertainty will continue to fall on creators and cultural markets. Platforms and distributors increasingly rely on disclosure, content-integrity, and monetization rules to manage synthetic or repetitive content; where a creator cannot clearly specify the scope of human authorship, the practical risk is not only copyright conflict but also reduced distribution or monetization under platform policies.<sup>13</sup> Licensing markets, meanwhile, are shifting toward risk-managed and indemnified channels: Getty Images barred third-party AI submissions in 2022 amid copyright uncertainty and later launched an in-house, indemnified generator, channelling demand into closed ecosystems that protect buyers but constrain creators.<sup>14</sup> Similar dynamics appear in book distribution. After a flood of low-quality AI titles, Amazon KDP capped daily uploads and now requires disclosure of AI-generated text or images—adding compliance burdens even when a work is largely human-authored.<sup>15</sup> The pattern is not a blanket prohibition of

---

<sup>12</sup> Jacob Noti-Victor, ‘Regulating Hidden AI Authorship’ (2025) 111 *Virginia Law Review* 139, 139; Firas Massadeh and others, ‘The Legal Protection of Artificial Intelligence Generated Work: The Argument for Sui Generis over Copyright’ (2024) 6(1) *Corporate Law and Governance Review* 49, 49.

<sup>13</sup> YouTube, ‘Disclosing use of altered or synthetic content’ (YouTube Help) <<https://support.google.com/youtube/answer/14328491>> accessed 2 January 2026 (“we require creators to disclose content that is meaningfully altered or synthetically generated”); YouTube, ‘YouTube channel monetization policies’ (YouTube Help) <<https://support.google.com/youtube/answer/1311392>> accessed 1 January 2026.

<sup>14</sup> James Vincent, ‘Getty Images bans AI-generated content over fears of legal challenges’ (*The Verge*, 21 September 2022) <<https://www.theverge.com/2022/9/21/23364696/getty-images-ai-ban-generated-artwork-illustration-copyright>> accessed 4 January 2026; Getty Images, ‘Getty Images Launches Commercially Safe Generative AI Offering’ (25 September 2023) <<https://newsroom.gettyimages.com/en/getty-images/getty-images-launches-commercially-safe-generative-ai-offering>> accessed 9 January 2026.

<sup>15</sup> Ella Creamer, ‘Amazon restricts authors from self-publishing more than three books a day after AI concerns’ (*The Guardian*, 20 September 2023) <<https://www.theguardian.com/books/2023/sep/20/amazon-restricts-authors-from-self-publishing-more-than-three>>

AI tools, but a default toward labels, limits, and exclusionary guardrails because decisionmakers lack a reliable way to identify which parts of a hybrid work are legally protectable. Without a clearer method, creators bear the cost in lost income, forfeited rights, and diminished trust.<sup>16</sup>

This Article offers a textually grounded, precedent-faithful solution: a fourfold Framework of Originality—*Conception, Orchestration, Curation, and Inscription*. Rather than ask the monolithic question ‘who controlled the details?’, the Framework locates human authorship in four *loci* that frequently appear in contemporary workflows and, crucially, are perceptible in the work: (1) Conception—the author’s specific expressive constraints and framing; (2) Orchestration—the documented procedure and iterative steering that predictably shape expression; (3) Curation—the selection, coordination, and arrangement of candidate materials; and (4) Inscription—the substantive edits, compositing, or layout decisions that fix the final expressive form of the work. The Framework traces and protects those increments that are both human and discernible. It thus reconciles *Feist*’s constant, low originality threshold<sup>17</sup> with *Burrow-Giles*’s focus on ‘original intellectual conceptions’,<sup>18</sup> honours *Thaler*’s human-authorship predicate, and implements § 103(b)’s instruction to confine rights to ‘material contributed by the author’.<sup>19</sup>

---

-books-a-day-after-ai-concerns> accessed 3 January 2026; Amazon, ‘Content Guidelines’ <[https://kdp.amazon.com/en\\_US/help/topic/G200672390](https://kdp.amazon.com/en_US/help/topic/G200672390)> accessed 2 January 2026.

<sup>16</sup> YouTube, ‘YouTube channel monetization policies’ (n 13); Amazon (n 15).

<sup>17</sup> *Feist Publications, Inc* 499 US 340, 345 (1991).

<sup>18</sup> *Burrow-Giles Lithographic Co v Sarony* 111 US 53, 58 (1884).

<sup>19</sup> 17 USC § 103(b).

This Article contributes to the literature in three respects. Practically, it yields an implementable registration rubric: structured applicant disclosures keyed to the four dimensions; an internal checklist for examiners; and standardized certificate annotations that publicly differentiate between what is protected (human increments) and what remains in the public domain (unedited machine detail). It guides creators in documenting and scoping claims *ex ante* and helps platforms to design intake, attribution, and provenance tools that align with those claims. Doctrinally, it unifies scattered strands into a single, administrable test. Theoretically, it replaces an all-or-nothing control narrative with a process-based account of authorship that is compatible with creativity scholarship and legal causation: It offers a stable vocabulary and an evidence-led inquiry that fits how people create today, in the age of human–AI collaboration.

Following the Introduction, Section 2 diagnoses the limits of the current originality doctrine in hybrid and tool-mediated workflows, showing how existing proxies like ‘control’ misfire in practice. Section 3 introduces this Article’s core contribution, the Fourfold Framework that identifies where human authorship appears in AI-assisted works and ties protection to perceptible, documented contributions. Section 4 shows how to apply the Framework in registration, licensing, and adjudication without altering existing law, thus offering creators, platforms, and courts a structured approach to protecting human authorship and preserving the public domain. Section 5 concludes.

## **2. The limits of the current originality doctrine in the face of AI and other nontraditional creativity**

### **2.1. Originality and human authorship in US law: the baseline**

Three propositions anchor US copyright law. First, a work is ‘original’ only if it is independently created and exhibits at least a minimal quantum of creativity—‘even a slight amount will suffice’.<sup>20</sup> Originality does not turn on novelty, effort or aesthetic merit, and copyright never extends to ideas, procedures or facts as such.<sup>21</sup> Second, protection attaches only to human authorship. Although neither the US Constitution nor the 1976 Act says ‘human’, courts have long read the terms ‘Authors’ and ‘Writings’ as referring to a human mind.<sup>22</sup> Third, originality must be expressed and fixed; the law does not protect abstract conceptions until an author has embodied them in concrete expressive form.<sup>23</sup>

These baseline commitments emerged well before machine learning systems, and have remained stable for more than a century. In *Burrow-Giles Lithographic Co. v. Sarony*, the Court protected a photograph because it embodied the photographer’s ‘original intellectual conceptions’—with proof of originality, intellectual production, thought, and conception—not because the camera ‘authored’ anything; the opinion even quotes an English authority that described the author as the ‘inventive or master

---

<sup>20</sup> *Feist Publications, Inc* (n 17) 345–46.

<sup>21</sup> *Ibid.* 345–46, 352–54; 17 USC § 102(b); *Bleistein v Donaldson Lithographing Co* 188 US 239, 251–52 (1903).

<sup>22</sup> US Constitution art I § 8 cl 8; *Trade-Mark Cases* 100 US 82, 94 (1879).

<sup>23</sup> 17 USC § 101.

mind’.<sup>24</sup> *Bleistein* cautioned courts against aesthetic gatekeeping, reinforcing the principle that the threshold of creativity is low and emphasizing that it is not a referendum on taste.<sup>25</sup> And in *Feist*, the Court distilled originality to its canonical formula—independent creation plus a modicum of creativity—while reiterating that facts (and, by statute, ideas and methods) remain free for all.<sup>26</sup>

On human authorship, both courts and the US Copyright Office are explicit: *The Trade-Mark Cases* framed copyrightable writings as the ‘fruits of intellectual labor’ founded in the ‘creative powers of the mind’,<sup>27</sup> and the Office’s Compendium states that it will register works created by a human being<sup>28</sup> while refusing works produced by a machine that ‘operates randomly or automatically’ without human creative input.<sup>29</sup> Consistent with that view, the Office declines claims to works that trace to animals or natural forces (eg, the monkey selfie).<sup>30</sup> The Ninth Circuit has held that nonhuman animals lack statutory standing under the Copyright Act.<sup>31</sup>

Modern digital tools did not unsettle these basics. The fundamental principles accommodate photography,<sup>32</sup> sound recording, and computer-assisted design<sup>33</sup> by identifying and protecting the human choices that shape expressive outcomes—the

---

<sup>24</sup> *Burrow-Giles Lithographic Co* (n 18) 58–60.

