



Why InterPositive Was Worth Up to \$600 Million to Netflix

A Strategic Analysis of Netflix's AI Filmmaking Acquisition
Comprehensive Review of 14 Patent Filings Covering a Full-Stack Cinematic AI System

Introduction

On March 5, 2026, Netflix announced having acquired InterPositive, LLC, a sixteen-person company with little conventional studio footprint: no library, no public product launch, and no disclosed commercial business. Bloomberg reported that total consideration could reach \$600 million, including earnout. If that figure is accurate, the acquisition is best understood not as a talent deal or a content deal, but as a strategic purchase of a new production-and-post infrastructure layer. The core question is not simply why Netflix bought InterPositive. It is which part of the media supply chain Netflix believed it could now internalize.

InterPositive's own description provides the first answer. This is not a prompt-to-video company, and it is not a general-purpose generative-video system repackaged for film. Affleck describes it as a production-specific model built from a project's own material: filmmakers shoot first, then train against their own dailies and related production assets, and only then use AI in post to improve editorial flexibility, relight or reframe shots, remove wires, enhance backgrounds, and generate missing coverage. For readers in production, post, and media-tech infrastructure, that distinction is the whole point. InterPositive is not trying to replace principal photography. It is trying to compress, automate, and reallocate selected tasks downstream of it.

That framing matters because the patent count, while substantial, is not by itself the valuation story. As of March 2026, the InterPositive portfolio reflected in the public family record comprises four granted U.S. patents (US 12,322,036; US 12,438,995; US 12,511,837; US 12,511,904), twelve published PCT applications (including WO2025255429A1), and one provisional application, US 63/657,756. The filings span LiDAR-assisted spatial capture, dataset construction, training methodology, captioning, language interfaces, style-controlled generation, and consumer extensions, suggesting a coordinated system and a deliberate attempt to claim multiple layers of a filmmaking stack. The enforceable core today is narrower than the headline count: four granted U.S. patents, with the balance still subject to national-phase entry and examination. For studios, facilities, and vendors, the practical takeaway is not that InterPositive owns "AI filmmaking" in the abstract. It is that Netflix now controls a serious, multi-layer attempt to protect the workflow architecture around production-specific model training and AI-assisted post.

What more plausibly explains the acquisition price is the combination of that IP with proprietary data, workflow integration, and operational readiness. InterPositive appears to have spent four years building controlled datasets on soundstages, organizing them around filmmaking metadata rather than generic web labels, and connecting the resulting models to a production-facing pipeline. On the public record, the strategic value is less the novelty of the underlying ML primitives than the integration: spatial data, controlled training, shot- and frame-level metadata, captioning, generation, and quality control combined in a form that can sit inside studio security requirements and feed existing post workflows. In

other words, Netflix did not just buy models. It appears to have bought a way to operationalize models against real production assets.

That distinction is also why Ben Affleck matters to the deal, though not mainly for celebrity reasons. In production technology, adoption is rarely blocked by feature demos alone; it is blocked by mistrust from filmmakers, post supervisors, line producers, VFX leadership, and labor. InterPositive's positioning is unusually disciplined on that front. Affleck's language is consistent: the system is trained on the material already shot, is intended to improve the technical and logistical parts of production and post, and is not presented as a tool for replacing performances or generating films from scratch. For this audience, that is not branding fluff. It is part of the deployment logic.

Netflix's decision to keep the technology exclusive sharpens the competitive implications. InterPositive's tools are not being sold into the broader market as a commercial platform. That means the transaction is not only about internal efficiency at Netflix. It is also about denying peers access to a potentially important new layer between traditional post-production tooling and heavier VFX or virtual-production workflows. If InterPositive works as described, the relevant question for the rest of the industry becomes immediate: who supplies the non-Netflix equivalent?

This is why the acquisition matters beyond the major studios. For post houses and technology vendors, InterPositive points to a likely reconfiguration of the stack. A production-specific AI layer that can sit between camera originals and final delivery — or, more precisely, between dailies, editorial decision-making, VFX prep, shot cleanup, and finishing — would not eliminate existing systems of record. It would sit on top of them and change where margin, leverage, and differentiation accrue.

This report proceeds in four parts. Part I examines the acquisition and the industrial logic behind it: what InterPositive appears to have built, how it fits within Netflix's broader AI posture, and how it alters the competitive landscape. Part II maps the architecture implied by the filings and the public record. Part III analyzes the patent portfolio itself: what is granted, what is merely published or pending, where the claim set appears strongest, and where obviousness, prior art, or design-arounds may limit exclusivity. Part IV assesses the strategic consequences for Netflix, competing studios, post houses, and technology vendors.

The purpose is not to argue that InterPositive is either revolutionary or overstated. It is to identify, with industry realism, what Netflix likely bought and why the rest of the market should care. The portfolio has genuine strengths: breadth across the stack, a coherent integration story, a closed-data posture more defensible than internet-trained generative systems, and a framing that aligns better with production realities than most AI film tools to date. It also has real limits: the base ML is largely commodity, the international position is not yet mature, individual layers may be designed around, and several implementation details remain inferential rather than confirmed on the public record. Netflix appears to have bought the only publicly described version assembled into a studio-safe, filmmaker-facing, patent-backed system — and paid accordingly.



(12) **United States Patent** (10) **Patent No.:** **US 12,322,036 B1**
Affleck-Boldt (45) **Date of Patent:** **Jun. 3, 2025**

(54) **LIDAR DATA UTILIZATION FOR AI MODEL TRAINING IN FILMMAKING**
(71) Applicant: **Benjamin Geza Affleck-Boldt**, West Hollywood, CA (US)
(72) Inventor: **Benjamin Geza Affleck-Boldt**, West Hollywood, CA (US)
(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **18/959,397**
(22) Filed: **Nov. 25, 2024**
Related U.S. Application Data
(60) Provisional application No. 63/857,756, filed on Jun. 7, 2024.

(51) **Int. Cl.** (2017.01) **G06F 7/521**
G06F 15/205 (2011.01)
(52) **U.S. Cl.** (2017.01): **G06F 15/205** (2013.01); **G06F 7/521** (2017.01); **G06F 2207/20081** (2013.01)
(58) **Field of Classification Search**
CPC: **G06F 15/205**, **G06F 7/521**, **G06F 2207/20081**
USPC: **382/154**
See application file for complete search history.


(56) **References Cited**
U.S. PATENT DOCUMENTS
11,288,864 B2 3/2022 George et al.
11,461,963 B2 10/2022 Mani-cangam et al.

OTHER PUBLICATIONS
Lin et al. "VideoDirectorGPT: Consistent Multi-Scene Video Generation via LLM-Guided Planning" 2024.
Hong et al. "TapVideo: Language pretraining for text-to-video generation via transformers" 2022.
* cited by examiner

Primary Examiner — Stephen P Coleman
(74) *Attorney, Agent, or Firm* — MARSHALL, GIERSTEIN & BERLIN LLP; Randall G. Raeth

(57) **ABSTRACT**
A method enhances AI model training in filmmaking using Lidar data to correlate 2D videos with 3D spatial data, process filmmaking metadata, and simulate professional techniques. A system includes processors and memory to correlate video data with Lidar data, process filmmaking metadata, and train AI models for realistic video content generation. A computer-readable medium contains instructions for using Lidar and metadata to train AI models in simulating professional filmmaking techniques, enhancing video content realism.

20 Claims, 26 Drawing Sheets



Front page of US 12,438,995 B1 — the first granted InterPositive patent, awarded October 7, 2025. Filed under inventor Benjamin Geza Affleck-Boldt and assignee Fin Bone LLC, it covers the integration of video language models with AI for filmmaking, including LiDAR-assisted spatial capture, metadata processing, and transfer learning applied to professional cinematography.

Table of Contents

PART I: THE ACQUISITION

- Chapter 1 — The Stealth Operation
- Chapter 2 — Public Skeptic, Private Builder
- Chapter 3 — What InterPositive Actually Does
- Chapter 4 — The Technical Reality
- Chapter 5 — Netflix's Broader AI Strategy
- Chapter 6 — The Competitive Landscape
- Chapter 7 — The Labor Question

PART II: THE ARCHITECTURE

- Chapter 8 — The Technology Stack
- Chapter 9 — The Closed-Loop System
- Chapter 10 — The Production Pipeline

PART III: THE PATENT PORTFOLIO

- Chapter 11 — Understanding the Portfolio: What '14 Filings' Actually Means
- Chapter 12 — Independent Claims and Patent Strength
- Chapter 13 — Patentability Assessment
- Chapter 14 — The Prior Art Landscape
- Chapter 15 — Patent Filing Strategy

PART IV: STRATEGIC ASSESSMENT

- Chapter 16 — What Netflix Actually Bought
- Chapter 17 — Who Is Locked Out
- Chapter 18 — IP Strength and Vulnerabilities
- Chapter 19 — The Consumer Tools Wildcard
- Chapter 20 — What Comes Next
- Chapter 21 — Conclusions

Patent Portfolio

Annex A: Individual Patent Analyses (B.1–B.16)

References

PART I

The Acquisition

- Chapter 1 — The Stealth Operation
- Chapter 2 — Public Skeptic, Private Builder
- Chapter 3 — What InterPositive Actually Does
- Chapter 4 — The Technical Reality
- Chapter 5 — Netflix's Broader AI Strategy
- Chapter 6 — The Competitive Landscape
- Chapter 7 — The Labor Question

CHAPTER 1

The Stealth Operation

In 2022, while publicly dismissing generative AI on podcasts and at investment conferences, Ben Affleck quietly incorporated Fin Bone, LLC in Los Angeles and began assembling a sixteen-person team of engineers, researchers, and creatives. The company was backed by RedBird Capital Partners, the investment firm led by Gerry Cardinale. It operated in near-total obscurity for almost four years. No press releases, no demo reels, no venture funding announcements. When Netflix acquired the company in early 2026 — by then renamed InterPositive, LLC — the industry learned about it from a Netflix blog post, not a press tour.

