digitcult @Scientific Journal on Digital Cultures

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

Nodal Prompting: rethinking authorial practice in Generative AI

Nodal Prompting: ripensare la pratica autoriale nell'IA Generativa

Giuseppe Bruno
0009-0009-4472-6683
Politecnico di Torino
ROR 00bgk9508
giuseppe_bruno@polito.it

Tatiana Mazali
0000-0002-2828-694X
Politecnico di Torino
ROR 00bgk9508
tatiana.mazali@polito.it

10.54103/2531-5994/31893

Abstract

[En.] The growing integration of Generative AI tools in audiovisual production has turned prompting into the primary interface between human intent and machine output. However, existing theoretical frameworks treat the prompt as a discrete textual artifact, overlooking the increasingly complex interaction models enabled by node-based workflow environments such as ComfyUI, where generative instructions are distributed across interconnected parametric structures.
This article proposes the concept of nodal prompting, a mode of human-AI interaction intrinsic to graph-based generative tools, in which authorial intent is wired across a network of nodes, each contributing partial specifications that jointly determine the generative output.
The argument is developed through qualitative analysis of a production case study: Il partigiano Libero, an audiovisual series on a partisan during the Italian Resistance, developed by Motion Pixel and ANPI Chiomonte, and situated within a critical review of existing literature on prompt engineering, parametric authorship, and human-AI co-creation.
The analysis identifies a constitutive paradox at the core of nodal prompting: the author delegates expressive execution to the machine while simultaneously exercising greater upstream control through system design. This redistribution of creative agency challenges linear models of AI-assisted authorship and suggests that efficiency gains in generative production are conditional on significant investment in workflow architecture. Ultimately, the figure of the author is not dissolved but fundamentally restructured.

Keywords: Generative AI, ComfyUI, Nodal prompting, Human-AI interaction, Creative agency, Audiovisual production.

[It.] La crescente integrazione degli strumenti di IA generativa nella produzione audiovisiva ha trasformato il prompting nell'interfaccia principale tra l'intento umano e l'output della macchina. Tuttavia, i quadri teorici esistenti trattano il prompt come un artefatto testuale discreto, trascurando i modelli di interazione sempre più complessi resi possibili da ambienti di workflow basati su nodi come ComfyUI, in cui le istruzioni generative sono distribuite su strutture parametriche interconnesse.
Questo articolo propone il concetto di nodal prompting, una modalità di interazione umano-IA intrinseca agli strumenti generativi basati su grafi, in cui l'intento autoriale

Received: 17/06/2026
Accepted: 16/07/2026
Published: 06/07/26

Milano University Press

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

è collegato attraverso una rete di nodi, ciascuno dei quali contribuisce con specifiche parziali che concorrono a determinare l'output generativo.
L'argomentazione viene sviluppata attraverso l'analisi qualitativa di un caso studio produttivo: Il partigiano Libero, una serie audiovisiva su un partigiano durante la Resistenza italiana, sviluppata da Motion Pixel e ANPI Chiomonte, e inserita in una rassegna critica della letteratura esistente su prompt engineering, autorialità parametrica e co-creazione umano-IA.
L'analisi individua un paradosso costitutivo al centro del nodal prompting: l'autore delega l'esecuzione espressiva alla macchina, esercitando al contempo un maggiore controllo a monte attraverso la progettazione del sistema. Questa ridistribuzione dell'agency creativa mette in discussione i modelli lineari di autorialità assistita dall'IA e suggerisce che i guadagni in termini di efficienza nella produzione generativa dipendano da un investimento significativo nell'architettura del workflow. In definitiva, la figura dell'autore non viene dissolta, ma radicalmente ristrutturata.

Parole chiave: IA generativa, ComfyUI, prompting nodale, interazione uomo-IA, agenzia creativa, produzione audiovisiva.

1. Introduction

Generative AI is now systematically integrated into audiovisual pipelines, redefining the timescales, costs, and visual grammar of filmmaking. Thus, it has altered the conditions under which visual content is conceived and realized (Banh & Strobel, 2023). High-end proprietary models such as Runway Gen-3 and Sora have opened the cinematic world to faster prototyping and previsualization workflows, but their closed, subscription-based nature and content moderation policies restrict both research freedom and the depiction of sensitive material (Singh, 2023). Text-to-Image and Text-to-Video models are now capable of generating high quality outputs from natural language descriptions and have shifted significant portions of the visual development process towards a paradigm in which one of the main instruments of creative expression is the prompt. Both in the industry discourse and academic research, the prompt has emerged as the main defining artifact of human-AI interaction in generative contexts; a prompt is now the unit through which human intention is translated into machine output.

The prompt is treated by the community as a fundamentally textual and linear entity, whether in the form of LLM instruction sequences and reasoning patterns (Brown et al., 2020; Wei et al., 2023), reusable prompt pattern catalogs modeled on software design (White et al., 2023), or automatic optimization strategies for diffusion-based text-to-image generation (Lian et al., 2024). In all these trajectories, the prompt remains anchored to the monodimensional paradigm of a textual sequence, however structured or compound, that is issued to a model and processed as a discrete input.

These assumptions leave a conceptual void regarding practices that node-based workflow environments such as ComfyUI provide (Deng et al., 2025; Kyaw & Sivalingam, 2025). In these, generative instructions are not organized as text sequences, but the prompt appears distributed across directed graphs of different nodes. Each and every

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

node is composed of different parameters, whose configuration jointly determines the generative output (Deng et al., 2025). In these contexts, the design of a system becomes the real authorial act, leaving behind the classical inscription of a textual instruction as the only control element. What node-based workflows expose is that a substantial and often dominant portion of creative control resides not in the text field of the positive prompt node, but in the topological choices, parameter weights, model hierarchies, and conditioning signals that compose the graph [1].

