

Human Generative Workflow

An invitation to conversation - 2026

Since November 2023, a small group of filmmakers and technologists has gathered with a simple objective: to ensure the future of filmmaking is shaped intentionally, not accidentally. Those conversations focused on building a shared understanding, a common vocabulary, and a collective responsibility for how AI is integrated into the creative process. This document is the result of those discussions and their valuable input.

The problem in plain language

One acronym is doing too much work.

The term “AI” is being asked to describe everything from features like noise reduction that have lived inside post-production software for years, to the new capabilities of generative models that didn’t exist two years ago. It is used in the same breath to mean a color grading assist in DaVinci Resolve and a prompted synthetic performance that replaces an actor. It applies to both the evolving CGI pipeline that artists spent a decade mastering and a deepfake made without consent.

That ambiguity is showing up in every contract negotiation, credit conversation, guild discussion, and audience disclosure question. ***It creates friction and confusion where the industry needs shared ground.***

This paper deliberately focuses on one challenge: developing a shared language around forms of AI use in production workflow. It does not attempt to resolve equally important questions around training data, licensing, labor transition, consent, or economic policy. Those issues deserve their own discussion and will ultimately require legal, commercial and social solutions. A common vocabulary, however, is a necessary foundation for all of them.

A map, not a switch

The binary that AI is either “used” or “not used” doesn’t hold. The industry needs a map for the distinctly different forms of AI use in production, rather than conflating such different uses together. Over time, these categories could form the basis of a shared classification system: something productions use internally, in credits, and in guild-facing documentation.

An industry-built classification system that distinguishes between these distinct AI use cases could serve productions, crews, guilds, and eventually audiences.

Three distinct forms of AI usage are emerging, each of which can be used for concept development and exploration, as well as for producing final footage. This document focuses on the latter: where the following three uses are used for final production footage.

Utility Techniques & Embedded AI - Copyrightable

This foundational category encompasses utility-driven machine learning, as well as the AI tools already inside media software – Premiere, DaVinci Resolve, Avid, After Effects – and its ecosystem of plugins. Many of these types of AI-enabled features have existed for years.

Importantly, workflows in this category generally help to answer a question that has a defined correct answer and is not subject to much creative interpretation. The following is a non-exclusive example list of techniques that fit in this category:

- **Denoising:** Removing noise/damage from image or sound
- **Upscaling and Sharpening:** Enhancing resolution of existing human-captured media
- **Audio Separation:** Isolating distinct audio tracks from a mixed source
- **Rotoscoping:** Automating matte generation for existing footage
- **Depth Information:** Extracting spatial data from flat images
- **Markerless Mocap:** Capturing human performance without physical tracking suits
- **Wire-removal & Inpainting:** Erasing production artifacts from human-directed scenes

Most productions already use these techniques without flagging it as “AI use,” because until recently, no one needed to make the distinction. ***The outputs from these techniques are generally copyrightable when used as listed above.***

Human Generative Workflow (HGW) - Copyrightable

Artist-controlled pipelines that integrate specialized generative models inside VFX or production tools. These are human authored, copyrightable, and the product of deliberate creative decisions. HGW works within professional tools like Nuke, Unreal, After Effects, Blender, ComfyUI, and DaVinci Resolve, and brings specialized generative models into that pipeline under the artist’s control, often alongside tightly trained “LoRA” (Low-Rank Adaptation) models, control nets, and structured inputs. ***The outputs are human crafted and copyrightable.***

Machine Generative (MG) - Not Copyrightable (likely not copyrightable)

Purely machine-generated likenesses, voices, performances, or stories that substitute for human creative work. A prompt-based generative tool is not precise enough and replaces human creative vision with a generative model. A prompt may include text, images, or other inputs. The key copyright question is not how creative those inputs are, but whether the user exercises sufficient control over the outputs specific expressive elements, which current generative systems may ultimately determine themselves. The distinction that these are *exclusively* “prompted generations” from machines is not a technicality: It is the difference between a set of tools and processes that serve human creative vision and a prompt-based generative tool that replaces it. ***These outputs are NOT copyrightable, as the machine solely generates the footage based on prompts.***

What we mean by Human Generative Workflow

Human Generative Workflow is a pipeline where, much like CGI, a human artist (the director and their team of human artists) is making decisions. The generative models integrated into these workflows serve the artist’s creative vision by preserving, if not enhancing, precision and control. The artist defines the parameters, provides the inputs, and evaluates what comes back at every step of a process that often involves multiple traditional and modern tools working together.

HGW takes different shapes in different crafts. The walkthrough and worked examples that follow draw from visual production. The same logic applies in sound, music, editing, and the writing room: artist-led, granular control, and iteration. It will extend to crafts and ways of making work we haven’t yet named or even imagined.

A workflow walkthrough

An HGW production might begin with an artist or team hand-painting reference art that establishes the look. That reference work becomes training data for a small custom model, fine-tuned on the team’s own artwork and original IP rather than the open internet. The artists generate outputs using the trained model, then paint over what isn’t right and retrain. This custom-trained model can then be used to generate elements that integrate with the rest of the production pipeline in various ways. Scenes – or keyframes used in the animation process – are assembled with intentional camera, lighting, and composition. Traditional compositing brings everything into the final frame. Color and finishing happen as on any other production.

