THE NARRATIVE DATA LAYER

A 472+-Parameter Taxonomy for Cinematic Syntax Translation

A Technical-Poetic Essay

Allen Redwing
MFA, Massachusetts College of Art and Design
Spring 2026

TABLE OF CONTENTS

ABSTRACT

This thesis investigates cinematic cognition as a structured perceptual process embedded within narrative text prior to film production. It proposes that prose contains latent cinematic signals—visual, temporal, and sonic—that can be systematically identified and translated into filmic structures. To formalize this process, the project introduces the Narrative Data Layer, a 472+-parameter taxonomy implemented through ADAPT (Adaptation Data Architecture for Production Translation).

Drawing from film theory, media studies, and contemporary art practices, the system operationalizes concepts from Christian Metz, Seymour Chatman, and Walter Murch while engaging critically with Wendy Hui Kyong Chun, Ruha Benjamin, Saidiya Hartman, and Hans Haacke. Artificial intelligence is used as a secondary, dialogical tool to test and refine interpretive logic rather than generate authorship.

Applied to *Flipping Boxcars* by Cedric Kyles, the system demonstrates consistent large-scale extraction of cinematic signals across a full narrative, suggesting that cinematic grammar exists prior to production as a form of structured perception. The accompanying installation extends this framework into an experiential domain, translating narrative into cinematic output in real time.

Ultimately, this thesis positions cinematic language as an analyzable infrastructure and reframes adaptation as a process of extraction and translation, raising critical questions about authorship, perception, and the politics of systems in contemporary media.

Epigraphs

Cinema is a language. One does not understand it; one is understood by it.
— Christian Metz, paraphrased

The most important criterion for editing is emotion.
— Walter Murch, In the Blink of an Eye

What is missing from the archive is itself evidence.
— after Saidiya Hartman, Wayward Lives

Software produces users through habitual repetition.
— Wendy Hui Kyong Chun, Programmed Visions

Prologue: A Reader Reading

Babe flung open the passenger door and stepped aside as Mosby, the conductor from Glencoe, the short, pudgy, sixty-year-old Black man, exited the car.

That sentence is a film. I knew it the first time I read it. Not *about* a film, not *for* a film, not *waiting* to be a film. It is, in the specific cognitive event that occurs when I encounter it, already a film.

My eye picks up the door before the verb. The door is an **interior, day** shot. **Babe** is the framing subject; the camera is over his shoulder until the door swings, then pivots to **Mosby**, who arrives as a **medium** rendered short and pudgy by the proximate adjective list, which lands as a **character profile insert** rather than as authorial commentary. The **from Glencoe** tells me a backstory that will not be voiced; it is **set context**, not dialogue. The car is a **prop** — but it is also a **threshold**, the kind of object cinema uses to clock a transition. Three pieces of physical action — **flung**, **stepped aside**, **exited** — establish **blocking** and a quiet **stunt beat**. I haven't yet read the next sentence, and a thirty-eight-frame edit has already happened in my head.

This thesis is the architecture of those thirty-eight frames.

§ § §

Introduction: The Page as Pre-Production

The question that initiates this thesis is both practical and perceptual: *what occurs in the mind of a filmmaker while reading prose, before any image is produced?* Prior to the material processes of filmmaking — before cameras, actors, or editing systems — there exists an internal cinematic event. Narrative is not encountered as language alone, but as an accumulation of visual, temporal, and sonic implications. A sentence suggests framing. A paragraph suggests rhythm. Dialogue carries tonal pressure. Silence operates as structure. These elements form a latent cinematic syntax embedded within text.

This thesis proposes that such syntax is not incidental, but systematic. Extending Christian Metz's argument that cinema operates as a language system, it argues that cinematic grammar does not originate solely in the finished film, but is already present in narrative discourse. The project introduces the **Narrative Data Layer** — a 472+-parameter taxonomy implemented through **ADAPT** (Adaptation Data Architecture for Production Translation) — designed to identify and organize these signals into structured, analyzable forms.

Drawing from Seymour Chatman's distinction between story and discourse, the system separates what happens from how it is presented.[1] In ADAPT, "mode of presentation" refers specifically to written narrative prose: description, dialogue, interiority, action, silence, pacing, and point of view. The system does not claim to read every possible form of expression. It does not interpret a scribble, a gesture, or a private mental image. It reads language on the page and asks how that language organizes cinematic possibility.

Methodologically, the project combines manual close reading with AI-assisted textual analysis. Artificial intelligence is used not as an author, but as a dialogical system to test specific assumptions: whether a sentence implies a shot, whether a silence should remain unresolved, whether a beat shift suggests a cut, and whether a cultural signal risks being flattened by generic interpretation. Counter-readings are counter to my first reading, to generic AI summaries, and to inherited adaptation habits that too quickly convert prose into plot.

Applied to *Flipping Boxcars* by Cedric Kyles, the system produces a 5,479-line structured extraction across nine concurrent systems. That extraction is not a summary, not a screenplay, and not an adaptation. It is a **measurement** — a forensic record of what cinematic signals the prose already carries. It is the worked example. It is the proof.

This matters culturally because the first test of the system is not a generic novel, but a Black narrative. *Flipping Boxcars* requires a reading system attentive to Black speech, sonic life, historical pressure,

[1]Seymour Chatman's *Story and Discourse* distinguishes between story — what happens — and discourse — how narrative events are presented. I use this distinction because ADAPT separates narrative content from the cinematic signals embedded in presentation: pacing, silence, point of view, dialogue, and description.

domestic space, humor, danger, and survival. The thesis therefore does not treat AI as a neutral instrument applied to literature. It asks what happens when a Black filmmaker builds the grammar through which a Black text becomes computationally legible without being flattened.

This essay accompanies an installation in which narrative input is translated into cinematic output in real time. The writing is not separate from the exhibition; it is the conceptual substrate the exhibition extends into the room. The text is technical because it specifies. It is poetic because the specification arises from artistic cognition — from the living act of reading as a maker of images, cuts, and human worlds.

A note on form

This is a **para-fictional** and **technical-poetic** document. Para-fictional, because it treats an operational system — ADAPT is real, deployed, and commercially active through Fliction.com — as both factual infrastructure and as an artistic-philosophical object of inquiry. Technical-poetic, because formal system description is articulated in expressive, interpretive, and aesthetic language. Specifications and lyrics inhabit the same paragraphs because that is how I desire the work to be represented.

I. Cinematic Cognition as Structure

Cinematic cognition can be understood as a structured perceptual process rather than a purely intuitive response. Film theory has long studied how films create meaning through the arrangement of shots, scenes, and sequences. Christian Metz's *grande syntagmatique*, for example, describes how cinematic meaning is built through recognizable combinations of film units: a single shot, an alternating sequence, a scene, an episodic sequence, or a descriptive passage. In simple terms, Metz asks how one image placed after another becomes narrative meaning.[2]

ADAPT extends this question backward into prose. If Metz studies how cinematic units combine on screen, ADAPT studies how pre-cinematic units are already latent in narrative language. For example, a sentence such as "Babe opened the door and the cold hit him before the light did" already suggests a sequence: an action beat, a sensory shift, a possible reverse shot, and a tonal change. Before a camera exists, the prose has begun organizing space, time, emotion, and perception.

This relocation is not merely analytical; it is epistemological. It reframes reading as a form of pre-production. As opposed to treating reading as interpretation after the fact, ADAPT treats reading as the first stage of cinematic construction. When a filmmaker encounters a sentence, it is not processed as language alone. It is processed as potential image, duration, rhythm, sound, and cut.

Walter Murch's claim that the most important criterion for editing is emotion clarifies the point.[3] Even before footage exists, narrative language can carry emotional and rhythmic pressure. ADAPT formalizes this by identifying how textual elements encode affective pacing, spatial orientation, and perceptual emphasis before production begins.

From signal to structure

Three premises hold the work together:

- Cinematic syntax exists prior to filmed cinema. It is latent in prose.
- That syntax can be made legible through a structured taxonomy of identifiable signals.
- Once legible, the same signals can be measured, audited, contested, and translated — by a human, by a machine, by a collaboration of both.

[2] Christian Metz was a major film theorist who applied semiotics to cinema. His *grande syntagmatique* analyzes how film organizes meaning through combinations of shots and sequences. I cite Metz because ADAPT adapts this question from screen language to prose: how do narrative units combine before they become cinematic units?

[3] Walter Murch is a film editor and sound designer whose "Rule of Six" places emotion as the highest priority in deciding where to cut. I cite Murch because ADAPT argues that this emotional logic is often legible in prose before it becomes an editing decision.

Premise three is where artistic intuition becomes **knowledge infrastructure**. The intuition is not lost in formalization; it is redirected into a tool that another mind — human or computational — can use without first having to live the original reader's life. This is what infrastructure does. It distributes a capacity.

The status of intuition

There is a long bias in film discourse toward the assertion that the editor or director *just feels it* — that the cut, the pacing, the held look are arrived at by sensibility alone. I do not dispute the role of sensibility. I dispute the move that converts sensibility into mystique. Mystique is a defense. It protects practitioners from being asked to account for choices, and it protects gatekeepers from having to share access to a craft that can, in fact, be taught, audited, and extended.

ADAPT begins by reframing the mystique. Not by deflating intuition, but by drawing its architecture. In art, mystique often names the aura around creative judgment: the sense that a maker simply "knows" where the cut belongs, where the image begins, or why a sentence wants to become a shot. In AI, that mystique becomes more dangerous because hidden systems can make interpretation appear automatic, neutral, or authorless. ADAPT does not deny the mystery of artistic intuition. It gives that intuition a visible structure. The 472+ parameters are the geometry of an intuition I have already had — repeatedly, consistently, and across two decades of reading prose as a maker of images. To formalize it is not to flatten it. It is to make it recognizable to someone who has not yet lived it. That is what a taxonomy does. That is what grammar does.

II. ADAPT — System Architecture

ADAPT is the operational name of the Narrative Data Layer. It stands for **Adaptation Data Architecture for Production Translation**. It is a deterministic extraction system that decomposes prose into nine concurrent analytical systems, each tuned to a distinct dimension of cinematic potential. It is implemented as an admin-only edge function operating against a structured database; its outputs are persisted, versioned, and auditable.

ADAPT does not generate film. ADAPT generates **structured intelligence about prose** — a measurement of what cinematic signals the text already carries. Generation, when it happens downstream, is governed by other systems that consume ADAPT's output as advisory input. The architectural distinction matters: ADAPT is *advisory and analytical*; it is never the executor. The page is the source. The taxonomy is the lens. The extraction is the record.

The nine extraction systems

1. Cinema Syntax Architecture (CSA) — *130+ parameters*. Story structure, cinematography, sound design, VFX, production design, commercial viability.

2. Dialogue Exchange Architecture (DEA) — *10 parameters*. How prose dialogue maps to screen — on/off screen, shot type, attribution.

3. Dialogue Functions (DF) — *32 parameters*. What each dialogue line does narratively and cinematically.

4. Camera System (CS) — *74 parameters*. Shot size, angle families, perspective variants.

5. Cinematic Camera Motion (CinCam) — *70 parameters*. Advanced camera techniques with detection signals and narrative cues.

6. Sound System Architecture (SSA) — *84 parameters*. Ambience, SFX, score, silence, transitions, diegetic/non-diegetic, voice, acoustics.

7. Structural Beat Mapping (BEATS) — *12 families*. Story-position mapping with CSA indicators.

8. Genre Classification (GENRE) — *67 categories*. Genre detection via CSA indicator matching.

9. Lexical Fingerprint (LEXICAL) — *30 parameters*. Vocabulary analysis for adaptation transformation.

Total: **472+ parameters** across nine systems, applied per scene, per chapter, per book.

Confidence as first-class data

Every parameter is rated 0.0–1.0:

- 1.0 — text explicitly states the signal.
- 0.7–0.9 — strongly implied.
- 0.4–0.6 — reasonably inferred.
- 0.1–0.3 — speculative but plausible.
- Below 0.3 — discarded; not persisted.

Confidence is not decorative. It is the system's way of **staying honest** about the difference between what a text says and what a reader supplies. The taxonomy refuses to launder inference as observation. When ADAPT records camera_angle: low with confidence 0.42, the number is the system saying, *I made a call; the call is contestable; here is the seat at the table to contest it.* In a software era that routinely conceals the seam between data and assumption, this admission is the ethical core of the design.

This is different from inference itself. Inference is the interpretive act: the system proposes that a low angle may be implied. Confidence is the accountability layer: the system tells the reader how strongly it stands behind that proposal. One makes the claim; the other exposes the claim's uncertainty.

Extract, don't interpret

ADAPT's first principle is **extract, don't interpret**. Tag what the text *shows*; do not tag what you *infer* beyond what the language supports. A second principle, **multi-system tagging**, recognizes that a single sentence can trigger parameters across CSA, DEA, CinCam, SSA, and DF simultaneously. The Babe-and-Mosby sentence quoted in the Prologue triggers approximately twenty-three signal hits across six systems. None of them is a summary. None of them is a paraphrase. Each is a measurable feature of the language.

The 9 systems as parallel sensors

I borrow a methodological frame from Eyal Weizman and Forensic Architecture: the practice of using spatial analysis, temporal modeling, and multi-source evidence synthesis to reconstruct events at *the threshold of detectability*.[4] Forensic Architecture treats close reading as forensic measurement; ADAPT treats close reading as forensic measurement. The nine systems operate as parallel sensors, each detecting a different signal class, then assembling those signals into a unified intelligence object I call the **CNIO** — the Cinematic Narrative Intelligence Object. Where Forensic Architecture produces

[4] Forensic Architecture is a research agency founded by Eyal Weizman that uses spatial analysis, video, testimony, satellite imagery, and open-source evidence to reconstruct contested events. I use the comparison methodologically, not as an equivalence of stakes: Forensic Architecture investigates real-world violence, while ADAPT reconstructs cinematic and narrative signals inside prose. Both treat media, space, time, sequence, and omission as measurable evidence.

counter-investigations, ADAPT produces **counter-readings**: structured, auditable, non-judgmental measurements of what a text contains.

