

Defining and Operationalising “AI Film” Within a Transmedia Storytelling Framework

Qiushi Wang

*Department of Journalism and Communication, The University of New South Wales, Sydney,
Australia*

z5463726@zmail.unsw.edu.au

Abstract. We advance an operational definition of “AI film” that integrates three lenses—computer-science systems and capabilities, communication/media studies on storyworlds and circulation, and industry/labour governance on contracts and credit. We formalize five dimensions—Generative Agency (GA), World Modularity (WM), Orchestrated Personalisation (OP), Algorithmic Distribution (AD), and Labour/Credit governance (LC)—and translate them into auditable coding criteria. Each dimension is specified with decision anchors (e.g., GA from assistive use to automated control; AD within a 7–14 day attention window) and evidence requirements (assets, logs, platform signals). Applying the framework to three contemporary cases (Marvel’s Secret Invasion title sequence, the indie horror *Late Night with the Devil* publicity stills, and TriStar’s *Here* with on-set face replacement), we map a gradient from peripheral augmentation to embedded, real-time generativity. We find that OP remains largely latent despite technical feasibility, while AD is frequently catalyzed by controversy or milestone coverage. The framework reduces false positives in classifying “AI film”, clarifies where ethical and legal risks concentrate (disclosure, consent, remuneration), and offers reproducible guidelines for empirical studies. We conclude with implications for credit attribution, contractual language, and cross-platform orchestration in transmedia environments.

Keywords: Ai film, transmedia storytelling, generative models, algorithmic distribution, labour and credit

1. Introduction

Generative AI has moved from peripheral utility to a visible force in screen production and circulation. Models that synthesise images, video, and sound now support concept art, previs, grading, and even final shots; at the same time, platform recommenders shape how works find audiences [1,2]. Contracts and guild guidance have started to codify these shifts, especially around authorship, credit, and digital replica [3-5]. Yet scholarship and practice still use “AI film” inconsistently. Some label any AI-assisted workflow as AI film; others reserve the term for model-generated footage. This inconsistency blurs comparisons across studies and complicates policy and credit.

We address that gap by advancing a concise definition of AI film that integrates three lenses—computer science (systems and capabilities), communication and media studies (storyworlds and circulation), and industry/labour (contracts and market infrastructure)—and operationalises it for empirical analysis. Specifically, we (1) formulate an interdisciplinary definition grounded in established literatures [6], (2) translate it into coding criteria that distinguish routine automation from creative control, and (3) demonstrate its utility through illustrative cases, examining implications for credit, consent, and distribution. Applying these criteria reduces false positives in classification and clarifies where ethical and legal risks concentrate. We first situate related work, then define and operationalise 'AI film', apply the criteria to three contemporary cases, and discuss implications for labour and audiences.

2. Background: industry, labour and policy context

Since early 2024, text-to-video systems moved from lab demos to tools with public previews. OpenAI announced Sora, which produces minute-long clips with complex scenes [7]. Google DeepMind introduced Veo, presented as a state-of-the-art video generator and now tied to the Gemini product line [8]. Runway released Gen-3 Alpha in June 2024, targeting professional workflows for pre-visualisation and short-form production [9]. These launches normalised generative video in creative pipelines and raised questions about authorship, labour, and standards.

Labour settlements in Hollywood formalised those questions. The Writers Guild of America (WGA) 2023 Minimum Basic Agreement (MBA) states that AI cannot write or rewrite literary material and that AI-generated text is not “source material” for credit [3]. The Screen Actors Guild–American Federation of Television and Radio Artists (SAG-AFTRA) 2023 Television (TV)/Theatrical framework and bulletins require informed consent and compensation when a performer’s digital replica is created or used [4,5]. They also distinguish digital replicas from fully synthetic performers. Although implementation remains contested, these provisions set default expectations at scale.

Policy has also shifted. The EU Artificial Intelligence Act (Regulation (EU) 2024/1689, adopted 13 June 2024) introduces transparency duties for certain AI systems, including synthetic media; Article 50’s obligations begin to apply in 2026 [10,11]. For cross-border releases and platforms, these provisions signal labelling and disclosure norms for generative content.

Public debate sharpened around concrete film and TV cases. Marvel’s *Secret Invasion* faced backlash for an AI-generated title sequence in June 2023, crystallising worries about replacing creative labour [12,13]. In March 2024, the indie horror *Late Night with the Devil* drew criticism for several AI-generated stills; the directors confirmed the use and defended it as limited [14]. These episodes show that audience reception and professional norms are still in flux.

Finally, distribution itself is algorithmic. Large-scale recommender systems shape what reaches whom and when, linking creative choices to market outcomes [1]. This environment makes AI film not only a production question but a circulation question.

We therefore define the term with care, align inclusion rules to that definition, and synthesises evidence across computer science, communication, and industry practice.

3. Defining “AI film”

Scholars and practitioners use “AI film” in different ways. Placing the definition at the start clarifies scope, evidence, and coding rules for this article.

3.1. Computer science lens

From a systems perspective, an AI film is a screen work in which machine-learning models synthesise or transform a material share of the audiovisual content, or decisively automate creative choices in pre-production, production, or post.

This view anchors the term in generative models that have enabled text-to-image and text-to-video pipelines—GANs [15], diffusion models [16], latent diffusion for high-resolution imagery [17], and early text-to-video systems [18].

3.2. Communication and media-studies lens

Communication and narratology foreground how storyworlds travel across media and how audiences meet them. In this lens, an AI film is a film whose storyworld, circulation, or audience address is shaped by algorithmic systems—during creation, cross-platform orchestration, or distribution.

The definition builds on work in transmedia storytelling and media-conscious narratology [19] and on studies of algorithmic culture and automated media. In short, computation is not only a tool; it is an actor in meaning-making and reach [20,21].