This article fills that void by proposing the concept of nodal prompting as a new form of authorial agency: a mode of human-AI interaction intrinsic to graph-based generative environments, in which authorial intent is wired across a network of interconnected nodes, each contributing partial specifications that jointly determine the generative output. With this new concept, the kind of prompt engineering also differs and changes from the classical definition. If the textual prompt is a speech act destined to a model defined by a message, the nodal prompt is instead an architecture that preconfigures the conditions of generation. The minimum unit of authorial intervention is no longer the string of text, but the coordinated ensemble of nodes, parameters, and connections composing a generative workflow.

The nodal prompting concept carries a constitutive paradox. In fact, the author who works within a node-based workflow delegates expressive execution to the machine, while simultaneously exercising a form of upstream control over the system (Goodfellow, 2024; Gunkel, 2025). This paradox is developed fully in Section 4.

These theoretical propositions were first documented in Bruno (2025) and were developed through qualitative analysis of a production case study: *Il partigiano Libero*, an audiovisual series depicting the story of a partisan fighter during the Italian Resistance in the Val di Susa, produced by Motion Pixel in collaboration with ANPI Chiomonte and realized entirely within an open-source ComfyUI pipeline incorporating Stable Diffusion XL, FLUX.1 [dev] (Black Forest Labs, 2024), CogVideoX (Hong et al., 2023; Yang et al., 2025), and custom-trained LoRA models (Hu et al., 2022). The case is analytically significant for several intersecting reasons. First, it demanded consistent visual characterization of a historically specific protagonist across different generative scenes. Second, the deliberate use of locally deployable open-source tools bypassed the content moderation policies of proprietary platforms, enabling the representation of wartime imagery that closed systems would have filtered. Third, the complexity of the resulting workflow, spanning image generation, video synthesis, character conditioning, and super-resolution upscaling, makes it an exemplary instance of nodal prompting in professional practice. The case is analyzed in conjunction with a critical review of existing literature on prompt engineering, parametric authorship, and human-AI co-creation.

[1] In ComfyUI, creative control is distributed across functional nodes (model loaders, samplers, ControlNet conditioners, LoRA injectors), each exposing specific parameters that jointly determine the output. The text prompt occupies a single node among many. See LearnOpenCV Team (2024), https://learnopencv.com/introduction-to-comfyui-for-stable-diffusion/; and Raphiki (2025), where the prompt is formalized as a JSON guide injected into specific graph nodes, effectively becoming a structured meta-prompt: https://dev.to/worldlinetech/automating-image-generation-with-n8n-and-comfyui-521p.

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

2. Literature review

2.1. Prompt engineering and its taxonomic limits

There is a foundational assumption in the literature on prompt engineering: the prompt is a discrete, textual, and linear artifact. It is the primary unit through which human intention is translated into machine output. This assumption, productive within a wide range of interaction contexts, has been elaborated into increasingly systematic frameworks. Brown et al. (2020) established the canonical taxonomy of zero-shot, one-shot, and few-shot prompting, demonstrating that pre-trained language models can be steered toward novel tasks by varying the number and structure of in-context examples provided in the input sequence. Building on this foundation, Wei et al. (2023) introduced chain-of-thought prompting, a technique that embeds explicit intermediate reasoning steps into the prompt, typically activated by formulas such as "let's think step by step", and demonstrated consistent performance gains on complex multi-step tasks. Both contributions share a structural commitment: the prompt is a string of natural language, however elaborate, issued to a model as a single discrete input.

White et al. (2023) systematized this approach into a Prompt Pattern Catalog, explicitly modeled on software design patterns. It documented reusable interaction structures such as *role prompting*, *output automater*, and *flipped interaction*. The catalog extends the notion of the prompt from a formulation to a design. It also acknowledges that practitioners increasingly work with families of prompts, sequences of instructions, and composable interaction templates. Yet even within this expanded conception, the fundamental unit of authorial intervention remains the text string, in which patterns are patterns of language, not configurations of heterogeneous computational elements.

With the diffusion of text-to-image models, the literature has largely focused on prompt optimization, i.e., the automatic rewriting or augmentation of user descriptions to improve alignment with model preferences. Lian et al. (2024) propose LLM-grounded Diffusion, a framework in which a large language model first generates structured spatial layouts (bounding boxes, positional annotations, scene graphs) that serve as intermediate representations guiding the diffusion model. This approach introduces a non-textual control dimension, but it does so within a pipeline where layout generation is itself treated as a prompt-level task, produced by another language model and injected as a structured string. The spatial and parametric dimensions of generation remain subordinate to a textual logic.

What these trajectories share, and what constitutes their collective limit, is the assumption that generative instructions are organized as text sequences, however structured, compound, or automatically generated. This assumption leaves a conceptual void regarding practices enabled by node-based workflow environments such as ComfyUI, where the prompt appears as distributed across directed graphs of different nodes, each contributing partial specifications that jointly determine the generative output. In these contexts, the text field of the positive conditioning node is one input among many: it operates alongside *LoRA weights*, *sampler configurations*, *guidance values*, *ControlNet signals*, and *seed*

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

states, each of which encodes a dimension of authorial intent that no single prompt string could express. Existing taxonomies of prompting offer little room for this mode of interaction, and workflow design itself has rarely been discussed as a form of authorial practice.