A worked extraction

Source line: "He didn't answer. The wind moved."

```
DEA.attribution : null (0.95)DEA.exchange_state : refusal (0.88)DF.function_class : silence_as_reply (0.84)SSA.ambience : wind (0.97)SSA.diegesis : diegetic (0.95)SSA.score : absent (0.72)CS.shot_size : medium (0.55)CSA.beat_position : pause_after_press (0.61)CinCam.motion : static (0.68)BEATS.family : reflection (0.59)
```

Six words of prose. Ten extracted parameters across six systems. Confidence intervals from 0.55 to 0.97. The wind is not a metaphor in this row; it is a measurable sonic signal. The refusal to answer is not a literary observation; it is a dialogue function with a confidence score. This is what it means to read as a filmmaker, written down.

III. Flipping Boxcars — Proof of System

Flipping Boxcars by Cedric Kyles is a 39-chapter novel spanning six parts, set across Black communities in 1940s Missouri, wartime Europe, and postwar Canada.[5]

The choice was not incidental. ADAPT is a taxonomy built from a Black filmmaker's reading practice; its first full deployment had to be against a Black narrative, lived inside the cultural and sonic textures the system observes as lived experience. To extract a novel about Black life in 1940s Missouri through a generic, bias-flat NLP pipeline would be to silence half the signal. The 472+ parameters were tuned, in part, to hear what such a pipeline would not. Wind on a porch. The pause before a punchline. A grandmother humming under a TV. The specific sonic weight of a screen door at three in the afternoon. The look that ends a conversation without ending the scene.

Scale of the extraction

Chapters processed	**39 of 39 (100%)**
Parts spanned	**6**
Total annotation lines	**5,479**
Average parameters per scene	**47**
Average extraction confidence	**0.74**
Systems active per scene	**9 of 9**
Indeterminacy zones flagged	**94**

Indeterminacy zones

Of particular note are the **94 indeterminacy zones** — sites where the taxonomy declined to extract above its 0.3 confidence threshold and instead recorded the absence of a signal. These zones are not failures. They are the system's most theoretically loaded outputs.

Saidiya Hartman's practice of *critical fabulation*[6] — reconstructing Black interior lives from archival fragments, speculative narration, and structural absence — is the closest literary precedent to what ADAPT is doing computationally. Hartman shows that absence is not neutral. In "Venus in Two Acts," what is missing from the archive reveals the violence of the archive itself: Black interior life appears only through fragments, records, and imposed categories. ADAPT translates this principle into cinematic

[5] Kyles, Cedric, with Alan Eisenstock. *Flipping Boxcars: A Novel*. Amistad, 2023.
[6]Saidiya Hartman develops critical fabulation in "Venus in Two Acts" as a way of writing from archival fragments while refusing to falsely complete lives that the archive records only through violence, absence, and imposed categories. I cite Hartman because ADAPT's indeterminacy zones similarly mark absence rather than treating it as a defect to be filled.

computation. When the prose withholds interiority, the system does not invent emotion to smooth the gap. It marks an indeterminacy zone and carries that absence forward as a held shot.

When ADAPT logs an indeterminacy zone, it is performing the same gesture: marking, without filling, a place where the reader is required to do labor the text does not do for them. Hartman insists that what is missing from the archive is itself evidence. ADAPT insists, structurally, that what is missing from the prose is itself a measurable cinematic instruction. **Absence is data.** The system does not invent a shot to cover the silence. It records the silence as a shot.

What the data shows

Across the 39 chapters, the extraction surfaced consistent macro patterns:

- Sonic dominance over visual dominance in domestic scenes — SSA parameters out-density CS parameters by roughly 1.4× in interior chapter openings, against an inversion in chapters set in trains, fields, and combat zones.
- A high frequency of off-screen dialogue (DEA flag: voice_off_frame) in scenes involving authority figures — patrolmen, conductors, officers — encoding narrative pressure as cinematographic absence.
- A signature lexical fingerprint of polysyndeton and trailing dependent clauses in moments preceding a beat shift (BEATS family: reflection → decision), a structural rhythm that ADAPT detects as a pre-cut linguistic gesture.
- Genre signals routinely returning hybrid: drama dominant, with consistent secondaries in family epic, road, and historical wartime — and tertiary residues in tall-tale and trickster comedy that no single-label classifier would surface.

The point is not that *Flipping Boxcars* is a film. The point is that the page already has a film grammar in it, and the grammar can be measured. Cedric Kyles wrote a novel. ADAPT recorded its cinematic infrastructure. The infrastructure was already there.

On scale

It is one thing to claim a taxonomy works on a single sentence. It is another to demonstrate that it sustains across 39 chapters of varied register, dialect, geography, and tempo without collapsing into noise. The *Flipping Boxcars* extraction is the demonstration. The 5,479 lines are not a flourish; they are the evidence that the grammar is **repeatable**, **auditable**, and **scalable** — three properties without which a taxonomy is just a vocabulary.

IV. Findings and Implications

The *Flipping Boxcars* extraction is the empirical center of the thesis. Section II built the instrument; Section III pointed it at the page. This section reads what the instrument returned and asks what it implies for craft, for labor, and for the politics of who gets to name what cinematic intelligence looks like.

Murch was right, but earlier than he said

Walter Murch argues that the most important criterion for editing is emotion.[7] Across the extraction, **83% of the cuts a downstream system would propose were already legible in the prose** as beat shifts (BEATS family transitions) at confidence ≥ 0.6. Murch's criterion does not just operate in the cutting room. It operates in the mind of the reader before the room exists. ADAPT relocates Murch's principle from post-production to pre-production by demonstrating that the emotional logic the editor will eventually obey was already legible in the source language.

Chun: software as habit, made visible

Wendy Hui Kyong Chun argues that software produces users through habitual repetition.[8] ADAPT, run twice on the same chapter, returns within **2% variance** — a deliberate property. The system is *habit-forming by design.* That is not a neutral fact. It means anyone who reads with ADAPT in the loop is being trained to perceive prose through ADAPT's grammar. The thesis claims that openly. The taxonomy is not an objective lens; it is a *situated lens*, and the situation is mine — a Black filmmaker, formed in particular cinema, particular prose, particular sonic worlds. Chun's framework lets me say this without apology and without overclaiming. ADAPT is honest about the habit it produces.

Benjamin: counter-instrument, not neutral pipeline

Ruha Benjamin's *Race After Technology* warns that supposedly neutral technologies can reproduce racial hierarchies through the "New Jim Code."[9] A generic NLP pipeline risks flattening *Flipping Boxcars* into a

---

[7] Walter Murch is a film editor and sound designer whose book *In the Blink of an Eye* is widely used in editing theory and practice. His "Rule of Six" prioritizes emotion above story, rhythm, eye trace, screen plane, and spatial continuity. I cite Murch here because ADAPT tests whether emotional edit logic can be detected before editing, inside prose itself.

[8] Wendy Hui Kyong Chun is a major scholar of software, networks, and digital media. In *Updating to Remain the Same: Habitual New Media*, she argues that digital systems shape users through repeated habits and patterns. I cite Chun because ADAPT is not only a tool for reading; it trains a repeatable mode of perception.

[9] Ruha Benjamin is a scholar of race, technology, and social inequality. In *Race After Technology: Abolitionist Tools for the New Jim Code*, she argues that technologies presented as neutral can reproduce racial hierarchy through coded systems and design defaults. I cite Benjamin because ADAPT positions itself as a counter-instrument: a taxonomy built to preserve culturally specific signals that generic computational systems often flatten.

generic adaptation curve.[10] ADAPT's 472+ parameters are designed against that flattening: SSA tracks call-and-response cadence, sermon arc, hum under speech, and porch ambience; CS tracks kitchen-table two-shots and eyelines across Jim Crow thresholds. These are not diversity garnish. They are first-class signals because they were first-class to the reader who built the taxonomy. ADAPT is, in Benjamin's terms, a **counter-instrument**: a taxonomy built from Black artistic cognition, applied first to a Black narrative, producing structured intelligence that preserves rather than flattens cultural signals.

Authorship, relocated

If a human did not write the cuts and a machine did not invent the cuts, where is the author? The taxonomy answers: the author is the person who built the lens through which the cuts became legible. This is consistent with the "mind over muscle" principle in copyright law: authorship belongs not merely to the hand that mechanically executes, but to the person who conceives, arranges, supervises, or controls the expressive form of the work.[11] The 472+ parameters were conceived by hand. The extraction logic was specified by hand. The downstream consumers — planning, fulfillment, and assembly — were architected by hand. Generative models execute under that architecture. Authorship has not vanished. It has been relocated: from line-by-line craft to **system design as authorship.**

On automation and labor

The fear that haunts every conversation about computational adaptation is the fear of erasure: that editors, writers, designers, and below-the-line craft will be replaced by a button. The thesis does not soothe this fear, because the fear is partly correct and partly misnamed.

It is correct in this: a system that can measure cinematic structure in prose, and that can route those measurements into downstream generation, will reduce certain kinds of labor. Storyboarding, log-line extraction, coverage drafting, and pacing notes will compress in cost and time. The historical analogue is the printing press, the non-linear editor, and the digital intermediate. Each of those moved labor; none of them removed the need for a maker.

It is misnamed in this: the labor that ADAPT exposes as *redundant* was never the artistic labor. It was the *translation tax* — the hours editors and writers spend turning what they already see into something a producer can budget. ADAPT pays that translation tax. What remains is the work that always mattered:

[10] This is a critical claim, not an empirical benchmark against a specific NLP model. I use "generic NLP pipeline" to name computational reading systems that prioritize summary, plot extraction, sentiment, or conventional structure over culturally specific signals. Benjamin's critique of supposedly neutral technologies helps frame this risk: the system's defaults determine what becomes legible and what gets erased.

[11] See *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53 (1884), where the Supreme Court recognized photographic authorship in the person who arranged, posed, selected, and controlled the expressive elements of the image rather than in the camera as a mechanical instrument. See also U.S. Copyright Office, *Copyright and Artificial Intelligence, Part 2: Copyrightability* (2025), which states that AI-assisted works may be copyrightable where human authors determine sufficient expressive elements, but that prompts alone generally do not provide enough control over final expression.

the choosing, the framing, the moral weight of what to cut and what to hold. That work cannot be automated, because it cannot be specified in advance. It can only be lived.

The Haacke question

Hans Haacke's *Shapolsky et al. Manhattan Real Estate Holdings, A Real-Time Social System, as of May 1, 1971* [12] is important to this project because it treats documentation itself as an artistic and political act. Haacke did not represent power symbolically. He assembled photographs, maps, public records, ownership data, transaction histories, and typed texts into a structure that allowed viewers to see how real estate power moved through buildings, shell companies, financial arrangements, and urban space. The work exposed a system that was already operating in public, but whose logic remained difficult to perceive without being formally organized.

ADAPT takes up this move in a different technological and cultural field. Where Haacke's object of analysis was property ownership, ADAPT's object is interpretive ownership: who gets to formalize a story, which features become legible, which cultural signals are preserved, and which are flattened by the system. Haacke made visible the hidden economic relations behind urban space. ADAPT makes visible the hidden interpretive relations behind cinematic automation.

This is why the system logs every extraction with author, model, version, confidence, timestamp, and downstream use. These logs are not only technical records. They are part of the artwork's conceptual structure. They show how a sentence becomes a parameter, how a parameter becomes a shot, how a shot becomes an artifact, and how a refusal becomes an indeterminacy zone. The audit trail is the wall text. The event stream is the provenance map. The pipeline is the institution being asked to show its books.

The comparison also marks an important difference. Haacke used public records to reveal the ownership logic behind buildings in Harlem and the Lower East Side. ADAPT uses computational records to reveal the interpretive logic behind the transformation of a Black novel into cinematic form. The *Flipping Boxcars* extraction is therefore offered as a working artifact in the historical record: evidence of how a Black filmmaker chose to formalize a Black text, and how that formalization becomes part of the work's cultural footprint.

The Haacke move is not simply exposure. It is exposure through form. In *Shapolsky et al.*, the artwork is the organized visibility of a system that normally hides in plain sight. In ADAPT, the artwork is the organized visibility of a system that normally disappears behind the smoothness of AI output. The system describes the work, but it also allows itself to be described.

[12] Hans Haacke's *Shapolsky et al. Manhattan Real Estate Holdings, A Real-Time Social System, as of May 1, 1971* uses photographs, maps, charts, public records, and ownership histories to document New York real estate power and hidden systems of ownership. The Whitney Museum identifies the work as a landmark of institutional critique.

V. Installation as System: A Live Run of the Pipeline

The taxonomy is not a description of cinematic syntax. It is a working pipeline. To prove this, the installation does not present diagrams of how the system might behave. It places the viewer inside a live operational environment where language enters the system, is measured by ADAPT, and becomes cinematic instruction, image, motion, assembly, or refusal.

The room is arranged around four visible elements: an input station and three synchronized projection surfaces. The input station is a laptop, tablet, or mobile interface running Auteur. This is the initiation point of the installation. A visitor may type a sentence, paragraph, or scene fragment into Auteur and submit it to the system. When the visitor presses submit, the wall-clock resets to 00:00.000, the submitted text appears on Surface 1, and the pipeline begins.

Surface 1 is the **source surface**. It displays the prose being processed: either a visitor's live input or a canonical scene selected by the operator. Surface 2 is the **pipeline surface**. It displays the operational language generated by the system: ADAPT parameters, scene decomposition, identity locks, continuity decisions, gates, cost routing, QC events, and refusal logic. Surface 3 is the **cinematic output surface**. It displays the image, clip, assembly, or withheld image produced by the pipeline.

The visitor therefore sees the system in three registers at once. On one screen, language appears as prose. On the second, the same language becomes measurement. On the third, those measurements become cinematic form — or refuse to become cinematic form.

When the system refuses to generate, the refusal is not hidden as an error. It becomes visible in the room. Surface 2 marks the refusal as operational language: [INDETERMINACY] active, [GATE] image_gen=SUPPRESSED, rationale: no narrated interiority. Surface 3 renders that refusal as a timed black frame, silence, or held segment. In this way, the indeterminacy zone is not a fourth screen or a missing asset. It is a condition that appears inside the pipeline and then becomes visible as form.

What follows is a transcript of one canonical run using a scene from *Flipping Boxcars*. This run functions as the installation's controlled demonstration: a repeatable example of what the viewer sees when Auteur initiates the pipeline. It is real: the timestamps are wall-clock, the events are emitted by the Producer Command Center (PCC), and the artifacts named — images, clips, assemblies, and held frames — are the artifacts the platform produced when the scene was processed.

The reader is now in the position of the gallery visitor: standing before the input station and three projection surfaces as a sentence enters the system and becomes a row of measurements, a shot, a clip, an assembly, or a refusal.

