Plausibility is All You Need:

New Film Restoration Techniques

Erica Andreose [orcid:0009-0003-7124-9639],
University of Bologna, Bologna, Italy
erica.andreose2@unibo.it

Mateusz Miroslaw Lis [orcid:0009-0009-1697-9149],
ML/AI Engineer, Orkidee - AI Laboratory, Padua, Italy
miro@orkidee.it

Massimo Toniato [orcid:0009-0000-3988-4243],
Independent videomaker; Teacher of Audiovisual and Multimedia Laboratory, IIS "G. Valle", Padua, Italy; Consultant and trainer in the audiovisual field
info@massimotoniato.com

DOI: 10.6092/unibo/amsacta/8622

Preprint version

Submitted to the

Conference on Ethics and Aesthetics of Artificial Images

Venice, May 8-10, 2025

Iuav University of Venice & Venice International University (VIU)

Acknowledgments: Frames from *The Rival Sculptor* (1911) are reproduced by courtesy of the Cineteca di Bologna for research purposes.

ABSTRACT. This research investigates the philosophical and technical implications of applying generative artificial intelligence to film restoration, focusing on the notion of *plausibility* as a new aesthetic and epistemological paradigm. The study explores how AI-based frame interpolation models can reconstruct missing or degraded sequences in early twentieth-century films, not by recovering lost information, but by generating perceptually coherent continuities between surviving frames. In this context, authenticity no longer coincides with material originality but with the consistency of visual experience. To assess the boundaries and potential of interpolation-based restoration, we introduce a three-tier framework comprising *Frame Complexity*, *Motion Complexity*, and *Damage Complexity* (and an additional one *Negative Frame Complexity*). A case study on *The Rival Sculptors* (1911) demonstrates how plausibility can operate as a guiding principle for preserving cinematic experience across time and technology. By bridging aesthetic theory and computational methodology, the research suggests that AI does not simply reproduce the past but actively constructs believable futures for visual memory.

KEYWORDS. AI, film restoration, plausibility, generative models, authenticity, digital heritage.

1. Introduction

In recent years, the integration of artificial intelligence into cultural-heritage practice has begun to challenge the traditional boundaries between authenticity, authorship, and simulation. Within the domain of moving-image preservation, this tension acquires a particularly vivid form: cinema has always inhabited the fragile space between what is materially real and what merely feels real. Our project, *Plausibility Is All You Need*, was conceived within this liminal territory, as a collaborative effort between researchers and practitioners coming from three complementary environments: the University of Bologna (Digital Humanities and Digital Knowledge), the University of Udine (Digital Philosophy and AI Master's Program), and Orkidee, an applied-research context dedicated to the use of AI in cinematic and audiovisual production.

Our starting question was apparently simple: what do we actually see when we watch a digitally restored film? Unlike the static arts, where a reproduction can be clearly distinguished from its original, the moving image complicates this relationship. Every screening, whether on celluloid or on a streaming platform, is already a re-production, a renewed performance of an underlying sequence. The material carrier may change, yet the perceptual identity of the film persists. This continuity of perception, rather than the persistence of matter, becomes the true locus of authenticity. To illustrate this paradox, during our talk in Venice, we referred to the canonical example of Magritte's *The Treachery of Images*. The painting's caption, "Ceci n'est pas une pipe," reminds us that representation is never the thing itself. When applied to film, this insight expands dramatically: a restored frame is not the historical frame; it is an image that feels like it could have been there. The core concept emerging from our experimentation is therefore plausibility, the ability of an artificial image to sustain the illusion of continuity without claiming ontological truth.

Our case study centres on *The Rival Sculptors*[1] (1911), an American short film preserved by the Cineteca di Bologna.[2] The footage, affected by severe mechanical and chemical degradation, offered an ideal terrain for testing generative methods aimed at reconstructing missing or unreadable portions. We deliberately limited our workflow to open-source tools, in order to evaluate what can be achieved without proprietary or high-cost resources. The resulting methodology re-frames restoration as a process of regeneration: rather than recovering lost information, AI can fabricate believable transitions between surviving fragments, thus restoring the perceptual rhythm of motion.

The following sections describe the conceptual and technical foundations of this approach, outlining the framework of three complementary metrics: Frame Complexity, Motion Complexity, and Damage Complexity, that emerged from practice. Through these metrics,

[1] IMDb page: https://www.imdb.com/it/title/tt1065493/?ref_=ext_shr_lnk [Accessed 10 October 2025]
[2] https://cinetecadibologna.it/ [Accessed 10 October 2025]

plausibility becomes not merely an aesthetic impression but an operational criterion for assessing AI-assisted film restoration.