3.3. Industry and labour lens

Industry practice makes the term operational through contracts, credits, and market infrastructure. For writers, the 2023 WGA MBA states that AI cannot write or rewrite literary material and that AI-generated text is not source material for credit [3,22].

For performers, SAG-AFTRA's 2023 TV/Theatrical agreements and related guidance require informed consent and compensation for digital replicas and govern certain digital alterations [4,23]. On the market side, large-scale recommender systems shape discovery and viewership; this affects how AI-assisted films find audiences [1].

3.4. Working definition (this study)

We define an AI film as a screen work in which machine-learning or allied algorithmic systems (i) generate or transform a material share of the audiovisual text, or decisively automate creative choices in pre-production, production, or post; or (ii) orchestrate cross-platform versioning or audience address beyond routine automation; or (iii) the work's labour, credit, or consent arrangements materially hinge on AI-specific provisions (e.g., digital replicas or synthetic performances) such that what is seen on screen or how it is billed would differ absent those systems.

Inclusion criteria and notes.

(1) Material may be established quantitatively (e.g., proportion of minutes or shots) or qualitatively (e.g., proximity to a principal performance, key narrative events, or a defining stylistic signature).

(2) Orchestrate extends beyond default utilities (such as automatic captioning) to include documented A/B or territory/segment variants, conditional asset delivery, or algorithmic assembly that shapes narrative address.

Exclusion criteria and edge cases.

(1) Routine utilities (upscaling, denoising, automatic captioning) are insufficient unless they displace creative control or alter meaning.

(2) Experimental assets that do not appear in the released text are excluded; paratexts integral to release (e.g., main titles) may be included.

(3) Purely algorithmic distribution effects (e.g., generic recommender exposure) are insufficient without on-text generation/transformation or orchestrated versioning.

While the technical capacity for content-layer versioning is clear, public evidence of audience-tailored on-text variants at scale remains limited. Deployed personalisation concentrates in packaging assets rather than narrative beats—for example, Netflix’s artwork/thumbnail personalisation with A/B or bandit-style selection [24,25]. Authorial judgement (methods note): On this basis, OP is emergent but not yet routine in 2024.

4. Theoretical bases of transmedia narrative

This section situates the study in established transmedia theory and adjacent literatures on participation and platformised circulation. It moves from core definitions to system properties and governance. The aim is to provide conceptual anchors that the next section can translate into operational criteria.

4.1. Origins and definitions

Kinder first described entertainment “supersystems” in children’s media, where film, television, games, and toys co-produce a shared world [26]. The point was not simple tie-ins. It was a designed ecology in which each medium extends the whole.

Jenkins popularised the term transmedia storytelling as the planned dispersion of integral story elements across media, with each medium adding something distinctive to the unfolding plot or world [6]. Ryan and Thon later argued for a media-conscious narratology. In their view, narrative is a cognitive construct that can migrate between modalities while keeping the storyworld coherent [19]. However, these approaches foreground different emphases: Jenkins’s cultural studies lens centres collective meaning-making, whereas Ryan’s narratological angle prioritises structural coherence—one treats transmedia primarily as cultural practice, the other as a formal system of narrative logic.

4.2. Storyworlds, seriality, and modularity

Robust storyworlds are the substrate of transmedia practice. Audiences navigate characters, places, and timelines that persist beyond any single text [19,27]. Coherence is a design task. It depends on stable rules for space, time, and identity.

Seriality sustains attention by segmenting experience and inviting return [28]. Production then follows with modular craft. Characters, settings, and assets are built for reuse and recombination across channels and moments [6,29]. In practice, this modularity underpins franchise economies: Marvel’s character arcs are engineered to migrate between cinema, streaming, and games. Outside the U.S., China’s *Three-Body Problem* illustrates a regionally specific, multi-platform build-out—audio drama on Ximalaya (Jan 2022), animation on Bilibili (Dec 2022), and a live-action TV series on CCTV/Tencent (Jan 2023)—showing how assets are recombined for distinct channels and publics.

Nonetheless, modular recombination alone does not guarantee coherence; the craft problem remains how to stage additions without fracturing the world’s internal rules.

4.3. Participation and spreadability

Participation is a native feature of transmedia, not an add-on. Fans interpret, annotate, and sometimes produce materials that travel with the official text [6]. These practices turn reception into a feedback engine—discussion threads, fan edits, and wikis add value to the “official” property.

Spreadable media reframes how texts move in networks when people choose to pass them on [30]. At the same time, “spreadability” was pitched against earlier metaphors of “stickiness,” which foregrounded platform retention. The question remains whether user agency or infrastructural capture should be analytically primary. In the Korean webtoon economy, for instance, fans’ transcreation through LINE Webtoon Translate has documented how platform affordances channel grassroots labour into IP globalisation—highlighting both empowerment and normalisation of unpaid/under-credited work.

4.4. Algorithmic culture and platformisation

Today, circulation is platform-shaped. At scale, recommendation systems rank, route, and time materials for users [1,2]. Visibility is no longer only a marketing outcome. It is an interface outcome. Netflix’s thumbnail testing and TikTok’s “For You” feed are emblematic. Recent research characterises TikTok publics as “sides”—loosely bounded, algorithmically assembled clusters co-produced by recommendation, vernacular tags, and creator practices—showing how curation routines and community discourse jointly shape visibility [31,32]. This supports our use of OP/AD as distinct yet interacting layers in platformised circulation.

Beyond TikTok, research on Douyin (the Chinese counterpart) details how “traffic rewards” and algorithmic visibility shape creator strategies and labour conditions; related studies trace Douyin’s role in national discourse formation—indicating that recommendation regimes are also governance regimes.