<sup>25</sup> *Bleistein* (n 21) 251–52.

<sup>26</sup> *Feist Publications, Inc* (n 17) 345.

<sup>27</sup> *Trade-Mark Cases* (n 22) 94.

<sup>28</sup> US Copyright Office, *Compendium of the US Copyright Office Practices* (3rd edn, 2021) § 306.

<sup>29</sup> *Ibid.* § 313.2.

<sup>30</sup> *Ibid.*

<sup>31</sup> *Naruto v Slater* 888 F3d 418, 426–27 (9th Cir 2018).

<sup>32</sup> *Ets-Hokin v Skyway Spirits Inc* 225 F3d 1068, 1073–74 (9th Cir 2000).

<sup>33</sup> *Meshwerks Inc v Toyota Motor Sales USA Inc* 528 F3d 1258, 1264–65, 1268–69 (10th Cir 2008).

choice of subject, framing, timing, selection among variants, editing, and arrangement—while treating the machine as an instrument.<sup>34</sup> The 1976 Act expressly limits originality in compilations and derivative works ‘only to the material contributed by the author’,<sup>35</sup> a textual foothold that protects human selection, coordination, arrangement, and editing but not the underlying materials (or tool-generated elements).

The established structure has faced challenges when it tries to accommodate generative AI. When a system autonomously determines the traditional elements of authorship—plot, phrasing, composition or pictorial details—no human has created those elements, and they are therefore not protected. The Office’s 2023 policy guidance instructs applicants to disclose AI-generated material and to remove it from any claim; registrable protection extends only to the human-authored contributions.<sup>36</sup> The prompts, even when elaborate, that guided the system are typically insufficient because users do not exercise ‘ultimate creative control’ over the specific expressive content that the AI system generates.<sup>37</sup>

This doctrine yields sensible results at the extremes. A fully human-crafted work—even if made with sophisticated software—is protected. A work produced entirely by a machine with no human creative input is not. But it provides only a blunt instrument for the increasingly common middle ground in which human contributions

---

<sup>34</sup> US Copyright Office (n 2) 16192.

<sup>35</sup> 17 USC § 101; 17 USC § 103(b).

<sup>36</sup> US Copyright Office (n 2) 16193.

<sup>37</sup> *Ibid.* 16192.

are indirect (conceptual guidance), process-oriented (method and parameterization), curation (selection and arrangement), or editorial (post-generation inscription).

Existing law can register some of these contributions, especially selection and arrangement,<sup>38</sup> but it lacks an analytic vocabulary that maps what, exactly, the human author contributed to the final expression.<sup>39</sup>

## **2.2. The struggle to apply the originality doctrine to AI-assisted and nontraditional works**

These difficulties recur in two settings: first, AI-registration disputes that force decisionmakers to separate human-authored expression from AI-generated material in hybrid works; and second, long-standing low-control artistic practices, where authorship is real but distributed across constrained or stochastic processes.

### **2.2.1. AI-generated and AI-assisted works**

In *Thaler v. Perlmutter*, the DC Circuit affirmed the Copyright Office’s refusal to register an image where Stephen Thaler had ‘listed the “Author” ... as the ‘Creativity Machine’” and described it as ‘created autonomously by [a] machine’.<sup>40</sup> The court held that ‘the Creativity Machine cannot be the recognized author’ because the Copyright Act ‘requires all eligible work to be authored in the first instance by a

---

<sup>38</sup> *Ibid.*; US Copyright Office (n 5) 1.

<sup>39</sup> US Copyright Office, *Copyright and Artificial Intelligence: Part 2—Copyrightability* (Report of the Register of Copyrights, January 2025) iii

<<https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf>> accessed 8 January 2026 (Part 2 Report).

<sup>40</sup> *Thaler* (n 3) 1043.

human being’, and on that basis it affirmed the denial.<sup>41</sup> Reading the Act as a whole, the panel explained that the statute ‘makes no sense if an “author” is not a human being’, pointing to provisions that hinge on an author’s ‘children’, ‘widow’, ‘domicile’, ‘mens rea’, ‘nationality’, and even ‘life’ and ‘death’.<sup>42</sup> The court expressly declined to rule on whether the Constitution independently requires human authorship, and did not consider Thaler’s late-raised claim that he was the author by virtue of making and using the system, noting it was waived.<sup>43</sup> As for works in which human and AI contributions mix, the panel emphasized that such ‘line-drawing disagreements’ were ‘neither here nor there in this case’, because Thaler identified the machine as ‘the sole author’.<sup>44</sup> The holding is narrow and orthodox: A work produced entirely by an autonomous system lacks a human author, and copyright ‘subsists ... in original works of authorship’ only—so no copyright attaches.<sup>45</sup>

But far more commonly, cases involve hybrids—works combining human authorship with AI material. Since March 2023, the US Copyright Office has instructed applicants to disclose any AI content, and, in the Author Created field, to ‘provide a brief explanation of the human author’s contribution’ in the Author Created field;<sup>46</sup> do not list an AI system as an author, and should ‘explicitly exclude’ AI content ‘that is more than de minimis should be explicitly excluded’<sup>47</sup> in the Material Excluded

---

<sup>41</sup> Ibid. 1044.

<sup>42</sup> Ibid. 1048.

<sup>43</sup> Ibid. 1051-52.

<sup>44</sup> Ibid. 1050.

<sup>45</sup> Ibid. 1051.

<sup>46</sup> US Copyright Office (n 2) 16193.

<sup>47</sup> Ibid.

field. Where ‘[t]he traditional elements of authorship[] are determined and executed by the technology—not the human user’, that material is ‘not the product of human authorship’ and ‘must be disclaimed’;<sup>48</sup> in contrast, applicants may register human selection, coordination and arrangement of AI material or human modifications ‘to such a degree’ that they become protectable.<sup>49</sup> The guidance compares prompts to ‘instructions to a commissioned artist’: users ‘do not exercise ultimate creative control’, because ‘the machine determines how those instructions are implemented’ and, for example, ‘will decide the rhyming pattern, the words in each line, and the structure’.<sup>50</sup>

Three decisions illustrate how the Office’s guidance operates. The first case is *Zarya of the Dawn*. In this case, although Kristina Kashtanova initially obtained a registration for her 18-page comic, the Office revisited the claim after learning that she had produced the images using Midjourney, a subscription-based text-to-image generative AI service. In a letter issued in early 2023, the Office concluded that Kashtanova is the author of the text and the ‘selection, coordination, and arrangement of the Work’s written and visual elements’, but that ‘the images in the Work that were generated by the Midjourney technology are not the product of human authorship’.<sup>51</sup> The Office therefore cancelled the original certificate and issued a new one ‘covering only the expressive material that she created’.<sup>52</sup> It also noted that any sufficiently

---

<sup>48</sup> Ibid. 16192.

<sup>49</sup> Ibid. 16193.

<sup>50</sup> Ibid. 16192.

<sup>51</sup> US Copyright Office (n 5) 1.

<sup>52</sup> Ibid.

creative edits to the intermediate AI images ‘could provide human authorship and would not be excluded from the new registration certificate’.<sup>53</sup>

In *Théâtre d’Opéra Spatial*, Jason Allen sought to register his Colorado State Fair-winning image. The record shows he ‘input numerous revisions and text prompts at least 624 times’ and then used Photoshop and Gigapixel AI to refine and upscale the image.<sup>54</sup> The Review Board affirmed the Copyright Office’s refusal because Allen declined to disclaim the AI-generated portions: ‘Because Mr. Allen is unwilling to disclaim the AI-generated material, the Work cannot be registered as submitted’.<sup>55</sup> Echoing the Office’s 2023 guidance, the Board emphasized that when a system receives only prompts, ‘the “traditional elements of authorship” are determined and executed by the technology—not the human user’.<sup>56</sup> The Board also cited the Compendium’s rule that ‘typical technical alterations... do not warrant registration’,<sup>57</sup> and invited Allen to resubmit an application limited to his human modifications in which he could ‘explain why [those] rise to the level of copyrightable authorship’.<sup>58</sup>

In *SURYAST*, Photographer Ankit Sahni applied to register a two-dimensional artwork and listed ‘RAGHAV Artificial Intelligence Painting App’<sup>59</sup> as a co-author. He described creating the work by supplying three inputs—‘a base image, a style image,

---

<sup>53</sup> Ibid. 12.

