



IBC 2026 ACCELERATOR MEDIA INNOVATION PROGRAMME

FRAMES: Federated Retrieval, Agentic Media Environments and Software-Defined Workflows

**An R&D Exploration of AI Ethics, Responsibility, Transparency &
Risk for Creative Media Production**

Part One [September 2026]

Supporting R&D for FRAMES Final POC Results

AUTHORS:

Editor: Muki Kulhan, IBC Accelerators Media Innovation Programme

Co-Editor: Ben Schofield (GC-SC - The Global Creative & Security Community)

Project Co-Leads & First Authors: Roberto Iacoviello (RAI), Paola Sunna (EBU), Chris Vienneau, Raymond Drury (MovieLabs)

Project Team- First Authors: Tim Deussen (XRBB- Extended Reality Berlin-Brandenburg), Kathryn Webb and Ahsan Mallick (AIMICI), Ben Schofield and Aaron Bhugobaun (GC-SC - The Global Creative & Security Community)

Project Team- Contributing Authors: John Canning (AMD), Erik Weaver (ETC/USC-Entertainment Technology Center at University of Southern California), Florian Reimann (MovieLabs), Clive Santamaria (ITV), Cedric LeJuene, INSA (National Institute of Applied Science), Aidan Fuhrer and Nicholas Frei (Enigmatic), John Maxwell Hobbs (Streamline Media), Nick Kennedy (BBC R&D), Nicholas Fry (Enigmatic), Shamir Alibhai (Press Play Labs/Eddie AI)

TABLE OF CONTENTS

PART ONE: Executive Summary	2
Executive Summary	2
Introduction	3
Context and Scope	4
Key IBC Accelerator Project Drivers	6
PART TWO: POC Results, Project Learnings & Insights	7
Ethics	7
Responsibility	7
Transparency	8
Risk	8
EU AI Act Interpretation for FRAMES	8
Further Industry insights	10
PART THREE: Creating best practice	11
Training	11
Transparency	11
PART FOUR: Conclusions	14
Definitions	17
References	18

PART ONE: Executive Summary

Executive Summary

AI is already an integral part of media production workflows, spanning everything from script development and screenwriting to production and post-production. Executives everywhere are under pressure to capitalise on the perceived 90% cost savings by deploying AI in production.

However, the rapid adoption of AI brings with it complex questions of AI Ethics, Responsibility, Transparency, and Risk. The question is not whether to use it, but whether a production company or individual can demonstrate responsible, transparent use of AI and quantify the potential impact of any risks.

This paper, stemming from the R&D work of the **IBC 2026 Accelerator FRAMES consortium (FRAMES: Federated Retrieval, Agentic Media Environments, and Software-Defined Workflows)**, offers an overview of how narrative-driven, semantically guided generative AI for modern, open studio pipelines produces content and media experiences, and begins to address these broader questions of Ethics, Responsibility, Transparency, and Risk. This initial work opens the discussion for further investigation.

The project was comprised of an international team of major media industry organisations, broadcasters, academia and vendors, including Champions: RAI, EBU, and MovieLabs as co-leads, alongside Co-Champion members BBC, XRBB, ETC-USC, ITV, GC-SC (The Global Creative and Security Community) and INSA (National Institute of Applied Science), along with Participant vendors AMD/HP, Press Play Labs/Eddie AI, Streamline Media, Enigmatic and AIMICI.

Since March 2026, this consortium has maintained curiosity and research around Ethics, Responsibility, Transparency, and Risk at every step of the process. The results are outlined in this paper, serving as part of their Final POC materials, and aimed at guiding the media production, film, television, and broadcasting industries by outlining tangible learnings, solid insights, and knowledge; suggesting best practice; and showing how neglecting these aspects could put an organisation on the wrong side of a very expensive conversation.

Introduction

Media industries are leveraging AI to transform key stages of the media value chain, ranging from production to distribution and monetisation, to reduce costs and deliver faster. As of September 2026, AI is being integrated throughout most of the pipeline, from scriptwriting and screenwriting to music composition, editing, animation creation and more. Much of the core AI skills and experience is currently concentrated in a small group of freelancers working on the frontier (often referred to as ‘solopreneurs’), while larger production companies and broadcasters still explore how to create the best content with the best teams. Therefore, AI skills are in short supply to match the high demand, with the supply chain and opportunities to upskill moving at a cautious pace due to huge, unknown risks, where the responsibility is normally being directly linked to that individual creator, producer or technical person, rather than a parent company.