Nonetheless, platformisation is not monolithic. Comparative work emphasises how markets, infrastructures, and governance vary across regions and industries, urging analysts to specify platform logics rather than assume a single global model.

4.5. Authorship, labour, and credit

Transmedia multiplies sites of authorship. Writing, direction, design, marketing, and data work all shape the world that audiences encounter [29]. Branding and licensing add further constraints and opportunities.

Governance therefore matters. Contractual language, disclosure norms, and credit taxonomies decide who is recognised and paid. These arrangements influence which assets can be reused, which variants are legitimate, and how audience trust is maintained. Disputes over authorship in major franchises show that credit allocation is part of the infrastructure of transmedia, not a peripheral afterthought.

4.6. Synthesis: implications for operationalisation

Three implications follow for the measurement task in the next section.

First, strong storyworld design and serial craft reduce friction for cross-media extension—an empirical claim testable via asset inventories and reuse evidence.

Second, platformised circulation requires a clear analytic split between supply-side versioning and demand-side ranking. One concerns how materials are assembled for contexts and locales; the

other concerns how platforms surface those materials to users. Importantly, these logics can reinforce or undermine each other: careful localisation can be muted by ranking suppression, while algorithmic amplification can elevate variants never designed for centrality. Region-specific regimes (e.g., TikTok vs. Douyin) further complicate outcomes and should be measured rather than assumed.

Third, labour and credit governance sets practical limits. Consent, disclosure, and compensation affect what can be made, what can be reused, and how audiences respond. These constraints are not external to storytelling. They shape the arc of a transmedia project. Taken together, world design, circulation logics, and governance frame the operational criteria developed in the next section. These conceptual anchors inform the layered framework below, which we then translate into scoring rules.

5. Theoretical framework: AI film within a transmedia system

We propose a layered framework that links narrative form, computational agency, and labour/market governance. The aim is explanatory: to account for how AI changes what is produced, how it circulates, and who is credited.

(1) Generative agency (GA).

Machine-learning systems can synthesise or transform audiovisual material and can automate creative choices. GA captures that capacity and its share in the final cut. It rests on advances in GANs and diffusion models [15-17]. Anchors: none → assist → co-create → automate. Evidence: tool logs, prompts/scripts, real-time overrides.

(2) World modularity (WM).

Transmedia worlds are modular. Plots, characters, and assets can be recombined across media. WM draws on multiplicity and seriality from transmedia studies [6]. Operational focus: reusable, parameterisable assets (models/rigs/prompts) with versioning/APIs. Evidence: asset graphs, documented cross-title/scene reuse.

(3) Orchestrated personalisation (OP).

Algorithms can assemble versions, cut trailers, and route audiences across platforms. OP names this layer of algorithmic orchestration that exceeds routine automation. It aligns with work on algorithmic culture and automated media [20,21]. Evidence: conditional delivery/A-B cuts, targeting flags, regionalisation logs.

(4) Algorithmic distribution (AD).

Recommender systems shape exposure and timing. AD denotes the distributional environment that filters discovery and viewing [1]. Measurement: attention dynamics within a 7–14-day window (media/article counts, trend curves, platform amplification signals, sentiment). Measurement (operational threshold). We treat the first 7–14 days after release as an operational threshold for attention dynamics, aligned with weekly reporting cycles on major SVODs [33], and with evidence that early view trajectories predict long-run popularity and that bursts decay with characteristic patterns [34,35]. We report primary counts in this 7–14-day window and run robustness checks at 3 days (opening-weekend sensitivity) and 28/91 days (long-tail). See Table 1 for robustness checks across 3/7/14/28/91-day windows.

Table 1. Robustness checks for AD attention window sizes (3/7/14/28/91 days)

Window (days)	Primary purpose	Operational justification (APA in-text)	Key indicators to compute	Interpretation guide	Reviewer checklist
3	Opening burst sensitivity; weekend/front-page amplification check	Captures launch shock; early popularity predictive power [34,35].	• 3-day share of 91-day cumulative views (%)• Spearman $\rho(\text{rank_3d}, \text{rank_91d})$ • $R^2: \log(\text{views_91d}) \sim \log(\text{views_3d})$ • Exposure attribution: recommender surfaces vs. search/direct	High 3-day share with rapid decay \Rightarrow front-loaded promo-driven titles; early flag but unstable for slow-burners.	Report CIs; stratify by title type (film/series) and region.
7	Align with weekly reporting cycles; early trajectory predictiveness	Matches Monday–Sunday cycles [33]; early views predict long-run outcomes [35].	• 7-day share of 91-day cumulative views (%)• Spearman $\rho(\text{rank_7d}, \text{rank_91d})$ • $R^2: \log(\text{views_91d}) \sim \log(\text{views_7d})$ • Day-of-week pattern (Mon–Sun)	Balances signal and noise; stable for cross-title comparisons in week one.	Explain anomalies (holiday effects, outages).
14	Two-week amplification window; captures second-week decay/word-of-mouth	Still within burst-and-decay regime; tests second-week effects [34].	• 14-day share of 91-day cumulative views (%)• Δ growth 7 \rightarrow 14 days (slope)• Retention ratio: 14d/7d• Rank stability: rank_14d vs. rank_91d	Separates slow-burners from front-loaded hits; checks rank stabilisation by week two.	Segment by release cadence (season drop vs. weekly episode).
28	Near-term tail; consolidates launch + early long-tail	Captures ~1-month accumulation; beyond initial burst, pre-quarter [34].	• 28-day share of 91-day cumulative views (%)• Δ growth 14 \rightarrow 28 days• Rank stability: rank_28d vs. rank_91d• Category-specific tail behaviour	Useful for quarter-end forecasting; flags durable discovery beyond week two.	Disclose catalogue effects (library vs. new release).
91	Quarter-scale reference horizon for cumulative attention	Biannual engagement horizons in streamer reporting [36,37].	• Cumulative views at 91 days (baseline)• Exposure share: recommender surfaces vs. search/direct• Final quarter rank• Burst class (endogenous/exogenous)	Reference horizon to evaluate early-window predictiveness and rank stability.	State coverage; align time zones; deduplicate bundles.