<sup>54</sup> US Copyright Review Board (n 6) 2

<sup>55</sup> Ibid. 1.

<sup>56</sup> Ibid. 7.

<sup>57</sup> US Copyright Office, Compendium of US Copyright Office Practices (3rd edn, ch 900 ‘Visual Art Works’, 28 January 2021) § 909.3(A) <<https://www.copyright.gov/comp3/chap900/ch900-visual-art.pdf>> accessed 9 January 2026 (“Typical technical alterations”).

<sup>58</sup> US Copyright Review Board (n 6) 5.

<sup>59</sup> US Copyright Review Board (n 7) 2.

and a “variable value determining the amount of style transfer.”<sup>60</sup> The Review Board affirmed the Office’s refusal, finding that the work ‘does not contain sufficient human authorship’, because ‘the new aspects of the Work were generated by the “RAGHAV app, and not Mr. Sahni—or any other human author”’.<sup>61</sup> The Board explained that Sahni ‘did not control where those elements would be placed ... and what colors would be applied to them—RAGHAV did’,<sup>62</sup> and even if he treated RAGHAV like a filter, ‘selecting the strength of a visual filter, by itself, is not sufficiently creative’.<sup>63</sup> Consistent with the Office’s guidance, the decision reiterates that in true hybrids ‘the Office will register the human’s contributions’.<sup>64</sup>

The current control proxy—asking *who determined the expressive details*—is sound at the poles but ill-suited to the hybrid middle. That frame usefully polices fully autonomous outputs, but it under-credits creative labour that is *iterative and supervisory*: parameter steering, rejecting dozens of variants, and final approval. The Office has long framed the crucial question as whether the machine was merely an ‘assisting instrument’ or whether the work’s expressive elements ‘were actually conceived and executed not by man but by a machine’, a binary that does not map cleanly onto today’s stochastic tools.<sup>65</sup> As Ginsburg has argued, copyright authorship rests on ‘creativity in conceiving the work and controlling its execution’,<sup>66</sup> and

---

<sup>60</sup> Ibid. 2.

<sup>61</sup> Ibid. 3.

<sup>62</sup> Ibid. 7.

<sup>63</sup> Ibid. 6.

<sup>64</sup> Ibid. 5.

<sup>65</sup> US Copyright Office (n 2) 16192.

<sup>66</sup> Jane C. Ginsburg, ‘The Concept of Authorship in Comparative Copyright Law’ (2003) 52 DePaul Law Review 1063, 1072 (2003).

crucially ‘the author’s intellectual conception of the work need not reflect a complete or even an accurate prediction of the resulting work’s contents’.<sup>67</sup> Copyright theory has long recognised that creativity is ‘complex, decentered, and emergent’,<sup>68</sup> produced by ‘dynamic interactions’<sup>69</sup> rather than a single act of perfect control, and ‘importantly shaped by the concrete particulars of expression’.<sup>70</sup> Read against that baseline, a control test coloured by determinism risks treating unpredictability as non-authorship even when a human’s plan and supervision shape the outcome in recognisable ways.

### **2.2.2. Non-AI antecedents**

The tension between authorship and control predates AI. For more than a century, artists have reduced direct control without forfeiting authorship. US law has long credited human intellectual conception that is perceptibly embodied in the finished work, even when tools or contingent methods supply details. The originality threshold is modest: Copyright requires independent creation plus a ‘modicum of creativity’, and that bar is ‘extremely low; even a slight amount will suffice’.<sup>71</sup> The law does not require predictability; the key is whether the work shows traceable links between a human plan and the final expression.<sup>72</sup>

---

<sup>67</sup> Ginsburg and Budiardjo (n 8) 352.

<sup>68</sup> Julie E Cohen, ‘Creativity and Culture in Copyright Theory’ (2007) 40 UC Davis Law Review 1151, 1153.

<sup>69</sup> *Ibid.*

<sup>70</sup> *Ibid.* 1154.

<sup>71</sup> *Feist Publications, Inc* (n 17) 345-46.

<sup>72</sup> *Burrow-Giles Lithographic Co* (n 18) 58–60.

This logic holds across a wide range of practices. In postwar abstract painting, Jackson Pollock, known for his large-scale poured canvases, described his technique as constrained rather than random: ‘[I]t seems to be possible to control the flow of paint to a greater extent ... I deny the accident ... I have a general notion of what I’m about and what the results will be ... I approach painting in the same sense one approaches drawing, that is, it’s direct ...’<sup>73</sup> And in post-tonal experimental music, John Cage, a pioneer of indeterminacy, wrote of composing with ‘chance operations, some derived from the *I-Ching*’, externalizing detail choices while preserving an artistic plan.<sup>74</sup> In the same spirit, the US Copyright Office has long framed computers as assisting instruments with which the human conceives and executes the traditional elements of authorship and reserves registration for the human-authored contribution.<sup>75</sup> Contemporary scholarship likewise favours a decentred account of creativity, and cautions that copyright’s role is ‘essential but limited’,<sup>76</sup> lest it impair the ‘mobility... indispensable’ to creative practice.<sup>77</sup>

This reasoning extends to assembly-based art forms as well. When a creator selects and arranges preexisting materials, whether in collage, montage, or compilation, the Copyright Act expressly protects the ‘selection, coordination, or arrangement’ so long as the result reflects original judgment, even if the inputs themselves are

---

<sup>73</sup> Bernice Rose, *Jackson Pollock: Drawing into Painting* (The Museum of Modern Art 1980) 19.

<sup>74</sup> John Cage, *Silence: Lectures and Writings* (Wesleyan University Press 1961) 17.

<sup>75</sup> US Copyright Office, Sixty-Eighth Annual Report of the Register of Copyrights for the Fiscal Year Ending June 30, 1965 (Library of Congress 1966) 5; US Copyright Office (n 2) 16192.

<sup>76</sup> Julie E Cohen, *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice* (Yale University Press 2012) 81.

<sup>77</sup> *Ibid.*

unprotectable.<sup>78</sup> *Feist* provides the model: Even where inputs are unprotectable facts (and, by analogy, other unprotected materials), the law may protect the author’s non-trivial selection and arrangement of them.<sup>79</sup> That framework applies to AI-assisted workflows: a creator may claim the curation and layout that are perceptibly embodied in the work, although not the component AI outputs.<sup>80</sup>

Photography doctrine shows how to credit creativity without demanding micromanagement of every pixel. Predictability is not required; *traceability* is. In *Burrow-Giles*, the US Supreme Court grounded protection for photographs in the photographer’s ‘original intellectual conceptions’, including posing, lighting, and arrangement.<sup>81</sup> Appellate courts likewise credit protectable choices such as posing, lighting, camera angle, composition, and the decisive moment of capture—confirming that authorship can reside in how a scene is lit and framed, when the shutter falls and how the subject is staged.<sup>82</sup> These principles transfer cleanly to AI-assisted practice: Where the work shows traceable human contributions, such as storyboards or reference sketches; masks and guides; non-trivial retouching; compositing and layout; or the creative selection and arrangement of multiple outputs, copyright can attach to those human increments while leaving machine-determined details unprotectable.<sup>83</sup>

This focus on traceability, meaning perceptible links between human conception and the expression fixed in the work, avoids an all-or-nothing inquiry into *who determined*

---

<sup>78</sup> 17 USC §§ 101,103(b).

<sup>79</sup> *Feist Publications, Inc* (n 17) 345; 17 USC § 103(b).

<sup>80</sup> US Copyright Office (n 2) 16191-94.

<sup>81</sup> *Burrow-Giles Lithographic Co* (n 18) 58–60.