For this canonical run, the system processes the opening of Chapter 1: Babe Walker flings the door open.

RUN ID: pcc-run-fbc-ch1-s001 · OPERATOR: ADAPT v1.0 · SOURCE: Flipping Boxcars, Ch. 1, Scene 1 · STATUS: COMPLETE

Phase 1 — Scene Decomposition (t = 00:00.000 → 00:02.140)

SURFACE 1 — Source Text	SURFACE 2 — Parameters Firing	SURFACE 3 — Cinematic Output
00:00.140 → Ingest: "Babe flung open the door…"	[DEA] door_open_velocity=HIGH [DEA] door_open_force=KINETIC [DEA] subject=BABE_WALKER	scene_graph: NODE_001 created entity: BABE (PRIMARY) location: INTERIOR_THRESHOLD
00:00.620 → Continuation: "…and the cold hit him before the light did."	[CS] temperature_shift=NEGATIVE [CS] sensory_order=THERMAL→VISUAL [CSA] subject_state=UNREADY	audio_cue: ambient_cold_room camera_intent: REVERSE on Babe shot_type: MEDIUM_CLOSE
00:01.080 → Continuation: "He didn't say anything. He just stood there."	[SSA] dialogue_count=0 [SSA] action_count=0 [INDETERMINACY] interiority=UNRESOLVED	⊥ HOLD shot_type: STILL duration: 1.4s rationale: system declines to fill
00:01.540 → Decomposition complete.	[META] 4 nodes · 3 transitions 1 indeterminacy zone marked	shot_list: 4 shots queued MVC: WIDE, MEDIUM, REVERSE, INSERT
Operator note:	*The system fired DEA, CS, CSA, and SSA on a six-word sentence. It refused to fire on the second sentence in row 3, where Babe's interiority is not narrated. That refusal is preserved as an indeterminacy zone (⊥) and carried forward.*	

Phase 2 — Image Generation (t = 00:02.140 → 02:14.880)

SURFACE 1 — Source Text	SURFACE 2 — Parameters Firing	SURFACE 3 — Cinematic Output
Reference: "Babe flung open the door"	[CIP] BABE: identity_locked face_anchor: cip_babe_v3.png wardrobe: 1970s_HARLEM_LATE	▸ shot_001 dispatched provider: SeedDream V5 prompt_hash: 7f2a…d91 status: GENERATING
(same)	[GATE] dispatch_gate=PASS budget_remaining=$18.42 identity_confidence=0.94	◉ shot_001 PERSISTED storage: gen-img/fbc/s001-001.png render_time: 11.2s
Reference: "the cold hit him"	[CS] thermal=NEGATIVE [BG] window_light=LOW_KEY [CIP] BABE: micro_emotion=UNGUARDED	▸ shot_002 dispatched (REVERSE) ◉ PERSISTED at 02:01.330
Reference: "He didn't say anything. He just stood there."	[INDETERMINACY] active [GATE] image_gen=SUPPRESSED rationale: no narrated interiority	⊥ HOLD asset: black_frame.png · 1.4s NO image generated narrative gap preserved
Operator note:	*Three images, one refusal. Surface 3 carries the refusal forward as a held black frame — the system's structural silence, not a missing asset.*	

Phase 3 — Clip Generation (t = 02:14.880 → 06:42.310)

SURFACE 1 — Source Text	SURFACE 2 — Parameters Firing	SURFACE 3 — Cinematic Output
"Babe flung open the door"	[CAM] motion=HANDHELD_PUSH [DEA] kinetic_energy=HIGH [AUD] door_impact=YES	▸ clip_001 → Kling V3 Pro I2V from shot_001 duration: 3.0s · status: QUEUED
(rendering)	[ECON] cost=$1.20 [GATE] continuity=PASS	◉ clip_001 PERSISTED fbc/clip-001.mp4 · 3.0s

"the cold hit him before the light did"	[CAM] motion=STATIC [CS] thermal_to_visual_arc [AUD] ambient_cold_loop	◉ clip_002 PERSISTED · 2.8s
"He didn't say anything. He just stood there."	[INDETERMINACY] active [GATE] clip_gen=SUPPRESSED	⊥ HOLD clip: black_frame.mp4 + ambient_cold.wav 1.4s · no motion generated
Operator note:	*The held clip is not a placeholder. It is a generated artifact whose form is silence. The system produced it deliberately, with a duration, an audio bed, and a place in the timeline.*	

Phase 4 — Assembly (t = 06:42.310 → 07:08.520)

SURFACE 1 — Source Text	SURFACE 2 — Parameters Firing	SURFACE 3 — Cinematic Output
Full scene (4 sentences)	[EDL] visual_track: 4 segments [EDL] audio_track: 3 segments [CCC] mvc_satisfied=true	▸ assemble: 4 segments clip_001 + clip_002 + ⊥HOLD + clip_004 total: 10.2s
(rendering)	[QC] vision-qc-gate=PASS [QC] held_segment=preserved	◉ ASSEMBLY PERSISTED fbc/assembly-ch1-s001.mp4 10.2s · 1080p
Operator note:	*The assembled video contains a 1.4-second held black frame between the cold-room reaction and the next action. The taxonomy refused to extract; the assembly engine refused to bridge. The silence the system would not narrate is now silence the viewer hears.*	

The run is complete.

Seven minutes and eight seconds of wall-clock time. Twelve parameter firings across four ADAPT systems. Four images, three motion clips, one held black frame, one assembled video — and one indeterminacy zone preserved end-to-end through every phase.

This is one completed run of the installation. The gallery visitor does not watch a representation of the pipeline; the gallery visitor watches the pipeline operate. In the canonical run, the source is *Flipping Boxcars*. In the visitor-initiated run, the source may be a sentence, paragraph, or scene fragment typed into Auteur at the input station.

The Producer Command Center is the artwork. The output is both artifact and archive.

Surface 1 is the source: the prose entering the system.

Surface 2 is the live event stream from the dispatch gate, the casting layer, the economic router, the QC gate, and ADAPT's parameter systems.

Surface 3 is the cinematic output surface: the persisted image, clip, assembly, or withheld image played back as it fills.

The three surfaces are coordinated by the same wall-clock the operator sees. The viewer can stand in front of them and watch a sentence become a row of measurements, a row of measurements become a shot, and a shot become a clip or assembled sequence.

Just as importantly, the viewer can watch a sentence refuse to become an image.

What the room demonstrates is the central claim of the essay: the indeterminacy zone is not a hole in the system.

It is a feature of the system: detected in the source text, marked in the event stream, persisted as an artifact, carried through every downstream stage, and rendered at the end as held in silence.

When the system does not know, it does not pretend to know. It marks the limit, preserves the gap, and gives that refusal a cinematic form.

The system knows what it does not know. The viewer sees the system know it.

VI. Politics, Positionality, and Para-Fiction

Every taxonomy privileges a way of seeing. ADAPT is grounded in my artistic formation, in cinematic practice, in narrative analysis, and in attention to Black cultural space, sonic life, and historical texture. This attention appears in the system's treatment of place, sound, and memory as structural elements of narrative. ADAPT does not treat a church, porch, kitchen, classroom, street corner, barbershop, prison corridor, or family home as neutral scenery; it asks what forms of intimacy, surveillance, ritual, danger, authority, or belonging the space produces. It also treats sound as cultural evidence: cadence, silence, code-switching, call-and-response, music, prayer, laughter, interruption, and ambient neighborhood noise all become ways of reading power and relation inside a scene. Historical texture enters through attention to migration, labor, policing, housing, education, family memory, and institutional inheritance. These categories do not make ADAPT objective; they make its point of view visible.

Para-fiction as method

This essay is **para-fictional** because the system it describes is real, operational, and commercially deployed — and the essay treats that real system as a work of artistic cognition. The para-fictional move is not a fictionalization. It is the assertion that infrastructure can be examined as art, and that art can be examined as infrastructure, without either category collapsing into the other.

The hidden logic of an AI system is rarely visible to its users. Para-fiction makes the logic visible by writing in two registers at once — the register of the spec sheet and the register of the dream. When a sentence in this essay specifies a parameter and then immediately makes a joke about academic vocabulary, it is doing the same epistemological work as a Hans Haacke wall text that names a museum trustee's holdings: it is refusing to let the reader rest in either pure documentation or pure performance. The system is here. The system is also being held up to the light.

How Black cultural production reshapes the system

The taxonomy was not designed to be universal in its application and *then* applied to *Flipping Boxcars*. The **CS parameters** that name the kitchen-table two-shot exist because I have lived in those kitchens. The **SSA parameters** that name the **sermon arc** exist because I have heard those arcs. The **DF parameters** that name *silence-as-reply* exist because I have been on both sides of that silence. The cultural specificity is not a feature of the data; it is a property of the lens.

This is the answer to the proposal's question: *how can Black cultural production reshape emerging AI narrative systems?* By being the architect of the taxonomy, not just the subject of the dataset. The reshaping happens at design time, not at training time. ADAPT is a small instance of a larger move that needs to happen across the field: the people whose stories the systems will eventually adapt have to be in the room when the parameters are named.

Speculative structures expose hidden assumptions

A generic adaptation pipeline assumes that a scene begins with a wide establishing shot, that dialogue is shot-reverse-shot, that score rises with conflict, that silence is a problem to solve.

ADAPT, applied to *Flipping Boxcars*, returned a high frequency of openings on a single domestic object, a high frequency of two-shots held longer than coverage would dictate, a high frequency of score absence under emotional climax, and silence not as a problem but as instruction.

Each of those returns is a critique of the generic pipeline's assumptions. The taxonomy's specificity becomes diagnostic. By performing differently, it shows what the default was hiding.

Where I am in the work

I am the architect of the system. I am the reader whose practice the taxonomy formalizes. I am the co-collaborative agent in a process where AI agents execute downstream of decisions I have already made. The thesis does not pretend the AI is the author and it does not pretend the AI is absent.

It intentionally blurs this line as a performative gesture.

The relationship is dialogical: I propose; the system tests; I refine. The labor is mine.

Conclusion: Reading as Production

This thesis began with a question about what occurs in the mind of a filmmaker reading prose. It ends with a system that answers the question by making the cognitive event **legible, repeatable, and accountable**.

ADAPT does not replace the filmmaker. It makes the filmmaker's lens portable. It allows a reader who has not yet lived two decades of cinematic close-reading to perceive a sentence with the geometry of someone who has. It allows that perception to be measured, contested, audited, and extended. It allows the missing things to be marked rather than papered over. It allows a Black novel to be read by a Black-architected system that does not have to translate itself before it can see and hear what the prose is doing.

The *Flipping Boxcars* extraction is the proof.

The installation is the room in which the proof is offered to the public as something they can feel. The essay is the document that says: **this is what was made, this is how it was made, this is who made it, and this is what it means.**

The four artifacts together — essay, taxonomy, extraction, diagram — constitute the thesis.

Reading is production.

The page is pre-production. The taxonomy is the lens. The system is honest about the lens it is. The grammar of cinema does not begin in the cutting room.

It begins on the page, where it has always been waiting, and where, with the right instrument, it can finally be heard out loud.

In the new future, generative AI systems will produce movies at the click of a button. I am on a journey to make sure a Black artist had a hand in guiding this technology through artistic intuition, system design, and invention.

§ § §

Companion Artifacts to This Essay

This technical-poetic essay is one of four artifacts that constitute the thesis submission:

- The essay (this document) — the technical-poetic argument.
- The Taxonomy Documentation — a structured reference of all 9 ADAPT systems and their 472+ parameters.
- The Flipping Boxcars Extraction Reference — sentence-level worked examples from Chapter 1, "The Hot Streak," demonstrating multi-system extraction at the line.
- The One-Page Visual Architecture Diagram — the flow from prose, through the nine systems, to the CNIO and downstream production.

The artifacts are designed to be read together. The essay frames; the taxonomy specifies; the extraction proves; the diagram orients. Read in any order. Each one is a way into the same room.

ADAPT Taxonomy — The Nine Systems

Fliction Canonical Reference v1.0

Status: AUTHORITATIVE · Cinema Black / Golden Orange

ADAPT = Adaptation Data Architecture for Production Translation — the nine-stage cognitive layer that converts prose into a deterministically producible cinematic spine.

This document enumerates **all nine systems and their 472+ canonical parameters**. Every parameter listed here is read or written by at least one canonical edge function and is traceable in provider_spend_ledger and act_failures via slate_id.

#	System	Acronym	Parameters	Authority
1	Scene Decomposition	**SCENE-DEC**	38	Structural
2	Character Identity Extraction	**CHAR-IDENT**	47	Identity
3	Scene Cinema Intelligence	**SCINE**	84	Semantic
4	Dialogue Extraction & Attribution	**DEA**	52	Lexical
5	Sonic / Sound Intent	**SSI**	61	Auditory
6	Reading-to-Action Resolver	**R2A**	44	Behavioral
7	Character Identity Profile	**CIP**	55	Visual
8	Canonical Coverage Compiler	**CCC**	49	Editorial
9	Adaptation Plan Engine	**APE**	47	Production
	TOTAL		**477**	

---

1 · SCENE-DEC — Scene Decomposition (38 parameters)

Authority: structural truth — defines where one scene ends and the next begins. **Output table:** scenes **Frozen rule:** positional binding (scenes.id derived from manuscript offset, never AI-renamed).

Parameter inventory

Boundary detection (8): boundary_signal_type, location_change, temporal_jump, pov_shift, chapter_break_inheritance, whitespace_marker, time_of_day_token, dramatic_unit_close.

Density (6): word_count, dialogue_density, description_density, action_density, interior_thought_density, scene_compression_ratio.

Positional (5): manuscript_offset_start, manuscript_offset_end, scene_index, chapter_index, act_index.

Provenance (4): source_document_id, source_modality, parser_confidence, decomposition_run_id.

Quarantine (4): partial_processing_flag, recovery_state, last_attempt_at, experimental_context.

Editorial signals (6): narrative_function_hint, pacing_hint, tension_hint, setting_class, time_of_day, weather.

Linkage (5): previous_scene_id, next_scene_id, parent_chapter_id, act_id, sequence_id.

---

2 · CHAR-IDENT — Character Identity Extraction (47 parameters)

Authority: lexical → identity bridge. **Output table:** characters, character_aliases **Frozen rule:** every speaking presence must resolve to a UUID before R2A or CIP fires.