Netflix did not disclose the financial terms of the acquisition. Bloomberg subsequently reported that the total deal consideration may reach as high as \$600 million, including performance-based earnout provisions — a figure that, if accurate, would make it one of the largest acquisitions in Netflix's history and among the most significant AI-related transactions by a major Hollywood studio. Bloomberg also reported that after working on the technology for several years, Affleck began soliciting investment in 2025, meeting with venture capital firms and Hollywood companies about using the technology — conversations that led to Netflix evaluating InterPositive as an in-house production tool rather than a commercial platform. The transaction followed Netflix's decision to walk away from its months-long pursuit of Warner Bros. Discovery Inc. In this context, the InterPositive acquisition represents a markedly different approach: rather than massive horizontal consolidation, Netflix chose a targeted investment in proprietary AI filmmaking capability.

The gap Affleck identified was specific: the generative-AI companies building film tools were led by engineers who had never run a set. They could generate pixels, but they could not tell you why a director might need three additional angles of coverage from a scene already in the can, or why matching color temperature across a sequence matters more than resolution. As Affleck put it on the Joe Rogan podcast in January 2026, the outputs of large language models are "really shitty" when it comes to producing meaningful creative work from scratch. His thesis was never about generation from nothing — it was about augmenting footage that already exists.

Netflix confirmed the acquisition through its technology blog on March 5, 2026, noting that all sixteen InterPositive team members would join Netflix and that Affleck himself would serve as a senior adviser. The blog post framed the acquisition in language that would become the company's official position: "InterPositive's mission — to use emerging technology in ways that protect and expand creative choice — is deeply aligned with Netflix's long-standing belief that innovation should serve storytellers and the creative process." The technology, Netflix made clear, would be offered to its creative partners but

would not be sold commercially in the marketplace.

But the most striking part of the story is not the acquisition itself — it is what Affleck was saying about AI in public while he was building it in private.

CHAPTER 2

Public Skeptic, Private Builder



Affleck during the Netflix acquisition announcement interview. The production environment behind him — rigging, studio lighting, director's chairs — is the world InterPositive was designed to serve.

The timeline of Affleck's public statements about AI, read alongside the InterPositive patent filing dates, reveals a deliberate pattern. At CNBC's Delivering Alpha summit in 2024, Affleck told the audience that AI "cannot write you Shakespeare" and that "nothing new is created" by large language models. These were not throwaway comments — they were a precise articulation of the boundary his company was designed to operate within.

On the Joe Rogan Experience in January 2026, weeks before the Netflix acquisition became public, Affleck elaborated: it is "bullshit" to think AI can produce meaningful movies from scratch. But he then drew the distinction that defines InterPositive's entire product thesis. "If you can shoot a scene in a studio and then make it realistically look like the North Pole using AI instead of actually going to the North Pole, that saves money, saves time, and lets you focus on the performances." That is not a man who dislikes AI. That is a man who has spent four years building a very specific version of it.

In the acquisition announcement interview, Affleck described his initial encounter with generative AI video tools: "It really scared me. I can't believe what this is capable of. But what I learned pretty quickly was that that was a little bit of an illusion. It would often fall apart." His alarm was real — but so was his recognition that the technology in its generic form lacked filmmaking knowledge: "There was this really deep engineering, math, science level of expertise associated with creating this — but no artistic, no filmmaking information whatsoever." That gap became InterPositive's founding premise.

Through his production company Artists Equity, co-founded with Matt Damon, Affleck had already signaled an interest in restructuring how films get made — shifting financial participation toward talent and crew. InterPositive extends that logic into the production pipeline itself: if AI can eliminate expensive location shoots and reduce post-production cycles, the savings become leverage in a profit-sharing model.

The public skepticism served a second function. By consistently framing generative AI as creatively inadequate, Affleck was implicitly validating InterPositive's approach — tools that enhance existing footage rather than replace the filmmaking process. Every interview was, in effect, a market-positioning statement for a product no one knew existed.

CHAPTER 3

What InterPositive Actually Does

InterPositive's system begins where traditional production ends: with the dailies. Raw footage from a soundstage — every take, every angle, every lighting setup — becomes the training dataset for a production-specific model. This is not a model trained on the open internet. It has never seen a stock photo or a scraped YouTube frame. It knows only what was shot on that set, by that crew, for that project.

As Affleck described the approach in the acquisition announcement: "It's not about text-prompting or generating something from nothing. You're building a model from your own material. That's how this works. You have to create your movie essentially first before you can really build your model around your movie using AI." The system takes what a director has already captured and extends it — generating synthetic coverage from angles that were never physically shot, adjusting lighting conditions across a sequence, or producing clean plates for visual effects compositing without requiring a dedicated VFX pass.

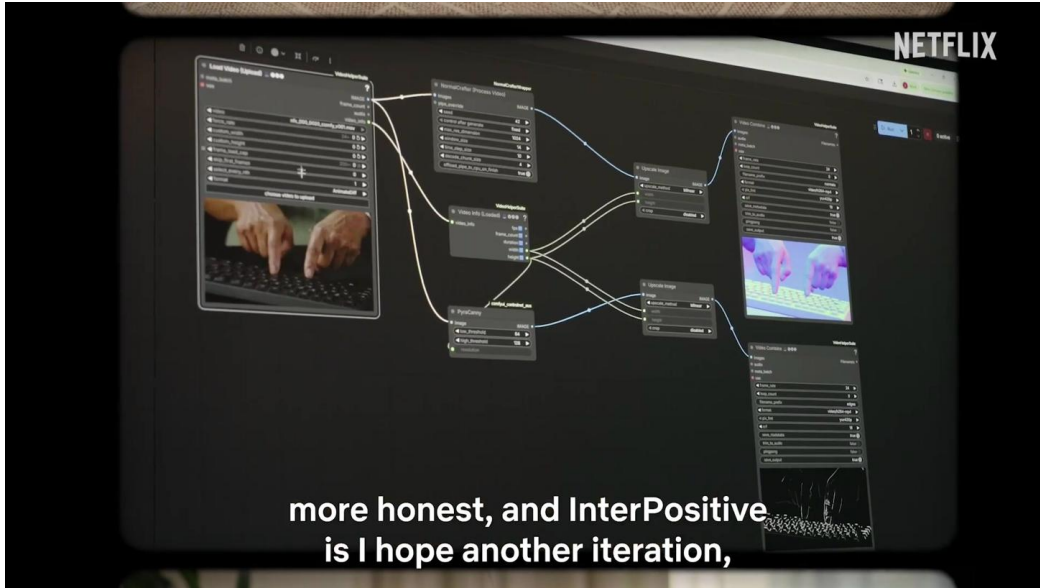
The practical applications, as Affleck described them concretely: "You can use your own model to remove the wires on stunts, reframe a shot, get a shot you missed, shape the lighting, enhance the backgrounds. You can take some of those problems out. Yes, you can do it more quickly, you can do it more easily. You're giving more choice, you're giving more opportunity."

The technical architecture rests on a captioning model (internally designated SamildAnach) that converts visual frames into structured text descriptions, and a generation model (Filmmaker) that produces new frames conditioned on those descriptions plus the original footage. The two models operate in a closed loop: the captioner evaluates the generator's output, and the generator refines based on that evaluation. This feedback cycle, protected across multiple patent filings, is what maintains visual consistency across generated frames.


In his official statement for the Netflix blog, Affleck elaborated on the development process: "Together with a small team of engineers, researchers and creatives, I began filming a proprietary dataset on a controlled soundstage with all the familiarities of a full production. I wanted to build a workflow that captures what happens on a set, with vocabulary that matched the language cinematographers and directors already spoke and included the kind of consistency and controls they would expect."

One additional point warrants emphasis: the tools focus exclusively on filmmaking technique, not on performances. Affleck draws this line sharply: "It's not a way of imposing a new set of reactions or something alien or foreign to the character. It can only understand this and only build this tool because


it's trained on the character that the actor has already built." InterPositive manipulates and refines environments, lighting, and camera logistics — everything surrounding the performance, but never the performance itself.



Internal ComfyUI-based node-graph workflow used to orchestrate captioning, enhancement, and generation steps within the InterPositive pipeline. This screen represents backend ML workflow infrastructure rather than the editor-facing review/compositing interface.



InterPositive review/compositing web app running in a Chromium browser. The purple-themed interface shows a production queue, preview player, and batch-processing controls, indicating a proprietary editor/reviewer-facing application layered on top of the underlying ML workflow stack. The secondary monitor displays a production-tracking integration panel.



The InterPositive platform mid-operation — job counter reads 50/50 as the system processes a batch of shots. The workflow supports relighting, reframing, wire removal, and shot generation trained on the production's own footage.



Output review stage within the InterPositive web app. The second monitor shows After Effects integration and production-tracking panel documentation, indicating direct connection to standard post-production delivery pipelines.

CHAPTER 4

The Technical Reality

What Netflix acquired is not a breakthrough in machine learning. Frame interpolation, view synthesis, neural radiance fields, Gaussian splatting, inpainting-based reframing, and diffusion model fine-tuning all exist as published research and, in some cases, as commercial products. The individual ML components in InterPositive's stack are implementations of known techniques. The technology is commodity. The value sits elsewhere.

The value sits in the integration layer — and that is where the patent portfolio concentrates its protection. The specifications describe a system connecting spatial capture (LiDAR scans tokenized into numerical sequences), dataset construction (locked to production footage), training methodology (single-parameter variation to isolate visual attributes), model inference, and output delivery through a unified pipeline. Netflix's announcement video shows this pipeline operating across two distinct interface layers.

The announcement video points to two distinct interface layers within the InterPositive pipeline. One brief screen shows a node-based workflow environment consistent with ComfyUI, suggesting that InterPositive uses a graph-based orchestration layer for internal machine-learning workflows such as captioning, enhancement, and generation. But the clearer operational footage shows a separate custom browser-based application used for shot review, compositing, queue management, batch processing, and output review. The architectural distinction matters: ComfyUI appears to function as backend workflow infrastructure, while the production-facing surface used by editorial and post-production staff is the proprietary web app.