Notes: 'N-day share of 91-day cumulative views (%)' = $100 \times [\text{cumulative views from release to +N days}] / [\text{cumulative views from release to +91 days}]$. Use platform day boundaries/timezone. Titles observed for <91 days should be marked as censored and reported as 'N-day share of D-day cumulative views'. If using watch-hours, replace 'views' accordingly. References: Crane & Sornette (2008) [29]; Szabo & Huberman (2008) [35]; Netflix Top-10 methodology (2021) [33]; Netflix Engagement Reports (2023/2024) [36,37].

(5) Labour and credit governance (LC).

Contracts and guild rules set limits on AI use and define credit, consent, and compensation, including digital replicas [3-5]. Evidence: contractual clauses, disclosure statements, consent records, credited roles, union compliance.

Core propositions.

P1. Higher GA increases WM by lowering the cost of producing coherent variants and inserts.

P2. WM and OP jointly raise AD exposure, controlling for spend, because modular worlds travel well and can be versioned for niches.

P3. LC moderates the effect of GA on OP: strong consent/credit rules constrain certain automated uses while legitimising others.

P4. Audience participation (from the transmedia tradition) mediates the link between AD and reception: algorithmic exposure converts to engagement when audiences can act on, circulate, or co-produce materials [6,29].

Scope conditions.

The framework applies to projects where AI affects either (i) the composition of shots/sequences, (ii) cross-platform orchestration, or (iii) labour/credit arrangements. It does not treat routine utilities (e.g., denoising) as GA unless they shift creative control.

Operationalisation (summary).

GA: % of shots/minutes or key creative decisions attributable to model outputs.

WM: count of distinct media nodes and unique reusable assets tied to the same storyworld.

OP: presence of algorithmic assembly/versioning beyond format defaults; documented A/B variants.

AD: share of initial views attributable to recommender surfaces; attention dynamics within a 7–14-day operational threshold (weekly SVOD cycles; early-views predictiveness; burst-decay), with robustness checks at 3 and 28/91 days.

LC: existence of AI-related consent/credit provisions (e.g., WGA/SAG-AFTRA compliance) and their scope.

This framework underpins our definition and coding criteria in the next subsection and structures the analysis that follows.

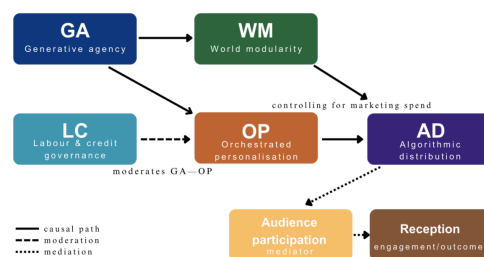


Figure 1. Propositional path model linking GA, WM, OP, AD, LC, and reception (via audience participation)

6. Production pipeline: from script to release

Building on the framework above, We follow the path of a project. We move from script, to picture, to release. At each step we note the mechanism, the constraint, and the variable it touches (GA, WM, OP, AD, LC).

6.1. Script generation

Work starts with words. Large language models help with loglines, beats, scenes, and dialogue. Dramatron shows how prompts can keep a story hierarchy across title, characters, locations, and scenes [38]. In user studies, professionals co-wrote short plays and screenplay fragments with the system. This is assistance, not wholesale authorship. It still sits under credit and disclosure rules. The WGA 2023 MBA says AI text is not “source material” and cannot replace human writing [3].

Thus, at the writing stage, GA is present but LC constrains its reach. This stage also seeds WM by turning ideas into modular beats that can travel across media.

Once words lock, the task shifts. Plans must become shots. The generative problem moves from discrete tokens to continuous space-time.

6.2. Video imagery

Text-to-video diffusion now produces short shots with consistent motion and style. Imagen Video uses a cascade of video diffusion and spatio-temporal super-resolution [18]. Make-A-Video learns text-conditioned clips without paired text-video data [39]. Lumiere renders an entire clip in a single space-time pass to improve temporal coherence [40]. Sora brings these ideas into a product context for high-definition short-form assets [7]. These models enable inserts, establishing shots, previs, and stylised passages. Here GA can be high. When creators reuse generated assets across platforms, WM rises. When systems help assemble variants for different audiences, OP enters the workflow. LC still matters if digital replicas or synthetic performances are involved.

Footage cut into the timeline is not the end. Attention is scarce. Exposure depends on surfaces that rank and route content.

6.3. Release and circulation

Distribution is algorithmic. On streaming platforms, recommenders decide what appears, to whom, and when. The Netflix stack links candidate generation and ranking to watch-time goals [1].

YouTube uses deep retrieval and ranking at very large scale [2]. These systems shape the early audience for any AI-assisted film. Policy can also alter trust and marketing. Under the EU AI Act, certain synthetic media require transparency measures [11]. In our terms, this stage foregrounds AD. It also feeds back into OP, since platform signals inform versioning and routing.

6.4. Pulling the stages together

Across the three stages, GA enables new assets, WM/OP organise them across media and audiences, AD governs discovery, and LC moderates what is allowed and how credit is assigned.

We next turn these observations into concrete inclusion rules and coding criteria.

