

HUMAN-AI COLLABORATION IN CREATIVE CONTENT GENERATION

Vinesh Tapre¹, Rohita Sharma², Nirupama Singh³

¹Lovely Professional University, Punjab, India.

²Lovely Professional University, Punjab, India.

Email: rohitasharma2009@gmail.com

³Department of Visual Arts; Former Dean, School of Performing and Visual Arts; Former Head, Department of Visual Arts, University of Himachal Pradesh, Dehra Gopipur, District Kangra, Himachal Pradesh, India.

Abstract: Generative AI is transforming the conceptual development and design of text, images, music and video ideas and variations – both in terms of real time efficiency and cost. Not in the creative ability of humans, though, but of DED. The outcome of human's activity and what they feed into their personal life, the assessment of them, responsibility, the decision to communicate each other with the scale changes and transformations made by the machine. For a miniproj, the human centric approach is suggested and proposed for the content generation: co-creation and creativity. Our stages are: Framing of the brief, independent/humanased seed, articulation of alternatives, AI, getting it to mature, getting it approved then released to the learning outcomes. Clearly identifies who's responsible for each and every step, and all elements of the provenance ledger –the sources, what the consents are, the model(s), model version(s), prompts, candidates, the editor, and who's doing the reviewing of the edit and ledger disclosure persons. Please use the design in writing, visual communication, campaign ideas and for production of educational media. Issues of originality, relevance, voice, factuality and accessibility/efficiency now come into the picture when it comes to quality and release gates. Controlled prototype – same prototype compared to Human (or non-human) (or both) prototype and same set of Briefs and same experts used and blinded during evaluation. Only illustrative results (not necessarily experimental result) indicated to show the propose analysis. The least successful creative engagement is to have the partners develop a project or if following the 'quick and post' route. It's a clear and incremental process in which each individual can maintain ownership of their own decisions about what to write and be ready to take responsibility for what they write, while the AI can offer them more avenues for exploration, faster iterations of revisions and test options.

Keywords: human-AI collaboration, generative AI, co-creativity, creative content, provenance, responsible AI, design workflow

1. INTRODUCTION

There are few stages that are being repeated and / or iterated in order to create creative content 'observing', 'framing', 'associating', 'drafting', 'critique and revision'. Adding another player: a probabilistic system capable of fairly quickly providing coherent extensions or changes of a prompt and context. The system doesn't see (not feel) the audience; doesn't have a stake in the message or is unresponsive to consequences. Despite this, it has its own benefits such as brainstorming, reframing/adapting, summarisation, visually exploring, repetitive production, etc. However, the question "can humans be creative in the same way as AI?" is not the question, but rather how their capacities can be matched up to create consistent teamwork.

Limited evidence of mixed pictures reasoning being used. This approach to language modelling also has reduced the time needed to do this and improved the average quality of the achieved professional task in writing [6]. Results: From the visual art works shown and short stories, it seems that the productivity improvements and risks of converging with each other exist individually or exist when they are created in conjunction. Moreover, a creative-problem-solving study showed the effect of novelty increase when human crowds is used, and value and quality increase when human-AI solutions are used as well as the explicitly differentiated one shows higher performance the



AI-assisted search [9]. Findings support the theory of ‘complementarity’ but results also reveal that speed/quantity of candidates is not suitable criteria to define creativity.

Four issues of practical problems that the project will be based upon are: The division of labour problem may be ambiguous (confusing); results may be achieved too early; the level of traceability may be low and measures to measure quality may be solely efficiency based, and not on originality, or harmfulness. It sets up a sustainable collaborative process that encompasses process scoring, governance and prototype review. Limited creative team overseeing the content and low to medium risk content. Other Domain Control Requirements – Requirements where high stakes are involved in communication (e.g. medical, legal, financial or political) will not be a part of the prototype.

2. CONCEPTUAL FOUNDATIONS AND PROJECT SCOPE

A. Human–AI Co-Creativity

Co-creativity – Co-production of a creative product between a computational system(s) and humans. Lubart's first categorisation was geared towards the possibility's computers could bring as nannies, pen pals, coaches or colleagues in the creative process [1]. Later on, a distinction is made between alternation and task division: The former means that the collaborators alternate with each other and the latter that they have specialised job roles [2]. A process of co-creation that takes place between the partners, the creative product and the co-creation process [4] is modeled by COFI. Collaboration: Considered as the activity of the project machine is not "literal" enters a word/words, searches for a word/words in another word/words, but it is interpreted, selected, and transformed, according to the results of the decision of a person.