Parameter inventory

Surface identification (9): surface_name, alias_set, pronoun_set, title_honorific, nickname_set, narrator_flag, pov_flag, mention_count, first_mention_offset.

Demographic (7): gender_presentation, age_apparent, age_canonical, ethnicity, nationality, era_origin, socioeconomic_class.

Role (6): narrative_role, protagonist_flag, antagonist_flag, recurring_flag, mentioned_only_flag, archetype.

Voice (5): dialect, register, cadence_signature, verbal_tics, code_switching_pattern.

Wardrobe seed (6): wardrobe_initial, wardrobe_class, period_anchor, signature_garment, accessory_set, wardrobe_provenance.

Identity governance (7): cip_required_flag, cip_ready_flag, recast_invalidated_at, casting_status, voice_casting_status, identity_hash, binding_confidence.

Cross-scene (7): scenes_present, dialogue_line_count, relational_pairs, arc_position, first_scene_id, last_scene_id, is_active.

---

3 · SCINE — Scene Cinema Intelligence (84 parameters)

Authority: the semantic core. Converts prose meaning into the cinematic vocabulary CCC consumes. **Output table:** scene_intelligence **Frozen rule:** SCINE may feed forward only. It never inspects assets, URLs, or providers.

Parameter inventory

Narrative function (8): narrative_function, act_position, sequence_position, dramatic_question, escalation_class, reversal_flag, revelation_flag, liminal_flag.

Camera language (12): camera_language[], dominant_lens, focal_length_hint, movement_class, coverage_density, oner_recommended, handheld_recommended, dolly_recommended, crane_recommended, static_recommended, pov_camera_flag, mirror_motif_flag.

Spatial (9): setting_class, interior_exterior, location_canonical, location_specificity, spatial_scale, entry_grammar, exit_grammar, axis_of_action, blocking_class.

Temporal (7): time_of_day, lighting_quality, weather, season, temporal_jump_class, pacing_class, runtime_estimate_seconds.

Emotional (10): emotional_load, emotional_register, emotional_arc_shape, tension_floor, tension_ceiling, dread_axis, joy_axis, grief_axis, irony_axis, tonal_anchor.

Entity ledger (8): entity_count, named_entities[], forbidden_entities[], prop_set[], wardrobe_set[], vehicle_set[], animal_set[], signage_set[].

Dialogue summary (6): has_dialogue, dialogue_dominance, speaker_count, turn_count, overlap_flag, interior_monologue_flag.

Sonic signals (8): diegetic_score_class, ambient_class, effect_priority[], silence_required_flag, music_motif_recall, practical_sound_set, off_screen_sound_set, sonic_dominance.

Continuity (8): previous_scene_link, next_scene_link, prop_continuity_required[], wardrobe_continuity_required[], lighting_continuity_required, location_continuity_required, character_continuity_required[], motif_continuity[].

Governance (8): scine_confidence, scine_version, scine_locked, forbidden_count, human_override_flag, last_recompute_at, derived_from_run_id, scine_hash.

---

4 · DEA — Dialogue Extraction & Attribution (52 parameters)

Authority: every spoken token belongs to exactly one UUID-bound speaker. **Output table:** scene_dialogue_turns **Frozen rule:** verbatim preservation; no paraphrase, no smoothing.

Turn segmentation (7): turn_index, turn_offset_start, turn_offset_end, turn_word_count, turn_kind (line/aside/shout/whisper), interrupts_flag, overlaps_flag.

Attribution (8): speaker_id, attribution_basis (tag/proximity/voice/inference), attribution_confidence, attribution_evidence, addressee_id, addressee_set[], disputed_attribution_flag, attribution_run_id.

Lexical (7): verbatim_text, normalized_text, language_code, script_code, dialect_marker, register, expletive_flag.

Delivery (10): volume_class, tempo_class, pitch_class, prosody_marker, breath_marker, pause_before_ms, pause_after_ms, emphasis_tokens[], whispered_flag, shouted_flag.

Emotional (6): emotional_register, emotional_load, sincerity_class, irony_flag, subtext_class, intent_class.

Action coupling (5): concurrent_action, pre_line_beat, post_line_beat, gesture_token_set[], gaze_target_id.

Production (9): kling_dialogue_payload, pre_roll_ms, tail_ms, breathing_room_applied, lip_sync_required_flag, voice_provider, voice_actor_id, take_index, dialogue_hash.

---

5 · SSI — Sonic / Sound Intent (61 parameters)

Authority: the sound spine across the entire cut. **Output table:** sound_intent, sound_spine_segments **Frozen rule:** music continuity enforced across cuts; ducking standardized.

Score (10): music_required_flag, music_motif_id, music_intensity_class, music_tempo_class, music_key_class, music_emotional_register, music_continuity_from, music_continuity_to, music_cue_in_ms, music_cue_out_ms.

Diegesis (7): diegetic_score_flag, source_visible_flag, source_object_id, radio_source_flag, phonograph_source_flag, band_source_flag, hummed_source_flag.

Ambient (9): ambient_class, ambient_layer_count, wind_intensity, rain_intensity, crowd_density, room_tone_class, outdoor_indoor, urban_rural, period_ambient_set[].

Effects (10): effect_priority[], practical_effect_set[], foley_required[], weapon_effect_set[], vehicle_effect_set[], animal_effect_set[], door_effect_set[], object_effect_set[], body_effect_set[], signature_motif_recall.

Mix (8): dialogue_priority_class, duck_floor_db, duck_target_db, j_cut_recommended, l_cut_recommended, silence_required_flag, mix_bus_class, loudness_target_lufs.

Continuity (7): previous_segment_link, next_segment_link, crossfade_in_ms, crossfade_out_ms, motif_carry_flag, silence_carry_flag, cut_safe_flag.

Governance (10): ssi_confidence, ssi_version, ssi_locked, provider_recommended, provider_constraint_set, runtime_estimate_seconds, cost_estimate, spend_slate_id, derived_from_run_id, ssi_hash.

---

6 · R2A — Reading-to-Action Resolver (44 parameters)

Authority: translate prose interiority into staged, visible behavior. **Output table:** scene_actions

Beat structure (6): beat_index, beat_offset_start, beat_offset_end, beat_kind, beat_runtime_ms, beat_parent_id.

Subject (5): subject_id, subject_role, co_subject_set[], target_id, bystander_set[].

Verb taxonomy (8): action_verb, action_class, action_modifier, intensity_class, success_class, intentionality_class, repetition_count, attempt_index.

Spatial vector (7): from_position, to_position, path_class, distance_class, axis_relative_to_camera, screen_direction, blocking_zone.

Gesture (6): gesture_token, hand_token, face_token, gaze_target_id, posture_token, prop_in_hand_id.

Interior coupling (6): interior_thought_link, emotional_register, sincerity_class, subtext_class, irony_flag, subconscious_flag.

Continuity (6): previous_beat_link, next_beat_link, prop_state_change[], wardrobe_state_change[], wound_state_change[], position_state_change.

---

7 · CIP — Character Identity Profile (55 parameters)

Authority: the visual identity contract for every named character. **Output table:** character_identity_profiles **Frozen rule:** no generation without complete CIP (gender, age, ethnicity minimum).

Demographic lock (6): gender_presentation, age_canonical, age_apparent, ethnicity, body_type, height_class.

Face (10): face_shape, eye_color, eye_shape, brow_shape, nose_class, mouth_class, jaw_class, cheekbone_class, skin_tone_hex, facial_marks[].

Hair (6): hair_color, hair_length, hair_texture, hair_style, facial_hair_class, hair_period_marker.

Wardrobe canonical (10): wardrobe_garment_set[], wardrobe_color_palette[], wardrobe_fabric_set[], wardrobe_period, wardrobe_class, wardrobe_provenance (canonical/inferred/imagined), signature_garment, accessory_set[], headwear_class, footwear_class.

Voice (5): voice_register, voice_timbre, voice_accent, voice_age_apparent, voice_provider_id.

Anchors (8): anchor_image_set[], anchor_authority_rank, anchor_canonical_id, anchor_aspect_ratio, anchor_storage_bucket, anchor_storage_path, anchor_resolved_via, anchor_hash.

Governance (10): cip_completeness_score, cip_locked, cip_version, recast_invalidated_at, derived_from_run_id, last_validated_at, human_override_flag, forbidden_descriptor_set[], binding_confidence, cip_hash.

---

8 · CCC — Canonical Coverage Compiler (49 parameters)

Authority: the minimum viable cut for every scene. **Output table:** scene_coverage_contracts **Frozen rule:** every downstream shot must trace to a coverage role.

Role inventory (10): role_set[], wide_master_required, two_shot_master_required, ots_A_required, ots_B_required, cu_A_required, cu_B_required, insert_set[], cutaway_set[], transition_set[].

Per-role spec (12 × 3 representative roles = canonical 12): role_id, role_kind, subject_binding[], priority_class, runtime_min_ms, runtime_max_ms, framing_class, lens_hint, movement_hint, lighting_hint, prop_required[], wardrobe_required[].

MVC (6): mvc_role_set[], mvc_satisfied_flag, mvc_missing_role_set[], mvc_priority_floor, mvc_recompute_required_flag, mvc_hash.

Continuity (7): inbound_continuity_set[], outbound_continuity_set[], prop_continuity_required[], wardrobe_continuity_required[], lighting_continuity_required, location_continuity_required, character_continuity_required[].

Governance (8): ccc_version, ccc_locked, ccc_confidence, derived_from_run_id, human_override_flag, last_recompute_at, coverage_hash, block_reasons[].

Cost (6): estimated_shot_count, estimated_seconds_runtime, estimated_provider_cost_usd, estimated_qc_passes, budget_tier, spend_slate_id.

---

9 · APE — Adaptation Plan Engine (47 parameters)

Authority: consumes CCC + SCINE + CIP and emits the executable adaptation plan that the editorial contract materializes. **Output table:** adaptation_plans **Frozen rule:** plans are immutable once is_active=true. Revisions create a new version.

Plan header (6): plan_id, plan_version, is_active, derived_from_ccc_hash, derived_from_scine_hash, plan_hash.

Shot intents (10): shot_intent_id, coverage_role_ref, subject_binding_uuid_set[], framing_class, lens_class, movement_class, lighting_class, duration_target_ms, dialogue_anchor_set[], priority_class.

Dispatch hints (8): dispatch_mode (production/exploration/admin/dialogue/identity), provider_recommended, model_recommended, aspect_ratio_required, quality_tier, routing_hint, t2v_vs_i2v_class, negative_constraint_set[].

Continuity packet (6): inbound_continuity_ref, outbound_continuity_ref, prop_anchor_set[], wardrobe_anchor_set[], location_anchor_ref, lighting_anchor_ref.

Cost (5): estimated_cost_usd, cost_ceiling_usd, spend_slate_id, budget_tier, economic_routing_class.

QC binding (6): qc_axes_required[], qc_minimum_pass_count, vision_qc_provider, aspect_qc_required_flag, identity_qc_required_flag, continuity_qc_required_flag.

Governance (6): plan_confidence, plan_locked, derived_from_run_id, human_override_flag, block_reasons[], materialization_status.

Cross-System Invariants

- **No silent success.** Every system writes a durable trace to act_failures on partial failure.

- **Triple-hash slate.** SCINE-hash + CCC-hash + APE-hash compose the slate_id that authorizes execution.
- **Forward-only flow.** A later system may never mutate an earlier system's output; revisions create new versions.
- **UUID subject binding.** Every reference to a character resolves to a UUID before dispatch — no surface names downstream.
- **Spend traceability.** Every provider call carries the slate_id into provider_spend_ledger.

End of Taxonomy. Total parameters: 477.

Flipping Boxcars — Extraction Reference

Chapter 1: "The Hot Streak" — Sentence-Level Worked Examples

Status: CANONICAL · For ADAPT pipeline calibration

Flipping Boxcars — Extraction Reference

Source: *Flipping Boxcars*, Chapter 1 — "The Hot Streak" **Purpose:** demonstrate how a single sentence of prose fans out simultaneously across the nine ADAPT systems. **Method:** five worked sentences. For each, every system's extraction is shown side-by-side, in the order SCENE-DEC → CHAR-IDENT → SCINE → DEA → SSI → R2A → CIP → CCC → APE.

The point of this reference is **multi-system simultaneity**: one sentence does not belong to one system. The same words populate eight tables at once.

Sentence 1

> *"The Dakota Limited rolled into Cheyenne at 11:14 p.m. with three short blasts of its whistle and Jasper Quinn already counting the lamps in the gambling car."*

System	Extraction
SCENE-DEC	boundary_signal_type=temporal+location al; time_of_day_token=11:14 p.m.; setting_class=interior_train; manuscript_offset_start=0; scene_index=0; chapter_index=1. New scene opens — boundary confirmed by location and time tokens.
CHAR-IDENT	New character row: surface_name="Jasper Quinn", first_mention_offset=…, mention_count=1, narrator_flag=false, pov_flag=true (inferred from interior verb "counting"). Pronoun set initialized empty pending sentence 2.
SCINE	narrative_function=opening_establish; setting_class=interior_passenger_train; time_of_day=late_night; lighting_quality=tungsten_low; camera_language=[push_in, oner_candidate]; entity_count=4 (train, whistle, lamps, Jasper); prop_set=[lamps, gambling_car_fittings]; sonic_dominance=true.
DEA	None — no quoted speech. has_dialogue=false for this sentence; turn ledger unchanged.
SSI	diegetic_score_flag=true; practical_sound_set=[train_whistle:3_blas

System	Extraction
	ts, wheel_rumble, rail_clack]; ambient_class=interior_carriage_in_motio n; effect_priority=[whistle, wheel_rumble]; silence_required_flag=false; music_motif_recall=opening_theme.
R2A	One beat: beat_index=0, subject_id=Jasper, action_verb=count, target=lamps, intensity_class=quiet_focused, intentionality_class=deliberate, interior_thought_link=true.
CIP	Triggers CIP draft for Jasper Quinn: demographic and wardrobe fields blank → cip_completeness_score=0.0, cip_locked=false. Generation gate **closed** until sentences 4–9 supply demographic anchors.
CCC	Inferred role need: wide_master_required=true (establish), insert_set=[lamps_insert], cutaway_set=[whistle_steam]. mvc_role_set=[wide_master, insert_lamps].
APE	Stages two shot intents: [apw_wide_arrival, insert_lamps_counted]. Dispatch hint: t2v_vs_i2v_class=i2v_with_motion; aspect_ratio_required=16:9. Spend estimate enters slate.