<sup>82</sup> *Rogers v Koons* 960 F2d 301, 307 (2d Cir 1992); *Leibovitz v Paramount Pictures Corp* 137 F3d 109, 116 (2d Cir 1998); *Ets-Hokin v Skyry Spirits Inc* 225 F3d 1068, 1073–74 (9th Cir 2000); *Rentmeester v Nike, Inc* 883 F3d 1111, 1119 (9th Cir 2018).

<sup>83</sup> *Feist Publications, Inc* (n 17) 348; 17 USC § 103(b); US Copyright Office (n 2) 16192.

*every detail* and aligns with long-standing doctrine that credits human conception and controlled execution without aesthetic gatekeeping.<sup>84</sup>

The through-line across non-AI domains is clear: US law does not demand omniscient control but requires that the final work discernibly embody identifiable human choices. The threshold is low; the contribution must be human, while the machine's function is that of an instrument. This baseline accommodates unpredictability: Creative practice often relies on processes that externalize execution without forfeiting authorship. And it furnishes an administrable grammar for hybrid cases—examiners and courts can ask, as they do in photography, where the *human author's contribution appears in the work*. That inquiry avoids both overclaiming and undercrediting and protects traceable human expression, however indirect.

### **2.3. The limits of proposed fixes**

Recognizing the strain that AI-assisted creation places on traditional doctrine, commentators have responded with competing supplemental frameworks. Each, however, falters once one asks how an examiner or court should identify and define what is protectable in a hybrid work.

Vision centred accounts prioritize human conception or narrative. The influential formulation by Ginsburg and Budiardjo insists on both 'a mental step (the conception of a work) and a physical step (the execution of a work)', and identifies the 'core concept' of authorship as 'creativity in conceiving the work and controlling its

---

<sup>84</sup> *Bleistein* (n 21) 251–52; Ginsburg and Budiardjo (n 8) 352.

execution’, while warning that ‘conception alone... does not suffice... [b]ut neither does mere execution’.<sup>85</sup> This is a salutary correction to determinist critiques, as Ginsburg emphasizes that the author’s conception ‘need not reflect a complete or even an accurate prediction of the resulting work’s contents’.<sup>86</sup>

But a vision-only test slides toward protecting unexpressed ideas. Pamela Samuelson anticipated the practical consequence decades ago: As user instructions become ‘increasingly brief or general’, the claim to authorship ‘becomes increasingly difficult to defend’.<sup>87</sup> And even the stronger suggestion that authorship may arise from initiation alone—that ‘a computer user who initiates the creation of computer-generated expression should be recognized as the author’<sup>88</sup>—cannot tell a registrar *which* increments in a work are human-authored or the scope in which to register them. In short, vision centred theories correctly refuse to treat unpredictability as disqualifying, but they do not yield an administrable method for tying protectable authorship to *traceable, perceptible* human expression in the record.<sup>89</sup>

A related approach analogizes the user to an art director, the ‘master’ who commands the machine. This perspective is familiar from collaborative practice: The person who sets the creative brief, iterates on drafts and gives final approval *feels* like the author.

But the analogy collapses under the basic predicate of US law: ‘The Creativity

---

<sup>85</sup> Ginsburg and Budiardjo (n 8) 343, 346-47, 352.

<sup>86</sup> Ibid. 352.

<sup>87</sup> Samuelson (n 8) 1201.

<sup>88</sup> Denicola (n 9) 251.

<sup>89</sup> *Feist Publications, Inc* (n 17) 345-46; *Burrow-Giles Lithographic Co* (n 18) 58-60.

Machine cannot be the recognized author’, because the Act ‘requires all eligible work to be authored in the first instance by a human being’.<sup>90</sup> Merely being ‘in charge’ does not identify which aspects of the resulting expression the applicant authored (and therefore may properly claim). US practice does recognise human authorship in perceptible increments (eg, creative selection, coordination and arrangement, or non-trivial edits that change the appearance and placement of elements), but the art-director analogy by itself does not generate the evidentiary vocabulary that will mark those increments and disclaim the rest.<sup>91</sup> Without that vocabulary, claims either balloon (treating all machine-determined pixels as directed by the ‘master’) or evaporate (because the submission never isolates the human-made parts).

Causation focused frameworks proceed differently: Instead of vision or direction, they ask who is legally responsible for the fixed expression. Shyamkrishna Balganesesh describes copyright as embodying a theory of ‘authorial causation’, and proposes an independent requirement of ‘copyrightable causation’.<sup>92</sup> Similarly, Dan L. Burk argues that causation, volition and intention should inform authorship analysis.<sup>93</sup> This lens resonates with the infringement doctrine’s volitional conduct requirement: The Second Circuit asked ‘who makes the copies?’ and held that the *customer*, not the remote-DVR service, did so;<sup>94</sup> the Ninth Circuit likewise required a defendant’s act to be volitional, and declined to impose direct liability on a provider that did not itself

---

<sup>90</sup> *Thaler* (n 3) 1041.

<sup>91</sup> 17 USC § 103(b); US Copyright Office (n 2) 16192-94.

<sup>92</sup> Balganesesh (n 10) 1.

<sup>93</sup> Burk (n 11) 263.

<sup>94</sup> *Cartoon Network LP, LLLP v CSC Holdings, Inc* 536 F3d 121 (2d Cir 2008) 131.

cause the reproduction.<sup>95</sup> Causation thus helpfully explains *why* a bare prompter is not the author of every detail. Yet as a registration tool it stalls, because it does not specify which human choices in a tool-mediated workflow count as legally cognizable (planning constraints? iteration logs? selection and arrangement across variants?) nor where to look for them in the work. Without a concrete measure tied to perceptible evidence, causation tends to devolve into case by case intuition, assigning responsibility somewhere along the chain, without yielding a standard method with which to carve out the human-authored increments and disclaim the machine-determined ones.

### **3. The fourfold framework of originality**

The preceding sections show a disconnect between the structure of originality doctrine and hybrid, tool-mediated creation. The core problem is not the originality threshold, but traceability: in complex workflows, we lack a disciplined way to identify which increments of the finished expression are perceptible as human-authored and traceable to contemporaneous human choices in the record (rather than to the tool’s autonomous generation). This Section supplies that method. It replaces the monolithic ‘control’ proxy with a Fourfold Framework—*Conception*, *Orchestration*, *Curation*, and *Inscription*—that locates human authorship in recurring loci and ties protectability to evidence that can be shown, preserved and examined.

---

<sup>95</sup> *Perfect 10, Inc v Giganews, Inc* 847 F3d 657 (9th Cir 2017) 666–67.

### 3.1. The four dimensions

*Conception* marks the starting point of human authorship in assistive workflows: the moment when a creator fixes a plan, whether visual, structural, or thematic, that the work later expresses. This includes the specific constraints the author intends to realize, such as a character model sheet, a colour script, a page plan, a compositional rule or camera schema. The classic photography opinion explains the point: The law protects photographs insofar as they embody an author’s ‘original mental conception’, realized through posing, lighting, and arrangement.<sup>96</sup> The camera is an instrument, not the author. In AI-assisted workflows, *Conception* counts when it is realized and traceable in the work: Dated storyboards, layout grids, or model sheets that match visible features in the registered work (for instance, a recurring silhouette or page rhythm) are probative. In contrast, generic prompts or mood boards that never materialize in the work remain unprotectable ideas.<sup>97</sup>

*Orchestration* captures the procedural middle—the ways through which a human guides execution through parameterized, traceable control that visibly shapes the outcome. This includes pipelines, control inputs, parameter regimes or tuned adapters that predictably shape the output’s expressive features. § 102(b) does not protect methods as such, but the features that the methods reliably realize, when visible in the work and traceable to human control materials, stand as evidence of human authorship. In today’s image pipelines, this can include using fixed seeds or latents for

---

<sup>96</sup> *Burrow-Giles Lithographic Co* (n 18) 60.