2. Methodology

The methodological core of our project lies in translating the theoretical notion of plausibility into a reproducible computational process. In designing our workflow, we compared several classes of generative models currently used in visual reconstruction, ranging from diffusion-based image synthesis to neural inpainting and motion transfer. Each of these approaches embodies a different philosophy of digital restoration: diffusion models invent missing content, inpainting models fill absence, while motion models connect what survives. We deliberately chose to explore the third approach, focusing on frame interpolation, as the most coherent with our conceptual premise.

Interpolation operates not by recreating what has been lost, but by predicting the transition between existing visual fragments. Within the context of moving images, this means predicting how motion, light, and texture evolve from one frame to the next. Traditional film restoration techniques often rely on static correction, repairing or replacing damaged frames through manual retouching or image compositing. In contrast, our AI-based interpolation methodology treats motion itself as the unit of restoration: rather than filling the gaps by duplicating or smoothing existing frames, the system learns to generate plausible transitions that preserve the temporal coherence of the sequence. The “restored” frame becomes a hinge between two authentic ones, a plausible image that sustains the illusion of motion without pretending to recover historical truth.

2.1 Choice of Tools

Our workflow relied on two open-source frame-interpolation models: RIFE[3] (Real-Time Intermediate Flow Estimation) and FILM[4] (Frame Interpolation for Large Motion), selected for their ability to handle complex motion dynamics and degraded visual input with limited computational cost. Beyond their conventional use for smooth motion synthesis, we hypothesised that these networks could be strategically repurposed as generative systems, capable of reconstructing plausible continuity across damaged sequences. An example of the application of such models can be seen in Figure 1.

[3] https://github.com/hzwer/ECCV2022-RIFE [Accessed 10 October 2025]
[4] https://film-net.github.io/ [Accessed 10 October 2025]