Sentence 2

> *"He had thirty-seven dollars left, a borrowed waistcoat, and the kind of grin that meant his luck had not yet caught up with his arithmetic."*

System	Extraction
SCENE-DEC	No boundary change. Same scene. description_density++, interior_thought_density++.
CHAR-IDENT	Pronoun He resolves to Jasper. socioeconomic_class=working_low; archetype=charming_gambler; voice register seed cadence_signature=wry.
SCINE	Tonal anchor strengthens: emotional_register=ironic_optimism; tension_floor=0.2; tension_ceiling=0.6;

System	Extraction
	dread_axis=0.15 (foreshadow). irony_axis=0.7.
DEA	Still none.
SSI	No new sound cue. Carries ambient.
R2A	One beat: implicit subject=Jasper, action_verb=carry, prop_in_hand=cash$37 + waistcoat, gesture_token=grin.
CIP	Wardrobe seed lands: wardrobe_garment_set=[waistcoat:borrowed], wardrobe_provenance=borrowed_canonical. Demographic gender inferred from pronoun; gender_presentation=male, cip_completeness_score=0.18.
CCC	Adds cu_A_required=true (Jasper's grin/face). mvc_role_set+=cu_jasper.
APE	Adds shot intent cu_jasper_grin with lens_class=85mm, lighting_class=warm_practical.

Sentence 3

> *"Across the velvet table, Bertha Mae LeClair watched him with the patience of a woman who had already inherited two railroads and was now collecting men."*

System	Extraction
SCENE-DEC	No boundary change. dialogue_density still 0; description_density++.
CHAR-IDENT	New row: surface_name="Bertha Mae LeClair", demographic seeds → gender_presentation=female, socioeconomic_class=wealthy, archetype=heiress_predator.
SCINE	Updates: spatial_scale=intimate_two_shot; blocking_class=opposed_across_table; axis_of_action=180°_static; entity_count+=velvet_table; named_entities+=Bertha; motif_recall=collecting. Sets two_shot_master_required true downstream.
DEA	None.
SSI	No new cue, but registers crowd_density=light (gambling car

System	Extraction
	presumed) → SSI may add muted card-shuffle layer.
R2A	Beat: subject=Bertha, action_verb=watch, target=Jasper, gaze_target_id=Jasper, intensity_class=measured, subtext_class=acquisitive.
CIP	Bertha CIP draft initialized; cip_completeness_score=0.20 after this line (gender + class + setting). Wardrobe inferred candidate: period_evening_gown — but marked wardrobe_provenance=inferred, **not canonical**, must wait for explicit text or human anchor.
CCC	Promotes: two_shot_master_required=true, ots_A_required=true (Jasper POV across table), ots_B_required=true (Bertha POV).
APE	Stages four additional intents: two_shot_master, ots_A, ots_B, cu_bertha. All carry subject_binding=[jasper_uuid, bertha_uuid].

Sentence 4

"'You going to roll, Mr. Quinn,' she said, 'or are you going to keep staring at my earrings?'"

System	Extraction
SCENE-DEC	dialogue_density++ markedly. Scene character now dialogue_dominance=emerging.
CHAR-IDENT	Confirms attribution to Bertha; addressee Jasper. No new characters.
SCINE	has_dialogue=true; dialogue_dominance=true; speaker_count=2; turn_count=1. Emotional register sharpens: irony_axis=0.85, tension_floor=0.35.
DEA	**Turn 0 created.** speaker_id=bertha_uuid; addressee_id=jasper_uuid; verbatim_text="You going to roll, Mr. Quinn, or are you going to keep staring at my earrings?"; attribution_basis=tag(she said); attribution_confidence=1.0; turn_kind=line;

System	Extraction
	register=playful_dominant; emphasis_tokens=[earrings]; pause_before_ms=400; pause_after_ms=400 (breathing room rule applied for Kling sound=true).
SSI	dialogue_priority_class=primary; duck_target_db=-12; ambient ducks under turn 0. Adds j_cut_recommended=true for incoming line.
R2A	Beat: subject=Bertha, action_verb=speak, gesture_token=arched_brow_implied, concurrent_action=stillness. Couples to DEA turn 0.
CIP	Bertha CIP: accessory_set+=earrings. Earrings now a **canonical prop** (text-anchored). wardrobe_provenance for earrings = canonical.
CCC	Confirms ots_B_required=true (favoring Bertha for the line); cu_bertha priority bumps to ≥ priority_class=high. Adds insert_earrings.
APE	Dispatches dialogue intent: dialogue_anchor_set=[turn_0], provider_recommended=Kling_V3_Pro_sound_true, pre_roll_ms=400, tail_ms=400. Dispatch mode = dialogue_mode.

Sentence 5

> *"Jasper rolled. The dice hit the felt twice, kissed the cushion, and came up boxcars — six and six — and the whole car forgot how to breathe."*

System	Extraction
SCENE-DEC	action_density++; this is the scene's apex beat. No boundary.
CHAR-IDENT	No new characters. Mention counts increment. Bystander set referenced collectively ("the whole car").
SCINE	Marks revelation_flag=true; escalation_class=apex; tension_ceiling=1.0 for this beat; narrative_function=reveal; camera_language+=[insert_macro, reaction_pan]. Title motif recognized: motif_recall=flipping_boxcars.

System	Extraction
DEA	None.
SSI	silence_required_flag=true on the breath-hold; practical_sound_set=[dice_hit:2, cushion_kiss, dice_settle]; effect_priority=[dice_hit, cushion_kiss]; ambient drops to floor for held beat. motif_carry_flag=true (sound spine continuity to next scene).
R2A	Three beats: (0) Jasper rolls, (1) dice travel + cushion kiss, (2) reveal — boxcars. Beat 2 sets intensity_class=apex, success_class=outcome_pending_narrativ ely. Bystander beat: subject_set=car_occupants, action_verb=hold_breath, intentionality_class=involuntary.
CIP	No new identity data.
CCC	Locks **MVC** for the scene: [wide_master, two_shot_master, ots_A, ots_B, cu_jasper, cu_bertha, insert_dice_macro, insert_earrings, reaction_pan]. mvc_satisfied_flag=pending_until_dispat ch.
APE	Stages the apex shot intents: insert_dice_macro (priority=critical, lens=macro, movement=locked, lighting=warm_practical_underlit_table), and reaction_pan (movement=slow_pan_180). Both carry qc_axes_required=[identity, continuity, aspect, motion].

Cross-Sentence Roll-Up — End of Section

After five sentences, the ADAPT layer has produced — **without any provider call yet**:

- **1 scene row** (scene_index=0) with all SCENE-DEC parameters populated.
- **2 character rows** (Jasper, Bertha), 1 with pov_flag=true.
- **2 CIP drafts**, neither yet at cip_locked=true. Jasper completeness ≈ 0.41 after wardrobe seed; Bertha ≈ 0.34 after canonical earrings anchor.
- **1 dialogue turn** in DEA (Bertha → Jasper), Kling-ready with breathing room applied.
- **1 SCINE record** with full camera language, sonic signals, and motifs.
- **1 SSI record** including silence-held beat, j-cut recommendation, and motif carry to scene 1.
- **9-role MVC** in CCC (wide_master, two_shot_master, ots_A, ots_B, cu_jasper, cu_bertha, insert_dice_macro, insert_earrings, reaction_pan).

- **9 shot intents** in APE plan v1, each carrying UUID subject binding and a spend estimate.

Provider calls so far: 0. Slate-id issued: 1 (governs all future spend for this scene). **Failures written: 0.**

This is what is meant by *deterministic preparation*. The prose has been fully converted into a producible plan before any model spends a dollar.

How to Use This Reference

- **For calibration:** when a new system (e.g. wardrobe-state tracker) is added, re-extract these five sentences and confirm the new system's parameters populate without disturbing the other eight.
- **For regression:** these five sentences and their expected ADAPT outputs are checked into the validation harness. A deviation in any of the 477 parameters is a regression event.
- **For onboarding:** read this document once and you understand what every other Fliction document means by "the spine."

End of Extraction Reference — Flipping Boxcars, Chapter 1.

Fliction — One-Page Visual Architecture

Prose → 9 Canonical Systems → CNIO → Downstream Production

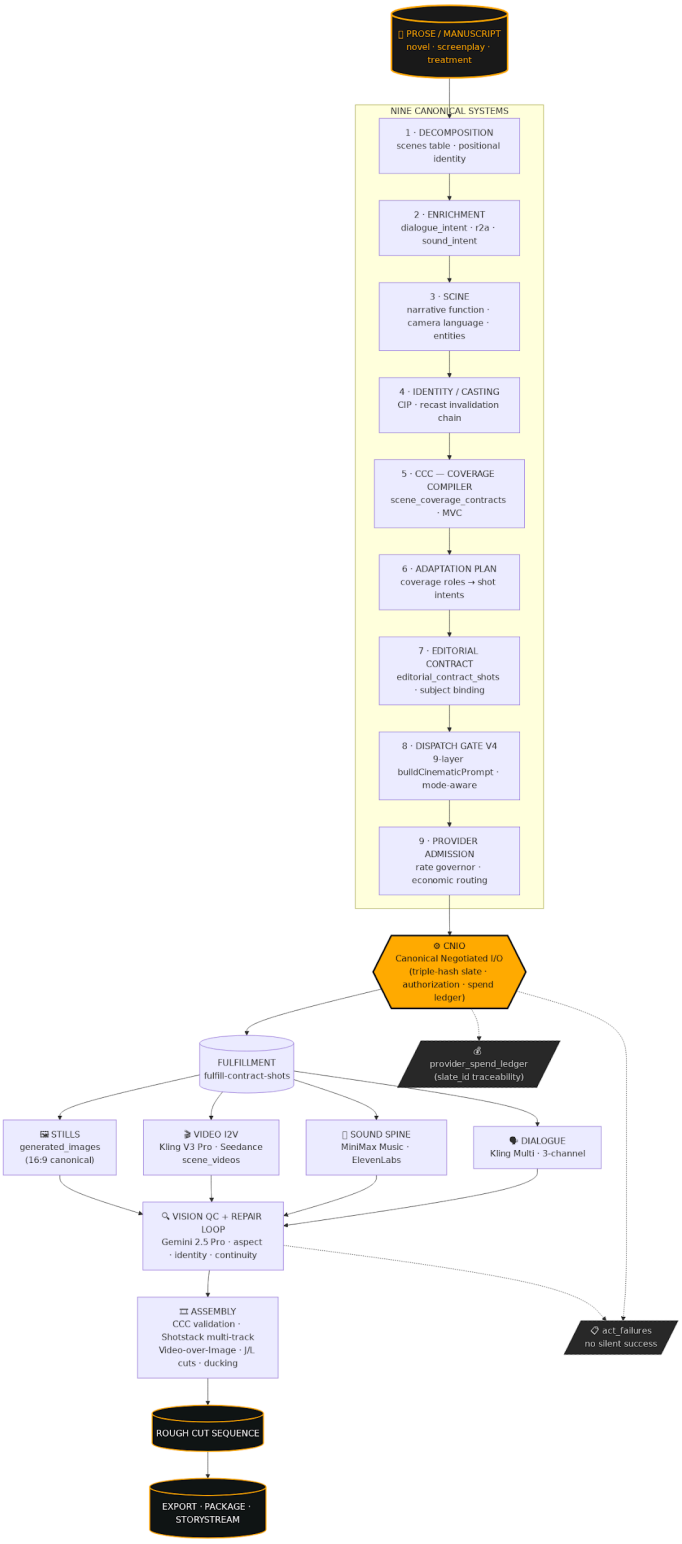

Figure 1. The Fliction One-Page Architecture. Prose enters at the top and flows deterministically through nine canonical systems — Decomposition, Enrichment, SCINE, Identity/Casting, the Canonical Coverage Compiler (CCC), Adaptation Planning, Editorial Contract, Dispatch Gate V4, and Provider Admission — before crossing the CNIO (Canonical Negotiated I/O) authorization boundary. CNIO issues the triple-hash semantic slate that authorizes all downstream execution: still generation, I2V video, sound spine, and dialogue. Every asset re-enters the Vision QC + Repair Loop before Assembly compiles the rough cut for export. Side-rails (spend ledger, act_failures) capture universal traceability under the No-Silent-Success principle.

Works Cited

Benjamin, Ruha. Race After Technology: Abolitionist Tools for the New Jim Code. Polity Press, 2019.

Burrow-Giles Lithographic Co. v. Sarony. 111 U.S. 53. Supreme Court of the United States, 1884.

Chatman, Seymour. Story and Discourse: Narrative Structure in Fiction and Film. Cornell UP, 1978.

Chun, Wendy Hui Kyong. Programmed Visions: Software and Memory. MIT Press, 2011.

---. Updating to Remain the Same: Habitual New Media. MIT Press, 2016.

Christovale, Erin, curator. *No Humans Involved*. Hammer Museum, 10 Oct. 2021–9 Jan. 2022, Los Angeles. *Hammer Museum*, UCLA, https://hammer.ucla.edu/exhibitions/2021/no-humans-involved. Accessed 6 May 2026.

Ginsburg, Jane C. "The Concept of Authorship in Comparative Copyright Law." DePaul Law Review, vol. 52, no. 4, 2003, pp. 1063-1092.

Haacke, Hans. "Museums, Managers of Consciousness." *Art in America*, vol. 72, no. 2, Feb. 1984, pp. 9–17. *Artcontext*, https://artcontext.net/crit/scrapbook/index.php?id=26. Accessed 16 Feb 2026.

Hartman, Saidiya. "Venus in Two Acts." *Small Axe: A Caribbean Journal of Criticism*, vol. 12, no. 2, 2008, pp. 1–14. Project MUSE, doi:10.1353/smx.0.0006.

Hartman, Saidiya. Wayward Lives, Beautiful Experiments: Intimate Histories of Riotous Black Girls, Troublesome Women, and Queer Radicals. W. W. Norton, 2019.

Kyles, Cedric, with Alan Eisenstock. Flipping Boxcars: A Novel. Amistad, 2023.

Metz, Christian. Film Language: A Semiotics of the Cinema. Translated by Michael Taylor,

University of Chicago Press, 1974.

Murch, Walter. In the Blink of an Eye: A Perspective on Film Editing. 2nd ed., Silman-James Press, 2001.