Three elements appear to distinguish this from academic implementations. First, the dataset is locked: models train only on authorized production footage, which eliminates the copyright exposure that plagues open-internet-trained systems. Second, the pipeline appears to embed provenance tracking — the patent specifications describe metadata identifying source footage, model versions, and processing steps. Third, the color science layer — preserving exposure, color temperature, and dynamic range consistency across original and generated frames — reflects production expertise that pure ML teams rarely possess. This candid assessment matters because it clarifies what Netflix actually acquired. Netflix did not acquire a platform. Netflix acquired a team that knew how to connect existing open-source ML infrastructure to existing professional media tools, train production-specific models within them, and deliver the output in a form that the existing post-production pipeline could consume. The barrier to building an equivalent system is not access to technology. It is the expertise to integrate it — to connect the ML layer to the media layer within a security perimeter that studios will trust.

CHAPTER 5

Netflix's Broader AI Strategy

The InterPositive acquisition does not exist in isolation. Netflix has been assembling a vertically integrated AI production stack through a series of moves that, taken together, suggest a systematic strategy to control every layer of the content pipeline.

Netflix's Q3 2025 shareholder letter described generative AI as a "significant opportunity" across content production, recommendations, and advertising. On the accompanying earnings call, Co-CEO Ted Sarandos went further, characterizing Netflix as "all in" on AI, and clarifying the company's position: Netflix is "not worried about AI replacing creativity" but "very excited about AI creating tools to help creativity." InterPositive is the most concrete manifestation of that stance.

Eyeline Studios, Netflix's in-house visual effects division built on the Scanline VFX acquisition, handles heavy visual effects — the kind of work traditionally outsourced to facilities like Industrial Light & Magic or Weta. InterPositive occupies a different position in the pipeline, handling the on-set and near-set augmentation that sits between principal photography and heavy post-production.

In early 2025, Netflix's production of *EI Eternauta* included AI-assisted VFX work — public reporting confirmed AI use for a Buenos Aires building-collapse sequence. Netflix has also invested in Ready Player Me for avatar and digital-human technology, adding another layer to what is becoming a comprehensive internal capability.

The pattern is consistent: rather than licensing external tools, Netflix builds or acquires proprietary technology and keeps it exclusive. The strategic logic is vertical integration in service of margin protection. Every production capability Netflix owns internally is a capability it does not pay a vendor markup on — and a capability competitors cannot access.

David Fincher has reportedly used InterPositive products on an upcoming film starring Brad Pitt — as reported by Bloomberg. If accurate, it would represent the highest-profile third-party validation of the technology to date, given Fincher's reputation for exacting technical standards.

CHAPTER 6

The Competitive Landscape

InterPositive is not the only company applying AI to film production, but its positioning is distinct — and the competitive picture is more nuanced than a simple comparison of feature sets suggests.

AI-powered post-production tools already exist in widely used software. DaVinci Resolve Studio ships with Relight FX for AI-driven relighting, Magic Mask for subject isolation, UltraNR for neural-engine noise reduction, and IntelliTrack for AI object tracking. Third-party plugins like Beeble's SwitchLight extend these capabilities further with advanced 3D relighting via normal maps and surface passes.

Flawless AI, founded by director Scott Mann, deserves particular attention because it occupies a complementary position to InterPositive. Flawless operates on the performance layer — its TrueSync technology modifies actors' lip movements and facial expressions for dubbing and localization. Flawless works on the human element of the frame. InterPositive explicitly avoids it. The distinction is clean and deliberate: one handles the performance layer, the other handles the technical layer, and neither attempts to generate content from nothing. The critical difference for the market: Flawless is available to any studio or production that wants to license it. InterPositive is not — it belongs to Netflix.

The prompt-to-video tools — Sora, Runway, Luma, Pika, Kling — operate in a fundamentally different space. They train on internet-scale video datasets and generate video from text prompts or images. InterPositive assumes filming has already occurred and enhances it with metadata-driven cinematography, spatial intelligence, and style control. It works downstream of principal photography, not as a generative replacement for it. Comparisons between InterPositive and these tools, while common in press coverage, are largely misplaced.

The business model difference is equally significant. Flawless licenses its technology. DaVinci Resolve is commercially available. Beeble sells plugins. InterPositive, under Netflix ownership, is exclusive — unavailable to competitors at any price. This creates an asymmetry: Apple, Amazon/MGM, Disney, Paramount, Sony, and every other major studio face the same production cost pressures and the same need to automate foundational post-production tasks. None of them have InterPositive. None of them can license it.

CHAPTER 7

The Labor Question

The InterPositive acquisition landed during a period of acute tension between Hollywood's creative workforce and the studios over AI. The SAG-AFTRA negotiations of 2023 produced contract language specifically addressing AI-generated performances, and the below-the-line unions — camera operators, lighting technicians, set builders, location scouts — have been watching the technology closely.

Affleck is among the Hollywood insiders who signed on to the Creators Coalition on AI, a group established in late 2025 describing itself as a central hub for cross-industry discussions about how AI is impacting the entertainment industry. The Coalition's position — that it represents a commitment to responsible, human-centered innovation rather than a rejection of AI — aligns closely with the messaging Netflix deployed around the acquisition.

InterPositive's pitch is that AI handles the "logistical difficult technical stuff," freeing directors to "focus on the performances." That framing positions the technology as efficiency-enhancing rather than labor-displacing. But the craft roles that InterPositive's tools can reduce — location crews, additional camera setups, certain VFX passes — represent real employment for real workers. Affleck acknowledged the economics bluntly on the Joe Rogan podcast: AI could replace offshore rendering work and redirect those budget savings toward actors and core creative elements. That is a direct statement about labor displacement, however it is framed.

Netflix's Chief Content Officer Bela Bajaria addressed this tension in the acquisition announcement: "For me it's really not about cheaper. It's really about better." Elizabeth Stone, Chief Product and Technology Officer, reinforced the position: "Our approach to AI has always been focused on meaningfully serving the needs of the creative community and our members." In her official blog statement, Bajaria was more expansive: "Our relationship with artists has always been grounded in trust: supporting the full range of their creativity and ensuring they have the power to decide how their films and shows are made."

Whether this distinction between "better" and "cheaper" holds as the technology matures and its cost-saving potential becomes clearer will depend less on the technology itself and more on how aggressively it is deployed, and whether the savings are shared with the workforce or captured entirely as margin improvement. What will likely determine that outcome is not intent but leverage — and leverage, in this case, begins with what the patents actually protect.

PART II

The Architecture

Chapter 8 — The Technology Stack

Chapter 9 — The Closed-Loop System

Chapter 10 — The Production Pipeline

CHAPTER 8

The Technology Stack

The InterPositive system is architected as a seven-layer stack, where each layer is protected by one or more patent filings and each is dependent on the layers beneath it. This layered architecture creates a competitive moat: a competitor cannot replicate just one layer — they must build or acquire the entire stack.

Layer	Technology	Patents / Apps	Function
Layer 1 Spatial Capture	LiDAR hardware, 3D positioning	US 12,322,036 WO2025255446A1	LiDAR sensors, laser locators, and robotic equipment (Techno crane, Panther Dolly, KUKA arm). Captures X, Y, Z coordinates, pitch, yaw, roll at frame level.
Layer 2 Dataset Construction	Soundstage filming, synthetic data, simulation	WO2025255439A1 WO2025255427A1 WO2025255441A1	Controlled soundstage footage combined with real production data and virtual simulation environments.
Layer 3 Training Methodology	Single-parameter variation	WO2025255428A1 WO2025255436A1	Isolating single camera parameters per training iteration; creating the metadata framework for comprehensive model training.
Layer 4 Core Models	Captioner + quality control	US 12,511,904 WO2025255432A1	The captioner model (SamildAnach) analyzes footage frame-by-frame and generates structured cinematographic metadata.
Layer 5 Language Interface	VLM integration	US 12,438,995 WO2025255437A1	Bridges natural language filmmaker instructions with AI parameter-driven generation. Allows filmmakers to describe intent in their own vocabulary.
Layer 6 Output Generation	Style control + technique simulation	US 12,511,837 WO2025255429A1 WO2025255425A1 WO2025255426A1	Style-controlled generation, technique simulation, and the video content generation engine.
Layer 7 Consumer Apps	Consumer tools	WO2025255433A1	Film stock simulation, processing techniques, consumer-accessible filmmaking tools.

The layered architecture means that interdependence runs downward: Layer 7 depends on Layer 6, which depends on Layer 5, and so on to Layer 1. A competitor could theoretically copy any single layer, but without the layers beneath it, the output would be inferior. This interdependence is genuine, but it should not be overstated. A well-funded competitor with deep expertise in both ML and cinematography could potentially build an equivalent system using different technical approaches at

each layer without infringing any individual patent. The portfolio's defensive strength is that doing so across all seven layers simultaneously would require reimagining the entire system.




FIG 2A from US 12,438,995: The Filmmaker Computing System architecture, showing all processing modules — from metadata processing and LiDAR integration through language model training and prompt processing — connected to camera hardware, cloud infrastructure, and the production user interface.

CHAPTER 9

The Closed-Loop System

The captioner model (US 12,511,904, internal codename SamildAnach) and the Filmmaker model (provisional US 63/657,756) form a closed feedback loop that is more sophisticated than a simple linear inference pipeline.

The Feedback Loop. The process operates as follows: (1) Raw footage enters the system from either a controlled soundstage capture session or real production footage. (2) The captioner model analyzes it frame-by-frame and shot-by-shot, generating structured metadata labels describing camera settings, movement patterns, lighting characteristics, composition, and technical parameters. (3) The Filmmaker model consumes these labels and generates enhanced output: color correction, relighting, wire removal, background enhancement, continuity correction, and parameter-controlled enhancement. (4) The enhanced output feeds back to the captioner for validation and refinement. (5) Both models improve iteratively as discrepancies between the captioner's description and the Filmmaker's execution are resolved.

How This Differs from Conventional Pipelines. In a conventional linear pipeline, an encoder takes raw input and produces a learned representation, which a decoder converts to output. This architecture has no feedback: the encoder cannot validate whether its encoding was correct by examining the decoder's output. The captioner-Filmmaker loop is different. The captioner teaches the Filmmaker what cinematography looks like through structured metadata. The Filmmaker demonstrates its understanding by producing output consistent with those labels. The captioner validates and refines its analysis based on the Filmmaker's output. Over time, both models improve in tandem.