2.2. From distributed agency to parametric and distributed authorship

Authorship was once seen as a singular, intentional, and expressive act, traceable to a discrete origin in the mind of a human creator. This assumption has been progressively eroded by practice, long before the emergence of generative AI, but it has come to surface again with node-based generative workflows.

Goodfellow (2024) articulates this condition through a systems-theoretic framework, proposing a spectrum of authorship that runs from pure human creation - "which draws directly from the interior world of affect, emotions and ideas" - through co-evolved works produced with tools and collective labor, to works that are largely devoid of human involvement. The critical contribution of this framework is not the spectrum itself, but the observation that AI-generated works do not occupy its edges: they are "the consequence of a highly distributed model of production" that sits across algorithms, training data, underlying information systems, and the artists who both contribute to and extract value from them. The expansion of AI, Goodfellow argues, "entangles the artist in ever more complex webs of production and dissemination, whereby the boundaries between the work of the artist and the work of the networked technologies are increasingly distributed and obscured." Within this condition, the role of the artist shifts from the creator of aesthetic objects to their "translator and curator", a figure who selects, conditions, and steers within generative infrastructures that exceed direct control.

Recent empirical studies have also supported these theories regarding node-based generative systems. In a workflow-based approach to game prop design in ComfyUI, Deng et al. (2025) have demonstrated that the node-based architecture enables "full-process parametric control of prop design from concept to detail" while substantially reducing production time. Their study reveals that creative control in such environments does not reside primarily in natural language description but in the configuration of the generative graph itself (the selection of base models, the injection of conditioning adapters, the calibration of sampling parameters), each of which functions as a discrete authorial decision encoded in a non-textual form. Kyaw & Sivalingam (2025) extend this logic to multimodal generation, presenting a node-based platform for integrated text, audio, image, and video production. Their system demonstrates that node-level architecture enables "targeted editing of individual nodes, automatic branching for parallel storylines, and node-based iterative refinement", capacities that require the graph structure to function as the primary site of creative control, rather than as a support layer for textual instruction.

While these studies document the technical affordances and operational efficiency of node-based environments, they remain system-oriented and do not treat the redistribution of control as a matter of authorship. The concept of nodal prompting moves beyond this operational dimension, treating workflow design itself as an authorial act

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

rather than a technical optimization, and identifying the paradox of authorship that motivates this article: intensified control at the level of system design coupled with delegated execution at the level of individual generation. This paradox, developed fully in Section 4, offers a way of framing what nodal prompting adds to purely system-oriented accounts of node-based generation.

The most important counterpart to these observations is given by Gunkel (2025), which argues that generative AI systems cannot be adequately understood through the instrumentalist theory of technology, which situates any tool as a neutral means to a human-determined end. The act of prompting "is not a command but an invitation to joint participation", an engagement with a system that responds with output that is often unpredictable and reformulatory. Moreover, recognizing AI systems as cooperative participants in the *semiotic economy of writing*[2] does not require granting them intentional agency, but does require acknowledging their contribution to the co-construction of meaning. Taken together, these perspectives converge on a shared diagnosis: in generative AI workflows, authorship is not dissolved but redistributed, relocated from the moment of textual enunciation to the upstream act of system architecture.

2.3. Human-AI co-creation and agency

One of the most important questions that this article raises and that has become central in professional communities is how agency is distributed between human and machine in co-creative contexts. As AI systems transition from tools performing static tasks to collaborative partners capable of proactive and potentially autonomous action, understanding the mechanisms through which agency is allocated and negotiated becomes a prerequisite for both system design and critical analysis.

Starting from the sociotechnical turn in media studies and more recent scholarship on the platform society, agency belongs not only to human beings but also to the technologies that participate in action (Latour, 2005). The relation between people and media is a process of co-construction, in which human action and sociotechnical configurations mutually shape one another. In this perspective, agency is not a fixed property of the individual subject, but a relational effect emerging from the interaction between actors, infrastructures, and affordances. The concept of agency thus evolves to include relationality with machines and their affordances, giving rise to an idea of distributed and relational agency (Couldry, 2012; Van Dijck et al., 2018).

The actor-network theory of Latour (2005) was fundamental in challenging the idea that action originates only in human intention, by showing how non-human entities also mediate, enable, and constrain action within heterogeneous assemblages. Couldry (2012) extended this insight to digital media, arguing that media are not neutral channels, but socially embedded infrastructures that organize what can be seen, said, and

[2] Gunkel refers to the distributed system of sign production and meaning-making in which human authors and generative models both participate, not as neutral tools but as participants contributing to the co-construction of meaning.

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

done. Van Dijck et al. (2018) further developed this sociotechnical view to platform ecosystems, where mechanisms of datafication, commodification, and selection encode normative configurations that distribute agency across users, interfaces, algorithms, and governance structures. In these three interpretations, the common thread is the rejection of both technological determinism and social determinism: the relationship between technology and users is the result of an ongoing negotiation between the parties.

Within generative AI environments, this sociotechnical perspective becomes directly operative. In a complex analysis, Zhang et al. (2025) revealed that agency in human-AI co-creation operates along a spectrum from passive tools to cooperative partners, and that its distribution is both dynamic, shifting across interaction stages, and granular, enacted through specific operational mechanisms at the input, action, output, and feedback layers of the interaction cycle. The review distinguishes agency as the intentional initiation of action and the subjective experience of directing outcomes, as distinct from control, that is, the concrete mechanisms and commands through which those intentions are executed (Zhang et al., 2025). This distinction is analytically central: it allows for a multi-level reading of co-creative interaction that separates goal-directed guidance from its operational realization. The review identifies a critical gap in the existing literature, i.e., the absence of a comprehensive synthesis mapping agency patterns to operational control mechanisms across diverse co-creative contexts and modalities.