B. Complementary Capabilities

The distribution of work in the framework has been performed based on the comparative strengths. Content relating to indigenous and social/cultural content and moral perception, taste, stakeholder relationships and responsibilities are introduced. Quick recombination, plenty of cheap choices, format conversions and consistency are some of the characteristics of models. People can't be objective, objective people can get tired, and objective people can look at objects newly, models can hallucinate, and unconsciously mimic stereotyped models, which can also camouflage the source of a concept. It's not a check nor a human measure one.

Table 1: Division of labour in a human–AI creative team

Activity	Human contribution	AI contribution	Final responsibility
Intent and brief	Purpose, audience, values, boundaries and success criteria	Questions, reformulations and brief completeness checks	Creative lead
Divergent ideation	Independent seeds, unusual associations and cultural context	Rapid alternatives, combinations and perspective shifts	Creative lead
Drafting	Narrative logic, meaning, voice and key original expression	First drafts, variants, adaptation and structural suggestions	Named author/designer
Evaluation	Taste, relevance, factual and cultural judgement	Rubric checks, comparison and inconsistency detection	Human reviewer
Release	Consent, rights, disclosure and acceptance of residual risk	Metadata and checklist support; never autonomous publishing	Accountable approver

C. Aim, Objectives, and Scope

Through need of their collaboration, developing and demoing an open source creative system using artificial intelligence (AI) and humans to build the creative artefacts. The objectives will be to: (1) maintain human intent/agency; (2) increase ease of use in the process of exploring while keeping diversity; (3) maintain versioned

provenance; (4) include quality rights/safety review; (5) develop a plan for an academic mini project; (6) compare the multiple-sides of collaboration using common protocol. It also considers that of text and image concepts and multi modal campaigns, but it doesn't claim to be in itself a measure for creativity, and that AI created content is always copyrightable.

3. REVIEW OF RELATED WORK

The criteria for human computer interaction is compiled as: Set expectations, give some assistance to correct, show uncertainties and efficiently discard/reconsider unwanted results [3]. So these are all some key factors to take into account when using creative technologies, the quality of the interaction that they provide and consequently the quality of the products produced. Vocabulary of initiative, turn taking, contribution and communication [4] is provided by COFI. They also believe that an effective partner should lead the leaden place - the use - to other solutions, go according to directions, and never give away his power to say no to an idea, without losing a job.

Table 2: Evidence synthesis and design response

Study	Relevant finding	Project response
Noy and Zhang [6]	Generative AI improved time and average quality in professional writing tasks	Measure efficiency and quality separately
Doshi and Hauser [7]	Individual creativity improved, but collective content became less diverse	Require independent seeds and diversity review
Zhou and Lee [8]	AI use increased productivity; average visual novelty declined while peak novelty could rise	Track average and best-candidate quality
Boussieux et al. [9]	Human–AI solutions gained value; differentiated search supported novelty	Prompt for distance from prior candidates
Amershi et al. [3]; COFI [4]	Interaction design affects trust, control and collaboration	Use visible stages, undo, critique and approval gates

4. DESIGN REQUIREMENTS

Brief, structured, stand-alone, human ideation, controlled access to models, batch/iterative creation, comparing candidates, in-place editing, feedback/citation/source checking, disclosure and post-release feedback. Non-functional requirements include such aspects, as usable, reproducible, privacy, exportable, accessible and tolerable latency/predictable cost, etc. For each “candidate” we have to make, we need to have a “block” that cannot be automatically substituted to the human-made “block” and is not published as an asset.