These advancements raise critical questions regarding authenticity, attribution, copyright, transparency, accountability, and ethical use. These present a new set of commercial risks that need to be quantified to drive executive focus; mitigation will depend on potential costs, especially when considering the use of IP and archive materials. It is becoming clear that these models and platforms present a new set of risks that require a different approach to current controls and governance.

Industry dynamics, organisational structures, and the skill sets of human professionals further complicate the situation. Global regulatory frameworks differ significantly, requiring companies to adapt responsibly to regional requirements while lacking the necessary knowledge.

To embed and implement responsibility into AI systems within the media sector, monitoring and evaluation efforts are also guided by several key concerns and priorities. This requires shared industry best practices and ongoing stakeholder engagement, which this paper outlines through the research conducted within this IBC 2026 Accelerator FRAMES.

The team also investigated a variety of solutions, as outlined throughout this paper, including introducing AI training across production companies and individuals, with a focus on competency and risk to comply with new and upcoming industry standards in the UK and EU. In the UK, progress has been made to streamline AI training standards through industry-recognised competency, namely via one of the FRAMES Participant solution vendors, AIMICI, and its Responsible AI Certificate, which provides training through courses being recognised in the EU via Ireland. This initiative will help set the standard for the knowledge required for AI use in the industry. This paper also includes a new working proposal for an industry-standard AI documentation and disclosure process.


Context and Scope

FRAMES is one of the nine projects selected for the IBC Accelerator 2026 programme; their challenge objective was to create an end-to-end media production framework based on AI agents adopting a structure provided by the "MovieLabs Ontology for Media Creation." The project's scope encompasses AI-assisted media production, focusing on pre-production and production phases, and involves key activities such as scriptwriting, character generation, shot and image creation, and

background music composition, all essential for producing a final video of professional quality.

While the market offers numerous AI tools and agents, AI cannot become a genuine competitive advantage for an organisation without extensive metadata and a robust framework covering Ethics, Responsibility, Transparency, and Risk. The scenarios analysed within the project provided a meaningful context for examining the challenges and considerations involved in defining a framework for responsible AI in the media sector.

Workflow Overview



This schema diagram describes the “knowledge-graph-centred” AI production pipeline the team worked on together. The team ingested scripts and archive materials into Google Drive. AI extracts metadata from archive material; AI agents build the shortlist from the script; and MovieLabs’ partner Yamdu breaks down the shortlist. The ontology for Media Creation (OMC) provides a shared semantic and data model for describing the production’s content, assets, characters, activities, rights, and tools. OMC underpins the Knowledge Graph.

The graph then distributes structured production data to INSA via Kitsu/HOLLM (our production brains), Enigmatic (to rename files according to the naming convention), and Streamline (a shot-review tool). Approved shots pass through an iterative review process before enhancement and final editing in Topaz and Resolve.

Parallel tasks include generating a soundtrack inspired by archive material and creating an automated rough cut of archive material using Eddie AI. We then compare this rough cut with manually selected content.

Key IBC Accelerator Project Drivers

The adoption of AI in the media production sector is driven by the need for innovation, productivity improvements, and enhanced human capabilities, whilst keeping creative integrity intact. With this in mind, and within their short five-month sprint since IBC's Kickstart Day in February 2026, this IBC Accelerator FRAMES team prioritised routinely exploring these aspects throughout their R&D process

This raised questions for the team about AI Ethics, Responsibility, Risk, and Transparency, which this paper explores. These aspects include, but are not limited to, questions such as:

- How can legitimate ownership and copyright regarding AI-generated content be determined?
- If content is generated exclusively by artificial intelligence without sufficient human creative input, can it be legally protected under copyright law, or are competitors free to copy it directly?
- Who is liable for due diligence and risk assessment, especially if something goes wrong? The prompt engineer? The production company? The technical team? The AI model provider or the creator platform? More importantly, at what stage of the production should this be considered?
- What exactly are the guarantees of the legal indemnities? These contractual safeguards protect users against copyright infringement lawsuits arising from generative AI output or training data. However, the scope and applicability of such protections vary significantly depending on the training architecture and the provider's terms of service. Under what conditions do such indemnities become completely void?
- What rights do original authors retain when their works are included in training sets?