The Captioner as the "Reader" and the Generator as the "Writer." The relationship between the captioner patent (US 12,511,904) and the style control patent (US 12,511,837) is best understood through this lens. The captioner is the reading function — it analyzes existing footage and articulates what cinematographic choices are present. The style control patent is the writing function — it uses the style understanding extracted by the captioner to constrain and condition video generation. Together, they create a closed-loop capability: analyze existing footage to extract style metadata, then use that metadata to generate new footage that matches.

The Production Interface Layer. InterPositive appears to expose its AI capabilities through a layered interface model. Beneath the surface, a ComfyUI-based node graph likely handles internal orchestration of captioning, metadata generation, enhancement, and output steps. At the surface, the working production UI shown in the announcement video is a custom browser-based application with a preview player, queue and status panels, batch job management, and output review. This makes the

system accessible to post-production staff without requiring them to work directly inside the ML orchestration graph. In other words: ComfyUI is part of the engine room; the proprietary web app is the production-facing bridge to editorial, VFX, and finishing workflows.

This layered architecture also provides transparency at the right level for each audience. Internal ML engineers and technical staff can examine and adjust the node-graph orchestration layer. Editors, colorists, and post supervisors interact with the browser-based review interface, examining output, managing queues, and approving or flagging results — without requiring knowledge of machine learning principles or direct access to the underlying workflow graph.

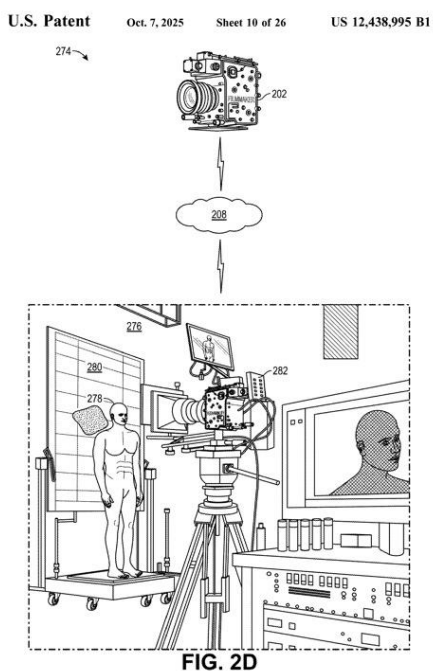


FIG 2D from US 12,438,995: A green screen compositing environment connected to the Filmmaker Computing System. A mannequin subject is filmed against the green screen (left), and the composited result with a generated background is visible on the monitor (right) — demonstrating the input-to-output pipeline.

CHAPTER 10

The Production Pipeline

While the public framing focuses on the dailies as the primary training input, the likely input dataset is substantially broader than raw video frames. A production generates a rich, interconnected set of assets that — embedded together — give the model a far deeper understanding of the project than footage alone.

The Full Input Set. A production-specific model likely ingests:

Video Dailies	Every take, every angle, every lighting setup — teaching visual logic, lens characteristics, set geography, actor appearance, blocking, and continuity.
Production Audio	Scene structure, take identification, emotional tone, pacing.
Text Materials	Script, shot lists, continuity notes, camera reports, lined script — teaching scripted intent, coverage plan, narrative structure.
Reference Images	Unit stills, reference photos, DP color references, LUT stills, lookbook — teaching target aesthetic, color intent, lighting references.
Production Documents	VFX breakdowns, storyboards, set blueprints — teaching planned vs. captured coverage, VFX intent, spatial layout.

U.S. Patent Oct. 7, 2025 Sheet 8 of 26 US 12,438,995 B1




FIG. 2B

FIG 2B from US 12,438,995: The controlled multi-camera environment for training data capture. Multiple cameras on wheeled dollies surround a mannequin figure, with crew on elevated platforms operating equipment at varying heights and angles.

U.S. Patent Oct. 7, 2025 Sheet 9 of 26 US 12,438,995 B1

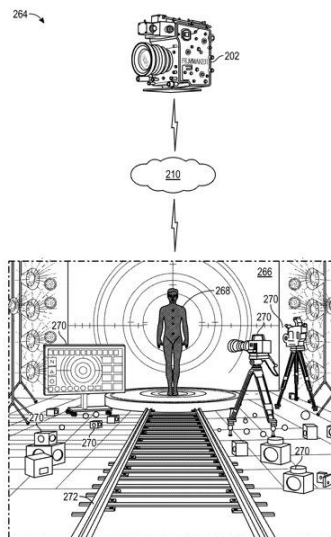


FIG. 2C

FIG 2C from US 12,438,995: Additional capture configuration from the patent, showing the multi-angle production data collection methodology.

The Per-Production Intelligence Layer. The trained model produces a comprehensive intelligence layer for that specific production. Capabilities include: automated baseline color consistency and shot matching; coverage gap identification, continuity checking, and take comparison; wire removal, background enhancement, reframing, and lighting correction; generation of missing shots and alternate framings; automated technical conforming and delivery spec compliance; and per-frame provenance metadata documenting every AI operation, its inputs, and where human review occurred.

The Provenance Layer. Any system that replicates InterPositive's capability must embed a rigorous provenance and documentation layer into the metadata of every output file. The system produces a per-shot, per-frame record tracking which shots were touched by AI tools and in what sequence, which work came from an automated system versus a human's creative decision, and where human review occurred at each stage. This is critical for ownership and chain of title, indemnification, and union compliance. When Netflix says InterPositive made the lighting better and didn't replace anyone, the provenance trail is what makes that statement defensible — or not.

PART III

The Patent Portfolio

Chapter 11 — Understanding the Portfolio: What '14 Filings' Actually Means

Chapter 12 — Independent Claims and Patent Strength

Chapter 13 — Patentability Assessment

Chapter 14 — The Prior Art Landscape

Chapter 15 — Patent Filing Strategy

Patent Portfolio

Complete InterPositive patent portfolio, filed under inventor Benjamin Geza Affleck-Boldt:

No.	Identifier	Title	Date
Provisional Application			Priority Date
—	US 63/657,756	Computing System for Generating Video with Professional Filmmaking Techniques (Filmmaker Model)	Jun. 7, 2024
Granted US Patents			Grant Date
1	US 12,438,995 B1	Integration of Video Language Models with AI for Filmmaking	Oct. 7, 2025
2	US 12,511,837 B1	AI-Based Video Content Creation with Predetermined Styles	Dec. 30, 2025
3	US 12,511,904 B1	Training a Captioner Model (SamildAnach)	Dec. 30, 2025
4	US 12,322,036 B1	LiDAR Data Utilization for AI Model Training in Filmmaking	Jun. 3, 2025
International (WIPO/PCT) Publications			Publication Date
5	WO2025255425A1	Video Content Generation via Metadata and LiDAR	Jun. 19, 2025
6	WO2025255426A1	Simulation of Professional Filmmaking Techniques	Jun. 19, 2025
7	WO2025255427A1	Synthetic Data Generation for ML in Filmmaking	Jun. 19, 2025
8	WO2025255428A1	ML Model Training with Parameter Variation	Jun. 19, 2025
9	WO2025255429A1	AI-Based Video Content Creation with Predetermined Styles (Intl.)	Jun. 19, 2025
10	WO2025255432A1	AI-Driven Generation Meeting Professional Film Standards	Dec. 11, 2025
11	WO2025255433A1	AI-Based Filmmaking Tools for Consumer Use	Dec. 11, 2025
12	WO2025255436A1	Cinematic Data Collection and Processing	Dec. 11, 2025
13	WO2025255437A1	Integration of VLMs (International Filing)	Dec. 11, 2025
14	WO2025255439A1	Training with Synthetic and Real-World Data	Dec. 11, 2025
15	WO2025255441A1	Simulation-Driven Training Environment	Dec. 11, 2025
16	WO2025255446A1	LiDAR Data Utilization for AI Model Training	Dec. 11, 2025

Total: 4 granted US patents, 12 WIPO/PCT publications, 1 provisional application. 329 claims (48 independent). 612 pages across all filings.

CHAPTER 11

Understanding the Portfolio: What '14 Filings' Actually Means

The InterPositive portfolio as reflected in the public family record comprises four granted US patents, twelve international PCT publications, and one provisional application. Not all filings are equal. They fall into three distinct categories — granted US patents, international PCT publications, and a provisional application — and the differences between these categories are fundamental to understanding what the portfolio actually protects, where, and how reliably.

Granted US Patents (4)

US 12,438,995, US 12,511,837, US 12,511,904, and US 12,322,036 are granted United States utility patents. This is the strongest category of intellectual property protection available. A granted patent means the USPTO examined the application — reviewed the claims against the prior art, assessed whether the invention was novel and non-obvious, evaluated whether it met the statutory requirements for patent eligibility — and concluded that the claims are valid. The patent holder can enforce these claims immediately: if a competitor practices the claimed invention in the United States, the patent holder can seek an injunction and damages in federal court.

Granted US patents carry a statutory presumption of validity under 35 U.S.C. § 282. This means a challenger bears the burden of proving invalidity by clear and convincing evidence — a high bar. That presumption is real, but it is not absolute. Granted patents can be challenged through inter partes review (IPR) at the Patent Trial and Appeal Board (PTAB), where the standard is preponderance of the evidence. Patent term: each patent is valid for 20 years from the earliest non-provisional filing date. Based on the 2024 filing dates, protection extends through approximately 2044–2045. These four patents are the enforceable core of the portfolio — the only filings that can be used today to stop a competitor.

International PCT Publications (12)

WO2025255425A1 through WO2025255446A1 (including WO2025255429A1) are publications under the Patent Cooperation Treaty (PCT), administered by the World Intellectual Property Organization (WIPO). A PCT application is not a patent. It does not grant any enforceable rights. It is a single international filing that preserves the applicant's right to seek patent protection in any of the 158 contracting states of the PCT. The twelve PCT publications in the InterPositive portfolio were filed on

June 6, 2025, with priority dates of June 7, 2024. They were published in 2025. As of this writing (March 2026), they are in the international phase. National phase entry, if pursued, would need to occur by approximately December 2026 to January 2027. Even after national phase entry, the path to a granted patent is long and uncertain — national examiners may narrow the claims, reject them entirely, or require significant amendments.