5. PROPOSED HUMAN–AI COLLABORATION FRAMEWORK

A. Seven-Stage Lifecycle

Figure 1 explains the suggested Life Cycle. Framing involves putting an appeal to a request in more ethical and measurable terms in the form of a brief. Seeding records to add “permited” idea and reference to what Human invented. This divergent (more interested in the model alternatives) is because of differing structure, perspective risk and media. Selection – no confusion using rubric (fluency and/or originality) change: refine – form of editing – human editing with some computer editing Approvals are for facts, similarity, rights, accessibility, disclosure, brand and/or author voice. A Learning Log is a record of achievement on the task for the current short term and/or any incidents that happen during the short term. Ms. Hayes teaches her students that they should use feedback to enhance their own current and future abilities and to accept those of others.




Figure 1: Proposed seven-stage human–AI creative collaboration lifecycle.

B. Collaboration Modes

Stage: 3 modes: 3 Options. Sometimes you may find it hard to do bound transformation, assist you will help you to change your layout sometimes, shortens or translates etc. In alternating mode, the "first-look" filters are used alternately by individuals and change models, who propose and challenge each other, respectively. The person (who has his own thoughts) and the person (who performs the production operations) - the model - the task division is determined. No independent candidates are allowed. The element of human rating throughout the obtaining a presentable product can be quite important and it should be documented how much human rating was used.

C. Governance Gates and Roles

Brief and Direction: the authors/designer has provided Brief and Direction; Original expression and editing: a Fact and Domain Reviewer has checked the claims made and the assumption that they are not similar or objects of which no (reproduction) rights exist has been extended with the aim of asking the right holders for reproduction rights; Verification of assets used in production: a Rights Reviewer has verified the legal status of these assets and extended the assumption that there are no reproduction rights to ask the right holders for rights.

6. CREATIVE AND TECHNICAL WORKFLOW

A. Brief-to-Release Architecture

All components of the architecture in Figure 2 are partitioned in four sections, such as, creative brief, generation workspace, human studio and release gata. Objects are now being created in each workspace (not one file), and are versioned. Explain about the purpose, target audience, voice, limitation and bibliography of the short stories. The generation workspace contains: Prompt templates, Model identifier and the candidates. This selection, combination and edits can be imported to the human studio. Review the evidence - when doing the regurgitation gate, it should be final approval and disclosure. Provenance ledger is important to "bind" things together but a ledger is not one that "says the truth of provenance.




Figure 2: Versioned content architecture and provenance ledger.

B. Operational Workflow

Table 3: Stage-level inputs, actions and outputs

Stage	Input	Human action	AI action	Output
1. Frame	Request, audience and policy	Define intent, constraints and rubric	Identify ambiguity; suggest questions	Approved brief
2. Seed	Brief and lawful references	Create concepts, examples and exclusions	No generation until the seed is recorded	Seed set
3. Diverge	Brief, seed and diversity plan	Choose perspectives and search directions	Generate contrasted candidates	Candidate pool
4. Select	Candidate pool and rubric	Score, compare and reject	Cluster, highlight differences and checks	Shortlist
5. Refine	Shortlist and critique	Rewrite, compose and decide meaning	Targeted variation and consistency support	Edited draft
6. Approve	Draft, ledger and evidence	Verify, disclose and sign off	Checklist and metadata support	Release package
7. Learn	Outcomes and incidents	Interpret feedback; revise practice	Summarise patterns under review	Lessons and new controls

C. Diversity-Preserving Prompt Strategy

They should attempt to think of ideas they know to be deviant rather than numerous paraphrases when they're generating. The prompt contains information about the difference (in terms of structure of how their representation of the narrative is constructed, range of emotions presented, cultural point of view, medium and the actions they want the audience to take) and every candidate has to describe its difference from previous prompts. Separate sessions of using context resets can be useful in minimising the local mimicry. The human examples are not modelling or prompts which will replicate a person creator. Directions that are rejected are recorded by the team as these form the creative history in accordance with the converged decisions.

D. Provenance Record

The minimum record will include the project/asset ID, short version, contributor and consent status, the reference asset's source, the reference asset's licence, date, prompt(s) and/or candidate's ID (some you will enter in the record), human edits, evidence of verification status and approver name, disclosure text, and disclosure channel. The CSS are cryptographically signed statement of edits and origin but don't contain the truthfulness of a claim. There's still a need for a human readable ledger, though.