- How can artists and publishers protect their works when billions of copyrighted assets are extracted without explicit consent, attribution, or financial compensation?
- What is the true value in ensuring works are protected? Especially for works that may include substantial AI use or the use of AI agents (eg., quantifiable benefits, creative friction, etc.)?
- How difficult is it to document AI use on a project, and can we improve on the process?
- How can we ensure that everyone in the industry has the skills and knowledge necessary to navigate these issues?

PART TWO:

POC Results, Project Learnings & Insights

The team completed an end-to-end workflow and produced a short film based on archive materials matched to generative environments, resulting in significant savings in production time and costs. Through this process, we raised questions about Ethics, Responsibility, Transparency & Risk and explored potential solutions.

Ethics

The EU AI Act is a binding regulation and reflects EU principles regarding human agency and oversight, privacy, and accountability. It provides a risk-based framework for assessing AI use.

AI can assist writers, directors, editors, composers, and visual effects teams, but final decisions must rest with responsible individuals, e.g., the creative or production team.

At FRAMES, when working with archives, we ensured that the material was not used to train the provider's models; we also verified where the videos are processed and stored, and who had access to them. We also applied its transparency requirements to distinguish between original RAI archive footage, AI-manipulated

material, and fully synthetic scenes, and to determine how each should be clearly labelled within the docufiction.

Responsibility

AI does not transfer responsibility to the software; the producing company remains responsible for the decisions made and the results obtained.

In the FRAMES project, we discussed which AI tools to use and examined the results to identify any errors, biases, or inappropriate similarities, as AI cannot replace responsible human judgment or oversight.

Transparency

The transparency requirements in [Article 50 of the EU AI Act](#) aim to ensure people can recognise when content has been generated or substantially manipulated by AI. These requirements are intended to reduce the risks of deception, impersonation, misinformation, and manipulation of individuals or public discourse.

The European Commission's voluntary [Code of Practice on Transparency of AI-Generated Content](#) translates these requirements into practical measures. These include clear audience-facing disclosures, machine-readable marking, and reliable provenance information. The Code promotes an approach that is accessible, consistent, and proportionate to the type of content and its context. Its application nevertheless requires producers to assess their specific use of AI and determine which obligations apply.

Risk

The EU AI Act adopts a risk-based approach (e.g., the exploitation of children, persons with disabilities, or the use of AI to generate non-consensual intimate material).

In the context of FRAMES, using AI exclusively for visual effects, editing, audio, colour grading, or creative assistance will not result in classification as a high-risk system. Further work will look at how to quantify risk in terms of impact and frequency to align mitigation with business priorities

In any case, Stella the snail gave her full consent, and no snails were harmed in the making of “One Small Crawl”.

EU AI Act Interpretation for FRAMES

- **Public-interest content** - FRAMES is an artistic docufiction centred on a fictional snail. It is not AI-generated text published to inform the public about matters of public interest, such as news or public-affairs reporting. We therefore concluded that the specific public-interest disclosure requirement does not apply. This conclusion is based on the purpose, format, and editorial context of the work—not simply on RAI's status as a public broadcaster.
- **Archive footage and AI-generated content** - The production combines original material from the RAI archive with AI-assisted, AI-manipulated, and fully synthetic content. The use of public-broadcast archive material does not determine the AI Act risk category on its own. The relevant questions are how the material is processed, whether it could mislead viewers about its authenticity, and whether its artificial origin should be disclosed.
- **Deepfakes** - We concluded that FRAME does not use deepfakes because its generated elements are not intended to be presented as authentic records of real people, places, or events. For example, an AI-generated background may recreate an image as a visual reference for framing and composition, but it does not reproduce or impersonate identifiable individuals or claim to document a real event. However, the AI Act's definition of a deepfake is not limited to representations of people. It can also cover AI-generated or manipulated representations of existing objects, places, entities, or events where the result could falsely appear authentic. Each scene should therefore be assessed according to whether viewers could reasonably mistake it for genuine documentary footage. As FRAME is an evidently artistic and fictional work, any applicable disclosure may be made in an appropriate manner that does not interfere with the display or enjoyment of the work.
- **Emotion and performance capture** - The artist manually animated the snail's emotions. No AI system analysed biometric data or inferred a person's emotional state; therefore, this process does not constitute emotion recognition under the AI Act. If a production used real actors and analysed their facial expressions, voices, gestures, or other biometric data to infer emotions and drive a generative-AI process, it could qualify as emotion recognition. The actors would need to be informed, and all applicable AI and data-protection requirements would have to be satisfied. Consent alone may not make such processing lawful: emotion recognition in the workplace is generally prohibited under Article 5 of the AI Act, except for narrowly defined medical or safety purposes. Capturing an actor's visible movements or expressions without using AI to infer emotions is not, by itself, emotion recognition.
- **3D models and digital twins** - The 3D models of real locations used in FRAME serve as visual references for framing and composition. They are not intended to be accurate, data-driven replicas that continuously represent or simulate the physical properties, current conditions, or behaviour of real locations. We therefore did not