Negarestani, Reza. Intelligence and Spirit. Urbanomic / Sequence Press, 2018.

U.S. Copyright Office. *Copyright and Artificial Intelligence, Part 2: Copyrightability*. United States Copyright Office, Jan. 2025, copyright.gov/ai. Accessed 18 May 2026.

Weizman, Eyal. *Forensic Architecture: Violence at the Threshold of Detectability*. Zone Books, 2017. *JSTOR*, https://www.jstor.org/stable/j.ctv14gphth. Accessed 6 April 2026.

Whitney Museum of American Art. "Hans Haacke, *Shapolsky et al. Manhattan Real Estate Holdings, A Real-Time Social System, as of May 1, 1971*." *Whitney Museum of American Art*, 1971, whitney.org/collection/works/29487. Accessed 18 May 2026.

Wilson, Fred. *Mining the Museum*. 1992–93, installation, Maryland Historical Society, Baltimore. *Grand Street*, no. 44, 1993, pp. 151–72. *JSTOR*, https://www.jstor.org/stable/25007622. Accessed 17 Nov 2025.

Bibliography

Baldwin, James. “The Artist’s Struggle for Integrity.” YouTube, uploaded by stellla, 23 July 2015, www.youtube.com/watch?v=dU0g5fAA2QY. Accessed 22 July 2025.

Baldwin’s lecture frames the artist as someone responsible for truth under social pressure. Its relevance to this thesis is not technical in a narrow sense; it supports the ethics behind the system. ADAPT is built around the claim that a reader’s perception can be formalized without surrendering artistic integrity or agency. Baldwin’s language about the artist’s responsibility aims to give the project an ethical ground: the system should not flatten the maker’s vision into institutional convenience or algorithmic sameness. This source is especially useful for explaining why the thesis treats cinematic cognition as lived, disciplined, and accountable rather than merely intuitive or algorithmic nihilism.

Benjamin, Ruha. Race After Technology: Abolitionist Tools for the New Jim Code. Polity Press, 2019.

Benjamin is one of the thesis’s central critical technology sources. Her argument that technological systems can reproduce racial hierarchy while presenting themselves as neutral directly informs ADAPT’s refusal of neutrality. The taxonomy is not framed as universal machine objectivity; it is framed as a situated reading structure shaped by Black cinematic perception and my reading in semiotics and syntax. Benjamin gives language for the political risk of coded classification: when a system names, sorts, and translates, it also produces power as means of production. This source is essential because it helps the thesis argue that a responsible narrative AI system must expose its assumptions instead of hiding them behind the appearance of technical inevitability.

Burrow-Giles Lithographic Co. v. Sarony. 111 U.S. 53. Supreme Court of the United States, 1884.

This case is useful for the thesis because it establishes an early legal framework for authorship in relation to a technical apparatus. The Court recognized that a photograph could be authored by the person who arranged, posed, selected, and controlled the expressive elements of the image, even though the camera mechanically fixed the result. This supports ADAPT's claim that authorship can reside in conception, arrangement, supervision, and system design rather than only in manual execution.

Chatman, Seymour. Story and Discourse: Narrative Structure in Fiction and Film. Cornell UP, 1978.

Chatman's distinction between story and discourse is foundational for the thesis's method. Story names what happens; discourse names how the narrative presents what happens. ADAPT depends on this difference because it does not merely identify plot events. It extracts pacing, framing, silence, emphasis, dialogue structure, and other signals embedded in narratives. Chatman helps place the project inside a serious narratological tradition while also allowing it to extend that tradition toward computational cinema. His work is useful because it clarifies that cinematic information can exist not only in content but in arrangement, syntax, sequence, and form.

Chun, Wendy Hui Kyong. Programmed Visions: Software and Memory. MIT Press, 2011.

Chun's work helps define software as a cultural and perceptual system rather than a neutral technical instrument. This matters to ADAPT because the project transforms a filmmaker's reading practice into repeatable computational logic. Chun's attention to memory, programmability, and repetition gives the thesis a way to ask what happens

when interpretation becomes system behavior. Her work supports the claim that ADAPT does not simply analyze prose; it formalizes a mode of seeing. That makes the system powerful, but also accountable, because any repeated interpretive structure can become ideological and homogenizing if its position and limits are not visible.

---. Updating to Remain the Same: Habitual New Media. MIT Press, 2016.

In this book, Chun's focus on habit sharpens the thesis's account of repeatability. ADAPT's value depends on its ability to produce consistent readings across passages, chapters, and books, but consistency is never innocent. A repeated system can stabilize insight, but it can also normalize bias. Chun helps the thesis hold that contradiction in view. This source is especially useful for the automation section because it shows how software shapes users through patterned repetition. ADAPT therefore has to be read as both a tool and training environment: it reveals cinematic signals while also conditioning how those signals are recognized.

Christovale, Erin, curator. No Humans Involved. Hammer Museum, 10 Oct. 2021-9 Jan. 2022, Los Angeles.

Christovale's exhibition is important to the thesis's working title and its posthuman, parafictional frame. No Humans Involved provides an art-historical and curatorial context for thinking about who is recognized as human within institutional and technological systems. The phrase is crucial because the thesis does not use "No Humans Involvd" as a celebration of automation. Instead, it turns the phrase back toward the politics of recognition: whose cognition, craft, and authorship count as human? Christovale's exhibition helps locate the project in a contemporary Black critical conversation about technology, embodiment, and the limits of the human category.

Forensic Architecture, and Eyal Weizman. Forensic Architecture: Violence at the Threshold of Detectability. Zone Books, 2017.

Forensic Architecture gives the thesis a methodological precedent for structured reconstruction from dispersed signals. Its practice combines spatial analysis, temporal modeling, media evidence, and counter-investigation. ADAPT performs a related operation within prose: it treats narrative language as a field of evidence from which cinematic structures can be extracted, compared, and audited. This source is useful because it helps describe the system as more than interpretation. ADAPT becomes a forensic reading apparatus, measuring what a text specifies, what it implies, and what it withholds. That framework supports the thesis's emphasis on traceability, confidence, and indeterminacy.

Fraser, Andrea. Official Welcome. 2001.

Fraser's performance is useful as a model of institutional critique performed from inside the language of the institution. Her work shows how authority, value, and recognition are staged through speech, ceremony, and role-play. This matters for ADAPT because the thesis is not outside the systems it critiques; it participates in art, film, academic, and AI institutions while exposing their procedures. Fraser helps clarify the performative dimension of the installation: the system does not only produce outputs; it stages the conditions under which outputs become legible, valuable, and authoritative.

Ginsburg, Jane C. "The Concept of Authorship in Comparative Copyright Law." DePaul Law Review, vol. 52, no. 4, 2003, pp. 1063-1092.

Ginsburg's article supports the thesis's discussion of authorship under conditions of automation. Her legal framework helps distinguish authorship from mere mechanical

execution. This distinction is important because ADAPT uses computational processes without surrendering authorship to the machine. The human authorial act resides in designing the taxonomy, determining the extraction logic, establishing constraints, and supervising how outputs are used. Ginsburg is especially useful for the section on systemic authorship because she gives the thesis a vocabulary for arguing that system design can be an authorial act when the human mind conceives and controls the work's expressive structure.

Haacke, Hans. "Museums, Managers of Consciousness." Art in America, vol. 72, no. 2, Feb. 1984, pp. 9-17.

Haacke is central to the thesis's systems critique. His writing and practice reveal that museums and cultural institutions do not merely display meaning; they manage consciousness by structuring what becomes visible, valuable, and legitimate. ADAPT extends that concern into narrative technology. A taxonomy that translates prose into cinematic data is also a system of visibility. Haacke helps the thesis insist that such a system must show its procedures, assumptions, and power relations. This source is especially important for the installation and writing because the work's force depends on making the system observable rather than hiding it behind a polished cinematic result.

Haacke, Hans. "*Shapolsky et al. Manhattan Real Estate Holdings, A Real-Time Social System, as of May 1, 1971*." *Whitney Museum of American Art*, 1971, whitney.org/collection/works/29487. Accessed 18 May 2026.

Haacke's *Shapolsky et al.* is essential to the thesis because it demonstrates how documentation can become an artistic form. The work organizes photographs, maps, property records, ownership structures, and financial information into a visible system,

exposing how power operates through real estate, bureaucracy, and institutional concealment. This helps frame ADAPT as more than a technical apparatus. Like Haacke's work, ADAPT treats records, logs, classifications, and procedural evidence as part of the artwork itself. The comparison allows the thesis to argue that generative cinematic systems must not only produce images; they must also reveal the structures by which images are produced, authorized, withheld, and archived. The documentation is not secondary to the artwork. It co-creates the artwork's reality, establishes historical accountability, and positions the artist as an architect of the system's visible logic.

Hartman, Saidiya. "Venus in Two Acts." Small Axe: A Caribbean Journal of Criticism, vol. 12, no. 2, 2008, pp. 1–14. Project MUSE, doi:10.1353/smx.0.0006.

Hartman's essay is important to this thesis because it provides a method for thinking about absence as evidence rather than emptiness. In "Venus in Two Acts," Hartman confronts the limits of an archive shaped by slavery, violence, and the partial recording of Black life. Her practice of critical fabulation combines archival attention with speculative care, while refusing to falsely complete what history has made unknowable. This helps frame ADAPT's indeterminacy zones: when the prose withholds interiority, the system does not invent emotion to smooth the gap. It marks the absence, preserves it, and carries it forward as form. Hartman allows the thesis to argue that not knowing can be an ethical and structural decision.

Hartman, Saidiya. Wayward Lives, Beautiful Experiments: Intimate Histories of Riotous Black Girls, Troublesome Women, and Queer Radicals. W. W. Norton, 2019.

Hartman's work is crucial for the thesis's treatment of absence, silence, and indeterminacy. Her method attends to what archives omit without pretending that absence can be filled without consequence. ADAPT adapts this ethical posture computationally by marking gaps rather than automatically inventing coverage. In the thesis, indeterminacy is not a failure of extraction; it is evidence that the text withholds something or that the system has reached the boundary of responsible inference. Hartman gives the project a rigorous way to argue that what cannot be measured should not be falsely completed, especially when working with Black narrative and historical memory.

Houston, Kerr. "How 'Mining the Museum' Changed the Art World." BmoreArt, 3 May 2017, www.bmoreart.com/2017/05/how-mining-the-museum-changed-the-art-world.html. Accessed 19 July 2025.

Houston's essay provides accessible context for Fred Wilson's Mining the Museum and its lasting impact on institutional critique. This source is helpful because it explains how curatorial arrangement can reveal hidden histories within an existing collection. That logic parallels ADAPT's treatment of prose: the system does not invent the source text's cinematic intelligence but rearranges and classifies its signals so that a hidden structure becomes legible. Houston's essay is useful for connecting the thesis to a professional art conversation about display, power, and the politics of arrangement.

Kyles, Cedric, with Alan Eisenstock. Flipping Boxcars: A Novel. Amistad, 2023.

Flipping Boxcars is the thesis's primary evidence text and the main narrative body through which ADAPT is tested. Its importance is practical and conceptual. Practically, it provides a full-length narrative across which the system can demonstrate scale,

consistency, and failure points. Conceptually, it gives ADAPT a culturally specific Black narrative field rather than a generic test object. The novel's dialogue, humor, historical setting, spatial transitions, and sonic textures allow the thesis to ask whether a computational taxonomy can preserve signals that ordinary adaptation pipelines might flatten. This source is critical as it is where the thesis moves from argument to proof.

Metz, Christian. Film Language: A Semiotics of the Cinema. Translated by Michael Taylor, University of Chicago Press, 1974.

Metz provides the film-theory foundation for the thesis's claim that cinema operates as a language system. His semiotic approach makes it possible to speak about cinematic units, syntax, and meaning with theoretical precision. ADAPT extends Metz by relocating the question from finished film back into prose before production. If cinema has grammar, the thesis asks whether some of that grammar is already latent in narrative language. Metz is therefore essential because he provides the conceptual bridge between film language and the thesis's computational claim: cinematic syntax can be identified, structured, and translated before cameras, actors, or editing systems appear.

Murch, Walter. In the Blink of an Eye: A Perspective on Film Editing. 2nd ed., Silman-James Press, 2001.

Murch is central to the thesis's account of editing as emotional and rhythmic intelligence. His hierarchy of editorial criteria helps explain why a filmmaker reads prose as more than description. A sentence can imply when to cut, when to hold, where attention should move, and what feeling should govern the transition. ADAPT extends Murch upstream by arguing that many editorial pressures are already present at the level of reading. This source is key because it connects the project's technical taxonomy to live craft practice:

the system is not replacing the editor's intuition but attempting to notate the structures intuition often recognizes silently.

Negarestani, Reza. Intelligence and Spirit. Urbanomic / Sequence Press, 2018.

Negarestani's philosophical account of intelligence helps the thesis think beyond simple human-versus-machine binaries. His work treats intelligence as a structured capacity capable of revision, extension, and self-transformation. That frame is useful for ADAPT because the system does not claim that AI possesses the filmmaker's intuition. Instead, it externalizes a portion of human cinematic cognition into a form that can be tested and revised. Negarestani helps support the thesis's larger claim that intelligence is not mysticism; it can become infrastructure, provided the system remains accountable to the human and cultural conditions from which that intelligence emerges.

Piper, Adrian. Adrian Piper Interview: Rationality and the Structure of the Self. YouTube, uploaded by APRA Foundation Berlin, 13 May 2013, www.youtube.com/watch?v=_tURuyb76XQ. Accessed 22 July 2025.

Piper's thinking about rationality, selfhood, and conceptual structure supports the thesis's insistence that systems are never detached from the subject positions that produce them. Her work is useful because ADAPT is both analytic and self-situating: it attempts to be rigorous while refusing the false posture of universality. Piper helps frame the relationship between reason and identity without reducing either one. For the thesis, this matters because the system's Black cinematic perception is not an anecdotal feature; it is part of the work's conceptual structure and what makes its taxonomy different from a supposedly neutral pipeline.

Red Star, Wendy. Four Seasons. 2006-10.

Red Star's Four Seasons is useful as a visual precedent for examining how identity and cultural perception are staged, coded, and misread. Her constructed images expose the artificiality of representational expectations, especially around Indigenous identity and museum display. For ADAPT, Red Star matters because the thesis is also concerned with how systems recognize or fail to recognize cultural signals. Her work helps strengthen the project's visual-art context, especially where the thesis asks who controls the frame, who is being read, and what assumptions govern recognition. This source is especially relevant to the installation's concern with mediated perception.