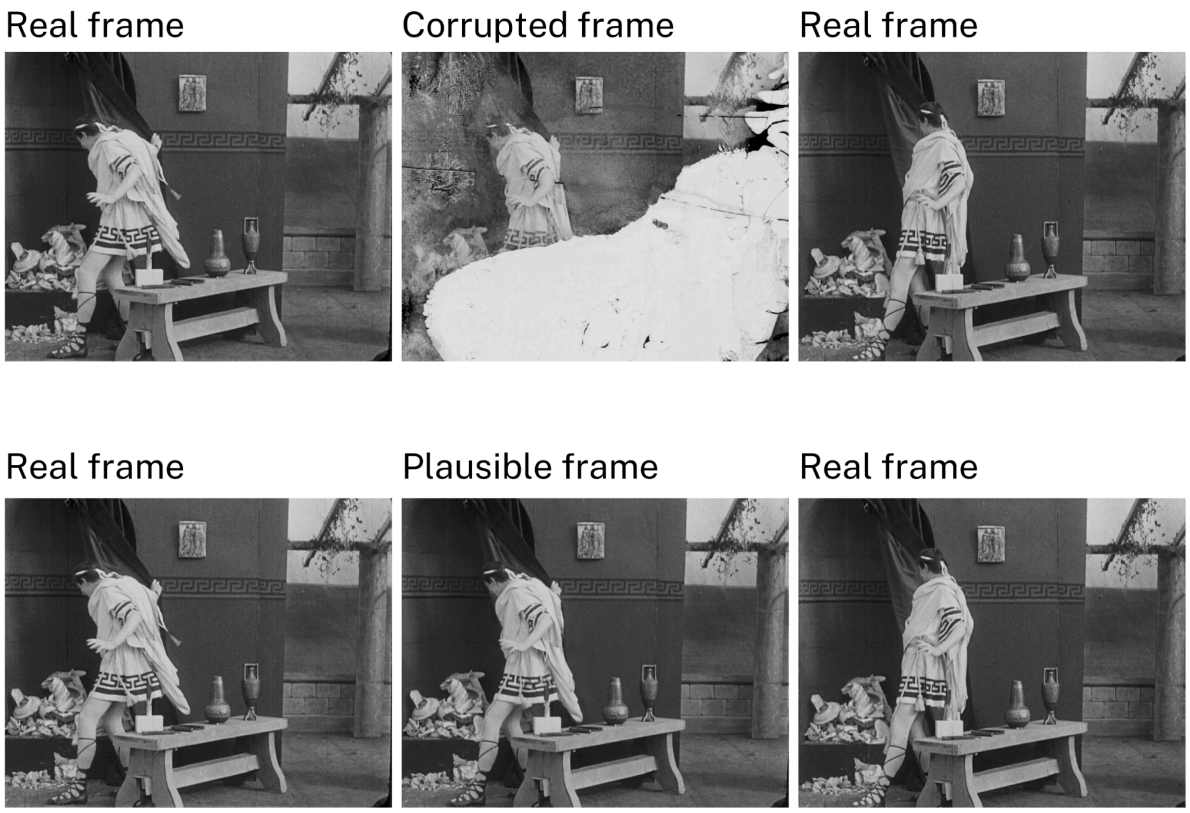

Figure 1. Exemplification of the concept of Plausibility applied to frame interpolation. Above, original frames from test footage. Below, the central frame has been reconstructed in a plausible way with interpolation models.

All operations were performed within *ComfyUI*[5], an open-source, node-based environment that allows modular integration of multiple AI components. This architecture enabled us to design and visualise each stage of the workflow: data preparation, frame removal, interpolation, post-processing, and visual comparison, while maintaining full transparency and reproducibility. One of the workflows implemented for this project is shown in Figure 2.

[5] https://www.comfy.org/ [Accessed 10 October 2025]

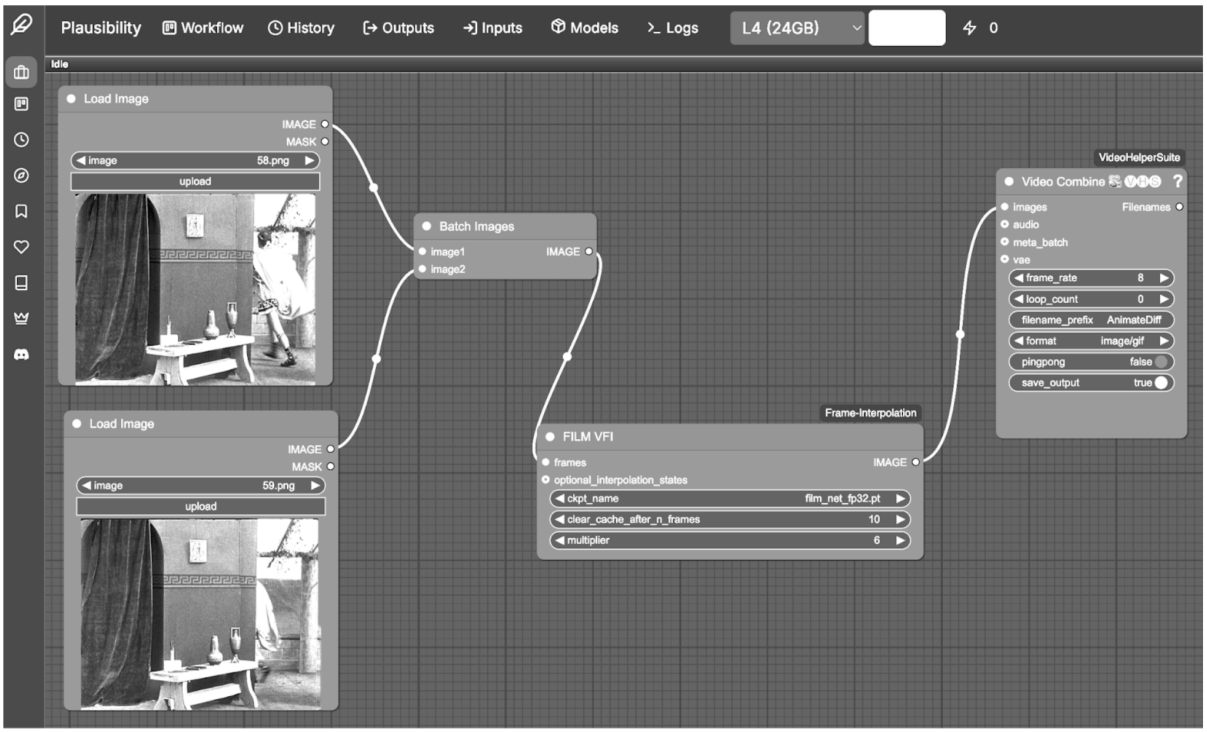

Figure 2. Example of a ComfyUI workflow used in the project. Each node represents a distinct processing step.