Provisional Application (1)

US 63/657,756 was a United States provisional patent application filed on June 7, 2024. A provisional application establishes a priority date but is not examined, never itself matures into a patent, and expires after twelve months unless a later-filed non-provisional or PCT application properly claims its benefit. The public record indicates that this occurred here. The U.S. applications that later matured into US 12,438,995, US 12,511,837, and US 12,322,036 were filed on November 25, 2024, claim the June 7, 2024 priority date, and trace back to provisional 63/657,756. The June 6, 2025 PCT filings should therefore be understood as international follow-on filings within the same family, not as speculative conversions. Multiple PCT publications expressly claim the benefit of both provisional 63/657,756 and corresponding November 25, 2024 U.S. non-provisional applications.

Category	Count	Status	Enforceable?	Certainty
Granted US Patents	4	Granted, enforceable	Yes — immediately	High (examined)
PCT Publications	12	Published, not granted	No — pending national phase	Low to moderate
Provisional	1	Expired; carried forward	No	N/A — served its purpose

The practical implication: when this report refers to InterPositive's "patent protection," the enforceable core is four US patents. The remaining filings represent varying degrees of potential future protection. The portfolio's strategic value includes all sixteen public filings — but its legal enforceability, as of March 2026, rests on four.

CHAPTER 12

Independent Claims and Patent Strength

The InterPositive portfolio comprises 329 total claims, of which 48 are independent claims. In patent law, independent claims define the broadest scope of protection — they stand on their own without reference to any other claim, and every dependent claim in a patent family must fall within the boundaries set by its corresponding independent claim. Under 35 U.S.C. § 112(d), a dependent claim incorporates all limitations of the claim it references and adds a further limitation; because a dependent claim is necessarily narrower than its parent, a finding of validity for the independent claim does not require separate validity analysis of the dependent claims beneath it. Independent claims are, in practice, the load-bearing walls of a patent portfolio.

InterPositive's 48 independent claims span the full technology stack, though what follows requires a careful distinction between what the independent claims expressly recite and what the broader specifications describe. At the capture layer, the independent claims in WO2025255446A1 and US 12,322,036 are framed around correlating 2D video with 3D LiDAR data, processing filmmaking metadata, and training AI models using the combined data. The specifications describe a more detailed implementation involving numerical tokenization of spatial coordinates, but the independent claims themselves are framed at a higher level of abstraction. At the training layer, the specifications across several filings describe the single-parameter variation methodology. At the inference layer, the specifications describe a closed-loop architecture in which the captioning model evaluates the generation model's output, but the independent claims of the granted US patents are centered on more specific procedural steps.

Each of the four granted US patents follows the same structural pattern: three independent claims covering the same core steps in three statutory categories — a method claim (computer-implemented method), a system claim (one or more processors plus memory with stored instructions), and a computer-readable medium claim (non-transitory medium with stored instructions). The substantive limitations are functionally identical across all three categories within each patent. This triple-coverage strategy provides redundancy: even if one claim category faces challenges, the other two may survive.

The broader specifications also describe pipeline-level integration — connecting spatial capture, dataset construction, model training, inference, and output delivery through a node-graph interface with embedded provenance metadata. Whether this system-level integration is captured in enforceable claim language, or resides primarily in the disclosure, is a question that would require claim-by-claim construction to answer definitively.

The breadth of independent claim coverage means that a competitor seeking to build a comparable system would need to design around not just one or two chokepoint patents but dozens of independent claims spanning multiple layers of the stack. Whether every claim would survive rigorous examination in all jurisdictions is an open question — patent eligibility standards for AI-related inventions remain in flux in both the United States and Europe — but the portfolio's structural redundancy provides substantial defensive depth. The core independent claim steps for each granted US patent are set out in Annex B.1–B.4.

CHAPTER 13

Patentability Assessment

The InterPositive portfolio navigates patent eligibility by anchoring its claims in specific technical implementations rather than abstract computational methods. In the United States, Section 101 of the Patent Act requires that claims directed to abstract ideas must demonstrate something "significantly more" — a concrete technical improvement. In Europe, the EPO's technical-effect doctrine requires that the inventive step produce a technical contribution beyond mere automation of a known process.

The portfolio's strongest eligibility position is at the spatial-capture and training layers, where the filings describe physical-world data acquisition (LiDAR scanning, on-set footage capture) tied to specific processing methods. The specifications detail numerical tokenization of spatial coordinates and single-parameter variation training protocols — concrete technical transformations that map well to both US and European eligibility frameworks.

The closed-loop feedback architecture between the captioneer and generator models, as described in the specifications, presents a credible technical-effect argument: the system produces measurably more consistent output than single-pass generation, and the feedback mechanism is a specific architectural choice rather than a generic instruction to "use AI."

The consumer-facing claims face the highest eligibility uncertainty, as they are closer to software method claims that have historically attracted scrutiny. However, to the extent these claims incorporate the locked-dataset and provenance-tracking elements from the production pipeline, they carry forward the technical specificity that strengthens the rest of the portfolio.

The overall portfolio structure — with redundant independent claims across multiple layers and jurisdictions — reflects a filing strategy designed to ensure that even if individual claims face narrowing or rejection in particular jurisdictions, the core technological protection remains intact.

CHAPTER 14

The Prior Art Landscape

Any credible assessment of InterPositive's patent portfolio must address the body of published work in the fields these patents touch. The individual ML techniques are not novel — a point the portfolio itself implicitly acknowledges by claiming not the algorithms but their application to filmmaking.

Style Transfer and Style-Conditioned Generation

Neural style transfer has been an active research area since Gatys, Ecker, and Bethge published their foundational work in 2015, demonstrating that convolutional neural networks could separate and recombine the content and style of images. Subsequent work extended this to video (Ruder, Dosovitskiy, and Brox, 2016). The claims in US 12,511,837 (style control) will need to be evaluated against this landscape. The patent's likely defense is specificity: rather than generic style transfer, the claims cover predetermined cinematic styles applied during generation using metadata extracted by a companion captioner system.

Video Captioning and Scene Understanding

Video captioning models have been published extensively — from early sequence-to-sequence approaches through modern transformer-based architectures. The claims in US 12,511,904 distinguish themselves by targeting cinematographic metadata rather than semantic descriptions. A general video captioner might describe "a person walking through a room"; the InterPositive captioner describes the shot as "35mm lens, 3:1 key-to-fill ratio, tracking camera movement at 2 feet per second." The novelty claim rests on this domain-specific application rather than on the captioning architecture itself.

LiDAR-Assisted Scene Understanding

The use of LiDAR data in computer vision is well established, particularly in autonomous driving (Waymo, Cruise) and robotics. The InterPositive claims focus specifically on integrating LiDAR data with filmmaking metadata and using it to condition video generation — a more specialized application. The risk here is that the claims could be characterized as applying known LiDAR processing techniques to a new domain without sufficient inventive step, particularly under European examination standards.

Single-Parameter Variation Training

The concept of isolating variables during training — changing one parameter at a time while holding others constant — is a well-established experimental methodology in science and engineering. In ML, ablation studies follow a similar logic. The patent claims apply this principle to cinematographic

parameters specifically, creating training pairs where the effect of each parameter can be learned in isolation. The question is whether applying a well-known experimental methodology to a specific domain constitutes a patentable inventive step. This area may face the most challenge from obviousness arguments.

Closed-Loop Training Architectures

The captioner-Filmmaker feedback loop described in the patents resembles established training paradigms including generative adversarial networks (GANs), reinforcement learning from human feedback (RLHF), and iterative refinement architectures. The InterPositive system's feedback loop is better characterized as iterative supervised learning with validation-based refinement than as adversarial training — the captioner validates the generator's output against structured metadata rather than competing with it. Prior art in iterative refinement and self-supervised training loops is extensive. The claims' defense rests on the specific application to cinematographic quality control and the integration with the captioner's domain-specific metadata vocabulary.

Net Assessment

The portfolio's novelty lies in the integration of these techniques into a coordinated filmmaking system, not in any individual technique. This is both a strength and a vulnerability. It is a strength because the specific combination and application to cinematography is arguably novel. It is a vulnerability because individual claims, examined in isolation against the prior art in their respective fields, may face narrowing. The four granted US patents survived examination, which provides a presumption of validity, but inter partes review (IPR) challenges could test the claims against a broader body of prior art than the examiner may have considered.

CHAPTER 15

Patent Filing Strategy

InterPositive pursued a deliberately sequenced patent filing strategy that reveals strategic thinking about IP protection, market positioning, and competitive dynamics. The strategy unfolded across four coordinated phases.

Phase 1: Provisional Foundation (June 2024)

On June 7, 2024, InterPositive filed provisional patent application US 63/657,756 for the core Filmmaker model. That filing secured the initial U.S. priority date, but it should not be described as a standalone patent asset or as a filing whose later conversion is uncertain on the present record. Before the provisional's June 7, 2025 expiration, InterPositive filed multiple U.S. non-provisional applications on November 25, 2024 claiming that same June 7, 2024 priority date, including the applications that later matured into US 12,438,995, US 12,511,837, and US 12,322,036. Strategic implication: InterPositive secured priority on the core technology while the company was still operating in stealth mode, then built a layered family of non-provisional and international filings around that priority foundation.

Phase 2: US Patent Filings (November 2024)

On November 25, 2024, InterPositive filed multiple US non-provisional patent applications claiming the June 7, 2024 priority date from the provisional. At least three of these have matured into granted patents: the integration of video language models with filmmaking AI (US 12,438,995), AI-based video content creation with predetermined styles (US 12,511,837), and LiDAR data utilization for AI model training in filmmaking (US 12,322,036). All applications were filed under the entity name Fin Bone, LLC. Strategic implication: InterPositive prioritized domestic protection across multiple layers of the stack simultaneously — the language interface, style control, and spatial foundation — well before the twelve-month provisional window expired.