The existing literature provides the conceptual tools to understand how agency, control, and authorship are distributed across sociotechnical environments, and to establish the conceptual infrastructure on which the analysis of nodal prompting depends. From prompt engineering, this article inherits the question of what constitutes the minimum unit of human-to-model communication, and the recognition that this unit has progressively expanded in structural complexity. From distributed and parametric authorship, the claim that creative agency in generative AI contexts is neither singular nor textual but architecturally encoded and distributed across the topology of the generative system rather than concentrated in any single inscription. Lastly, from research on co-creation, it inherits the distinction between agency and control as analytically separable dimensions of interaction.

The gap at the intersection of these three strands is precise: current frameworks do not appropriately recognize workflow design in node-based generative environments as a form of authorial practice, or the mode of human-AI interaction it instantiates. This gap motivates the proposal of nodal prompting, developed in Section 3 through the case study.

3. Case Study

The case selected is *Il partigiano Libero*, an audiovisual project produced by Motion Pixel in collaboration with ANPI Chiomonte, documenting the story of a partisan who participated in the Italian Resistance during World War II. The specific constraints of the project (historically sensitive imagery, the requirement for visual character consistency across episodes, and a zero-budget infrastructure) made the adoption of a fully open-source,

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

node-based pipeline a structural necessity rather than an aesthetic preference. The final result is conceived as a 90-second audiovisual output, designed to combine a cinematic look with realism and stylistic coherence across characters, environments, and scenes. Within this case, nodal prompting is a necessary and visible need. Its high-level workflow architecture has given a degree of parametric control that would otherwise have been delegated to opaque infrastructure.

The production pipeline comprised: Stability Matrix as the local environment manager; ComfyUI as the node-based generative interface; Stable Diffusion XL and FLUX.1 [dev] for image generation; CogVideoX for video synthesis; a custom LoRA model trained via Kohya_ss on archival reference photographs of the protagonist; ElevenLabs for voice generation.

3.1. Methodology

This study adopts a qualitative single case study design as its primary methodological framework. The research focuses in depth on a single instance to examine, through close analysis, the possible contours of nodal prompting, rather than pursuing statistical generalizability across a broad sample of AI-assisted productions.

The empirical basis of the analysis consists of materials generated during the production of the case study: approximately 70 ComfyUI workflow files documenting the pipeline's evolution across production stages (60 Text-to-Image workflows, from initial Stable Diffusion tests to the final FLUX.1 [dev] configuration; 8 Image-to-Video workflows, based on CogVideoX and AnimateDiff; 2 video upscaling workflows); production documentation including the detailed scene-by-scene script, voice-over text, and prompt list, the table of fixed seed values (Table 1), and the LoRA training dataset (55 archival photographs and their 100 × generic counterpart set, described in detail in Bruno (2025)); and a broader corpus of approximately 250 generated images and 150 generated videos.

The parameters of the analysis are the following: LoRA loading, describing the construction of the protagonist's visual identity, and discussed in Section 3.2.1; LoRA strength values (`strength_model`, `strength_clip`), representing the modulation of its influence within the generation process, discussed in Section 3.2.2; seed selection and fixation, which describe the scene-by-scene continuity of the visual style, in Section 3.2.3.

The three parameters analyzed here were selected as the most consequential for stylistic control among those available in the workflow, and the ones on which the most extensive practical experimentation was conducted during production. Their calibration followed an iterative, practice-based process of trial and adjustment rather than a systematic variation testing protocol: values were adjusted scene by scene until a satisfactory visual outcome was reached. The observations were then interpreted through the theoretical framework developed in Section 2, with particular attention to how creative agency is distributed across the nodal structure rather than concentrated in a single prompt string.

There were other parameters, used in the study but not specifically analyzed for the purpose of this article. These were related to base model choice, positive and negative conditioning, sampler configuration, guidance scaling via FluxGuidance, scheduler configuration, and VAE decoding.

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

3.2. The pipeline analysis and three nodal acts

In conventional filmmaking, authorial decisions are distributed across a chain of specialized agents: screenwriter, director of photography, editor, colorist. Each of these figures operates within a defined phase of production, the workflow is sequential, and responsibility is legible through role attribution. In *Il partigiano Libero*, this principle becomes concrete: LoRA weights, scheduler configurations, guidance values, and seed states each encode a dimension of authorial intent that the text prompt alone does not capture. Figure 1 illustrates the ComfyUI workflow used in the study, showing the node graph that encodes each authorial decision. Before a single image is generated, the author has already made many decisions that will constrain and direct every subsequent output: which base model to load, which LoRA to inject and at what strength, how to configure the sampler, which seed to fix or release, how to route the latent image through conditioning layers. These are not preparatory choices that precede the creative act. They are the creative act.