The choice of an open and reproducible framework was, for us, not merely a technical convenience but an ethical commitment consistent with the broader aims of *Plausibility Is All You Need*. Within the field of film heritage, the credibility of a digital restoration depends as much on how an image is reconstructed as on what the viewer ultimately sees. Proprietary or closed-source systems often conceal the decision-making processes encoded in their architectures, turning restoration into a form of algorithmic opacity, a "black box" whose inner logic remains inaccessible. By contrast, adopting open-source tools allowed us to align our methodology with our theoretical stance: if plausibility is to be understood as a shared, perceptual contract rather than a hidden algorithmic effect, then the process that generates it must also be open to inspection and replication. Through a reproducible workflow, we sought to ensure that both the methods and the results of our experiments could remain accessible, verifiable, and extendable by future researchers and institutions.

2.2 Operational Pipeline

The practical sequence followed four steps:

1. Identification and Removal of Damaged Frames: Sections exhibiting severe physical or chemical deterioration were isolated and excised. The resulting discontinuities defined the regions in which synthetic frames would later be generated, allowing controlled experimentation with different degrees of visual absence.

2. Interpolation and Generation: Using RIFE developed by Huang et al. (2022) and FILM introduced by Reda et al. (2022), we produced new intermediate frames bridging the removed segments. Each network was tested independently to observe qualitative differences in motion prediction, spatial coherence and texture reconstruction.

3. Evaluation and Comparison: Generated sequences were visually compared against adjacent original frames, focusing on coherence of motion and absence of artefacts such as ghosting or warping.

4. Iterative Refinement: Multiple passes were applied when residual temporal discontinuities were detected, adjusting node parameters directly in *ComfyUI*.

Through this process, interpolation becomes a controlled act of creation rather than a blind algorithmic correction. The researcher actively selects, tests, and validates each plausible solution.

2.3 Three Metrics for Assessment

Empirical work quickly revealed that the success or failure of interpolation could not be predicted by a single numerical measure. We therefore developed a qualitative table of metrics, each addressing a different dimension of difficulty:

Frame Complexity (FC)

Defined as the number and distribution of missing frames within a given shot. A higher FC generally increases computational effort and perceptual risk, but the correlation is not linear; occasionally, low FC sequences yield poor results due to other factors.

Motion Complexity (MC)

The relative amount and velocity of movement between adjacent frames. High MC sequences, such as rapid pans or gestures, tend to generate “ghosting” artefacts, while static or low-motion scenes interpolate smoothly. Although our evaluation was primarily qualitative, MC could in future be quantified by measuring pixel-change ratios.

Damage Complexity (DC)

A descriptor of the morphological variety and severity of visible damage (scratches, flicker, chemical blots). While DC plays a secondary role in interpolation, since corrupted frames are usually removed entirely, it becomes crucial when hybrid workflows integrate in-painting techniques.

Negative Frame Complexity (NFC)

Referring to situations where the number of missing frames cannot be determined because whole fragments are lost. In such cases, temporal inference becomes speculative, pushing interpolation toward the boundary between restoration and creative reconstruction.

The methodological design thus oscillates between scientific control and artistic judgement. Each generated frame represents a hypothesis about how motion might have continued. Validation cannot rely solely on pixel-level similarity but must appeal to human perceptual coherence,the viewer's ability to accept the motion as natural. This criterion of acceptance anchors the broader notion of plausibility that guides the entire project.

3. Results

The experimental phase focused on a 19 seconds fragment of *The Rival Sculptors* (1911), selected in collaboration with the Cineteca di Bologna. This sequence exhibited both mechanical and photochemical deterioration, including scratches, frame loss, and irregular motion. Its heterogeneous composition made it a paradigmatic test for evaluating the generative plausibility of interpolation models under real archival conditions.

After isolating and excising all unusable frames, we applied the workflow described above. Each tool (RIFE and FILM) was tested independently, followed by comparative analysis of their outputs. The criteria were visual continuity, absence of artefacts, and perceptual acceptance when viewed in real-time playback. An example of the results is shown in Figure 3.