This walkthrough of an HGW behind a set of shots shows the human hand at every step.

An HGW example: Dear Upstairs Neighbors

Dear Upstairs Neighbors [\[Trailer\]](#) [\[BTS\]](#), the six-minute animated short directed by Connie He and produced by Márcia Mayer, working with a crew of dozens of artists and a broader production staff, and with support from engineers at Google DeepMind, premiered at Tribeca Festival in June 2026. It was made by a team of roughly forty-five people consisting of animation veterans from Pixar and DreamWorks working alongside DeepMind researchers and engineers. Every frame on screen is generated by a series of different fine-tuned machine learning models alongside traditional editing tools. This document advocates that every frame is also entirely human-authored.

The team hand-painted concept art that established the visual language of the film. They fine-tuned small custom models on their own original artwork. They developed video-to-video workflows in which animators provided rough animation as the structural backbone for the generated output. When the results didn't match the team's intent, they painted over, retrained, and iterated.

Connie and her team made this film using Human Generative Workflow. The pipeline was artist-led. Control was granular, built on custom training on the team's own work. Iteration allowed for multiple cycles of creative exploration and discernment. Generative components contributed to a larger production process where human judgment governed every decision. Human Generative Workflow gave Connie language to describe her work and process. HGW proposes that creative authorship is human even when rendering is generative. The team's artistic decisions are visible in every frame even though no frame was hand-drawn. This is another incremental step along the continual evolution of computer-generated imagery rather than a departure from it.

The more useful question is not "Did you use AI?" but "Where, and how, did you use AI in the process?" The moment a tool moves from exploration into the final pipeline is where the most consequential decisions get made, often quickly and informally. That transition point is where consent, credit, and copyright begin to apply.

What using Human Generative Workflow is not: a tool for replacing writers and actors, automating performances, or generating story without a human author.

Copyright law defines and protects HGW

The US Copyright office has published¹ a document with interpretations and recommendations to help define human authorship, which could serve as a useful reference to distinguish between HGW and MG content. Here are a few notable excerpts:

- Copyright does not extend to purely AI-generated material, or material where there is insufficient human control over the expressive elements.

¹ [Report on Copyright and Artificial Intelligence, Part 2: Copyrightability](#). January 2025

- Whether human contributions to AI-generated outputs are sufficient to constitute authorship must be analyzed on a case-by-case basis.
- Based on the functioning of current generally available technology, **prompts do not alone provide sufficient control**.
- Human-authored expression that remains perceptible in an AI-assisted output may be protected by copyright, as may sufficiently original human modifications or the human selection, coordination, and arrangement of human-authored and AI-generated material; however, protection extends only to those human contributions, not to AI-generated elements standing alone, and must be assessed case by case.
- The inclusion of elements of AI-generated content in a larger human-authored work does not affect the copyrightability of the larger human-authored work as a whole.

The market is already moving

The classification conversation is not hypothetical. Filmmakers and studios are already staking positions that will shape the formal frameworks to come. A24's *Heretic* ended its 2024 credits with "No generative AI was used in the making of this film." In 2025, Vince Gilligan's *Pluribus* closed with "This show was made by humans." DreamWorks' *Bad Guys 2* used its credits to prohibit AI training on the work. And yet, many sub-contractors and production partners on projects like these are attempting to navigate the usage of new AI-enabled features in their standard tools (the "Embedded AI in Traditional Tools" category) and fast emerging HGW on a project by project basis. There is significant ambiguity and inconsistency in both policy and parlance. Sales companies and rights advocates are beginning to push for industry-wide certification standards. We need more granularity and clear language.

The industry is still stuck in the binary question, "was AI used, or not?" HGW responds to a different and more useful question about its use: "where in the process, by whom, and with what degree of control?" The future of entertainment will be shaped not by AI alone, but by the choices we make about how AI is integrated into the creative process.

What could happen next

None of this requires waiting for a regulatory mandate. Several of these steps could begin now, in parallel, with the people already in the room.

- **Publish this vocabulary and invite a response.** A shared industry glossary, even a provisional one, gives everyone something concrete to react to. Reaction is how a draft becomes a standard.
- **Start the consent conversation on Machine Generative (MG) outputs of actors.** That's where agreement is closest. Separating it from HGW in early negotiations allows progress on both fronts simultaneously.

- **Engage with market-driven labeling efforts.** Filmmaker-led disclosures and emerging certification schemes reflect real buyer and audience demand. A shared classification framework should be legible alongside them, not in opposition.
- **Think about governance as a living conversation.** The tools will keep changing. A multi-stakeholder body, modest in scope and flexible enough to revisit its own definitions over time, may serve the industry better than any framework written to be permanent.
- **How does this work across international co-productions?** A framework developed around US copyright rulings may not translate, subject to EU, UK, or Chinese law. And, China's regulatory approach to AI is moving on its own timeline.
- **Extend the conversation to music, sound, and the writing room.** Before those communities find themselves reacting rather than shaping.

*"The way forward must involve collaboration,
and the value of the human in the A.I. Age will be essential."*