consider them digital twins. This distinction does not remove the need for a transparency assessment. A generated model of a real location could still fall within the AI Act's deepfake definition if it were presented in a way that could reasonably be mistaken for authentic footage of that location.

- **Machine-readable marking and provenance** - Article 50 requires providers of generative-AI systems to make generated or substantially manipulated outputs detectable through machine-readable marking, as far as technically feasible. During the FRAME accelerator, this information was not always embedded by the AI tools or accessible within the production workflow. This represents an important practical gap when several AI systems are combined and their outputs are modified across different stages of production. Combining or orchestrating AI tools does not automatically transform a deployer into a provider. Provider responsibilities may nevertheless arise where an organisation develops or substantially modifies an AI system, or places it on the market or into service under its own name. Assess this distinction for each production workflow. To mitigate the lack of accessible machine-readable information, the FRAMES team manually documented the tools and processes used for each asset. The team classified every asset as original archive footage, AI-assisted or AI-manipulated material, or fully synthetic content. Where technically possible, AI-generated and substantially manipulated assets should also retain machine-readable metadata or provenance information throughout production and distribution. The internal asset register supports traceability and enables consistent disclosure to audiences, although manual documentation does not replace any machine-readable marking legally required from an AI-system provider.

Further Industry insights

While larger studios continue to integrate AI technologies into their operations, a growing number of independent, small-to-midsize, and AI-native studios are emerging. These entities frequently operate independently of broader industry initiatives, executing productions with minimal centralised studio oversight while rapidly navigating evolving AI governance standards. Consequently, any ethical best practices and framework developments must account for sole proprietors and AI-native studios amid ongoing industry disruption.

With multiple laws from different jurisdictions affecting multiple components of the production cycle, from deployment of AI, terms and conditions from suppliers, to trade union rules and various varying studio and insurance policies, the approach to AI transparency and reporting is fragmented, blocking progress with uncertainty. This project has sought to lay the groundwork for a universal AI transparency approach, enabling records to be

obtained for multi-stakeholder purposes.

Furthermore, there is a clear need for a global industry standard for evaluating AI supply chain risk, for example through an established risk-rating and tool-assessment approach or an approved directory of AI tools that evaluates vendors on operational practices, commercial viability, ethical track record, and transparency.

PART THREE: Creating Best Practice

Responsible AI Best Practices & the Transparency Gap

AIMICI has partnered on the project and has also been working on standards for responsible AI best practice in industry over the last two years. In their approach, they have identified two key levers to help the industry navigate AI concerns at scale:

Training

Not to be confused with technical AI model training, this educational effort is about ensuring everyone working on productions, especially key decision-makers, has an understanding of what AI can do along the pipeline, its common risks, and baseline knowledge of how AI governance and escalation processes work, and the role each of us plays within this due diligence process.

The team dug deep into this research, drawing on tangible learnings about what training has worked so far, especially the work of FRAMES Participant AIMICI and its establishment of affordable, standardised responsible AI e-learning courses released earlier this year. In the UK and EU, several unions, guilds, and training bodies have already backed the training and, with continued traction, could be well on its way to becoming a recognised industry standard that can be adapted globally and plugged back into further research modules.

Ensuring this kind of training is completed and understood beyond AI skills or tools training can help build a safer foundation for AI-led innovation as tools continue to progress and improve our productions.

Transparency

With a training foundation in place, the industry can also address the lack of transparency in AI usage across productions. Productions must accurately document and appropriately disclose AI use, both internally and externally. So far, the industry has not agreed on a shared standard for how to approach this.

Preview: A proposed standard for AI Transparency & Disclosure on Productions

AIMICI has designed a proposed new standard and begun testing it as part of the project, focused on bringing transparency and understanding of copyright implications to projects with hybrid AI and human workflows such as this.