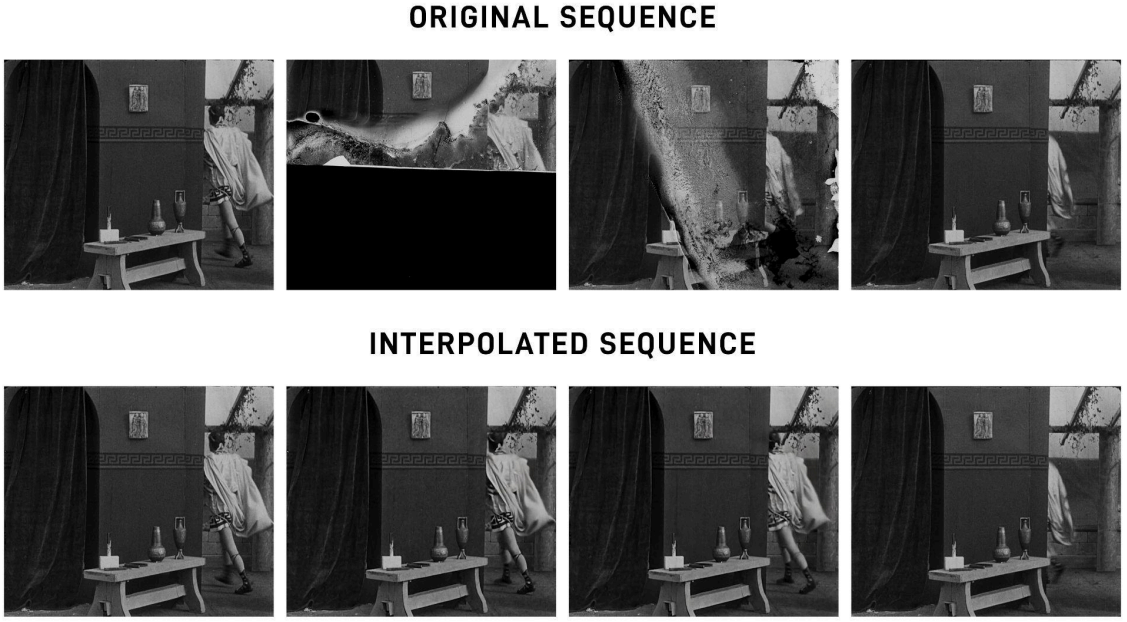

Figure 3. Example of an edited sequence via interpolation models. Above is the four-frame original sequence, below the same sequence including the regenerated frames (second and third position) to ensure perceptual continuity.

3.1 Observed Behaviour Across Complexity Levels

A pattern of correlation emerged observing the four metrics mentioned above:

Frame Complexity (FC)

As expected, the probability of plausible reconstruction decreases with the number of consecutive missing frames. Yet this relationship proved non-deterministic. Some low-complexity gaps failed completely, while certain high-complexity gaps generated surprisingly coherent results. The determinant appeared not to be the quantity of missing data but the predictability of motion within the scene.

Motion Complexity (MC)

Motion emerged as the dominant factor in determining success. When camera or subject movement was smooth and unidirectional, interpolation succeeded even with several missing frames. In sequences with erratic gestures or depth shifts, both RIFE and FILM produced ghosting and edge duplication. These effects, though algorithmically logical, betrayed the perceptual illusion of continuity. A particularly ghosting-affected example is shown in Figure 4.

Damage Complexity (DC)

Because damaged frames were removed before interpolation, DC had limited immediate impact. However, in adjacent regions where residual flicker or contrast loss remained, the models tended to over-smooth details, erasing part of the film's texture. This points to the need for adaptive pre-processing that respects historical grain and luminance variation rather than suppressing them.

Negative Frame Complexity (NFC)

Revealed itself during experimentation. In portions where entire fragments were missing, interpolation could not infer temporal structure. The generated frames, while visually smooth, lacked narrative alignment. Such cases marked the outer boundary of plausibility: a visually convincing illusion detached from historical continuity.