7. Case 1: Secret Invasion (Disney+, 2023) and AI-generated title design

Marvel's limited series *Secret Invasion* premiered on Disney+ on 21 June 2023. On premiere day, the show's director and executive producer, Ali Selim, confirmed that the main title sequence was created with AI tools by Method Studios. Selim framed the decision as thematically motivated by the series' story of shape-shifting Skrulls infiltrating human institutions [41]. Method Studios later issued a statement saying that a "custom AI tool" was used as one component in a pipeline that also relied on conventional art direction, 2D/3D animation, and compositing, and that "no artists' jobs were replaced" [41]. The disclosure triggered an immediate public backlash from artists and fans who criticised Marvel for using generative systems during an industry-wide labour dispute over AI's role in creative work [13].

GA — Generative Agency (medium).

The AI system influenced the visual form of an entire credit sequence. Selim described an iterative prompt-and-refine process ("we would talk to them about ideas and themes and words, and then the computer would go off and do something"), indicating non-trivial generative agency in

producing imagery [41]. Yet the sequence sits outside narrative action; it does not generate story beats or dialogue. It therefore registers as medium GA: significant authorship of a paratextual asset, not the diegesis itself.

WM — World Modularity (low).

Despite being a distinct asset, the credits' painterly, morphing aesthetic did not propagate as a reusable module across trailers, character posters, or ancillary Marvel materials in the way some title-design systems do. Publicly available coverage offers no evidence of a modular "asset pack" derived from the AI process for cross-media re-use. WM is low.

OP — Orchestrated Personalisation (none/low).

There is no evidence of audience-specific or territory-specific variations in the title sequence driven by algorithmic personalisation. Disney+ did not advertise any A/B tested or personalised versions. OP is none/low.

AD — Algorithmic Distribution (medium, indirect).

The controversy itself operated as an unplanned distribution engine. Entertainment and tech outlets quickly amplified the story; social platforms re-circulated clips of the credits with hostile commentary [42,43]. This earned attention likely altered recommendation dynamics on video platforms and news feeds during the release window. Although not an in-product algorithmic intervention, the event shows how controversy can function as an "attention router" in the attention economy. AD is medium.

LC — Labour/Credit governance (medium to high).

Method Studios' statement emphasising that no artists lost jobs and that AI was "just one tool" addresses key LC concerns: displacement, crediting, and consent in training data [41]. The timing placed the incident against the Writers Guild and SAG-AFTRA negotiations on AI protections; public discourse connected the opening titles to broader anxieties about synthetic labour and digital replicas [13]. While no formal grievance emerged from unions, *Secret Invasion* became an archetype in debates over disclosure standards and whether AI participation should entail distinct on-screen credits. LC is medium to high.

What this case contributes to this study?

Secret Invasion demonstrates how paratextual generative design can reframe public reception of a work, even when the core narrative remains human-crafted. It suggests that, within franchise storytelling, relatively small generative insertions can trigger outsized governance debates, particularly when the insertion intersects with labour negotiations. Analytically, it sharpens the boundary between use (AI as a pipeline tool) and credit (recognition, disclosure), and shows how AD can be catalysed indirectly by LC flashpoints.

8. Case 2: *Late Night with the Devil* (IFC/Shudder, 2024) and "three stills" of AI artwork

The independent found-footage horror film *Late Night with the Devil* opened in US theatres on 22 March 2024 and streamed on Shudder from 19 April 2024. As the film gained acclaim and notable box office for an indie release, viewers highlighted the presence of AI-generated images in a few interstitial title cards styled as 1970s broadcast graphics. Directors Cameron and Colin Cairnes issued a statement, via *Variety* and reported by outlets including *Entertainment Weekly*, clarifying that they had "experimented with AI for three still images," which were subsequently edited by the production team and appear only briefly [44].

GA — Generative Agency (low).

The use is narrowly bounded: three stills, used as transitional cards, not narrative shots. While this qualifies as generative agency, it is low in scope and influence on story meaning. The directors'

framing—AI as a minor element in a larger analogue aesthetic—supports this classification [44].

WM — World Modularity (low).

Interstitial cards are not designed as portable world-modules; they mimic broadcast ephemera. There is little evidence of re-use across promotional media beyond screenshots circulating on social platforms. WM is low.

OP — Orchestrated Personalisation (none).

There is no evidence of personalised variants or algorithmic tailoring. OP is none.

AD — Algorithmic Distribution (medium, event-driven).

The AI revelation created a discrete controversy cycle during opening weekend; coverage in genre press and social feeds amplified discourse around the film’s ethics even as reviews remained positive [45,46]. That cycle likely acted as an exogenous shock to visibility, functioning as controversy-as-distribution similar to Case 1, though at indie scale. AD is medium.

LC — Labour/Credit governance (medium).

Here the crux is disclosure and proportionality. The directors publicly enumerated the extent of AI use (three stills), credited their “amazing graphics and production design team,” and contextualised the shots as minor experiments [44]. This explicit quantification aids governance: it lets auditors test claims against the finished artifact (time-code the interstitials; compare to generative artefacts). At the same time, the backlash shows that even de minimis uses can trigger normative pushback in communities committed to human-made craft. From a governance perspective, the case tests whether materiality thresholds—e.g., below a certain duration, different disclosure or credit rules might apply—are acceptable to audiences and guilds. LC is medium.

What this case contributes to this study?

Late Night with the Devil is a clean micro-dose example. It reveals a non-linear relationship between the amount of AI and the intensity of public reaction: minimal use can provoke maximal discourse when the symbolic stakes are high. For our framework, the case is valuable because it decouples GA from LC and AD: low GA can still produce medium AD via an LC-triggered debate about credit, consent, and artistic authenticity. It also offers a practical blueprint for evidence-based disclosure (precise counts, brief durations) that future productions can emulate.