The transparency standard, in summary, looks to capture the following scope of attributes

- **WHAT** - Details of the project and its key phases
- **WHO** - Who is responsible for overseeing and reporting on AI usage across those stages
- **HOW** - What are the AI use cases and tools used
- **WHEN** - When were tools used, and what decisions were made on which use cases or tools can proceed or be stopped
- **WHERE** - Where did AI actions take place across a workflow, and what assets were affected

The key to applied transparency is ensuring appropriate documentation and records are in place so disclosure is possible at different levels. The AIMICI standard proposes the following **three levels of records**:

Level 1: Production AI Statement - *For viewing audiences*

A short, plain-language set of one to three statements, appropriate to show on screen and/or include in the credits, that the production creates according to the importance of the aspects of the AI used in the production and/or its compliance needs. Suggested format includes 1) Where/how AI was not used 2) Where/how AI was used 3) Human accountability in the process

Level 2: Production AI card - *For industry & stakeholders (Eg. internal staff, financiers, commissioners etc.)*

One-to-two-page reference detailing AI use cases and tools, and relevant attributes around risk, visibility and significance of the AI use. This is organised into production stages and includes a high-level record of human oversight over the governance process, and the businesses involved in record-keeping across production. This could be made publicly available as a website or hosted document.

Level 3: Auditable AI record - *For auditors, legal, insurers etc.*

Full record including signed declarations, decisions and a high-level actions log, and further details about tools and licensing at a minimum. Can also include records of consent and granular actions & asset-level data, such as the machine-readable AI marks required by the EU AI Act.

Applied Research: Putting the standard into practice on FRAMES

Within the FRAMES project, there was an opportunity to attempt creating a production AI card/Level 2 record for an AI-centric production process for the first time. We also explored how the team's proposed workflow could feed into a version of Level 3 records. We explored these two approaches in tandem.

Outcomes

- The auditable record (Level 3) design evolved throughout the project and remains a work in progress. With the EU AI Act requirements coming into force during the project, we began to see machine-readable AI marks emerge on some of the AI tools used. However, ultimately we had to rely mostly on manual tagging of human vs AI-driven assets and workflows within the Streamline platform
- We successfully produced a Production AI card (Level 2) based on this information & conversations with the team. (See Figure 1)
- A Level 1 Production AI Statement has also been drafted based on this information, but not yet added into the pre-roll or credits in the final edit. (See Figure 2)

Learnings

Copyright

- The biggest overall tension when designing the best human-AI hybrid workflow on the project was between achievable quality and ensuring copyright claimability.
- Truly knowing and understanding the copyrightability of a final piece of media and its content (eg., characters) remains complex. Variables could depend on the geographical location of content creation and publication, depending on the legal interpretation of that local law. This could include IP consent, provenance, tracking, detection, chain of title, and a record/audit trail of which AI systems are used, etc.

Transparency

- Substance - A simple, human-readable record explaining how AI has been used on the project is really helpful as an internal reference, building awareness and understanding of how and why tools were considered in the process. It's difficult to make judgements on how "important" or "significant" AI usage is for any particular use case without digging further into specific usage actions (inputs, outputs, etc.), not just looking at the final product.
- Process - Although it is best to establish a process for transparency documentation from the start of a project, it can be achieved retroactively. On projects where communication happens mostly online, you can capture some decision-making and tool-usage details from meeting transcripts and chat conversations, rather than relying solely on proactive self-reporting or deep technical documentation. Using AI tools responsibly can also support this effort. However, ultimately these approaches are still error-prone, and information should always be reviewed and confirmed by the team before finalisation. Even if information can't be documented in real time from start to end of a project, it is important to at least have a responsible person or people who know the project and process well enough to help capture the info retroactively with confidence

Other

- Tradeoffs were necessary between different responsible AI aspects of the project to achieve acceptable quality in the final film, in line with professional-grade animation standards. For example: Some AI tools have safeguards in place for flagging the use of identifiable copyrighted material as inputs; however, when you already own the rights to the content, this kind of functionality is obstructive rather than useful. This can lead to choosing tools that are less protective to achieve (still legitimate & ownable) outputs.

Future thoughts: An industry-standard, globally recognised Responsible AI stamp

Agreeing on a standardised format for a credit statement / on-screen disclosure is a must-do for the industry. However, a written statement alone still requires legal sign-off and interpretability for audiences who may still have lower AI literacy. For example, judging whether it is "good" or "bad" that a production has used AI to support VFX activities is not clean-cut- the use cases and tools are the deciding factor in how responsibly this has been done.