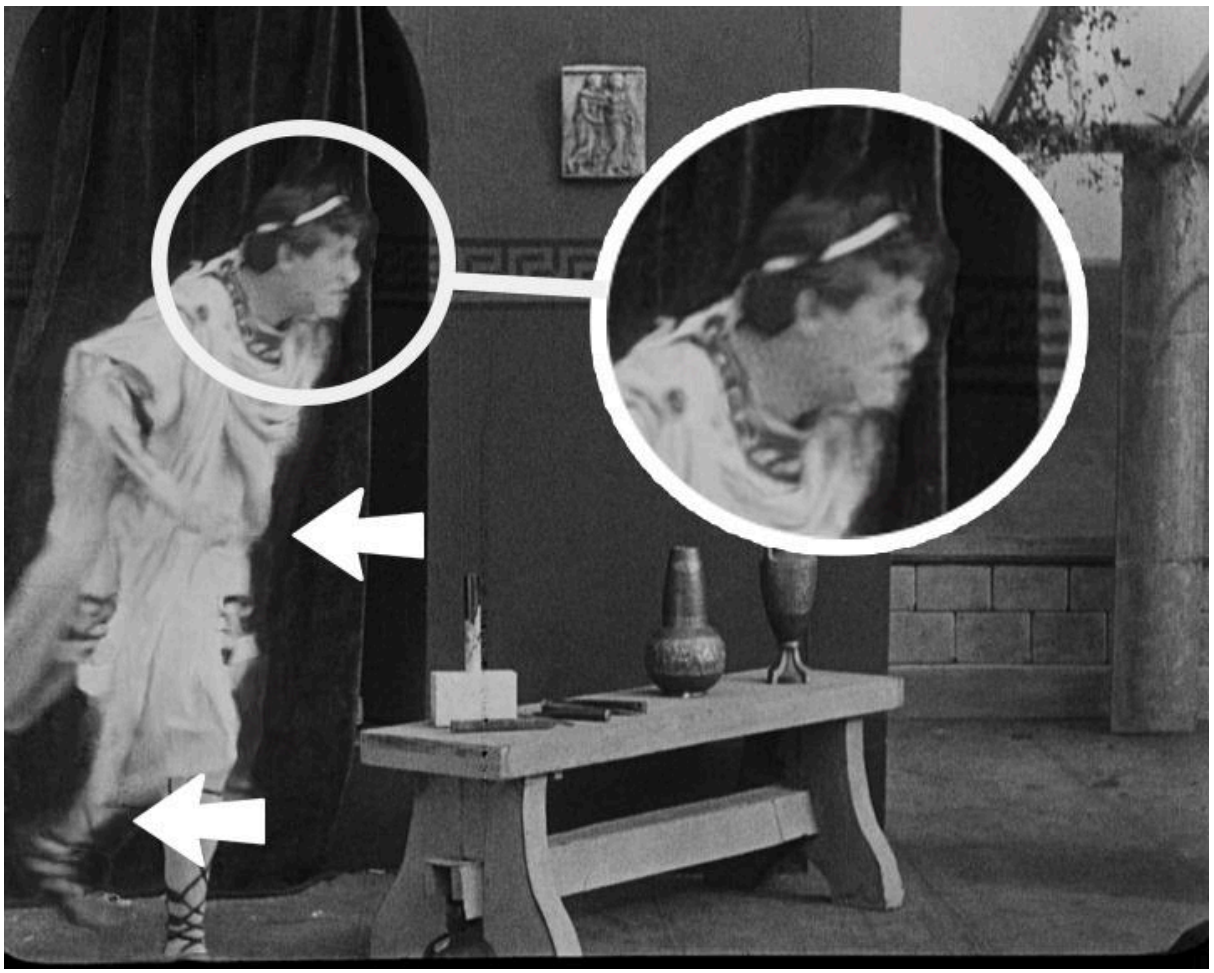

Figure 4. Example of a frame heavily affected by interpolative ghosting. The protagonist's facial features are so distorted as to appear unrecognizable; moreover, the arm and heel display a liquid-like deformation (see arrows).

In Figure 5, the matrix visualises the relationship between Frame Complexity (x-axis) and Motion Complexity (y-axis), illustrating how the interplay between these two factors determines the plausibility of AI-based restoration. Regions with low FC and low MC correspond to sequences that interpolate smoothly, while higher values along either axis increase the likelihood of artefacts such as ghosting or temporal drift. The upper-right area of the matrix thus marks the threshold where interpolation alone becomes insufficient and additional generative methods, such as diffusion-based[6] synthesis, may be required. The diagonal gradient expresses the transition from feasible restoration to speculative regeneration, emphasising that motion, more than frame count, is the decisive constraint on plausibility.

[6] A class of generative models that iteratively reconstruct an image or video by reversing a noise process. Starting from random noise, the model gradually "denoises" the signal, learning to generate new frames consistent with the statistical structure of the training data. Unlike interpolation networks, which predict motion between existing frames, diffusion models can invent entirely new visual transitions, enabling plausible but non-deterministic motion reconstruction.