9. Case 3: Here TriStar Pictures, 2024) and real-time generative face transformation

Robert Zemeckis’ Here adapts Richard McGuire’s graphic novel set in a single living room across centuries. Released by TriStar Pictures in November 2024 with a reported \$50 million budget, the film used Metaphysic’s real-time AI face-transformation system to present Tom Hanks and Robin Wright across roughly six decades of life within shots, with results visible on set during filming [47]. Metaphysic describes its approach as high-resolution, photorealistic face-swap and de-ageing “live and in real time,” reducing the need for extensive post-production compositing [48]. The production trained custom models on prior footage of the actors to achieve fast, frame-coherent transformations [47].

GA — Generative Agency (high).

Unlike the previous cases, AI here participates in principal photography, altering the visible performance continuously rather than substituting a title card or a few interstitials. The system outputs a synthetic visage that audiences read as “Hanks” or “Wright,” but at different ages, fused to live performances. Although blocking, line delivery, and scene construction remain human-driven, the visual identity of the character is co-authored by the generative model. That constitutes high GA in our framework.

This case also expands the scope of agency: the model makes per-frame choices to maintain identity coherence across lighting, lensing, and motion. The Here pipeline therefore functions as an algorithmic cinematography layer, not merely a post-production filter.

WM — World Modularity (medium).

The film’s de-ageing assets are potentially modular: publicity materials can showcase character life-stage variants, and the trained models could—subject to consent—be reused for ancillary marketing, behind-the-scenes reels, or future projects. Public sources do not confirm such cross-media reuse, but the capacity is evident. WM is medium.

OP — Orchestrated Personalisation (low, potential).

There is no public evidence that TriStar generated audience-specific promos where character ages are tailored to segments. But the existence of trained, controllable models implies a future capacity for versioning (e.g., different age beats for different trailers). For now, OP is low.

AD — Algorithmic Distribution (medium, tech-news amplification).

Coverage in *Wired* (via *Ars Technica*) and other outlets framed Here as a milestone in “AI-powered visual effects” at studio scale, which likely increased discovery among tech-interested audiences and within recommender systems [47]. This is a milder, more affirmative form of the AD dynamic seen in Cases 1–2: technology news acts as a distribution channel, but without a scandal frame. AD is medium.

LC — Labour/Credit governance (high salience).

By necessity, Here intersects with the most sensitive LC topics: digital replicas, training on past performances, and the scope of consent. Since late 2023, SAG-AFTRA’s TV/Theatrical agreement and subsequent union guidance have codified informed consent and compensation for digital replicas and AI-altered performances [4,49]. The core LC questions for Here are: Did performers consent to the scope of training on archival footage? What is the compensation framework for AI-altered screen time? How are credits assigned across VFX, AI engineering, and performers when a synthetic face model materially shapes perception? Public pieces indicate that on-set real-time previews shaped creative decision-making, which strengthens the case for transparent crediting of the AI VFX team as co-craft rather than invisible infrastructure [47]. LC salience is high.

What this case contributes to this study?

Here is an exemplary case of embedded generativity. It shows how AI can become a cinematic instrument used live, altering the acting image as it is recorded. That materially shifts the locus of agency in screen performance and intensifies the need for robust consent frameworks and credit taxonomies. It also demonstrates a plausible viability model: a mid-budget studio film using AI to expand narrative time without recourse to multiple casts or months of de-ageing post-work.

10. Cross-case synthesis on GA / WM / OP / AD / LC

GA gradient: From *Late Night with the Devil* (low GA: three stills), to *Secret Invasion* (medium GA: a full title sequence), to *Here* (high GA: live transformation of principal actors), the cases map a clear gradient of how deeply generative systems can permeate a screen work. This supports your argument that GA should be measured not only by minutes of AI content but also by proximity to core story functions (title paratext vs. screen performance).

WM potential vs. practice: All three cases indicate untapped modularity. Even when AI yields reproducible styles or controllable models, public evidence shows limited re-deployment across transmedia artefacts. That suggests that studios remain cautious about overtly marketing AI-derived modules, aware of reputational risks (Case 1 and 2), or that production agreements constrain reuse (Case 3).

OP as a future frontier: None of the cases deployed audience-tailored variants. Yet Case 3 makes OP technically feasible at scale: controllable, actor-specific models could enable personalised trailers or artwork under appropriate consent. This gap between capability and practice is an ethical “holding pattern” created by evolving union rules [4] and unsettled audience norms [13,44].

AD as controversy or milestone effect: In Cases 1–2, AD is catalysed by backlash; in Case 3, by milestone tech coverage. Both are algorithmic attention events that alter visibility without any change to the underlying text. That is analytically useful: it shows your AD construct can be operationalised via media-event coding (volume, tone, platform spread) even when the product itself does not algorithmically personalise content.

LC as the keystone dimension: Each case converts quickly into questions of consent, compensation, and credit. Case 1 concerns disclosure and job protection narratives. Case 2 stress-tests de minimis thresholds for disclosure. Case 3 collides directly with digital replica rules: if a live system trains on an actor’s past screen corpus, LC demands clear informed consent, option scopes, and pay formulas that reflect the ongoing economic value of the trained model. Recent SAG-AFTRA agreements and FAQs explicitly foreground consent, control, and compensation for digital replicas; those provisions become the primary interpretive lens for Case 3 [4,23,49].

11. Methods note

For each case, we can operationalise the five dimensions with simple, auditable measures:

GA: % of total runtime directly altered or generated; qualitative proximity to narrative core (paratext vs. performance).

WM: Count of distinct assets (trailers, posters, social shorts) that reuse AI-derived modules; evidence of reusable pipelines.

OP: Evidence of versioning/personalisation (territory or audience segment variants).

AD: Media-event coding (number of articles within first 7 days; social engagement); identification of controversy vs. milestone frames.

LC: Presence of disclosure statements; credit lines; alignment with SAG-AFTRA provisions on digital replicas/informed consent.