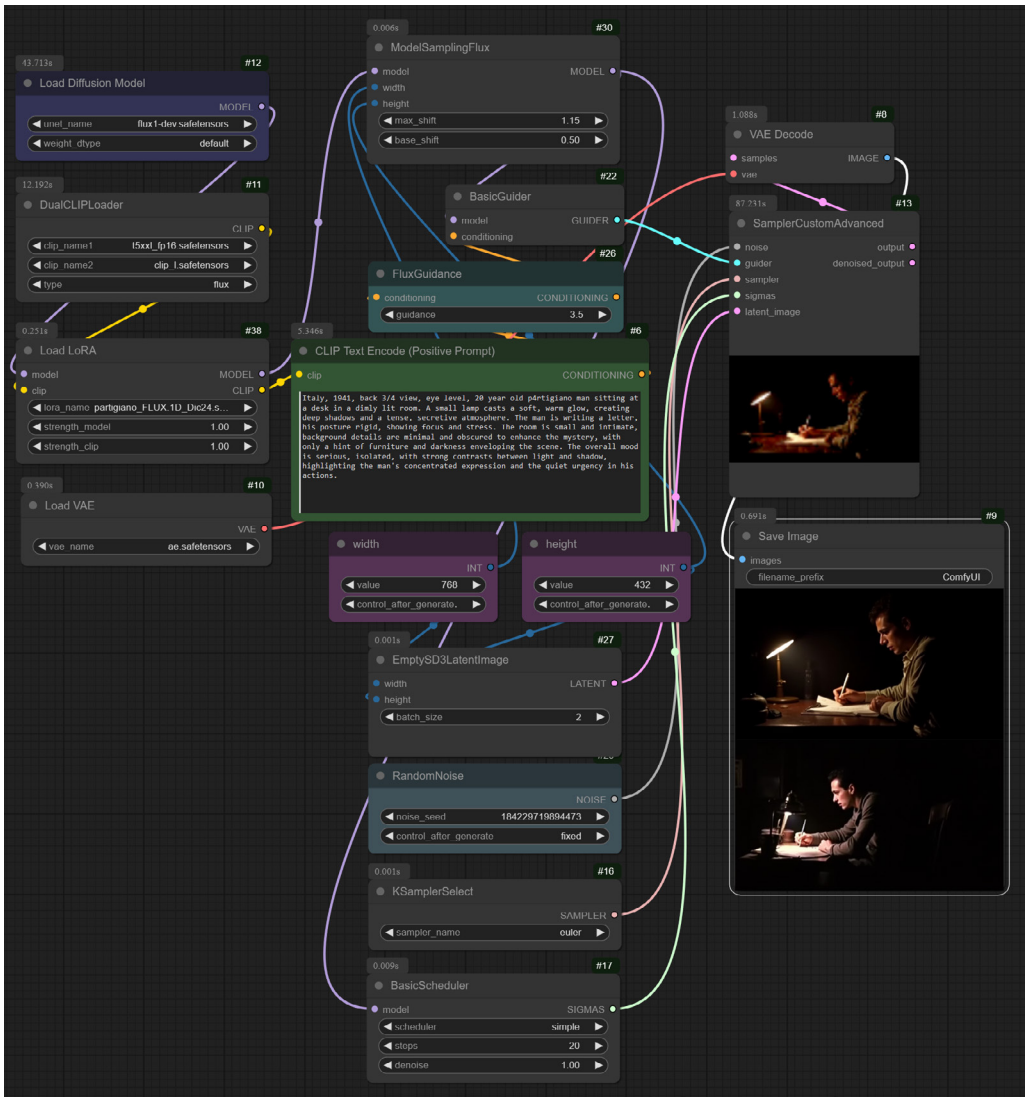

Figure 1. ComfyUI workflow for image generation using the FLUX.1 [dev] model and the custom LoRA trained on the partisan character

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

This observation is central to the concept of nodal prompting. The text prompt entered into the positive conditioning node is, in this pipeline, only one of many inputs that co-determine the output. It operates alongside LoRA weights, scheduler configurations, guidance values, and seed states, each of which encodes a dimension of authorial intent that a single prompt string alone could not express. The pipeline is not a tool the author uses; it is a structure the author writes, and each subsequent generation is a performance of that structure.

What follows is an analysis of three nodal configurations in the *Il partigiano Libero* workflow that exemplify this logic; three moments where a technical parameter functions, on closer inspection, as an act of authorial writing.

3.2.1. *The LoRA as Casting Decision*

The most consequential single decision in the production of *Il partigiano Libero* was the construction of a custom LoRA model trained on archival reference photographs of the protagonist, a partisan fighter whose face and bearing had to remain consistent across scenes set in different locations, different lighting conditions, and different emotional registers.

A LoRA (*Low-Rank Adaptation*) is a fine-tuning technique that modifies the weight space of a base diffusion model by injecting a learned low-rank matrix trained on a specific dataset. In practical terms, training a LoRA on 55 curated photographs of the protagonist, with a 100 × factor of generic male portraits (5,500 images generated with Stable Diffusion XL) and then injecting it into the FLUX.1 [dev] pipeline, meant that a particular visual identity was embedded into the generative model itself, not described to it, but constituted within it. The trigger word **`p4rtigiano man`**, required to activate the LoRA during generation, functions less like a prompt keyword and more like a proper name: it calls a specific entity into existence within the latent space. Figure 2 shows a generated image only using the base model (Figure 2a), and an image generated with the custom LoRA model (Figure 2b).

Read as an authorial act, LoRA training is the equivalent of casting. In traditional filmmaking, casting precedes shooting: the director selects a body, a face, a physical presence that will carry the narrative across scenes. Here, that selection happens upstream of any individual generation, encoded in the model weights themselves. The author is not describing a character; they are instantiating one. And critically, this instantiation is durable; it persists across every node in the graph that inherits the LoRA-modified model, establishing a kind of authorial commitment that no per-scene prompt could replicate with comparable consistency. The LoRA is the character before the character appears on screen.

This has implications beyond consistency. The choice of which photographs to include in the training dataset (their framing, their lighting, their historical specificity) shapes the visual register of the character in ways that are difficult to override through prompting alone. The author's curatorial decisions during dataset preparation (55 images selected from a larger archive, cropped and labeled with specific captions) are already expressive choices, encoding an interpretation of the protagonist's visual identity that will inflect every subsequent generation. The dataset is, in this sense, the first draft of the character.