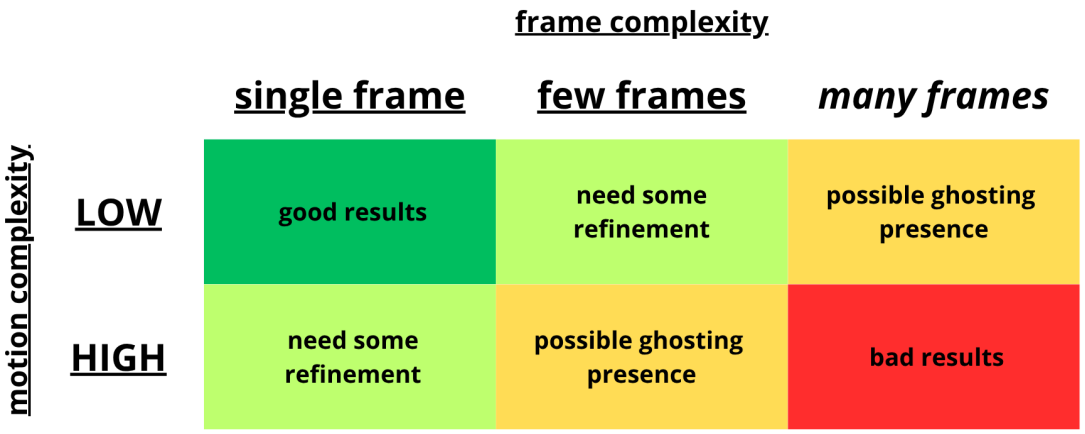

Figure 5. Interaction between Frame Complexity and Motion Complexity showing how increased motion amplifies restoration difficulty.

3.2 Comparative and Perceptual Evaluation of Models

When comparing the performance of RIFE and FILM, we observed complementary strengths rather than a clear hierarchy of results. RIFE produced smoother and more stable global motion, especially in scenes involving large displacements, but tended to blur fine-grained details and soften the film's texture[7]. FILM, conversely, preserved local sharpness and grain but occasionally introduced geometric distortion or flickering edges in fast-moving sequences.

The most convincing outcomes were achieved through selective compositing, combining the global coherence of RIFE with the local fidelity of FILM. The modular flexibility of *ComfyUI* proved essential for testing these hybrid configurations, enabling precise parameter control and rapid visual comparison.

From a phenomenological standpoint, plausibility could be felt rather than measured. When the regenerated segment played seamlessly within the surviving footage, the viewer's attention remained engaged; when artefacts appeared, the illusion broke instantly. The

[7] The term film's texture refers to the material and granular quality of the image, encompassing the fine visual details that contribute to the aesthetic character of the source material. It includes elements such as film grain (or its digital emulation), micro-contrast, and subtle variations in brightness and color that convey depth and realism. The loss or smoothing of this texture, as occasionally observed in RIFE interpolations, can result in an overly polished appearance, diminishing the perceived cinematic look of the footage.

restoration's value, therefore, resided not in recovering an "original" frame but in sustaining the viewer's belief in continuity, the perceptual contract at the heart of cinema itself.

4. Discussion

The term *plausibility* emerged from our practice, not from pre-existing theory. Traditional restoration vocabulary, *authenticity*, *fidelity*, *integrity*, proves inadequate once AI systems intervene to create content that never existed in the source material. What these systems generate is neither true nor false; it is convincing. The goal becomes not to reproduce the past but to preserve the experience of seeing a coherent motion through time. In this sense, the passage from restoration to regeneration is central. Restoration implies recovery of something that once existed[8]. Regeneration acknowledges that loss is irreversible and proposes to fill the gap with plausible invention. The regenerated frame is a fiction that serves truth; the truth of perception. In practice, this means accepting a controlled degree of artificiality as a condition for maintaining experiential authenticity.

Using AI to recreate missing imagery inevitably raises ethical questions. Should future viewers be informed that certain sequences are algorithmically generated? In our view, transparency should accompany technical innovation, yet disclosure need not undermine the aesthetic experience. What matters is that the reconstructed passage does not distort historical context or misrepresent intent. Within these boundaries, AI acts as a form of interpretive mediation, similar to translation: it enables access without claiming originality. As Zhou and Ramli (2025, p. 39) observe, "AI not only improves the visual effects of films but also enhances the audience's ability to empathize through modern audio-visual means, giving these historical images a new lease of life." Cinematic continuity is not only visual but emotional. Even minor temporal discontinuities can fracture immersion and displace the viewer from the film's affective rhythm. Unlike static works of art, where cracks or missing fragments may coexist with appreciation, moving images require seamless motion to sustain meaning. Hence plausibility becomes not merely a technical goal but an ethical commitment to preserve the feeling of the work[9].

The operator of generative restoration systems occupies an ambiguous position between archivist, artist, and engineer. Decisions regarding frame removal, interpolation parameters, and evaluation criteria are interpretive acts. The researcher's expertise thus becomes the guarantor of plausibility. Human judgement, not the algorithm alone, determines when a regenerated sequence "feels right." The human eye remains the final arbiter of continuity.

---

[8] "Restoration represents the methodological moment of recognizing the work of art in its physical consistency and in its dual aesthetic and historical nature, with a view to its transmission to the future." (Brandi, 1963) (translation by the authors).