7. QUALITY, ETHICS, COPYRIGHT, AND PROVENANCE

A. Candidate Generation and Scoring

Now let us consider an accepted "brief" b , an accepted "context" c and say p_k is a "prompt" for the search it is we intend to make. Now with a model we are able to find parameters θ and then sample our y_k from a conditional distribution. A set of candidates is many prompts and many runs, NOT quality/author.

$$y_k \sim P_\theta(y | b, c, p_k) \quad (1)$$

Each candidate is judged against the each of the criteria of originality O , relevance R , voice V , fidelity (factual or structural) F and risk of harm H with set weights, with the total of all the positive weights adding up one. Score can be compared and there is a numeric (based on performance) score for a failed rights/safety gate.

$$U(y) = w_o O + w_r R + w_v V + w_f F - w_h H, \Sigma w = 1 \quad (2)$$

The positive side of this will be after set level of diversity the team will not be able to pick up near duplicate. The farther apart and the more different the similarity $\text{sim}(\cdot)$ used, the larger the mean distance (D) of the pairs of candidates. For text, embedding / comparison is possible as is structural / perceptual / compositional for images.

$$D(Y) = 2 / [m(m-1)] \cdot \sum_{i < j} [1 - \text{sim}(y_i, y_j)] \quad (3)$$

B. Responsible-AI Controls

The four phases of the NIST AI Risk Management Framework include Govern; Map, Measure, and Manage [11] (which includes in the profile of risks for generative-AI, confabulation, harmful bias, information integrity, privacy and intellectual property human-AI configuration [12]). UNESCO is also worried about such problems as human dignity, fairness, transparency, responsibility to oversee and environmental responsibility [13]. These were used today in the creative studio below:

Table 4: Creative risks, preventive controls and release evidence

Risk	Possible harm	Preventive control	Release evidence
Confabulation	False claims appear credible	Source-grounded drafting; claim-by-claim review	Verified source list and reviewer sign-off
Bias or cultural harm	Stereotype, exclusion or appropriation	Diverse review; cultural context and exclusion rules	Review notes and corrected version
Privacy/likeness	Personal data, face or voice used without authority	Minimisation, consent and restricted model inputs	Consent record or documented removal
Copyright/similarity	Unlicensed source use or overly close output	Rights inventory, originality review and similarity search	Licences, source map and rights sign-off
Deceptive media	Audience mistakes synthetic content for authentic evidence	Visible disclosure and provenance metadata	Approved label and signed release record
Automation bias	Fluent output bypasses human judgment	Independent human seed, rubric and mandatory rejection round	Score sheet and named approver

C. Copyright and Disclosure

The laws regarding copyrights differ from place to place. In its Annual Report for 2025, the U.S. Copyright Office shares this theory that this current copyright system can also accommodate works created with AI's assistance. That's where in this process is not related to the legal part, but rather to the significant aspect of Human selection, arrangement and editing. Also, the content of the contracts, local law, publicity/personality rights and training data licences must be considered in the project's policy. Information should be made available to make the audience aware about what has been created, what has been human-produced/verified and what is available as proving of the work of origin.

8. PROTOTYPE METHODOLOGY

A. Research Design

The prototype features a comparison of the three modes, human-only, AI-only and human-AI collaboration, within a single brief. A few of the creative is different for every brief; the speculative fiction story is capped at 500 words, the public interest/ 3 posts is capped at that number of posts and the key visual/ caption idea wording is capped at 250 words and/ or caption idea (min 3 posts). The Purpose, Audience, Constraints and Time limit are identical to

those from all modes, and the Reference pack is the same for all modes. There is a mode change, re-shaping of the skills being practised and the order of the modes is cross-compensated to minimize any learning and fatigue effects.

Table 5: Prototype conditions and controls

Condition	Allowed process	Recorded evidence	Key comparison
Human-only	Standard digital tools; no generative model	Time, drafts, sources, revisions and final asset	Baseline originality, voice and effort
AI-only	Brief submitted to model; minimal formatting only	Prompts, model/version, candidates and output	Machine speed, fluency and unsupported claims
Human-AI	Full seven-stage framework with human final edit	Complete ledger, scores, changes and approval	Complementarity, control and traceability

B. Evaluation Procedure

The outputs are presented to a panel of at least three assessors, who have creative, domain and audience understanding and knowledge, and are judged - in a randomized sequence. They are assessed on a scale of anchored 1 – 5 in Originality and Relevance, Voice and Craft and Factualness and Accessibility. So, time taken and number of candidates taking the test, number of test iterations to number of candidates failing the test based upon the edit distance is collected. These are reported prior to averaging and attained an interrater agreement. Without a doubt, there are two types of binary gates; similarity check gate and right check gate. They subsequently fill in a brief survey on agency, cognitive load, calibration of trust and satisfaction. Qualitative comments provide an explanation of what evidence of change in score indicates.