<sup>97</sup> US Copyright Office, Part 2 Report (n 39) 18–19.

reproducibility; ControlNet-style conditioning to lock pose or edges;<sup>98</sup> or parameter-efficient adapters (LoRA)<sup>99</sup> that enforce a stable character outline or lensing, as documented in iteration grids. The Copyright Office’s rule is straightforward: On ‘current generally available technology’, prompts alone do not provide sufficient human control, and users ‘do not exercise ultimate creative control’ merely by issuing them; therefore, creators must disclaim machine-determined material.<sup>100</sup> Accordingly, *Orchestration* counts only when the applicant shows a documented, reproducible procedure that predictably imprints human-controlled features on the product, for example, a ControlNet edge map combined with fixed seeds that repeatedly yields the same low-angle silhouette, and the corroborating parameter logs and iteration grids that appear in (and explain) the work. In contrast, *ad hoc* prompt tinkering without more, generally will not qualify as authorship.<sup>101</sup>

*Curation* reflects downstream human authorship: the act of choosing, sequencing, and arranging material into a coherent, expressive whole. It concerns editorial judgment: choosing among variants, sequencing panels, pairing text and images, and designing spreads and grids. The statute expressly protects a compiler’s ‘selection, coordination, or arrangement’ when, as a whole, it is original, while leaving underlying materials unprotected.<sup>102</sup> The Copyright Office applies this template to hybrid works. For example, in *Zarya of the Dawn*, it recognised the author’s creation of the comic’s text

---

<sup>98</sup> Lvmin Zhang, Anyi Rao and Maneesh Agrawala, ‘Adding Conditional Control to Text-to-Image Diffusion Models’ (26 November 2023) 1 <<https://arxiv.org/pdf/2302.05543>>.

<sup>99</sup> Edward J Hu and others, ‘LoRA: Low-Rank Adaptation of Large Language Models’ (16 October 2021) 1 <<https://arxiv.org/pdf/2106.09685>>; Hugging Face, ‘LoRA’ (*Diffusers*) <<https://huggingface.co/docs/diffusers/en/training/lora>> accessed 9 January 2026.

<sup>100</sup> US Copyright Office (n 2) 16192; US Copyright Office Part 2 Report (n 39) 18.

<sup>101</sup> US Copyright Office (n 2) 16192.

<sup>102</sup> 17 USC §§ 101, 103(b); *Feist Publications, Inc* (n 17) 357.

and in her selection, coordination, and arrangement of text and images, but excluded the stand-alone Midjourney-generated images.<sup>103</sup> Two clarifications improve administrability. First, the scope is ‘thin’. Protection attaches only to the compiler’s specific selection and arrangement, not to the components themselves; the copyright covers the pattern of inclusion and order.<sup>104</sup> Second, there is still a threshold. The Compendium cautions that arranging only ‘two or three’ elements is generally *de minimis*;<sup>105</sup> selecting twelve images from 180 variants and sequencing them to create a diagonal visual cadence across a spread would typically clear the bar, but leaving a single AI output to stand alone, or placing side-by-side two images that lack an evident editorial relationship generally will not.

*Inscription* captures the final layer of human input, specifically the perceptible modifications, edits and compositing that transform the raw material into authored expression. This includes non-trivial retouching, repainting, compositing, typographic, and layout interventions, which are choices that modify the form, emphasis, or structure in ways that are ‘distinguishable in some meaningful way’. Courts apply the ordinary originality standard to ‘new matter’ in derivative works. The Seventh Circuit in *Schrock* held that there is no heightened test, and asks only whether the derivative contribution shows ‘sufficient nontrivial expressive variation ... to make it distinguishable from the underlying work in some meaningful way’. Here protection is ‘thin’ and confined to the incremental expression added.<sup>106</sup> The Second Circuit in

---

<sup>103</sup> US Copyright Office (n 5) 1.

<sup>104</sup> 17 USC § 103(b); *Feist Publications, Inc* (n 17) 348, 350–351.

<sup>105</sup> US Copyright Office (n 28) § 312.2.

<sup>106</sup> *Schrock v Learning Curve International, Inc* 586 F3d 513, 521-22 (7th Cir 2009).

*Batlin* likewise required ‘some substantial, not merely trivial originality’ and rejected mechanical copying and mere changes of medium as inadequate.<sup>107</sup> The Compendium lists ‘typical technical alterations’ that ‘do not warrant registration’, including routine sharpening, colour balancing, and the like.<sup>108</sup> In contrast, masking and relocating figures to change the focal hierarchy, repainting highlights and shadows, or compositing elements from multiple variants to create a new narrative relationship are the kind of human-perceptible edits that the Office registers, within § 103(b)’s limits to the new expression contributed by the author.<sup>109</sup>

Together, the Fourfold Framework provides an administrable method for locating and scoping authorship: Where, exactly, in the record, does human expression reside? In registration, the question is no longer *who controlled every pixel*, but *where is the human in the work and its contemporaneous record*. This includes *Conception* that is traceably realized (eg, storyboards or layout files that match the work’s look); *Orchestration* with a documented, reproducible pipeline that leaves a visible trace (eg, logs correlating steps to features); *Curation* that meets *Feist*’s minimal creativity (non-trivial selection/arrangement across variants); and *Inscription* consisting of non-trivial expressive edits (before and after frames). Modern provenance tools (the

---

<sup>107</sup> *L Batlin & Son, Inc v Snyder* 536 F2d 486, 490-92 (2d Cir 1976).

<sup>108</sup> US Copyright Office, Compendium of US Copyright Office Practices (3rd edn, 2021) § 909.3(A).

<sup>109</sup> US Copyright Office (n 2) 16192-94.

C2PA ‘Content Credentials’ standard) can embed tamper-evident edit histories and tool metadata that make these linkages legible at intake and in litigation.<sup>110</sup>

### 3.2. Applying the fourfold framework to difficult cases

The Fourfold Framework transforms the vague and often-overbroad ‘control’ test into a concrete, evidentiary standard. It aligns the existing structure—particularly § 103(b), *Feist*, and the human-authorship predicate that *Thaler* reaffirmed—with modern, tool-mediated creative practice. The result is not a theoretical overlay, but a method for adjudicating mixed human–AI works in a way that is doctrinally faithful, administratively viable, and materially fair. This Section applies the Framework across representative cases—AI image generation, heavily prompted art, style-transfer tools, stochastic methods, photography and collage—and shows how outcomes become more consistent, more precise, and more explainable when courts and the Copyright Office ask the right questions: where is the human in the work, and what evidence proves it?

We’ll start with the extreme case. In *Thaler v. Perlmutter*, the DC Circuit upheld the Copyright Office’s refusal to register an image attributed solely to an AI system, the ‘Creativity Machine’. As the court held, ‘the Copyright Act of 1976 requires all eligible work to be authored in the first instance by a human being’.<sup>111</sup> Dr. Thaler confirmed that ‘the present submission lacks traditional human authorship—it was

---

<sup>110</sup> Coalition for Content Provenance and Authenticity, C2PA and Content Credentials Explainer (v 2.2, 22 April 2025) 10-11 <[https://spec.c2pa.org/specifications/specifications/2.2/explainer/\\_attachments/Explainer.pdf](https://spec.c2pa.org/specifications/specifications/2.2/explainer/_attachments/Explainer.pdf)> accessed 12 January 2026.

<sup>111</sup> *Thaler* (n 3) 1041.

autonomously generated by an AI'.<sup>112</sup> That ended the inquiry. In this example, the Framework makes the reasoning explicit: There were no traceable human *Conception*, no documentable *Orchestration*, no human *Curation* among variants, and no *Inscription* visible in the work. No human increments mean no human authorship, and therefore no copyright protection. Applying the Framework does not change the outcome. Rather, it explains it more transparently and supplies a pertinent doctrinal vocabulary.

Consider next the hybrid scenario that *Zarya of the Dawn* addresses. Kristina Kashtanova created a comic book using Midjourney to generate images, but she authored the textual content and assembled the layout herself. The Copyright Office ultimately cancelled the original registration but issued a new one, protecting the comic's text and the 'selection, coordination, and arrangement' of the text and images—while excluding the raw Midjourney-generated panels.<sup>113</sup> This ruling aligns precisely with the Fourfold Framework: The text and layout reflect Kashtanova's *Conception* and *Curation*, but the office required her to disclaim the AI images that did not have human-controlled features or documented edits. What the Framework adds here is scoping clarity. If Kashtanova had provided contemporaneous materials—such as character model sheets, panel sequencing plans or edit-layer screenshots documenting post-generation modifications—these would have supported her additional claims under *Orchestration* and *Inscription*. But if the only human

---

<sup>112</sup> Ibid. 1043.