[9] "The digital realm seemingly allows one to do (almost) anything; therefore, the most difficult task for those engaged in digital restoration is not to overcome limitations, but to set them—maintaining a balanced assessment between the vast possibilities of manipulation and the respect for the originals and the actual needs of the project". (Dagna, 2014) (translation by the authors).

5. Conclusion

This study deliberately focused on open-source interpolation models, privileging transparency and reproducibility over raw generative power. Yet several directions emerge for future development. Recent progress in video diffusion models suggests the possibility of synthesising motion transitions directly, rather than inferring them through optical flow prediction. Integrating diffusion into our workflow could mitigate ghosting artefacts and enable controlled extrapolation beyond existing frames, though current computational demands remain significant. A complementary path involves domain-specific training: fine-tuning networks on curated datasets of archival film would allow them to internalise historical characteristics, like grain, contrast, and framing conventions. The resulting systems would act as context-aware restorers, embedding cultural knowledge into generative inference.

Similarly, hybrid workflows that combine interpolation and in-painting could balance temporal coherence with spatial reconstruction, using *ComfyUI*'s modular structure to orchestrate both processes within a unified pipeline. Beyond technical evolution, these directions converge toward a broader paradigm of living archives, in which preservation becomes an ongoing, regenerative dialogue between heritage and innovation. Rather than a static act of recovery, digital collections could evolve through periodic re-generation, updating obsolete formats and recovering usability as technologies change. In this sense, generative plausibility becomes not a one-time intervention but an ongoing dialogue between heritage and innovation.

The experiments conducted on *The Rival Sculptors* have demonstrated that plausibility can serve not only as an aesthetic criterion but also as a methodological bridge between archival conservation and computational creativity. In reframing restoration as an act of regeneration, the project challenges conventional definitions of authenticity and proposes a more dynamic, perception-based understanding of film heritage. Artificial intelligence does not resurrect the past; it constructs credible continuities that allow the past to remain visible. Each interpolated frame is a hypothesis, one among many possible, about how movement might have unfolded. By accepting this hypothesis as plausible, the spectator re-enters the filmic world without noticing the fracture that once separated presence from absence. The illusion of continuity thus becomes a form of historical empathy: a way of caring for images by allowing them to move again.

This approach also reveals the evolving role of the restorer. The human operator is no longer a mere conservator of matter but a designer of perception. Their task is to negotiate between evidence and imagination, calibrating algorithms until the regenerated sequence "feels right". In doing so, they reaffirm that the essence of cinema lies not in chemical emulsions or digital codecs, but in the sustained belief that what moves before our eyes is, if not true, at least possible. The philosophical stakes extend beyond archival practice. Plausibility suggests a paradigm for how we might coexist with generative technologies more broadly: accepting synthetic artefacts not as deceptions but as tools for maintaining continuity of memory, of culture, of meaning across the discontinuities of time. In this sense, AI does not threaten the

ontology of cinema; it re-enacts it. From the very beginning, film has relied on illusion, on the rapid succession of still frames that deceive the eye into perceiving motion. Generative restoration simply renews that ancient pact between perception and invention. Ultimately, the project affirms that what matters in preservation is not what is real, but what feels alive. The plausible image, born from collaboration between algorithm and archivist, becomes the medium through which lost movements are imagined back into being. By embracing plausibility as both a practical and philosophical principle, we open a path toward future-oriented film heritage, one in which authenticity is measured not by material origin, but by the continuity of experience it enables.

References

Brandi, Cesare (1963), 'Teoria del restauro', con L. V. Borrelli & J. R. Serra, Roma: Edizioni di Storia e Letteratura.

Dagna, Serena (2014), 'Perché restaurare i film?', Pisa: Edizioni ETS (Collana Vertigo).

Huang, Zhewei, Zhang, Tianyuan, Heng, Wen, Shi, Boxin & Zhou, Shuchang (2022), 'Real-Time Intermediate Flow Estimation for Video Frame Interpolation', *arXiv preprint*, https://arxiv.org/abs/2011.06294.

Reda, Fitsum, Kontkanen, Janne, Tabellion, Eric, Sun, Deqing, Pantofaru, Caroline & Curless, Brian (2022), 'FILM: Frame Interpolation for Large Motion', *arXiv preprint*, https://arxiv.org/abs/2202.04901.

Zhou, Linjuan & Ramli, Sharulnizam. (2025). *From Photochemistry to AI: The Evolution and Multiple Values of Film Cultural Heritage Restoration Technology. Scientific and Social Research*. 6. 37-41. doi: 10.26689/ssr.v6i12.9258.