Table 6: Metrics and acceptance criteria

Dimension	Measure	Acceptance rule
Creative quality	Blind mean for originality, relevance, voice and craft	Human-AI is non-inferior to the stronger baseline and improves at least two dimensions
Factual/structural fidelity	Supported-claim rate and constraint violations	No critical false claim; $\geq 95\%$ supported checkable claims
Diversity	Pairwise distance plus human concept-cluster count	No more than half of the shortlisted ideas in one concept cluster
Efficiency	Active minutes to release-ready asset	Meaningful reduction without a lower gate pass rate
Agency	Self-report plus proportion of human-originated final decisions	Creators can identify, reject and revise machine contributions

C. Analysis and Validity

Mode, median and CI's for the brief & participant reports. Mode can be either fixed or creator/b uniform or random, brief/reviewer can be either brief or random. The study needs to tell us what changes are needed to the model, and where knowledge and information is lacking. Threats to validity are sample size, evaluator taste, effects of practice, change in the model(s) and unexpected TD familiarity with tools, and blinding the AI style. The prototype is not used in absolute, there is no attempt to grade the creativity of human and computers.

9. IMPLEMENTATION PLAN

Table 7: Eight-week mini-project plan

Stage	Weeks	Activity	Evidence
1. Discovery	1	Briefs, stakeholders, policy and baseline workflow	Approved scope and risk register
2. Design	2	Rubric, roles, prompt patterns and ledger schema	Design specification
3. Study	6–7	Run controlled tasks and blind assessment	Dataset, score sheets and consent record
4. Report	8	Analysis, limitations, handover and archive	Final report and reproducible package

10. ILLUSTRATIVE RESULTS AND ANALYSIS

A. Simulated Comparison

To show the analysis proposed in this paper, a few simulated scores (from figure 3 and table 8) are provided. Have not been based on facts. But here, it's still compared to basically almost as human-ish as possible; in performance and almost as original as possible, except for the fact that it has been on the slower side of things. AI generation is whizz quick but isn't strictly distinctive or original in speech. The combination of human and AI – iteration and human selection and rewrite – is most efficient, and most complete profile. The pattern is reasonable, but isn't considered to be a conclusion—it will have to be verified with the protocol in Section VIII.




Figure 3: Illustrative simulated comparison of creative workflows; not experimental findings.

Table 8: Illustrative simulated scores for framework analysis

Mode	Originality	Relevance	Voice	Efficiency	Gate status
Human-only	82	78	88	46	Pass
AI-only	71	74	64	94	Fail: provenance/voice
Human-AI	88	91	90	86	Pass

B. Interpretation

This is an example of a composite score that's easily confused. But the idea of AIO can be the only option when efficiency is foremost on people's list while addressing the 'from' and voice issues. Without checking many duplicate outputs, expenses of checking saved if it is determined it is original as per the best candidate. Overall compromise is a full report, diversity of the set of candidates, and time/gate failure. Once the governance effort has been incorporated it will be more likely that there are added value benefits that if verified would indicate the framework is well-placed to be adopted by a team.

11. DISCUSSION, LIMITATIONS, AND FUTURE SCOPE

The framework has the advantage that there is no overlap between the different phases of the exploration-evaluation-release process and that there is a very structured approach. They can be created from a direction, which can undergo some changes by generative models to result in another direction in a different format. Only things that can't be done by a human being are selection of problems, understanding of the situation, finding out the damage and taking responsibility. Eliminate some of the risks of a lot of creators using the same generator with the same prompts by differentiating the prompts using human seeds. Informal turns to explicit approbation (with validation), provenance.