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

(a) Image without LoRA (b) With the created LoRA model

Figure 2. Image generation results using FLUX.1 [dev] for scene 3: (a) without LoRA and (b) using the created LoRA model

3.2.2. *LoRA Strength as Identity Negotiation*

If LoRA training instantiates the protagonist as a stable visual entity, and seed fixation controls the compositional grammar of the scene, the LoRA strength parameters (`strength_model` and `strength_clip`) govern the relationship between that entity and its environment at the moment of generation. These two scalar values, both ranging from 0 to 1 and located in the Load LoRA node (as shown in Figure 3), determine how strongly the LoRA's learned weight modifications influence the base model and the CLIP text encoder, respectively. At maximum strength, when both values are 1, the LoRA dominates: the protagonist's visual identity overrides the base model's tendencies, producing a highly consistent but potentially rigid output. At reduced strength, in this production frequently `strength_model = 0.8` and `strength_clip = 0.7`, the LoRA and the base model enter into a weighted negotiation, with the base model contributing a degree of contextual flexibility that the LoRA alone could not provide.

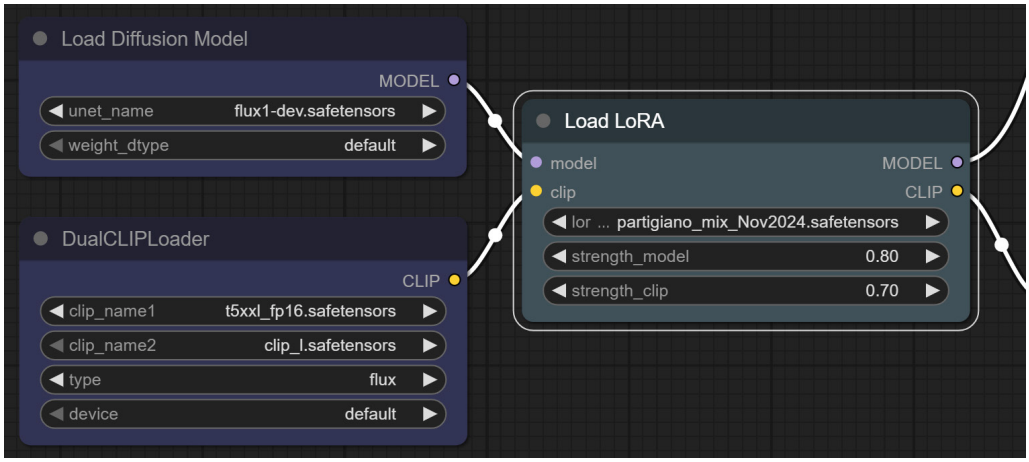

Figure 3. ComfyUI workflow detail for image generation showing the FLUX.1 [dev] model loading node, the dual clip loader, and the custom LoRA loader of the specific model trained on the partisan character, with `strength_model = 0.8` and `strength_clip = 0.7`.

This negotiation is not merely technical. Lowering `strength_model` means allowing the base model's broader visual knowledge, i.e., its understanding of period

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

lighting, architectural textures, atmospheric conditions, to inflect the output more strongly, at the cost of some identity rigidity. Lowering **`strength_clip`** means relaxing the constraint on how the text prompt is interpreted through the LoRA's modified embedding space, giving the CLIP encoder more latitude to respond to the textual description of the scene rather than anchoring everything to the LoRA's learned associations. The choice of specific values for each scene is therefore a decision about the relative weight of character identity versus scene expressivity, a tension that any narrative director must manage, and that here is encoded directly into the node parameters.

In the specific case of scenes where the protagonist appears in historically precise environments (wartime interiors, mountain paths, partisan gatherings), the authors found that default strength values produced outputs that felt visually over-determined, the character too plastically stable against backgrounds that demanded period specificity. Reducing strength allowed the base model's historical visual knowledge to surface, producing images where the protagonist inhabits the scene rather than being composited onto it. This micro-calibration, repeated across each scene type, constitutes a form of continuous directorial adjustment that in conventional production would be distributed across lighting design, costume, and set dressing departments. Here, it is condensed into two floating-point values.

3.2.3. *Seed Fixation as Narrative Control*

In diffusion models, the seed is the numerical value that initializes the random noise from which an image is generated. Changing the seed produces a different image from the same prompt and model configuration; fixing the seed guarantees that the same starting noise will be used across multiple generations, producing outputs that are structurally similar even when other parameters vary. In standard usage, seed management is discussed as a reproducibility mechanism, a way to recover a result that happened to be satisfactory. In the *Il partigiano Libero* workflow, it functions as something considerably more deliberate.

The production documentation records a set of fixed seeds assigned to specific scenes: the partisan's face across scenes 3 and 7 shares a single seed value; the letter-writing scene has its own; the mountain landscapes another (see Table 1 for the list of all seeds used; Figure 4 for an example of how seed value is entered into the system). This assignment is not arbitrary. Fixing a seed for the protagonist's scenes means that the spatial distribution of noise (and therefore the broad compositional logic of the generated image) remains stable across variations in prompt and lighting description. The character occupies roughly the same region of the frame; the posture and general atmosphere rhyme across scenes. The seed, in this reading, is not a technical parameter but a continuity decision: it performs the function that a shot list performs in conventional production, establishing a recurring visual grammar that the viewer can track across scenes.