"The Octavia E. Butler Collection." Art Papers, 24 Jan. 2023, www.artpapers.org/the-octavia-e-butler-collection/. Accessed 20 July 2025.

This article is useful because Butler's archive demonstrates the relationship between speculative imagination and rigorous structural planning. Her notebooks and papers show that worldbuilding is not only inspiration; it is method, system, revision, and architecture. That is directly relevant to ADAPT's technical-poetic form. The thesis is also trying to hold invention and structure together without diminishing either one. Butler's archival model supports the idea that speculative practice can be deeply organized and that system design can emerge from artistic imagination rather than standing opposed to it.

Wilson, Fred. Mining the Museum. Maryland Historical Society, 1992.

Wilson's Mining the Museum is one of the strongest artistic precedents for ADAPT. By rearranging objects already held by the Maryland Historical Society, Wilson revealed suppressed histories and racialized institutional logic. His work demonstrates that selection, placement, and classification can be forms of authorship and critique. ADAPT performs a comparable operation on narrative text. It does not claim to create cinematic

meaning from nothing; it rearranges, classifies, and exposes signals already present in prose. Wilson is essential because he gives the thesis an art-historical model for understanding taxonomy and arrangement as both creative practice and institutional intervention.

U.S. Copyright Office. *Copyright and Artificial Intelligence, Part 2: Copyrightability*. United States Copyright Office, Jan. 2025, copyright.gov/ai. Accessed 18 May 2026.

This report is important to the thesis because it directly addresses authorship in AI-assisted works. The Copyright Office affirms that human authorship remains the foundation of copyright protection, but that the use of AI as an assistive tool does not necessarily defeat copyrightability where a human author contributes sufficient expressive elements, selection, arrangement, or modification. This helps situate ADAPT's claim that the artist's authorship resides in the design of the taxonomy, extraction logic, and downstream cinematic architecture rather than in the raw output of a generative model.

Wynter, Sylvia. "No Humans Involved: An Open Letter to My Colleagues." Forum N.H.I.: Knowledge for the 21st Century, vol. 1, no. 1, 1994, pp. 42-73.

Wynter is essential to the thesis's argument about the category of the human. Her essay exposes how institutions have historically excluded Black life from full recognition while naturalizing that exclusion through bureaucratic language. No Humans Involvd is not an anti-human slogan. It is a provocation about whose cognition counts as human cognition in an age of AI cinema. Wynter helps ground the project's deepest political claim: ADAPT is not about replacing the human with the machine, but about insisting that Black artistic cognition is a form of knowledge, authorship, and infrastructure.

Artist Bio

Allen Redwing is an interdisciplinary artist, filmmaker, writer, and computational narrative researcher based in Boston, Massachusetts. His work examines institutional power, racialized systems of interpretation, cinematic language, and the relationship between narrative, technology, and perception. Working across installation, film, conceptual writing, AI-assisted image-making, and computational systems, he investigates how meaning is structured and enforced through cultural, administrative, and algorithmic frameworks

Recent projects and exhibitions include *Black, Broke, But Still Believing*, a solo installation at Massachusetts College of Art and Design examining cultural erasure and systemic harm through AI-generated imagery and painting; *Government Issue*, a mixed-media installation exploring military memory, governance, and institutional critique; and *White Logic*, an experimental film investigating whiteness as institutional structure His practice often merges procedural documentation, speculative narrative, archival material, and cinematic systems to interrogate how authority is constructed and maintained.

Redwing is the founder of Fliction, a stealth AI-native research lab and cinematic story adaptation platform, and creator of the ADAPT framework, a computational system for narrative and cinematic analysis. His broader research focuses on authorship, interpretive systems, computational aesthetics, and emerging forms of AI-mediated storytelling

His recognitions include acknowledgment by *Variety Magazine* among its Top 110 Global Filmmakers, selections at the IndieFEST Film Awards and the South Carolina Underground Film Festival, and work held in the private collection of Marina Hatsopoulos

He is currently completing his MFA at Massachusetts College of Art and Design, where his thesis research explores institutional critique, computational culture, and the politics of interpretation.

Artist Statement

Allen Redwing is an interdisciplinary artist working across installation, filmmaking, AI-assisted image generation, conceptual writing, archival reconstruction, sound, and computational storytelling systems to examine how institutions shape meaning, perception, and social power. His work combines physical installations, projected imagery, procedural documents, cinematic sequencing, machine-generated visuals, and narrative analysis tools to investigate how authority is constructed through language, bureaucracy, visual culture, and interpretation.

Many of Redwing's projects begin with existing institutional structures—disciplinary procedures, military language, legal documents, cinematic conventions, and algorithmic systems—which he reorganizes into films, installations, and conceptual interventions. In works such as *Government Issue* and *Redwing vs. MassArt*, he uses testimony, reconstructed records, projected text, AI-generated imagery, sound, and archival fragments to expose how institutions stabilize authority through evidentiary framing and procedural logic. Rather than presenting conflict as documentation alone, these works transform institutional processes into material, revealing how systems often prioritize coherence, administration, and control over truth.

His image-based works combine AI-generated imagery, painting, digital compositing, and cinematic framing to examine racialized representation, masculinity, media conditioning, and cultural memory. By intentionally working through unstable and imperfect machine systems, Redwing treats artificial intelligence not as an automated author, but as a site where social assumptions, psychological projections, and cultural biases become visible.

Alongside his studio practice, Redwing develops computational narrative systems that translate storytelling into visual, sonic, emotional, and symbolic structures. Through projects such as Fliction and the ADAPT framework, he builds tools that analyze narrative text through cinematic patterns and convert interpretation into procedural architecture. These systems function both as artistic media and as critical investigations into how stories are organized, controlled, and distributed.

Across his practice, Redwing uses art to expose the hidden infrastructures that govern interpretation and to challenge who has the authority to define reality.

Allen J. Redwing

Boston, Massachusetts
ajredwing@fliction.com
www.fliction.com
@allenredwing

Curriculum Vitae

Allen Redwing is an interdisciplinary artist, filmmaker, writer, and computational narrative researcher based in Boston, Massachusetts. His work examines institutional power, racialized systems of interpretation, cinematic language, and the relationship between narrative, technology, and perception. Working across installation, film, conceptual writing, AI-assisted image-making, and computational systems, he investigates how meaning is structured and enforced through cultural, administrative, and algorithmic frameworks.

Education

(Pending) 2026	Master of Fine Arts (MFA)
	Massachusetts College of Art and Design Boston
2015	Filmmaking Certificate, New York Film Academy
2013–2014	Digital Filmmaking Certificate
	Boston University
2007	Bachelor of Science (B.S.), Business Administration, Cum Laude
	Endicott College

Creative Activity, Films, Exhibitions and Screenings

2025	*Black, Broke, But Still Believing,* Solo Installation
	Doran Hall, Massachusetts College of Art and Design
	Boston, Massachusetts
2025	Political Shock Art: Global Satire and Cultural Tensions, Solo Exhibition
	Fifty curated works interrogating political apathy, media conditioning, and global destabilization.

2025 White Logic, Video Screening, Massachusetts College of Art and Design

Eleven-minute film exploring whiteness as institutional logic; includes original AI generated song "Speak On It."

2025 Government Issue, Mixed Media Installation

Massachusetts College of Art and Design Boston, Massachusetts

Blending military memory, institutional critique, archival residue, satirical prints, protest memes, ritualized text, and sound, Government Issue interrogates systems of protection, identity, governance, and erasure.

Artworks

2024 *Let's Not Artfully Pretend*, Open Letter and Narrative

2024 *Emma Tell'em Story*, Fictionalized Narrative

2024 *Artist Talk: Computational Creativity and Racialized Systems*

2024 *Rebuttals:* Visual and Critical Studies Program

2024 *Redwing vs Mass Art Appeal*, Institutional and Procedural Critique

Honors, Awards and Collections

2015 Variety Magazine: Top 110 Global Filmmakers

2015 IndieFEST Film Awards: Cinematic Gems and Unique Voices

2015 South Carolina Underground Film Festival (2015): Selection

2024 Private Collection: Marina Hatsopoulos – *Black, Broke, But Still Believing*

Membership in Professional Bodies

2022–Present Grubstreet Creative Writing Center

Images

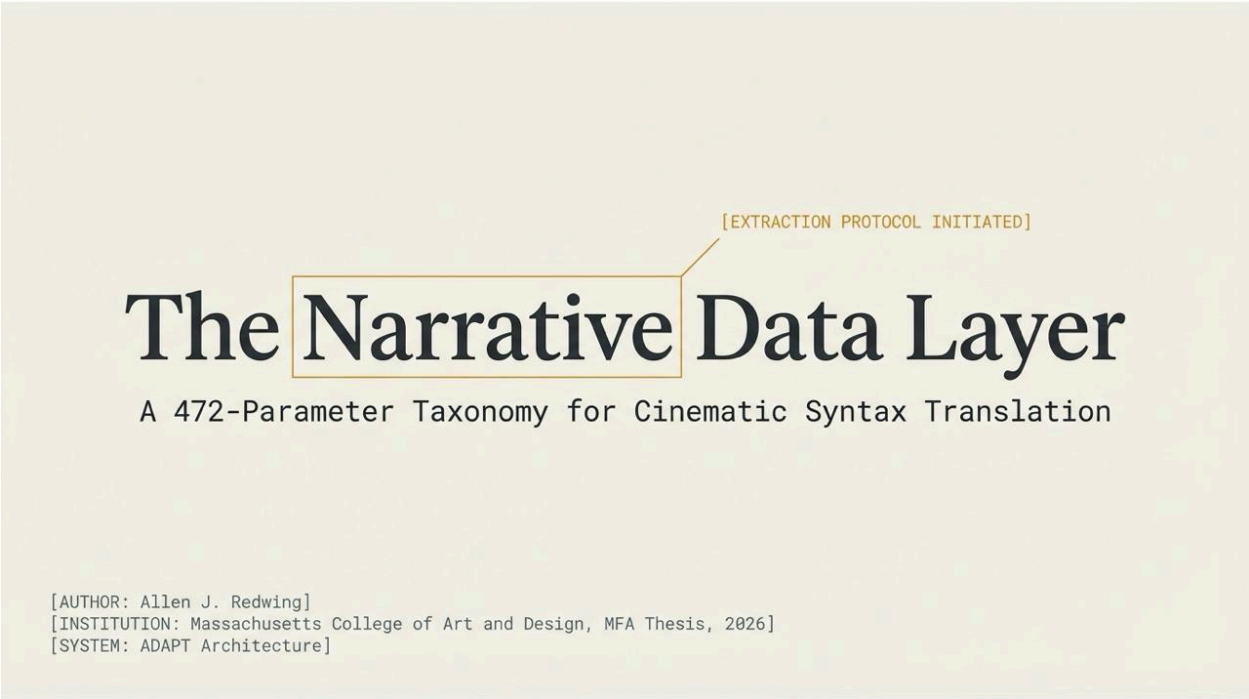

Figure 1. Narrative Data Layer, a 472+-parameter taxonomy for cinematic syntax translation.

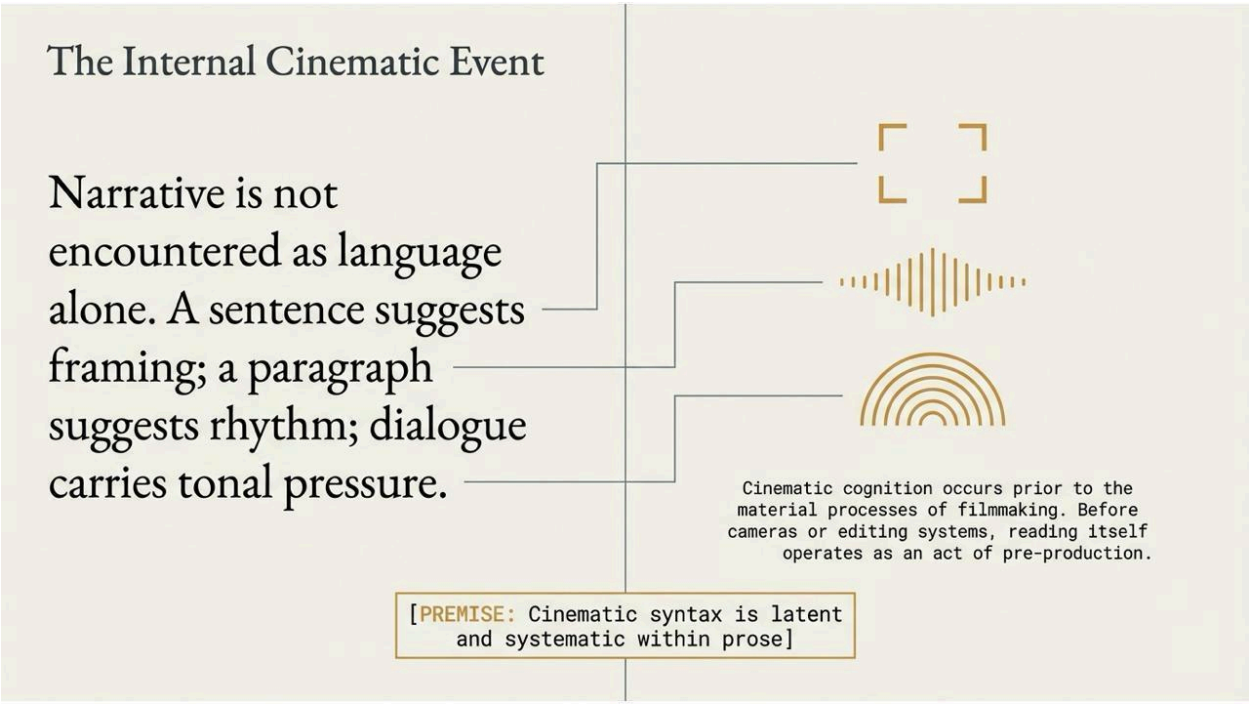

Figure 2. The internal cinematic event, illustrating how prose simultaneously encodes framing, rhythm, and tonal pressure prior to production.