All five are implementable through desk research and simple content analysis.

12. Conclusion

This paper defines AI film with an interdisciplinary lens and makes it operational. We link narrative form, computational agency, and labour governance to explain how works are made, how they circulate, and who is credited. The focus is explanatory rather than promotional. Our aim is to move debate from slogans to evidence.

We draw on transmedia theory to foreground storyworld design, serial craft, and modular assets. We then translate these ideas into measurable criteria that separate routine automation from creative control. The result is a compact framework that analysts can score with consistency and replicate across cases.

When applied to recent examples, the criteria yield stable judgements. They reduced false positives in classification and exposed where ethical and legal risks concentrate. Small generative insertions can have large distribution effects. Governance around consent, disclosure, and credit often decides whether those insertions are legitimate and accepted.

We offer three contributions. First, a precise definition that integrates computer science, media studies, and labour perspectives. Second, a set of coding rules that produce comparable

measurements across titles and pipelines. Third, a bridge between scholarship and practice that makes audit and review feasible for editors, producers, and regulators.

There are practical implications. Creators and firms can use the criteria to plan asset reuse, versioning, and localisation with clearer boundaries. Unions and standards bodies can adapt the indicators for credit taxonomies and disclosure norms. Platforms can align content labelling and ranking policies with transparent evidence rather than ad hoc rules.

Our study has limits. It focuses on publicly verifiable materials and may understate private workflows. Genre and market effects vary, especially outside English-language contexts. Future work should test the criteria longitudinally, expand cross-regional samples, and include audience studies on reception and trust.

Even with these limits, the central claim stands. Clear definitions and operational criteria make AI film research cumulative. They also make industry choices legible. By grounding analysis in transmedia craft and measurable signals, we offer a starting point for evidence-based dialogue about credit, consent, and distribution in the next wave of screen production.

References

- [1] Gomez-Uribe, C. A., & Hunt, N. (2015). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, 6(4), Article 13. <https://doi.org/10.1145/2843948>
- [2] Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)* (pp. 191–198). ACM. <https://doi.org/10.1145/2959100.2959190>
- [3] Writers Guild of America. (2023a). 2023 MBA: Contract changes—AI FAQ. Retrieved from <https://www.wga.org/contracts/contract-2023/contract-changes>
- [4] SAG-AFTRA. (2023a). TV/Theatrical contracts 2023—Summary of tentative agreement. Retrieved from https://www.sagaftra.org/sites/default/files/sa_documents/TV-Theatrical_23_Summary_Agreement_Final.pdf
- [5] SAG-AFTRA. (2024). Terms to know. Retrieved from <https://www.sagaftra.org/terms-know>
- [6] Jenkins, H. (2006). *Convergence culture: Where old and new media collide*. New York University Press.
- [7] OpenAI. (2024, February 15). Sora: Creating video from text. OpenAI Blog. Retrieved from <https://openai.com/index/sora/>
- [8] Google DeepMind. (n.d.). Veo. Retrieved September 18, 2025, from <https://deepmind.google/technologies/veo/>
- [9] Runway. (2024, June 17). Introducing Gen-3 Alpha: A new frontier for video generation. Runway Research. Retrieved from <https://runwayml.com/research/introducing-gen-3-alpha>
- [10] European Union. (2024). 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). *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- [11] European Commission. (2024). AI act – Application timeline. *Shaping Europe’s Digital Future*. Retrieved from <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [12] Forristal, L. (2023, June 21). Artists are upset that 'Secret Invasion' used AI art for opening credits. *TechCrunch*. <https://techcrunch.com/2023/06/21/marvel-secret-invasion-ai-art-opening-credits/>
- [13] Watercutter, A. (2023, June 21). Marvel’s Secret Invasion used AI for its title sequence, and people are not happy. *WIRED*. Retrieved from <https://www.wired.com/story/marvel-secret-invasion-ai-credits/>
- [14] Lammers, T. (2024, March 22). 'Late Night With The Devil' filmmakers acknowledge using AI in “three still images”. *Forbes*. Retrieved from <https://www.forbes.com/sites/timlammers/2024/03/22/late-night-with-the-devil-filmmakers-acknowledge-using-ai-in-three-still-images/>
- [15] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems 27 (NIPS 2014)*. https://proceedings.neurips.cc/paper_files/paper/2014/file/f033ed80deb0234979a61f95710dbe25-Paper.pdf
- [16] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*. <https://arxiv.org/abs/2006.11239>
- [17] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*