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

Table 1. Seed values used for image generation

Seed	Scene	Value
partisan	3, 7	184229719894473
letter	4	716242073457185
letter delivery	5	143001186633932
mountain landscapes	6	331831932758698
city landscapes	1	927715408442632

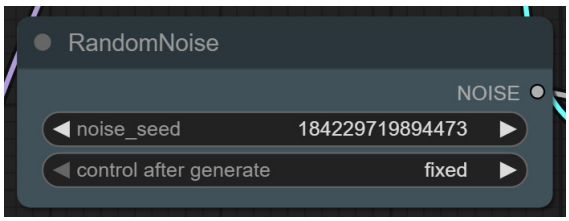

Figure 4. RandomNoise node showing the fixed seed for scenes 3 and 7. "`control after generate`" value can be **`fixed`**, **`increment`**, **`decrement`**, or **`randomize`**.

Conversely, releasing the seed, i.e., setting the RandomNoise node to generate a new value for each run, is an equally deliberate choice that signals a different authorial posture: an openness to variation, a willingness to let the model improvise within the constraints set by the other nodes. In scenes depicting landscapes or crowd backgrounds, where exact consistency is less narratively critical, the seed was left free. The author is not simply accepting randomness; they are allocating it, deciding which elements of the visual world are subject to authorial control and which are delegated to the model's generative latitude. Seed management, understood this way, is a form of selective authorship; a continuous negotiation between intention and contingency that runs through the entire production.

4. Discussion

The three nodal acts analyzed above converge on a structural paradox that is, we argue, constitutive of nodal prompting as a mode of authorship. Each decision described (training a LoRA, fixing a seed, calibrating strength values) represents an act of upstream control: the author intervenes at the level of the system before the generative act occurs, encoding intent into architecture rather than into instruction. In this sense, nodal prompting represents an intensification of authorial control relative to single-prompt interaction; the author is not asking the model to produce something but designing the conditions under which production becomes possible.

And yet this intensification of upstream control is inseparable from a corresponding delegation of downstream execution. Once the workflow is built and the parameters are set, generation proceeds without further authorial intervention. The model performs the work of image-making within the constraints the author has established, and its outputs are not fully predictable even within those constraints. The author does not

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

draw the image; they construct the system that draws it. This is not a diminishment of agency but a transformation of its locus, from the moment of execution to the moment of architecture.

What emerges from the case of *Il partigiano Libero* is that nodal prompting is not just a different way of writing prompts or an expanded version of prompt engineering, but a different form of authoring specific to node-based generative environments. Unlike Deng et al. (2025) and Kyaw & Sivalingam (2025), who document these parametric mechanisms as system affordances, this case reads them as authorial acts: casting, continuity, and directorial negotiation performed through nodes rather than through text. The analysis conducted in Section 3 clearly shows that the most consequential creative decisions were distributed across the workflow architecture, rather than concentrated in a single point. These decisions acted before individual generations, shaping the conditions under which images could be produced and interpreted. In this sense, the case study conducted supports the central claim of this article: in node-based generative systems, authorship is relocated from the prompt to the structured composition of a more complex workflow.

The concept of nodal prompting becomes useful in contexts like this, naming a mode of creative practice in which the author designs the environment that makes that particular generation possible, rather than simply asking a model to produce one image. The workflow becomes the site of composition, in which the prompt is just one of the main nodes in the graph, and it is no longer sufficient as the main unit of generation, although it remains relevant. In fact, in the case examined, the prompt was only one of the elements inside a broader system that already encoded decisions about character identity, visual continuity, and stylistic direction. This is the core difference between traditional prompt-based interaction and nodal prompting: the change of centrality from one to another. Prompt engineering treats the prompt as the central interface between intention and output; in nodal prompting, the workflow graph is itself the operative form of authorship.

The literature reviewed in Section 2 helps clarify why this distinction matters. Classical taxonomies of prompting, including more advanced approaches, remain grounded in a textual understanding of human-machine interaction. Even more recent proposals that extend prompting towards structured layouts or reusable patterns still preserve the prompt as the fundamental unit of intervention. The most consequential creative decision in the production of *Il partigiano Libero* was not expressed as text but encoded parametrically in the architecture of the graph itself. Existing taxonomies offer little room for this type of workflow, and available optimization strategies remain largely limited to text-based approaches. This observation coincides with the account of the artist's role shifting from the creator of aesthetic objects to the "translator and curator" within generative infrastructures (Goodfellow, 2024). Nodal prompting extends this shift because the author is not only a curator of outputs but an architect of the system that produces them.

At the same time, the case study confirms that nodal prompting is not a purely technical matter, but it extends to the concept of creative agency, as it defines a new form of it. The author is present in the depth of the process, and the agency is distributed

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

throughout the infrastructural design of the process, exercised well before generation occurs. If prompting is not a command but an invitation to joint participation (Gunkel, 2025), the model proposed by this case study requires the author to participate within a generative space in which their role is best understood as architectural rather than merely instrumental.

Seen from this perspective, *Il partigiano Libero* case shows that nodal prompting can be identified as a new form of authorial practice because it redefines where creative intention resides and how it becomes operative. The author is the designer of a generative system whose structure conditions the meaning of the result. Moreover, nodal prompting names a shift from prompt as instruction to workflow as authorship, which is clearly visible in the way the case study organizes creative choices in the graph.

The concept of nodal prompting is, however, conditioned by the specificity of the context described in Section 3, where workflow design was a condition of possibility rather than a preference. This singularity is analytically productive, because it renders visible the architectural logic of authorship that commercially mediated productions might obscure, but it also means that the concept, as formulated here, is grounded in a context where workflow design was not merely a technical preference but a condition of possibility.