<sup>113</sup> US Copyright Office (n 5) 1.

actions are prompt-only generation and routine upscaling, the claim stops at *Curation*, and she may only register that increment.

A more contested middle ground emerges in *Théâtre d'Opéra Spatial*, after Jason Allen entered a Midjourney-assisted image in a state fair art competition, reported that he had ‘input numerous revisions and text prompts at least 624 times’,<sup>114</sup> and then used Photoshop and Gigapixel AI in post-processing. The Review Board affirmed the Copyright Office’s refusal to protect the work ‘as submitted’<sup>115</sup> because ‘Mr. Allen is unwilling to disclaim the AI-generated material’.<sup>116</sup> It explained that, on the record before it, the Office ‘lacks sufficient information’ to determine whether his Photoshop adjustments ‘rose to the level of copyrightability’, while noting that ‘typical technical alterations’ like sharpening or balancing ‘do not warrant registration’.<sup>117</sup> The Board further quoted the AI Guidance that when an AI tool only receives a prompt, it is the technology that determines and executes ‘the “traditional elements of authorship”... not the human user’.<sup>118</sup> It invited a new application limited to Allen’s human modifications, if he could show how his edits met the originality standard. In the Framework’s terms: prompt churn alone does not establish *Orchestration*, but documented, perceptible edits (eg, compositing that changes focal hierarchy or light and shadow structure) can qualify as *Inscription*, if the claim disclaims the unedited AI portions and supplies evidence that makes the human increments

---

<sup>114</sup> US Copyright Review Board (n 6) 2.

<sup>115</sup> Ibid. 3.

<sup>116</sup> Ibid. 1.

<sup>117</sup> Ibid. 5.

<sup>118</sup> Ibid. 7.

traceable—yielding, at most, a registration limited to Allen’s contribution under § 103(b).

A similar analysis applies to *SURYAST*, where the applicant used a style-transfer tool to blend a base image with a style image at a specified strength value.<sup>119</sup> The Copyright Office refused protection, reasoning that the app—not the user—had determined the expressive elements.<sup>120</sup> The outcome is consistent under the Framework. If, however, the user had curated outputs with a discernible narrative sequence into a spread, or had composited and retouched results to introduce new relationships among visual elements, the resulting compilation and edits could support registration—again, limited to the new human expression. A raw style-transfer result with no human trace, by contrast, remains unregistrable.

The Framework also accommodates chance-based or low-control practices. In such cases, as when composers like John Cage or visual artists using stochastic processes, the question is not whether the artist had omniscient control, but whether the expressive plan is traceably realized in the work. If the artist documents the rule-set (*Conception*), process (*Orchestration*), and editorial selection (*Curation*), and the final work bears perceptible traces of the artist’s fixing hand (*Inscription*), she will satisfy the originality requirement. This approach mirrors the longstanding doctrine of *Burrow-Giles*, which protects photographs ‘so far as they are representatives of

---

<sup>119</sup> US Copyright Review Board (n 7) 2.

<sup>120</sup> *Ibid.* 7.

original intellectual conceptions of the author’,<sup>121</sup> and shows how process-based or contingent creativity can still yield registrable work. The Framework eliminates the metaphysical concern of ‘was the process too random?’, replacing it with a practical one: ‘Is the author’s contribution visible in the record?’

The same logic governs collage and graphic design. The Copyright Office has long treated collage as a compilation under § 103(b), extending protection to the author’s selection, coordination, or arrangement, but not to the source materials themselves. The Fourfold Framework formalizes this logic. If a designer selects, sequences, and juxtaposes images into a spread or poster—and substantiates his work through contact sheets or layout files—that satisfies *Curation*. If the designer then masks elements, paints transitions, or creates new relationships among parts, that’s *Inscription*. A conceptual plan or thematic schema underlying the arrangement might constitute *Conception*. Again, each dimension relates to what the record shows and what the work expresses. The more structured and traceable the author’s hand, the more clearly the work meets *Feist*’s modest bar.

Finally, the Fourfold Framework’s administrative value grows as platforms adopt provenance tooling. The C2PA Content Credentials standard provides tamper-evident provenance that ‘records and displays’ who produced a file and which tools and processes he used; the manifest travels with the asset and is cryptographically signed,

---

<sup>121</sup> *Burrow-Giles Lithographic Co* (n 18) 58, 60.

so improper changes are detectable.<sup>122</sup> In practice, this means that applicants can attach edit histories and tool metadata to works, allowing examiners to verify *Orchestration* (tool paths, control inputs), *Inscription* (layered edits), and *Curation* (assembly sequences) without mini-trials; in litigation, the same signatures and audit trails help to narrow disputes about who contributed what and when.<sup>123</sup> Adoption of these tools is already underway: Adobe has rolled out Content Credentials across Creative Cloud and Firefly,<sup>124</sup> Microsoft was a founding member of C2PA,<sup>125</sup> and camera makers are building in credentials at capture, for example, Leica M11-P (built-in Content Credentials)<sup>126</sup> and Nikon Z6III (C2PA-compliant firmware).<sup>127</sup> The Copyright Office has not announced any binding provenance standard, but it is actively modernizing its systems, through the Copyright Public Records portal and the e-Recordation program, which suggests administrative receptivity to standardized provenance that makes human authorship traceable in the registration record.<sup>128</sup>

### 3.3. Potential challenges to the fourfold framework and responses

---

<sup>122</sup> Content Authenticity Initiative, ‘How it works’ <<https://contentauthenticity.org/how-it-works>> accessed 2 January 2026; Coalition for Content Provenance and Authenticity (n 110) 3, 5, 10.

<sup>123</sup> Adobe, ‘Content Credentials overview’ (*Adobe Help Center*, 2 September 2025) <<https://helpx.adobe.com/creative-cloud/apps/adobe-content-authenticity/content-credentials/overview.html>> accessed 9 January 2026; Coalition for Content Provenance and Authenticity (n 110) 7, 11.

<sup>124</sup> Andy Parsons, ‘Seizing the moment and driving adoption for Content Credentials in 2024’ (Adobe Blog, 26 January 2024) <<https://blog.adobe.com/en/publish/2024/01/26/seizing-moment-content-credentials-in-2024>> accessed 19 January 2026.

<sup>125</sup> Microsoft, ‘Project Origin’ <<https://www.microsoft.com/en-us/research/project/project-origin/>> accessed 2 January 2026.

<sup>126</sup> Leica Camera, ‘Content Credentials’ <<https://leica-camera.com/en-US/photography/content-credentials>> accessed 16 January 2026; Leica Camera AG, ‘New: Leica M11-P—The world’s first camera to create a seamless chain of authenticity from capturing to publication with Content Credentials’ (26 October 2023) 1 <[https://leica-camera.com/sites/default/files/2023-10/press\\_release\\_leica\\_m11p\\_october\\_2023.pdf](https://leica-camera.com/sites/default/files/2023-10/press_release_leica_m11p_october_2023.pdf)> accessed 11 January 2026.

<sup>127</sup> Nikon, ‘Nikon develops firmware that adds a function compliant with C2PA standards to the Nikon Z6III full-frame mirrorless camera’ (14 October 2024) <<https://www.nikonusa.com/press-room/nikon-develops-firmware-that-adds-function-compliant-with-cp2a-standards-to-z6iii>> accessed 13 January 2026.

<sup>128</sup> US Copyright Office, ‘Search Copyright Records: Copyright Public Records Portal’ <<https://www.copyright.gov/public-records/>> accessed 8 January 2026; US Copyright Office, ‘Recordation System’ <<https://www.copyright.gov/recordation/pilot/>> accessed 15 January 2026.

A natural concern is that crediting *Conception* could blur the boundary between idea and expression. The Framework is structured to prevent that slippage. *Conception* counts only when it is *realized*, when contemporaneous materials (storyboards, scripts, layout grids, model sheets) let an examiner trace specific, pre-existing expressive constraints to features that are *visible in the work*, and only when they are coupled with *Orchestration*, *Curation*, or *Inscription*. This is not a doctrinal novelty; it is how US law has long tied authorship to ‘original intellectual conceptions’ that are embodied in expression rather than floating abstractions.<sup>129</sup> In short, the Framework uses *Conception* as evidence of authorship only when the artist perceptibly realizes that plan in the work.