Therefore, on top of a baseline statement, AIMICI has also proposed a further evolution: an easily understandable Responsible AI stamp that can be added to productions.

There have been attempts at this in a number of ways already across industry, from self-declared statements (eg. "No generative AI was used in the making of this film" - Heretic, 2024) to simpler icons and stamps (eg. "ProudlyHuman" mark - Slapper, 2016); however, they have been typically limited to the "no AI" stance, and there has not yet been anything that considers responsible AI working practice at large on a production.

By adhering to the combined requirements of appropriately trained staff and standardised transparency documentation, AIMICI's idea for a Responsible AI mark would look to help signify best practices across a range of productions including film & TV, scripted & unscripted, and AI-first, hybrid and traditional approaches.

PART FOUR: Conclusions

A set of emerging standards could help drive best practice in using AI in production. The pace of innovation in models, commercial pressures on traditional media distribution and advertising, and rapid inflows of investment capital have made this technology extremely fast-moving. The baseline set of metadata includes the identity of content, people, devices, and software, together with time and location, that need to be captured to support proper attribution and audit of new workflows. Where possible, this should rely on existing initiatives such as OpenTimeline, SMPTE RIS-OSVP, C2PA/CAWG. The experience gained through this project suggests summary best-practice guides that could help build new AI skills for current creative talent.

The "solopreneur" in content production: actual best practices on the frontier for major studios, advertisers and content creators

As rapid technology innovations disrupt the industry, we are seeing the rise of a new business paradigm fueled by evolving workflows and new roles. Generative AI tools are transforming freelancers into independent production houses (solopreneurs). However, the shift presents a significant challenge. Newly formed solopreneurs and small studio businesses often lack the resources, tools, and legal expertise to navigate AI policies, ethics, and governance, as well as the current standards and rules around media.

While this paper outlines foundational steps for transparency and ethical AI use in media production, it highlights the urgent need for a wider industry conversation and unified best practices/standards. Currently, the industry lacks a cohesive approach, with individual studios developing disjointed ethics and transparency frameworks and uncoordinated expectations for how others will adopt them. The industry needs to establish a baseline, holistic approach that provides resources for freelancers, solopreneurs, SMEs, and large studios. This can take the form of a curated registry of AI tools/providers that provides due diligence on training data, transparent business practices, ethics, and security measures.

By equipping productions with the insights they need to make sound technology and creative workflow choices, this project serves as a stepping stone toward a larger goal of responsible implementation and transparent innovation that drives the industry forward.

Closing remarks from Tim Deussen, Champion member of IBC FRAMES and Chairperson XRBB (Extended Realities Berlin-Brandenburg)

Five months of building FRAMES taught us that the hard problem in AI-assisted production is not generation. It is proof.

Almost every difficulty we met had the same shape. We could push for the quality a professional animation demands, or we could keep a clean claim to copyright, and holding both at once meant slowing down and writing things down. Machine-readable AI marks began appearing on our tools while the project was still running, but never on all of them, and never in a form our workflow could gather up, so in the end we tagged human and AI-driven work by hand. Safeguards designed to protect an anonymous user got in the way of a production that already owned the rights to its own material. And none of the record could be handed to software: what was made, by whom, and why was written down by people who knew the work.

That is not a complaint about the technology. It describes a missing layer. Regulation asks a production to account for itself at the level of the finished work. Our tools mark at the level of a single output. Authorship and rights live at the level of the individual asset. Nothing today carries a claim across all three, so people carry it instead.

None of us can close that gap alone, and a market of incompatible answers helps no one, least of all the artists whose work is at stake. So this paper ends with an invitation rather than a conclusion. We ask broadcasters, studios, producers, vendors, standards bodies and creator organisations to work with us on common standards for ethical AI, production transparency and artistic provenance: a disclosure record that holds good in more than one jurisdiction, provenance that survives post-production, and attribution that reaches the asset and its author, not only the finished master. Build it once, together, and international co-production stays possible. Leave it, and every territory will answer differently, and the people who pay for the difference will be the ones who made the work.

Appendices

Definitions

Within this paper, we will refer to the following terms - this is what they mean in the context of media production when AI is involved

Copyright - The legal right to claim creative ownership over a particular asset or piece of media, whether produced with AI or not. Specifically for the purposes of control, reproduction, distribution and display.