- (CVPR) (pp. 10684–10695). https://openaccess.thecvf.com/content/CVPR2022/html/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.html
- [18] Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., ... & Fleet, D. J. (2022). Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv: 2210.02303. <https://arxiv.org/abs/2210.02303>
- [19] Ryan, M.-L., & Thon, J.-N. (Eds.). (2014). *Storyworlds across media: Toward a media-conscious narratology*. University of Nebraska Press.
- [20] Napoli, P. M. (2014). Automated media: An institutional theory perspective on algorithmic media production and consumption. *Communication Theory*, 24(3), 340–360. <https://doi.org/10.1111/comt.12039>
- [21] Striphas, T. (2015). Algorithmic culture. *European Journal of Cultural Studies*, 18(4–5), 395–412. <https://doi.org/10.1177/1367549415577392>
- [22] Writers Guild of America. (2023b). 2023 Minimum Basic Agreement (MBA) – Summary. Retrieved from <https://www.wga.org/contracts/contract-2023>
- [23] SAG-AFTRA. (2023b). Contract bulletin: A.I. synthetic performer vs digital replicas. Retrieved from <https://www.sagaftra.org/contract-bulletin-ai-synthetic-performer-vs-digital-replicas>
- [24] Krishnan, G. (2016, May 3). Selecting the best artwork for videos through A/B testing. Netflix TechBlog. <https://netflixtechblog.com/selecting-the-best-artwork-for-videos-through-a-b-testing-f6155c4595f6>
- [25] Chandrashekar, A., Amat, F., Basilico, J., & Jebara, T. (2017, December 7). Artwork personalization at Netflix. Netflix Technology Blog. <https://netflixtechblog.com/artwork-personalization-at-netflix-c589f074ad76>
- [26] Kinder, M. (1991). *Playing with power in movies, television, and video games: From Muppet babies to Teenage Mutant Ninja Turtles*. University of California Press.
- [27] Herman, D. (2002). *Story logic: Problems and possibilities of narrative*. University of Nebraska Press.
- [28] Mittell, J. (2015). *Complex TV: The poetics of contemporary television storytelling*. New York University Press.
- [29] Scolari, C. A. (2009). Transmedia storytelling: Implicit consumers, narrative worlds, and branding in contemporary media production. *International Journal of Communication*, 3, 586–606. <https://ijoc.org/index.php/ijoc/article/view/477>
- [30] Jenkins, H., Ford, S., & Green, J. (2013). *Spreadable media: Creating value and meaning in a networked culture*. New York University Press.
- [31] Maddox, J., & Gill, F. (2023). Assembling “sides” of TikTok: Examining community, culture, and interface through a BookTok case study. *Social Media + Society*, 9(4). <https://doi.org/10.1177/20563051231213565>
- [32] Gerbaudo, P. (2024). TikTok and the algorithmic transformation of social media publics: From social networks to social interest clusters. *Information, Communication & Society*. Advance online publication. <https://doi.org/10.1177/14614448241304106>
- [33] Netflix. (2021, November 16). Top 10 things to know about our weekly Top 10. <https://about.netflix.com/news/top-10-things-about-Netflix-top-10>
- [34] Crane, R., & Sornette, D. (2008). Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences*, 105(41), 15649–15653. <https://doi.org/10.1073/pnas.0803685105>
- [35] Szabo, G., & Huberman, B. A. (2008). Predicting the popularity of online content. arXiv preprint arXiv: 0811.0405. <https://arxiv.org/abs/0811.0405>
- [36] Netflix. (2023, December 12). What we watched: A Netflix engagement report. <https://about.netflix.com/news/what-we-watched-a-netflix-engagement-report>
- [37] Netflix. (2024, May 23). What we watched: The second half of 2023. <https://about.netflix.com/news/what-we-watched-the-second-half-of-2023>
- [38] Mirowski, P., Mathewson, K. W., Pittman, J., & Evans, R. (2022). Co-writing screenplays and theatre scripts with language models: An evaluation by industry professionals. arXiv preprint arXiv: 2209.14958. <https://dl.acm.org/doi/10.1145/3544548.3581225>
- [39] Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., ... & Taigman, Y. (2022). Make-A-Video: Text-to-video generation without text-video data. arXiv preprint arXiv: 2209.14792. <https://arxiv.org/abs/2209.14792>
- [40] Bar-Tal, O., Chefer, H., Tov, O., Herrmann, C., Paiss, R., Zada, S., ... & Wolf, L. (2024). Lumiere: A space-time diffusion model for video generation. arXiv preprint arXiv: 2401.12945. <https://arxiv.org/abs/2401.12945>
- [41] Millman, Z. (2023, June 21). Secret Invasion credits were made with AI, and its director defends the choice. Polygon. Retrieved from <https://www.polygon.com/23768740/secret-invasion-credits-ai-method-studios>
- [42] DiPlacido, D. (2023, June 23). Marvel fans aren’t happy with Secret Invasion’s AI-generated opening credits. Forbes. Retrieved from <https://www.forbes.com/sites/danidiplacido/2023/06/23/marvel-fans-arent-happy-with-secret-invasions-ai-generated-opening-credits/>

- [43] Redmond, M. (2023, June 21). The 'Secret Invasion' opening credits are AI-generated and people are not happy. UPROXX. Retrieved from <https://uproxx.com/tv/secret-invasion-opening-credits-ai-marvel/>
- [44] Holub, C. (2024, March 21). Late Night with the Devil directors respond to controversy over AI-generated images. Entertainment Weekly. Retrieved from <https://ew.com/movies/late-night-with-the-devil-directors-address-ai-controversy/>
- [45] Turner, G. (2024, March 24). Late Night With The Devil's AI images controversy explained. ScreenRant. Retrieved from <https://screenrant.com/late-night-with-the-devil-ai-images-controversy-explained/>
- [46] Schneider, J. (2024, March 21). Directors of 'Late Night With the Devil' acknowledge the use of AI in film. PetaPixel. Retrieved from <https://petapixel.com/2024/03/21/directors-of-late-night-with-the-devil-acknowledge-the-use-of-ai-in-film/>
- [47] Edwards, B. (2024, February 1). Robert Zemeckis will use AI to de-age Tom Hanks and Robin Wright in Here. Ars Technica. Retrieved from <https://arstechnica.com/information-technology/2024/02/robert-zemeckis-will-use-ai-to-de-age-tom-hanks-and-robin-wright-in-here/>
- [48] Metaphysic. (2023, January 31). Metaphysic announces strategic partnership with CAA; named generative AI provider for Miramax's "Here" [Press release]. PR Newswire. <https://www.prnewswire.com/news-releases/metaphysic-announces-strategic-partnership-with-creative-artists-agency-named-generative-ai-provider-for-miramaxs-here-301734767.html>
- [49] Cobb, K. (2023, June 22). Marvel used AI to create "Secret Invasion" opening credits: "It felt explorative and inevitable". TheWrap. <https://www.thewrap.com/marvel-secret-invasion-ai-opening-credits/>