A second critique might lie in administrability: Will this framework invite evidentiary mini-trials? In practice, it leverages that machinery the Copyright Office already operates. Since March 2023, it has required applicants to disclose AI-generated material that is ‘more than de minimis’, provide ‘a brief explanation of the human author’s contributions’, and, where appropriate, disclaim machine-determined content in the ‘Material Excluded’ field.<sup>130</sup> The Copyright Office warns that knowingly omitting essential information risks cancellation or judicial disregard of the registration under § 411(b).<sup>131</sup> The Office’s 2025 report states that human contributions are registrable when they are ‘perceptible in AI-generated outputs’.<sup>132</sup> The Framework simply standardizes which aspects of their work applicants must

---

<sup>129</sup> *Burrow-Giles Lithographic Co* (n 18) 58-60.

<sup>130</sup> US Copyright Office (n 2) 16193.

<sup>131</sup> *Ibid.*

<sup>132</sup> US Copyright Office, Part 2 Report (n 39) iii.

preserve, where examiners look (the four loci), and how artists present their claims (§ 103(b)). The Framework does *not* demand perfect reproducibility, as not all creators can maintain sophisticated logs; it asks only for *traceability*—a plausible, documented link between human choices and visible features in the work, consistent with the Office’s emphasis on what is ‘perceptible’ in the work.<sup>133</sup>

A related practical concern is the gaming or fabrication of process records. Here too, the existing regime already provides tools: § 409 requires applicants to submit sworn information;<sup>134</sup> the Copyright Office may cancel registrations obtained by material misstatements; and, under § 411(b), courts may disregard registrations procured with knowing inaccuracies.<sup>135</sup> Provenance tools like C2PA’s ‘Content Credentials’ embed cryptographically verifiable edit histories into media files, and Adobe, Microsoft, and hardware makers such as Leica and Nikon are adopting them.<sup>136</sup> None of this imposes a mandatory technological hurdle on registrants; it simply makes honest documentation easier to verify when available.

A third line of criticism might urge us to side-step authorship altogether, either by relying on disclosure<sup>137</sup> or by creating new *sui generis* rights for AI outputs.<sup>138</sup>

Disclosure is necessary but not sufficient: it tells examiners *that* AI was used, not *what* the human authored. As for new exclusive rights, the Copyright Office’s 2025

---

<sup>133</sup> Ibid.

<sup>134</sup> 17 USC § 409.

<sup>135</sup> 17 USC § 411(b)(1)(A)–(B).

<sup>136</sup> Coalition for Content Provenance and Authenticity (n 110) 5; Adobe (n 123); Microsoft (n 125); Leica Camera AG (n 126); Nikon (n 127).

<sup>137</sup> Noti-Victor (n 12) 139.

<sup>138</sup> Massadeh and others (n 12) 49.

report squarely concludes that ‘[t]he case has not been made for additional copyright or *sui generis* protection for AI-generated content’,<sup>139</sup> warning that overprotection of machine-determined material could ‘undermine rather than advance’ copyright’s aims.<sup>140</sup> The Framework therefore hews to orthodoxy: It keeps autonomous outputs in the public domain while crediting human increments in hybrids—precisely the balance that the Office endorses.

Finally, some might call this ‘control’ talk by another name. It is not. The status-quo proxy asks who controls the details. The Fourfold test refines this, asking, *where in the work is the human?* This shift aligns with longstanding doctrine: Courts have long grounded photographic authorship in identifiable human choices.<sup>141</sup> The Copyright Office has now extended that recognition to AI, emphasizing that copyright attaches to the human authorship ‘perceptible in AI-generated outputs’,<sup>142</sup> to creative selection and arrangement, and to creative modifications. The Framework operationalizes those premises with a stable, evidence-led vocabulary.

In sum, the Fourfold Framework is both principled and practical. It guards the idea/expression line; fits the Copyright Office’s registration architecture; discourages gamesmanship; rejects *sui generis* detours; and replaces monolithic ‘control’ talk with a granular inquiry tied to evidence visible in the work. The result is a more faithful

---

<sup>139</sup> US Copyright Office, Part 2 Report (n 39) iii.

<sup>140</sup> *Ibid.* 36.

<sup>141</sup> *Burrow-Giles Lithographic Co* (n 18) 60.

<sup>142</sup> US Copyright Office, Part 2 Report (n 39) iii.

implementation of existing law that produces clearer records and narrower, more defensible claims.

#### **4. Implications for doctrine and policy in the age of human–AI creativity**

##### **4.1. Copyright Office practices: registration and record-keeping**

The Fourfold Framework requires no change to substantive law. The Copyright Office can implement it immediately as an administrative rubric that operationalizes what it already does when examining hybrid works. The Copyright Office has set two clear baseline requirements: applicants must describe the human contributions they seek to claim, and, in the ‘Limitation of Claim’ or ‘Material Excluded’ fields, they must disclose, and expressly disclaim, any more-than-*de minimis* AI-generated material.<sup>143</sup> The Office’s 2023 policy statement even provides model language and explains where in the on-line form it belongs. The Framework simply supplies a standardized vocabulary for the disclosures that the Office already expects, thus improving consistency without shifting doctrine.

This standardization yields three concrete, low-friction adjustments. First, for applicants, the Framework simplifies the ‘Author Created’ and ‘Material Excluded’ fields. Rather than offering *ad hoc* narratives, an applicant can scope claims using the Framework’s shorthand, eg, ‘Author Created: (*Curation*) selection, coordination and arrangement of 50 images from 1,000 variants; (*Inscription*) compositing and repainting; and (*Conception*) page grid realized in final layout’, paired with ‘Material

---

<sup>143</sup> US Copyright Office (n 2) 16193.

Excluded: unedited AI-generated base renders’. The Copyright Office’s own guidance invites precisely this kind of structured scoping (‘brief explanation of the human author’s contributions’, and the explicit exclusion of AI portions in ‘Material Excluded’). Had applicants in *Théâtre d’Opéra Spatial* and *SURYAST* led with Framework-scoped claims, each limited to documented human increments, they would have been more likely to have avoided ‘as-submitted’ refusals tied to undisclaimed AI material or insufficiently specified human authorship.

Second, for examiners, the Framework functions as an internal checklist that organizes the points that the 2023 policy directs examiners to evaluate, specifically ‘how the AI tool operates and how it was used’<sup>144</sup>, but channels that evaluation into the four loci in which human authorship is most often visible. In practice, that means asking: Is *Conception* traceably realized (eg, storyboards or model sheets that match features visible in the work)? Does *Orchestration* show a documented, reproducible procedure that leaves a visible trace (eg, seeds, masks, control inputs, iteration logs)? Does the submission claim a level of *Curation* (selection/coordination/arrangement) that clears *Feist*’s modest bar? And do the edit layers show *Inscription* that is ‘distinguishable in some meaningful way’, as derivative-work doctrine requires? By framing the questions in this way, the Fourfold Test makes the Copyright Office’s existing rules—disclose, disclaim, and register only the human increments—easier to apply consistently across examiners and art units.

---

<sup>144</sup> Ibid. 16192.

Third, for the public record, Framework-inflected annotations make certificates more informative. The Office already cancels protection and, where necessary, reissues it with narrowed scope.<sup>145</sup> Using uniform Framework terms on certificates (for example, ‘Registration limited to human-authored Curation and Inscription; AI-generated pictorial material disclaimed’) would generalize that practice and improve downstream notice. The Copyright Office is currently modernizing its record systems—implementing the Public Records portal and the e-Recordation pilot, which will make richer, standardized scope statements both administratively feasible and publicly searchable.<sup>146</sup>

Feasibility is not hypothetical; cognate agencies already rely on structured, machine-readable disclosures to improve accuracy, reduce cost, and enhance transparency. The USPTO requires claim-by-claim, marked-up amendment listings under 37 CFR § 1.121<sup>147</sup> and corresponding MPEP provisions,<sup>148</sup> standardized forms that make incremental contributions and exclusions apparent at a glance. The SEC’s Inline XBRL rule likewise imposed standardized tagging for public company disclosures, demonstrating how a modest format change can increase data quality while lowering preparation cost over time.<sup>149</sup> The Copyright Office’s move to require

---

<sup>145</sup> US Copyright Office (n 5) 1.