Furthermore, the primary author's dual role as practitioner and researcher, while enabling privileged access to the decision logic embedded in the workflow, introduces a reflexivity risk. These limitations define the epistemic perimeter of the findings and indicate the axes along which the concept must be tested to acquire broader explanatory purchase.

5. Conclusion

This article has shown that authorship in node-based environments is relocated from the prompt to the architecture of the workflow itself. The analysis of *Il partigiano Libero* has demonstrated that the most consequential creative decisions (LoRA training as a casting act, seed fixation as a continuity instrument, strength calibration as directorial negotiation) were enacted at the level of workflow architecture, upstream of any individual generation. This constitutes a form of authorship that is infrastructural: the graph is the text, and each generation is its interpretation.

These findings resonate with existing frameworks on distributed authorship, parametric co-creation, and human-AI agency, and suggest a mode of interaction that may be missing from current taxonomies of prompt engineering. Where classical frameworks treat the prompt as the fundamental unit of human-to-model communication, nodal prompting locates authorial practice in the design of the generative system itself, an upstream act of architecture that preconfigures the conditions under which production becomes possible, and whose expressive signature persists across every downstream generation.

The epistemic scope of these findings, grounded in a single qualitative case study, can be generative for further research. Future work should test the concept across commercial and institutional pipelines, extend it to multimodal and collaborative generative configu-

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

rations, and develop methodological instruments adequate to the specificity of the object, such as version-controlled workflow repositories and node-level logging protocols. What the case of *Il partigiano Libero* has made visible, above all, is that the generative system does not produce results independently of the architectural choices that precede it. The author is not dissolved by the machine; they are relocated upstream.

Acknowledgements

We would like to thank Stefano Sburlati, whose research and work at Motion Pixel inspired the use of an open-source generative workflow for the series, and began the study on nodal prompting, ultimately making this article possible.

References

Banh, L., & Strobel, G. (2023). Generative artificial intelligence. *Electronic Markets*, *33*(1), 63. https://doi.org/10.1007/s12525-023-00680-1

Black Forest Labs. (2024). *FLUX.1: A Family of Flow Matching Text-to-Image Models*. Technical release. https://bfl.ai/announcing-black-forest-labs/

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., … Amodei, D. (2020). *Language Models are Few-Shot Learners*. https://arxiv.org/abs/2005.14165

Bruno, G. (2025). *Generative AI in cinema: A case study on the series "Il partigiano Libero"* [Master's Thesis, Politecnico di Torino]. https://webthesis.biblio.polito.it/35421/

Couldry, N. (2012). *Media, society, world: Social theory and digital media practice*. Polity.

Deng, S., Ye, Z., & He, J. (2025). An Empirical Study on Game Prop Design: A Node-Based Generative Workflow in ComfyUI. *2025 IEEE/IEIE International Conference on Consumer Electronics-Asia (ICCE-Asia)*, 1–6. https://doi.org/10.1109/ICCE-Asia67487.2025.11263674

Goodfellow, P. (2024). The distributed authorship of art in the age of AI. *Arts*, *13*(5), 149. https://doi.org/10.3390/arts13050149

Gunkel, D. J. (2025). Prompted by me. Generated by ChatGPT. *Human-Machine Communication*, *10*(1), 2. https://doi.org/10.30658/hmc.10.2

Hong, W., Ding, M., Zheng, W., Liu, X., & Tang, J. (2023). CogVideo: Large-scale Pretraining for Text-to-Video Generation via Transformers. *International Conference on Learning Representations*. https://arxiv.org/abs/2205.15868

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2022). LoRA: Low-Rank Adaptation of Large Language Models. *International Conference on Learning Representations*. https://arxiv.org/abs/2106.09685

Kyaw, A., & Sivalingam, K. (2025). *Node-Based Editing for Multimodal Generation of Text, Audio, Image, and Video*. https://doi.org/10.48550/arXiv.2511.03227

Latour, B. (2005). *Reassembling the social: An introduction to actor-network-theory*. Oxford university press.

DigitCult
ISSN 2531-5994
VOL. 11, NO. 1
(2026)

Lian, L., Li, B., Yala, A., & Darrell, T. (2024). LLM-grounded Diffusion: Enhancing Prompt Understanding of Text-to-Image Diffusion Models with Large Language Models. *Transactions on Machine Learning Research*. https://arxiv.org/abs/2305.13655

Singh, A. (2023). *A Survey of AI Text-to-Image and AI Text-to-Video Generators*. 32–36. https://doi.org/10.1109/airc57904.2023.10303174

Van Dijck, J., Poell, T., & De Waal, M. (2018). *The platform society: Public values in a connective world*. Oxford university press.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2023). *Chain-of-Thought Prompting Elicits Reasoning in Large Language Models.* https://arxiv.org/abs/2201.11903

White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., Elnashar, A., Spencer-Smith, J., & Schmidt, D. C. (2023). *A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT*. https://arxiv.org/abs/2302.11382

Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., Yin, D., Zhang, Y., Wang, W., Cheng, Y., Xu, B., Gu, X., Dong, Y., & Tang, J. (2025). *CogVideoX: Text-to-Video Diffusion Models with An Expert Transformer*. https://arxiv.org/abs/2408.06072

Zhang, S., Wang, H., & Yi, X. (2025). Exploring Collaboration Patterns and Strategies in Human-AI Co-creation through the Lens of Agency: A Scoping Review of the Top-tier HCI Literature. *Proc. ACM Hum.-Comput. Interact.*, *9*(7). https://doi.org/10.1145/3757594