Phase 3: Captioner Model — The Core Algorithm (March 2025)

On March 12, 2025, InterPositive filed the captioner model patent application. This filing was distinct from Phase 2 in several respects: the resulting granted patent, US 12,511,904, lists InterPositive, LLC as assignee (the Phase 2 patents list Fin Bone, LLC), though the entity was not formally renamed from Fin Bone, LLC to InterPositive, LLC until November 17, 2025. The priority date is November 27, 2024 — roughly the same timeframe as Phase 2, but the filing was delayed by over three months, suggesting additional development or testing. The resulting patent is 68 pages — substantially longer and more detailed than the other three. Strategic implication: The delayed filing and greater technical

depth suggest that the captioner model represented newer or more refined technology. The decision to file this patent separately with a later priority date was deliberate — it protected the more recent innovation while maintaining separate claims strategies.

Phase 4: International Blitz — Complete Coverage (June 2025)

On June 6, 2025, InterPositive filed twelve WIPO/PCT international patent applications in a single coordinated filing day. Those filings are best understood as the international extension of a family that already had U.S. non-provisional support, not as evidence that the company waited until the last day to decide whether to convert the provisional. Multiple PCT documents expressly claim the benefit of provisional 63/657,756 together with corresponding November 25, 2024 U.S. non-provisional filings. The sequence appears to have been provisional first, U.S. non-provisionals second, and coordinated international filings third. The Phase 4 filings were published in December 2025, revealing the full scope of InterPositive's patent strategy to the industry for the first time. This public revelation may have accelerated Netflix's acquisition decision, which was announced just three months later.

The four-phase strategy accomplished three goals: maximizing coverage across multiple jurisdictions, protecting the most commercially important innovations with the strongest US patent positions, and ensuring that the portfolio in aggregate is difficult to design around.

PART IV

Strategic Assessment

Chapter 16 — What Netflix Actually Bought

Chapter 17 — Who Is Locked Out

Chapter 18 — IP Strength and Vulnerabilities

Chapter 19 — The Consumer Tools Wildcard

Chapter 20 — What Comes Next

Chapter 21 — Conclusions

CHAPTER 16

What Netflix Actually Bought

Netflix acquired six things:

1. Exclusive technology

InterPositive's tools will not be sold, licensed, or made available to any other studio. Every competitor will need to build or buy their own system from scratch. Netflix stated explicitly that it will offer access to its creative partners but does not plan to sell commercially.

2. Patent protection through 2045

The earliest patents in the portfolio won't expire for nearly two decades. During that time, competitors face either designing around the claims or negotiating licenses — neither of which Netflix is obligated to grant. Four US patents are granted and enforceable now; twelve WIPO applications are moving through national phase examination.

3. Proprietary training data

The purpose-built datasets captured on controlled soundstages over four years of development. This data is not available anywhere else and cannot be replicated without significant time and capital investment. And critically, every Netflix production now generates additional training data — the models improve continuously as more production data enters the system.

4. A specialized team

Sixteen people with deep expertise in both AI engineering and professional filmmaking. InterPositive lacks the tangible assets of a traditional studio — sound stages, physical libraries, or real estate. It has, however, a proven capability to attract, organize, and deploy talented engineers and filmmakers with success. This combination of skills is extremely rare and takes years to develop.

5. A production-ready system

InterPositive's technology is not a research prototype. The system appears to combine an internal ComfyUI-based orchestration layer with a proprietary browser-based review/compositing interface designed for use by post-production professionals. Netflix can deploy it on actual productions immediately. With Netflix's production infrastructure already in place, InterPositive's capabilities can be integrated into existing workflows without major additional capital expenditure.

6. The consumer tools patent

WO2025255433A1 opens the door for subscriber-facing features — a potential competitive advantage in the streaming wars that extends beyond professional production.

The reported total consideration — up to \$600 million including earnout, per Bloomberg — would reflect Netflix's assessment that the combined value of these assets justifies one of its largest acquisitions. The deal structure (upfront cash below the reported ceiling plus performance-based earnout) suggests Netflix is paying for demonstrated results, not just potential.

CHAPTER 17

Who Is Locked Out

Netflix has gone proprietary. The entire 16-person team, the trained models, the custom workflows, and the production-integration methodology are inside Netflix's walls. No other studio, streamer, or production company can access them.

Every major competitor — Apple, Amazon/MGM, Disney, Paramount, Sony — faces the same production cost pressures, the same union documentation requirements, and the same need to automate foundational post-production tasks without replacing human creative judgment. None of them have InterPositive. None of them can license it.

What they need is an equivalent pipeline — one they can deploy inside their own security perimeters — that ingests the full multimodal production dataset; builds a per-production model trained exclusively on that production's own material; automates the foundational pass across dailies grading, editorial, VFX cleanup, and conform; produces a clean, embedded provenance audit trail with every deliverable; and runs on infrastructure they control. This system does not currently exist as a product in the market. The individual components are all available. The integration is what's missing.

The barrier to building an equivalent is not access to technology. It is the expertise to integrate it, the production knowledge to make it work within real workflows, and the willingness to invest the capital and time. Whoever builds it first — or acquires it — owns the rest of the industry that InterPositive can no longer serve.

CHAPTER 18

IP Strength and Vulnerabilities

Strengths

(1) **Breadth.** Fourteen filings covering the complete pipeline from spatial capture through consumer applications. This is not a single point of IP; it is a coordinated moat covering every functional layer.

(2) **Coordination.** The filings interlock, protecting the system rather than individual components. Each patent addresses a critical component, and together they form a system that is substantially harder to design around than any individual patent.

(3) **Specificity.** The claims cover particular methodologies — soundstage protocol, numerical tokenization, multi-task loss, single-parameter variation — not vague AI concepts. Specific method claims are harder to design around than abstract functional claims.

(4) **Granted US patents.** Three filings are already granted and enforceable, carrying a presumption of validity under US law. They provide immediate defensibility in the world's largest IP market.

(5) **International coverage.** Eleven WIPO publications signal intent to seek global protection, creating awareness among potential competitors that infringement risk extends beyond the United States.

(6) **The data moat.** Even if every patent were invalidated, the proprietary training datasets — purpose-built on controlled soundstages — would remain impossible to replicate without the patented methodology, equivalent capital investment, and years of development.

Vulnerabilities

(1) **ML components are commodity.** The underlying machine learning techniques (transformers, attention mechanisms, loss functions, diffusion models) are not novel. The novelty lies in their application to filmmaking. Competitors could use the same underlying architectures with different integration approaches.

(2) **Individual layer design-around risk.** A competitor could potentially design around any single patent by using alternative approaches to the same functional layer. Monocular depth estimation instead of LiDAR. A different metadata schema. A different training protocol. Each substitution degrades the system somewhat but does not eliminate the capability entirely.

(3) Prior art exposure. The individual technical domains these patents touch — style transfer, video captioning, LiDAR-assisted scene understanding, single-parameter training — all have substantial bodies of published prior art. The three granted US patents have not been tested through inter partes review or litigation.

(4) International uncertainty. WIPO publications are not granted patents. National examiners may narrow or reject claims based on different prior art searches and different eligibility standards.

(5) The provisional application. US 63/657,756 (the Filmmaker model) has not been converted to a granted patent. If the provisional lapses without a corresponding full application, the core video generation model loses direct patent protection.

(6) Enforcement costs. Defending 14 patent filings across multiple jurisdictions is expensive. Strategic enforcement would likely focus on the three granted US patents against the highest-value targets.

(7) The presumption of validity is not invulnerability. Granted patents carry a presumption of validity, but inter partes review (IPR) proceedings at PTAB have invalidated claims in approximately 60-70% of instituted reviews.

Net Assessment

The portfolio is strong because of its coordination. Designing around any single patent is possible; designing around the entire coordinated system is substantially more difficult and would likely result in a materially different technology. The real defensive depth comes from the combination of patent protection, proprietary data, specialized talent, and production integration — not from any single element in isolation.

CHAPTER 19

The Consumer Tools Wildcard

WO2025255433A1 deserves separate strategic consideration because its commercial implications extend far beyond professional filmmaking. Netflix has 283 million subscribers as of early 2026. The consumer tools patent covers AI-based filmmaking tools for consumer use — specifically, capturing scenes in various formats and applying processing techniques to affect visual outcomes, including film stock simulation and film processing techniques.

Potential applications include: a Netflix mobile feature letting subscribers apply cinematic looks to personal videos using authentic film stock emulations; a Netflix Creator program offering simplified AI filmmaking tools to independent filmmakers; integration with Netflix's existing content creation pipeline for user-generated content; and licensing to device manufacturers for native camera app integration.

The film stock simulation capability is distinct from conventional filter applications. Rather than applying a color grade overlay after the fact, the patent describes AI-powered reproduction of the actual photochemical characteristics of specific emulsions — the characteristic color rendition, contrast curves, grain structure, and highlight/shadow behavior of particular film stocks. This is a meaningful technical distinction, though whether the claims as filed are specific enough to survive examination and distinguish themselves from existing film simulation products (such as those offered by VSCO, Halide, or Fujifilm's own film simulation modes) would require analysis of the actual claim language.

Filed only as an international WIPO publication with no corresponding US patent filing, this application has not been granted anywhere and faces the full uncertainty of national phase examination. Its commercial potential is real but speculative — it depends on the claims surviving examination, on Netflix choosing to build the consumer product, and on the technology delivering the claimed fidelity at consumer scale.

CHAPTER 20

What Comes Next

Based on the patent portfolio and the strategic trajectory of the acquisition, several developments seem likely:

Real-time on-set AI assistance. The system currently works in post-production. Extending it to real-time on-set use would allow directors and cinematographers to preview AI-enhanced footage during principal photography — essentially a real-time 'what if' system for creative decisions.

Automated dailies processing at scale. Every Netflix production generates thousands of hours of dailies. Processing them through InterPositive's system could automate color correction, continuity checking, and technical quality assessment across the entire production slate.

Custom models per production. Training a production-specific model on a show's existing footage, so the AI learns the specific visual language of that particular project — its lighting style, lens choices, color palette, editorial rhythm.

Pre-visualization. Using the system in pre-production to generate previsualization footage that reflects actual cinematographic parameters rather than simplified 3D animation.

Consumer product launch. The consumer tools patent suggests a subscriber-facing feature is in development. This could be a significant competitive differentiator in the streaming market.

Extended patent filings. The captioner patent (US 12,511,904) is currently US-only. International filings seem likely. Additional patents covering new capabilities developed by the team at Netflix are also probable. Continuation applications from the existing patents could further extend coverage.