<sup>146</sup> US Copyright Office, ‘Search Copyright Records: Copyright Public Records Portal’ <<https://www.copyright.gov/public-records/>> accessed 6 January 2026; US Copyright Office, ‘Recordation System’ <<https://www.copyright.gov/recordation/pilot/>> accessed 4 January 2026. (“allows users to electronically submit transfers of copyright ownership”); Library of Congress, ‘Copyright Online Recordation System Opens for Public Use’ (1 August 2022) <<https://www.loc.gov/item/prn-22-068/>> accessed 9 January 2026.

<sup>147</sup> 37 CFR §§ 1.121(c)-(d).

<sup>148</sup> US Patent and Trademark Office, Manual of Patent Examining Procedure (MPEP) § 714 ‘Amendments, Applicant’s Action’ <<https://www.uspto.gov/web/offices/pac/mpep/s714.html>> accessed 8 January 2026.

<sup>149</sup> Securities and Exchange Commission, ‘Inline XBRL Filing of Tagged Data’ (Final rule) (28 June 2018) 1 <<https://www.sec.gov/files/rules/final/2018/33-10514.pdf>> accessed 5 January 2026; Securities and Exchange

uniform Fourfold Framework-based scoping language is less complicated than either precedent, yet it will offer the same benefits: fewer errors, more comparable filings, and a clearer public record.

In short, the Framework refines, but does not rewrite, current practice. It translates the Copyright Office’s existing guidance (disclose, disclaim, and claim only human authorship) into a shared vocabulary that applicants can use *ex ante* and examiners can apply *ex post*, with certificates and public records that clearly mark the line that § 103(b) already draws between ‘material contributed by the author’ and everything else. The payoff is immediate: more precisely scoped claims, fewer reconsiderations, and a public record that shows, in plain terms, where the human work ends and the machine-determined details begin.

#### **4.2. Empowering creators and platforms: best practices and expectations**

The Fourfold Framework’s practical methodology allows creators and platforms to manage authorship in hybrid works, aligning creative practices with the Copyright Office’s legal requirements. This approach reliably lets creators secure protection for their contributions, and encourages platforms to build more transparent ecosystems.

For creators, the Framework translates directly into a workflow for documenting and asserting authorship. The first step is to maintain a contemporaneous record of human contributions corresponding to the four dimensions: storyboards or model sheets for

---

Commission, ‘SEC Adopts Inline XBRL for Tagged Data’ (29 June 2018)  
<<https://www.sec.gov/newsroom/press-releases/2018-117>> accessed 4 January 2026.

*Conception*; parameter logs and control inputs for *Orchestration*; selection notes and layout files for *Curation*; and version histories or layered files for *Inscription*. This documentation provides the evidentiary basis for a properly scoped copyright application, allowing creators to move beyond reliance on prompts alone.

Platforms and toolmakers can use the same discipline to build product and submission flows. Intake forms can ask users, in plain language keyed to the four dimensions, to identify their contributions and to upload corroborating artefacts (eg, a ControlNet mask or an edit-layer screenshot). Public-facing displays can mirror the registration record with scoped notices such as ‘Human authorship: Curation & Inscription; AI-generated base renders disclaimed’, improving transparency for licensors and audiences. Furthermore, by integrating open provenance standards like C2PA’s ‘Content Credentials’, platforms can allow creators to embed tamper-evident, machine-readable proof of their creative process directly into their files.

The adoption of Framework-style disclosures will also harmonize with emerging transparency regimes outside copyright. The EU AI Act imposes general transparency obligations, including labelling requirements for AI-generated or AI-altered content; structured, machine-readable provenance records make those duties administrable across the content supply chain.<sup>150</sup> In the United States, YouTube now requires creators to disclose their alterations of realistic content and additions of synthetic

---

<sup>150</sup> Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) (OJ L, 2024/1689, 12.7.2024) arts 50(2), 113.

media,<sup>151</sup> while US policy regarding AI accountability emphasizes provenance and authentication as necessary bases for trust.<sup>152</sup> In particular, NTIA’s 2024 AI Accountability Policy Report identifies ‘opacity around provenance of AI outputs’<sup>153</sup> as a barrier to accountability, and urges coordinated adoption of technical standards, precisely the role C2PA and Content Credentials can play. Taken together, these developments show that Framework-aligned records are more than registrable evidence; they are compliance assets across jurisdictions and platforms.

The result is a more predictable, lower-friction market. Creators know what to document and claim, platforms can align user interface and metadata with legal reality, and the public record more sharply demarcates the line between protectable human authorship and the public domain.

### **4.3. Safeguarding the public domain and avoiding overprotection**

The Fourfold Framework’s central public-law function is to protect the public domain by preventing overreach in hybrid works. It achieves this by clarifying the boundary between human and machine at each stage of the copyright life cycle. By tying scope to the author’s contribution, the Framework’s structural logic forecloses the familiar

---

<sup>151</sup> The YouTube Team, ‘How we’re helping creators disclose altered or synthetic content’ (YouTube Official Blog, 18 March 2024) <<https://blog.youtube/news-and-events/disclosing-ai-generated-content/>> accessed 7 January 2026.

<sup>152</sup> National Telecommunications and Information Administration, ‘AI Output Disclosures: Use, Provenance, Adverse Incidents’ (AI Accountability Policy Report, 27 March 2024) <<https://www.ntia.gov/issues/artificial-intelligence/ai-accountability-policy-report/developing-accountability-inputs-a-deeper-dive/information-flow/ai-output-disclosures>> accessed 19 January 2026.

<sup>153</sup> National Telecommunications and Information Administration, Artificial Intelligence Accountability Policy Report (March 2024) 16 <<https://www.ntia.gov/sites/default/files/publications/ntia-ai-report-final.pdf>> accessed 25 January 2026.

problem of *copyright creep*, where an artist uses his modest human contribution to claim *de facto* control over an entire AI-assisted work.

At registration, the Framework demands precision. By requiring applicants to articulate their claims in terms of specific, verifiable contributions, it forces a clear disclaimer of the machine-generated remainder. Instead of a vague assertion of authorship over a hybrid work, the public record will show a narrow claim limited to, for example, the *Curation* of images and the *Inscription* of textual overlays. This affirms the public domain status of the underlying AI-generated elements from the outset, keeping that material free for all to use.

In the marketplace, this clarity prevents the over-licensing of rights. A registration certificate that clearly circumscribes the scope of protection ensures that transactions are based on a precise understanding of what the author owns. Authors can license their limited rights in selection/arrangement or their new matter edits, but cannot purport to control unedited AI left in the public domain—precisely the allocation *Feist* envisioned when it protected a compiler’s choices while leaving underlying facts (and, by analogy, the purely machine-generated details) free.<sup>154</sup> This level of precision helps to maintain a competitive market for creative works, and prevents the inadvertent privatization of machine-generated content.

In adjudication, a dimension-tagged registration record streamlines infringement analysis. It allows a court to isolate the protectable human increments at issue, and to

---

<sup>154</sup> *Feist Publications, Inc* (n 17) 350.

determine whether the artist copied those specific elements. This prevents litigation from devolving into a metaphysical debate about ‘control’ by focusing the inquiry on the traditional question: *What protectable, human-authored expression was copied?*

By methodically separating human from machine contribution, the Fourfold Framework faithfully implements existing law, while delivering market and adjudicatory benefits that prevent overprotection. It protects what should be protected (traceable human expression) and keeps the remainder in the public domain, where it belongs.

## **5. Conclusion**

The copyright system must remain responsive to the realities of contemporary authorship while preserving the legal and conceptual boundaries that give it coherence. In the age of generative AI, the central risk is not the erosion of human authorship *per se*, but the drift toward abstractions, like ‘ultimate creative control’, that misfire in hybrid contexts. The Fourfold Framework asks a narrower, more administrable question: Where in the work is the human contribution perceptible and documented? By organizing evidence around *Conception, Orchestration, Curation* and *Inscription*, the Framework helps applicants to frame claims, examiners to scope registrations, and courts to confine relief to ‘the material contributed by the author’, consistent with § 103(b). It does not expand subject matter or raise originality; it offers a way to apply existing doctrine to the hybrid middle with greater discipline, greater clarity, and fewer disputes.