Limitations are substantial. Rubrics can only reveal a partial information concerning creativity – always cultural. There is a spectrum of modellers' behaviours ranging from the ability to change some of the terms in the equation up to similarity tools – some variants (Maybe there are conceptual borrowings that have not been identified). Not guaranteed of the source and originality / consent. An illustrative score is NOT a score! Larger multi-modal studies that provide comparison of novices to experts and evaluation of skills acquisition over a longer time frame would have been beneficial for future studies and would also have allowed for assessment of (ineffectively) cumbersome diversity controls to enhance the acquisition of skills.

12. CONCLUSION

By utilizing human agency in the entire creative process and governance, human-AI combination can create better creative pieces. Composed along a human-led inventive and governance process, Human-AI collaboration can aid in improving creative content. The seven stages suggested begin with intent and independent human seeds, and work through to the AI, human selection and substantive refinement, to finally arrive at verified approval, provenance and learning. It is a measure of originality, relevance, voice, fidelity, diversity, efficiency/harm NOT speed of production. Our concept would be to let AI create and edit the ideas, however, we would publish the ideas (created by a human), and the evidence/rights and consequences. It is collaboration itself (not so much a literal 'human author' to 'machine' proposal) that is touted in this (and in other cases) as design.

References

1. T. Lubart, "How Can Computers Be Partners in the Creative Process? Classification and Commentary on the Special Issue," *International Journal of Human-Computer Studies*, vol. 63, nos. 4–5, pp. 365–369, 2005, DOI: 10.1016/j.ijhcs.2005.04.002.
2. A. Kantosalo and H. Toivonen, "Modes for Creative Human-Computer Collaboration: Alternating and Task-Divided Co-Creativity," *Proceedings of the Seventh International Conference on Computational Creativity*, pp. 77–84, 2016.
3. S. Amershi et al., "Guidelines for Human-AI Interaction," *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2019, DOI: 10.1145/3290605.3300233.
4. J. Rezwana and M. L. Maher, "Designing Creative AI Partners with COFI: A Framework for Modeling Interaction in Human-AI Co-Creative Systems," *ACM Transactions on Computer-Human Interaction*, vol. 30, no. 5, article 67, 2023, DOI: 10.1145/3519026.
5. Z. Epstein et al., "Art and the Science of Generative AI," *Science*, vol. 380, no. 6650, pp. 1110–1111, 2023, DOI: 10.1126/science.adh4451.
6. S. Noy and W. Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023, DOI: 10.1126/science.adh2586.
7. A. R. Doshi and O. P. Hauser, "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content," *Science Advances*, vol. 10, no. 33, eadn5290, 2024, DOI: 10.1126/sciadv.adn5290.
8. E. Zhou and D. Lee, "Generative Artificial Intelligence, Human Creativity, and Art," *PNAS Nexus*, vol. 3, no. 3, pgae052, 2024, DOI: 10.1093/pnasnexus/pgae052.
9. L. Boussieux, J. N. Lane, M. Zhang, V. Jacimovic, and K. R. Lakhani, "The Crowdless Future? Generative AI and Creative Problem-Solving," *Organization Science*, vol. 35, no. 5, pp. 1589–1607, 2024, DOI: 10.1287/orsc.2023.18430.
10. V. Choudhary, A. Marchetti, Y. R. Shrestha, and P. Puranam, "Human-AI Ensembles: When Can They Work?" *Journal of Management*, vol. 51, no. 2, pp. 536–569, 2025, DOI: 10.1177/01492063231194968.
11. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, 2023. Available: <https://doi.org/10.6028/NIST.AI.100-1>.
12. C. Autio et al., "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile," NIST AI 600-1, 2024. Available: <https://doi.org/10.6028/NIST.AI.600-1>.
13. UNESCO, "Recommendation on the Ethics of Artificial Intelligence," adopted 23 November 2021. Available: <https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence>.
14. U.S. Copyright Office, "Copyright and Artificial Intelligence, Part 2: Copyrightability," January 2025. Available: <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf>.

15. Coalition for Content Provenance and Authenticity, “C2PA Technical Specification,” version 2.4, 2026. Available: https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html.
16. U.S. Copyright Office, “Copyright and Artificial Intelligence, Part 1: Digital Replicas,” July 2024. Available: <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-1-Digital-Replicas-Report.pdf>.