The Fincher validation. Bloomberg reports that David Fincher has used InterPositive products on an upcoming film starring Brad Pitt. If confirmed in the finished product, it would represent a critical proof point. If the technology performs at the level Fincher demands — and Fincher is known for technical precision — it will establish a benchmark that other productions and competitors will measure against.

CHAPTER 21

Conclusions



Affleck on the soundstage. The tools were designed by someone who understood production constraints firsthand.

The InterPositive patent portfolio represents the most comprehensive intellectual property position in AI-assisted filmmaking assembled to date. Fourteen filings cover every functional layer of a full-stack cinematic AI system, from physical LiDAR measurement through virtual simulation environments, structured metadata generation, language-model integration, style-controlled video generation, and consumer-accessible tools.

The technology is built on a thesis that has proven commercially viable: rather than trying to generate video from scratch using general-purpose AI trained on internet content, InterPositive built a system that understands the specific craft of cinematography and enhances real footage created by real filmmakers. The system works downstream of principal photography, augmenting rather than replacing the creative process. As Affleck concluded in the acquisition announcement: "I went into technology and did this because I love it. I'm passionate about it. And because I want to sustain and protect this thing that we all love."

Netflix's acquisition of this portfolio gives the streaming giant an exclusive, production-ready AI filmmaking system with no near-term competitive equivalent. The patent protection extends through

2045. The proprietary training data cannot be replicated. The specialized team is now embedded within Netflix's production infrastructure.

But the analysis must be honest about what the portfolio is and is not. It is a strong, coordinated IP position built on specific technical implementations and anchored by three granted US patents. It is not an impregnable fortress. The underlying ML techniques are commodity. Individual layers can be designed around. The WIPO applications are unexamined. The prior art landscape includes substantial published work in every domain the patents touch. The provisional application for the core Filmmaker model may or may not have been converted to a full patent. The consumer tools opportunity is speculative.

The strategic value lies not in any single patent or innovation, but in the systematic coordination of the entire pipeline combined with the proprietary data, specialized talent, and production integration that Netflix now controls. Each element reinforces the others. The patents protect the methods. The data makes the methods work. The team knows how to connect the methods to real production workflows. The production-ready pipeline makes the technology immediately deployable.

For Netflix, the acquisition delivers: exclusive production technology that competitors cannot access; production cost reduction through intelligent post-production automation; quality differentiation through AI-enhanced cinematographic consistency; a potential consumer product extension; a specialized team that understands both domains; and two decades of patent protection to defend the lead.

For everyone else in the industry, the acquisition represents an urgent signal. The individual components needed to build an equivalent system are available. The integration expertise is scarce. Whoever assembles it first — or acquires a team capable of doing so — captures the competitive advantage that InterPositive can no longer provide to the rest of the market.

Annex A: Individual Patent Analyses (B.1–B.16)

B.1

US 12,438,995: Integration of Video Language Models with AI for Filmmaking

Patent Number	US 12,438,995 B1
Filing Date	November 25, 2024
Grant Date	October 7, 2025
Inventor	Benjamin Geza Affleck-Boldt
Assignee	Fin Bone LLC
Independent Claims	1 (method), 7 (system), 14 (CRM) — 19 total claims

US 12,438,995 is the foundational interface patent of the InterPositive portfolio. It covers the integration of video language models (VLMs) with AI systems specifically designed for filmmaking. The patent protects the method by which filmmakers communicate creative intent using natural language and structured metadata rather than direct technical machine learning parameters. This is the critical bridge between human creative expression and AI generation — it solves the problem of how a director describes what they want and translates that description into machine-readable parameters that guide the AI system.

The patent covers the specific methodology of integrating VLMs with filmmaking-specific AI, including metadata processing, LiDAR data integration for spatial understanding, and the overall filmmaking pipeline architecture. The core innovation is not the VLM itself (which is a general-purpose model) but rather the adaptation of VLM outputs to control filmmaking parameters: camera position, lens choice, movement speed, lighting configuration, and compositional elements.

Claim breadth note: The independent claims require the system to receive LiDAR data — including spatial coordinates, distance measurements, and positional relationships — and integrate that with filmmaking metadata. A competitor generating AI video without a LiDAR capture step could argue they do not infringe, which narrows the claim's reach against systems that rely on monocular depth estimation or other non-LiDAR spatial understanding methods.

Independent Claims

Independent Claims: 1 (method), 7 (system), 14 (CRM)

- 1. Interfacing with an existing video large language model (LLM)
- 2. Receiving detailed metadata related to professional filmmaking techniques, including camera settings, shot composition, and lighting setups
- 3. Processing the received metadata to adapt the existing video LLM to generate video content that simulates professional filmmaking techniques
- 4. Receiving LiDAR data captured from a LiDAR sensor, including spatial coordinates, distance measurements, and relative positional information of objects within a scene
- 5. Integrating the LiDAR data with the processed metadata by combining spatial coordinates and distance measurements with filmmaking metadata to enhance the generated video content
- 6. Applying transfer learning techniques to the existing video LLM based on the processed metadata and LiDAR data to refine its video content generation capabilities

B.2

US 12,511,837: AI-Based Video Content Creation with Predetermined Styles

Patent Number	US 12,511,837 B1
Filing Date	November 25, 2024
Grant Date	December 30, 2025
Inventor	Benjamin Geza Affleck-Boldt
Assignee	Fin Bone LLC
Independent Claims	1 (method), 8 (system), 15 (CRM) — 20 total claims

US 12,511,837 protects style-controlled video generation — the capability to generate video content that faithfully reproduces a predetermined visual aesthetic. This patent represents a fundamental shift in how AI approaches video generation. Rather than creating video that simply meets functional requirements — correct camera movement, proper lighting exposure, appropriate composition — it teaches systems to generate video that conforms to predetermined aesthetic styles. This is the difference between generating a shot that has correct exposure and generating a shot that looks like it was shot by a particular cinematographer using particular aesthetic choices.

In practice, "predetermined styles" encompasses multiple concrete dimensions of visual identity: color grading consistency, lighting signatures, lens characteristics (focal length, iris characteristics, aberrations, depth of field), era-specific aesthetics (visual markers associated with particular periods of

cinematography), and motion characteristics (how the camera moves or doesn't). The technical approach relies on encoding style as actionable parameters that condition video generation.

Claim breadth note: The independent claims include a step requiring capture of "test images of the scene using different film formats to document visual effects." This ties the claim to a specific physical capture workflow. A competitor using purely synthetic or digitally-originated training data could argue this step is not met.

Independent Claims

Independent Claims: 1 (method), 8 (system), 15 (CRM)

- 1. Capturing one or more control images corresponding to a scene using standard digital video as a baseline
- 2. Capturing one or more test images of the scene using different film formats to document visual effects
- 3. Applying post-production alterations to the captured footage
- 4. Constructing a training dataset that includes a variety of shots captured under varied lighting conditions
- 5. Training an AI model with paired comparisons to enable it to learn specific visual signatures
- 6. Reviewing footage generated by the AI model to assess its authenticity and using feedback to refine the model
- 7. Optimizing learning cycles to enhance an efficiency of the training

B.3**US 12,511,904: Training a Captioner Model (SamildAnach)**

Patent Number	US 12,511,904 B1
Filing Date	March 12, 2025
Grant Date	December 30, 2025
Inventor	Benjamin Geza Affleck-Boldt
Assignee	InterPositive, LLC
Priority Date	November 27, 2024
Independent Claims	1 (method), 10 (system), 19 (CRM) — 26 total claims
Notes	68 pages — the most detailed patent in the portfolio

US 12,511,904 is, at 68 pages, the most detailed patent in the portfolio. It covers the training methodology for the SamildAnach captioner model — the system that analyzes footage and generates structured cinematographic metadata. This is the "reader" of the captioner-Filmmaker pair: it looks at video frames and articulates what it sees in the language of professional cinematography.

The captioner model is trained on a unique dataset specifically designed for cinematographic analysis. The training methodology uses multi-task learning — simultaneously training the model on multiple related tasks — to build a comprehensive understanding of visual production. The model learns to analyze shot types and compositions (establishing shots, close-ups, POV shots, over-the-shoulder), camera movements and technical parameters, lighting conditions and quality, film stock and processing characteristics, and color grading and exposure.

Independent Claims

Independent Claims: 1 (method), 11 (system), 21 (CRM) — 30 total claims

- 1. Organizing a dataset comprising raw video clips and corresponding metadata, wherein each video clip represents a specific shot varying one cinematic parameter at a time selected from the group consisting of focal length, camera movement, and framing
- 2. Extracting frames from the raw video clips at a consistent frame rate and storing the frames in a structured format
- 3. Associating each frame with corresponding metadata detailing the cinematic elements present in the frame, wherein the metadata includes information on focal length, camera movement, object distance, and framing style
- 4. Segmenting the raw video clips into shots and frames, wherein a shot comprises a continuous sequence captured without cuts, and frames are extracted at regular intervals from each shot

- 5. Applying frame-level labels to each frame based on the corresponding metadata, wherein the frame-level labels include focal length used during the shot, camera movement details, object distance from the camera, and framing style
- 6. Aggregating the frame-level labels to generate shot-level labels, wherein the aggregating includes calculating average focal length, determining predominant framing style, and smoothing camera movement data across the shot
- 7. Training the captioner model using the frame-level labels, the frames, and the aggregated shot-level labels to recognize and predict the cinematic elements in unseen video content
- 8. Iteratively refining the captioner model based on feedback from validation datasets to improve accuracy of cinematic element prediction
- 9. Deploying the trained captioner model to process and label a large video database, wherein the model generates metadata for new video content based on learned cinematic elements
- 10. Post-processing the generated labels to ensure consistency and accuracy, including performing outlier detection, confidence scoring, and manual quality control

B.4

US 12,322,036: LiDAR Data Utilization for AI Model Training in Filmmaking

Patent Number	US 12,322,036 B1
Filing Date	November 25, 2024 (US 18/959,397)
Grant Date	June 3, 2025
Priority Date	June 7, 2024 (via provisional US 63/657,756)
Inventor	Benjamin Geza Affleck-Boldt
Assignee	Fin Bone LLC → InterPositive, LLC
Independent Claims	1 (method), 8 (system), 15 (CRM) — 20 total claims
Notes	Foundation layer — every other patent depends on this layer

US 12,322,036 is the foundation layer of the InterPositive portfolio and the earliest granted US patent in the family (granted June 3, 2025). It covers the use of LiDAR sensors to correlate 2D video frames with precise 3D spatial data — giving the AI system a geometric understanding of the scene that no pixel-only approach can replicate. Every other patent in the portfolio depends on this layer: without the spatial understanding LiDAR provides, the entire system collapses into conventional pixel-based AI video generation.