Disclosure - The action of explaining to external parties how AI has been used on a project. This is usually best supported by accurate AI transparency records, although it can also be a broader agreed statement (e.g., AI has been used on this project; no AI has been used on this project, etc.).

Transparency - Typically, transparency is discussed in the context of a record of what kind of data AI tools have been trained on. However, in this paper we are primarily focused on the transparency around AI usage and decision-making on the project - production-level AI transparency, not tool-level. To be transparent about this, a record must be created containing accurate information.

Responsible AI - The broad practice of putting effort towards ensuring that AI is used within a production environment in a responsible way. Common themes include Copyright, data security, impact on jobs, etc.

Frameworks, standards & papers - Multiple frameworks have been both considered and proposed as part of the work on this project. A framework is typically a formalised approach to putting ideas or theory into action. Please see the references page as needed for the full list:

Solopreneur - Acts as a traditional business owner who works alone. Focuses on building automated assets and using freelancers to produce deliverables.

References

- [IBC Accelerators: A Framework for Generative AI - Responsible AI Framework paper \(2025\)](#): An output of last year's IBC accelerator project, this framework

proposes a high-level approach to operationalising responsible AI processes in production

- [Production Identity Whitepaper](#)
- [EIDR.org](#) content identity
- [C2PA CAWG](#)
- MovieLabs:
 - Ontology for Media Creation (CC BY 4.0): <https://movielabs.com/ontology-for-media-creation/>
 - Spec docs <https://mc.movielabs.com/docs/ontology/>
 - RDF/JSON <https://github.com/MovieLabs/OMC>
 - The 2030 Vision: <https://movielabs.com/the-2030-vision/>
 - Building Foundations for Trusted and Reliable AI in Media Creation Workflows: <https://movielabs.com/building-foundations-for-trusted-and-reliable-ai-in-media-creation-workflows/>
- [AIMICI Transparency Whitepaper \(2026\)](#): Released earlier this year, this self-published whitepaper lays out the thesis for why the industry needs improvements in AI transparency.
- [AIMICI Responsible AI Training / Training Standard](#): A training curriculum standard & live training courses that put them into action. They cover the key responsible AI-related topics the industry should be trained on, and support operationalising AIMICI's industry-informed responsible AI best practices. Accredited by a UK-based skills body for the production industry - ScreenSkills.
- [AIMICI AI Transparency & Disclosure Standard](#) (New, proposed) AIMICI's new proposed working standard suggests how to best gather and report information about how AI has been used on a project, including key information relating to responsible use & compliance. The result is a proposed format for documentation and disclosure.

Production AI Cards example

Figure 1: Level 2: Production AI Card for One Small Crawl, V1.0

STAGE	USE CASES	TOOLS	CLASSIFICATION	
Pre-production	Script generation, drafting & ideation	ChatGPT (OpenAI)	Invisible	Significant
	Archive footage search & selection (including via rough-cut generation)	Intellisearch (Resolve); Eddie AI	Invisible	Significant
	Shotlist generation from archive metadata	ChatGPT (OpenAI)	Invisible	Not significant
	AI metadata extraction & analysis of archive footage	RAI metadata pipeline	Invisible	Significant
	AI 3D asset generation	Rodin; Meshy;	Visible	Significant
	LoRA training on the character for posed reference images	Custom LoRA (via ComfyUI)	Visible	Significant
Production	Image(s)-to-video shot generation	Veo 3.1 (Google), Kling 3.0 / Kling 3 Omni, Seedance; via Higgsfield	Visible	Significant
Post-production	SDR-to-HDR conversion	Beeble AI	Visible	Not significant
	Music-focused story-arc generation	ChatGPT (OpenAI)	Invisible	Significant
	Music generation	Suno	Visible	Significant
	Sound effect generation	Mirelo SFX 1.6 text-to-audio	Visible	Significant
	AI enhancement (upscale / denoise / in-out painting) (exploratory)	ComfyUI; Topaz	Visible	Not significant
	Logging & rough cut generation	Eddie AI	Invisible	Not significant
All stages	AI-assisted coding for pipeline & tooling	Claude Code	Invisible	Not significant
	Agentic pipeline orchestration (agent piloting Kitsu & DaVinci Resolve via MCP)	Hermes Agent; n8n	Invisible	Not significant

Figure 2: Level 1: Production AI Statement for One Small Crawl, V1.0



Production AI Statement

- No AI was used for character design or deepfakes
- AI was used to create some footage & all SFX
- Final edits were curated & signed off by humans