The patent describes capturing spatial data using professional filmmaking equipment — Techno cranes, Panther Dollies, and KUKA robotic arms — to achieve sub-centimeter positioning accuracy. Critically, the spatial data is encoded via numerical tokenization rather than text descriptions: each coordinate is stored as a precise numerical value, preserving spatial accuracy and enabling the model to learn fine-grained geometric relationships between camera position, subject, and environment.

Claim breadth note: Every independent claim expressly requires LiDAR. A competitor using monocular depth estimation, stereo cameras, or structured light instead of LiDAR could argue non-infringement on the literal claim language — but would sacrifice the sub-centimeter spatial precision that enables the system's downstream capabilities. Google Patents notes one "cited by" publication: an OpenAI application titled "Generative video engine capable of outputting videos in a variety of durations, resolutions, and aspect ratios."

Independent Claims

Independent Claims: 1 (method), 8 (system), 15 (CRM) — 20 total claims

- 1. Receiving, by a computing system, LiDAR data captured by a LiDAR sensor on a film set, wherein the LiDAR data includes spatial coordinates, distance measurements, and relative positional information of objects within the scene
- 2. Processing the LiDAR data to generate a three-dimensional spatial map of the film set, wherein the spatial map includes positions of objects, lighting equipment, and camera positions
- 3. Correlating, by the computing system, the processed LiDAR data with video frames captured on the film set to create a spatially annotated video dataset
- 4. Generating, by the computing system, training data for an AI model using the spatially annotated video dataset, wherein the training data includes paired video frames and corresponding spatial metadata
- 5. Training, by the computing system, the AI model using the generated training data to enable the AI model to understand and generate video content that accurately reflects spatial relationships and depth cues within a scene

B.5 — B.16: International PCT Applications

The following twelve PCT applications were all filed on June 6, 2025, published in 2025, with priority dates of June 7, 2024. All are filed under Fin Bone LLC as applicant unless otherwise noted.

B.5

WO2025255425A1: Video Content Generation via Metadata and LiDAR

Publication Number	WO2025255425A1
Publication Date	Jun. 19, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

Covers the end-to-end filmmaking AI pipeline — the broadest patent in the portfolio. Describes how the system integrates all layers: capturing footage with spatial metadata, training the models, generating enhanced content. This is the primary international filing for the Filmmaker model.

B.6

WO2025255426A1: Simulation of Professional Filmmaking Techniques

Publication Number	WO2025255426A1
Publication Date	Jun. 19, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

The international counterpart of US 12,511,837, extending style-controlled generation patent protection to international jurisdictions. Protects the ability to simulate professional filmmaking techniques with controlled parameters and deterministic results.

B.7

WO2025255427A1: Synthetic Data Generation for ML in Filmmaking

Publication Number	WO2025255427A1
Publication Date	Jun. 19, 2025

Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

Covers the synthetic data component — generating artificial training data that supplements real production footage. Includes virtual camera rigs, lighting setups, and scene compositions that mirror real filmmaking. Enables continuous improvement of AI models between productions.

B.8

WO2025255428A1: ML Model Training with Parameter Variation

Publication Number	WO2025255428A1
Publication Date	Jun. 19, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

Covers the single-parameter variation training methodology — changing one cinematographic variable at a time while holding all others constant. This enables AI models to learn the precise effect of each parameter: how changing aperture affects depth of field, how changing focal length affects perspective.

B.9

WO2025255429A1: AI-Based Video Content Creation with Predetermined Styles (International)

Publication Number	WO2025255429A1
Publication Date	Jun. 19, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

The international counterpart of US 12,511,837, providing PCT coverage for style-controlled video content creation. Protects the methodology of applying predetermined visual styles — film stock emulations, lighting signatures, lens characteristics — to generated video using metadata extracted by a companion captioner system.

B.10

WO2025255432A1: AI-Driven Generation Meeting Professional Film Standards

Publication Number	WO2025255432A1
Publication Date	Dec. 11, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

Covers quality control validation for AI-generated content. Includes automated testing against professional filmmaking standards, human review workflows, and quality scoring. The patent ensures generated content meets the technical specifications required for professional film and television production.

B.11

WO2025255433A1: AI-Based Filmmaking Tools for Consumer Use

Publication Number	WO2025255433A1
Publication Date	Dec. 11, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

The consumer wildcard. Covers AI-based filmmaking tools for consumer use — capturing scenes in various formats and applying processing techniques, including film stock simulation. Opens potential subscriber-facing Netflix features. Filed only as international application with no corresponding US patent.

B.12

WO2025255436A1: Cinematic Data Collection and Processing

Publication Number	WO2025255436A1
Publication Date	Dec. 11, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025

Applicant	Fin Bone LLC
------------------	--------------

Covers the comprehensive data collection methodology for training AI filmmaking models. Includes multi-camera capture setups, synchronization protocols, metadata recording, and quality control procedures. The patent protects the specific methodology for creating high-quality training datasets for filmmaking AI.

B.13

WO2025255437A1: Integration of VLMs (International Filing)

Publication Number	WO2025255437A1
Publication Date	Dec. 11, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

The international counterpart of US 12,438,995, extending VLM integration patent protection beyond the United States through the WIPO/PCT system. Suggests it received special strategic attention as the most commercially important filing to protect worldwide.

B.14

WO2025255439A1: Training with Synthetic and Real-World Data

Publication Number	WO2025255439A1
Publication Date	Dec. 11, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

Covers the hybrid training approach — combining controlled synthetic data with real-world production footage. Synthetic data provides perfectly labeled ground truth; real-world dailies provide authentic variation. The patent covers combining both sources with a unified metadata framework and using LiDAR data as cross-validation. Every Netflix production generates additional real-world training data that continuously improves the models.

B.15

WO2025255441A1: Simulation-Driven Training Environment

Publication Number	WO2025255441A1
Publication Date	Dec. 11, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

Covers the virtual simulation environment — augmenting physical soundstage data with virtual scenes rendered using Unreal Engine or similar technology with accurate lens models, virtual lighting rigs, and proper 3D geometry. Virtual scenes include simulated LiDAR data in the same format as real-world captures. The advantage is scalability: virtual rendering is essentially free once built, enabling a single scene to be rendered thousands of times with different parameters.

B.16

WO2025255446A1: LiDAR Data Utilization for AI Model Training (International)

Publication Number	WO2025255446A1
Publication Date	Dec. 11, 2025
Priority Date	June 7, 2024
Filing Date	June 6, 2025
Applicant	Fin Bone LLC

The international counterpart of US 12,322,036 B1 (see B.4). Extends LiDAR-based spatial training protection to international jurisdictions via the PCT system. Subject to national phase examination. Given that US 12,322,036 is the foundation layer on which all other patents depend, this is strategically the most important international filing in the portfolio. Google Patents notes one 'cited by' publication: an OpenAI application titled 'Generative video engine capable of outputting videos in a variety of durations, resolutions, and aspect ratios.'

References

- [1] Netflix, "Innovation for Filmmaking, By Filmmakers: Why InterPositive Is Joining Netflix," [about.netflix.com](https://about.netflix.com/en/news/why-interpositive-is-joining-netflix), March 5, 2026. <https://about.netflix.com/en/news/why-interpositive-is-joining-netflix>
- [2] Netflix, Inc., "Q3 2025 Shareholder Letter," October 2025. https://s22.q4cdn.com/959853165/files/doc_financials/2025/q3/FINAL-Q3-25-Shareholder-Letter.pdf
- [3] Netflix, Inc., "Q4 2025 Shareholder Letter," January 2026. https://s22.q4cdn.com/959853165/files/doc_financials/2025/q4/FINAL-Q4-25-Shareholder-Letter.pdf
- [4] Netflix, acquisition announcement video featuring Ben Affleck, Bela Bajaria, and Elizabeth Stone, February 2026. <https://www.youtube.com/watch?v=o06LsOgSiv4>
- [6] Ben Affleck, interview on The Joe Rogan Experience (Episode #2440), January 2026. <https://www.youtube.com/watch?v=cWfKpFrEMn8>
- [7] Ben Affleck, remarks at CNBC Delivering Alpha, 2024. <https://www.youtube.com/watch?v=SVfGrncHF4>
- [8] United States Patent & Trademark Office, patents assigned to Benjamin Geza Affleck-Boldt / InterPositive, LLC / Fin Bone, LLC. US 12,438,995 B1 (granted Oct. 7, 2025); US 12,511,837 B1 (granted Dec. 30, 2025); US 12,511,904 B1 (granted Dec. 30, 2025). <https://www.uspto.gov>
- [9] World Intellectual Property Organization, PCT publications WO2025255425A1 through WO2025255446A1, published 2025. <https://www.wipo.int/pct/en/>
- [10] Bloomberg, "Netflix to Pay as Much as \$600 Million for Ben Affleck's AI Firm," March 11, 2026. <https://www.bloomberg.com/news/articles/2026-03-11/netflix-to-pay-as-much-as-600-million-for-ben-affleck-s-ai-firm>
- [11] Gatys, L.A., Ecker, A.S., and Bethge, M. "A Neural Algorithm of Artistic Style," 2015. arXiv:1508.06576.
- [12] Ruder, M., Dosovitskiy, A., and Brox, T. "Artistic Style Transfer for Videos," 2016. arXiv:1604.08610.
- [13] Press reporting on Netflix's El Eternauta AI-assisted VFX (Buenos Aires building-collapse sequence), 2025.

Netflix - InterPositive

3 Granted US Patents · 11 PCT Applications · 289 Claims
Patent Protection Through 2045