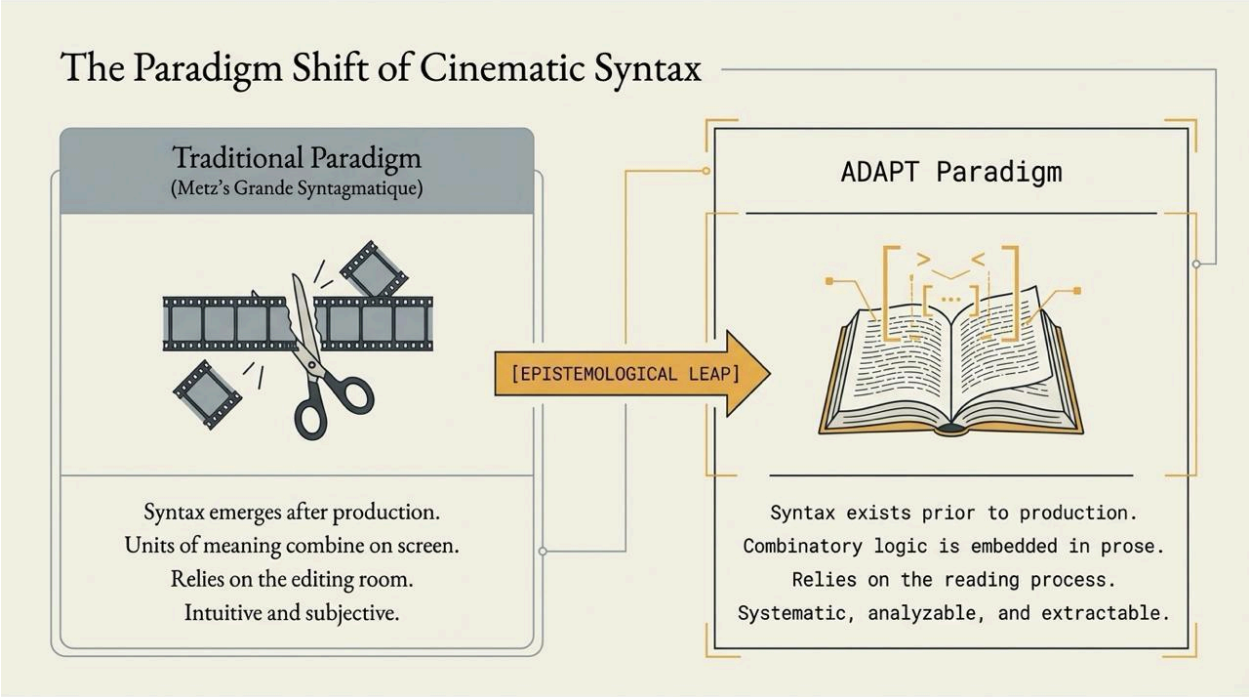

Figure 3. Paradigm shift from Metz's post-production Grande Syntagmatique to ADAPT's pre-production extraction of combinatory logic embedded in prose.

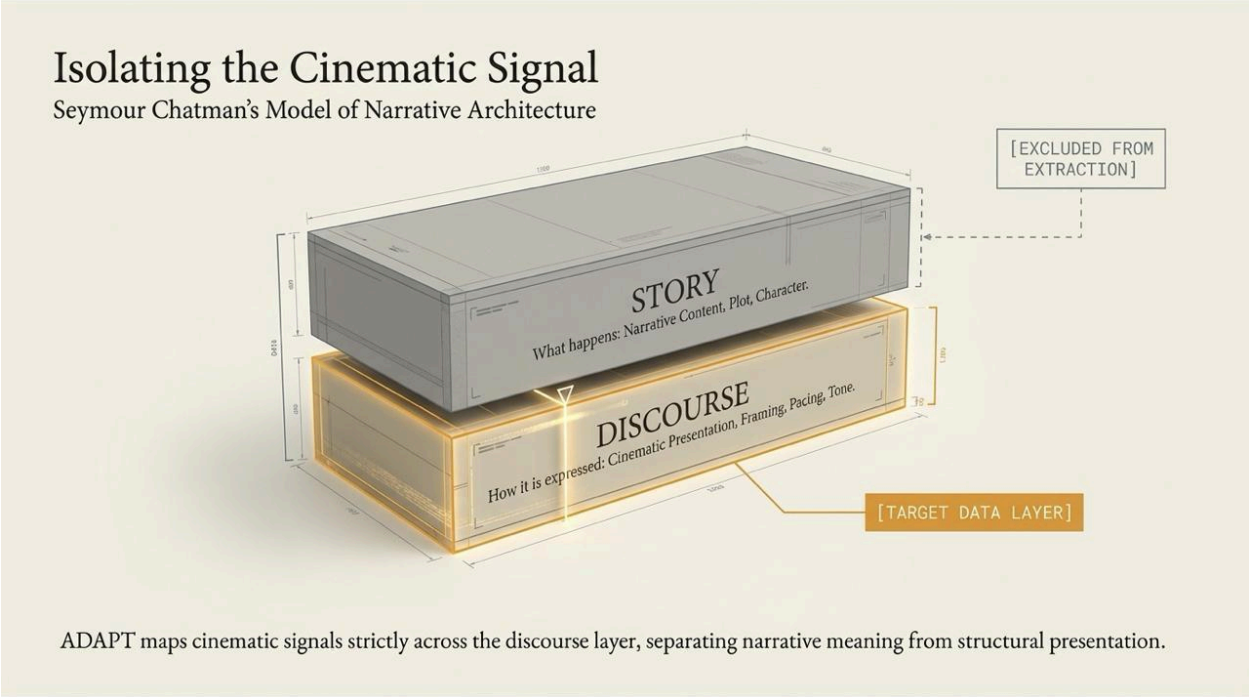

Figure 4. Chatman's story/discourse model adapted to isolate the discourse layer as ADAPT's target data surface for cinematic signal extraction.

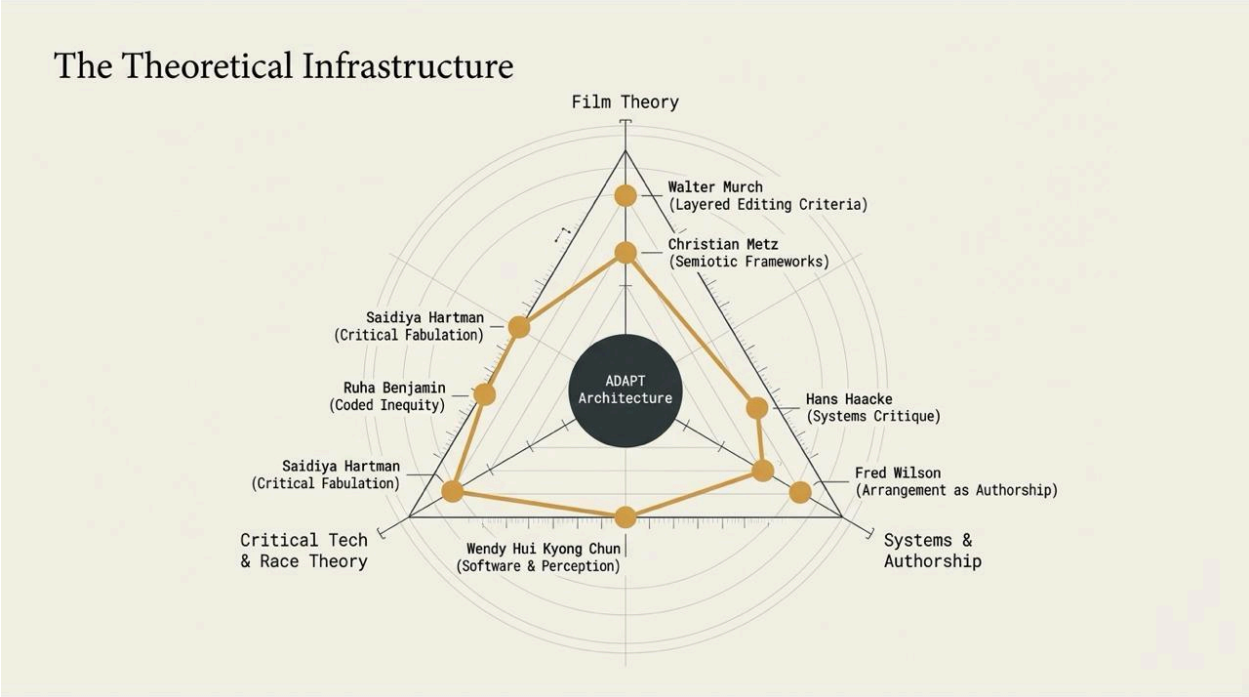

Figure 5. Radar diagram mapping ADAPT's theoretical infrastructure across three axes: film theory (Murch, Metz), critical tech and race theory (Chun, Benjamin, Hartman), and systems and authorship (Wilson, Haacke).

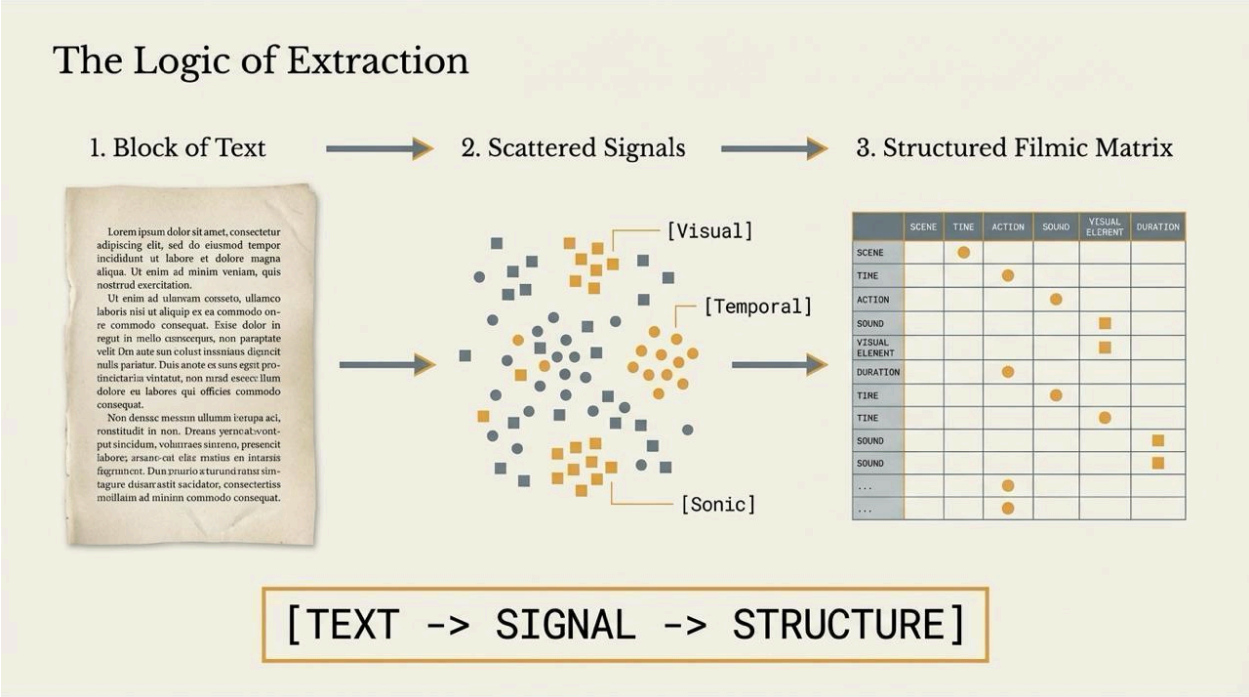

Figure 6. The extraction pipeline: narrative text decomposed into scattered visual, temporal, and sonic signals, then organized into a structured filmic matrix.

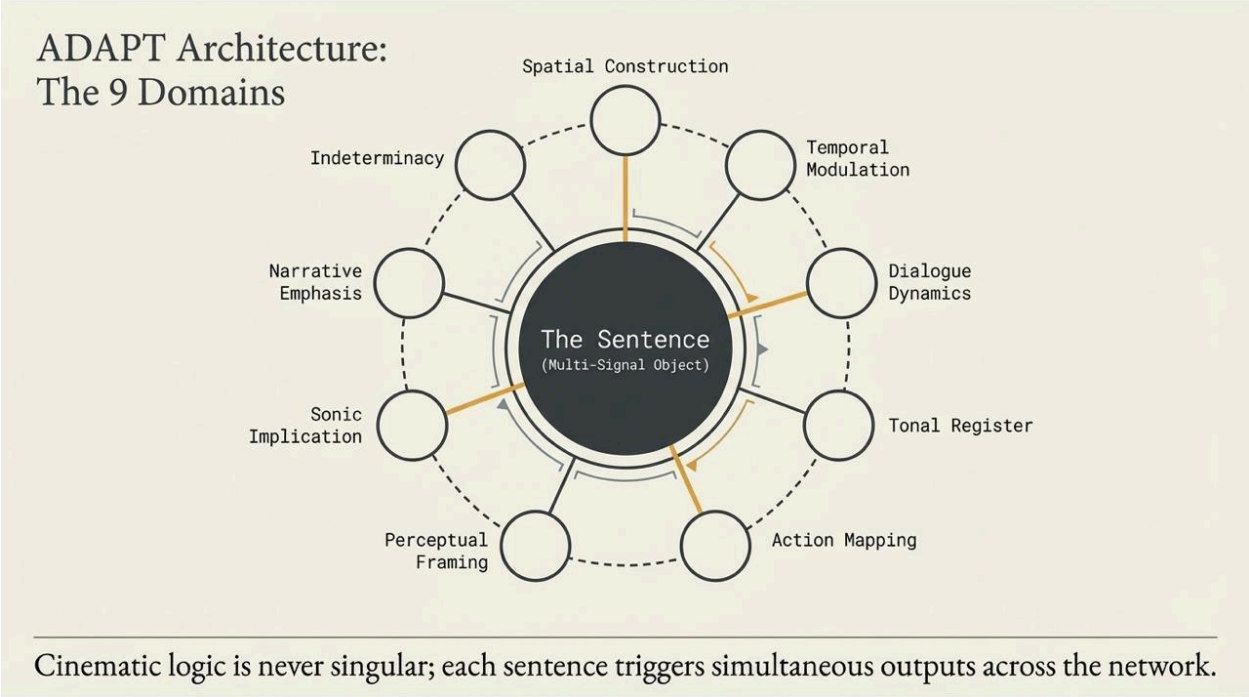

Figure 7. Radial schematic of ADAPT's nine extraction domains with the sentence positioned as a multi-signal object at center.

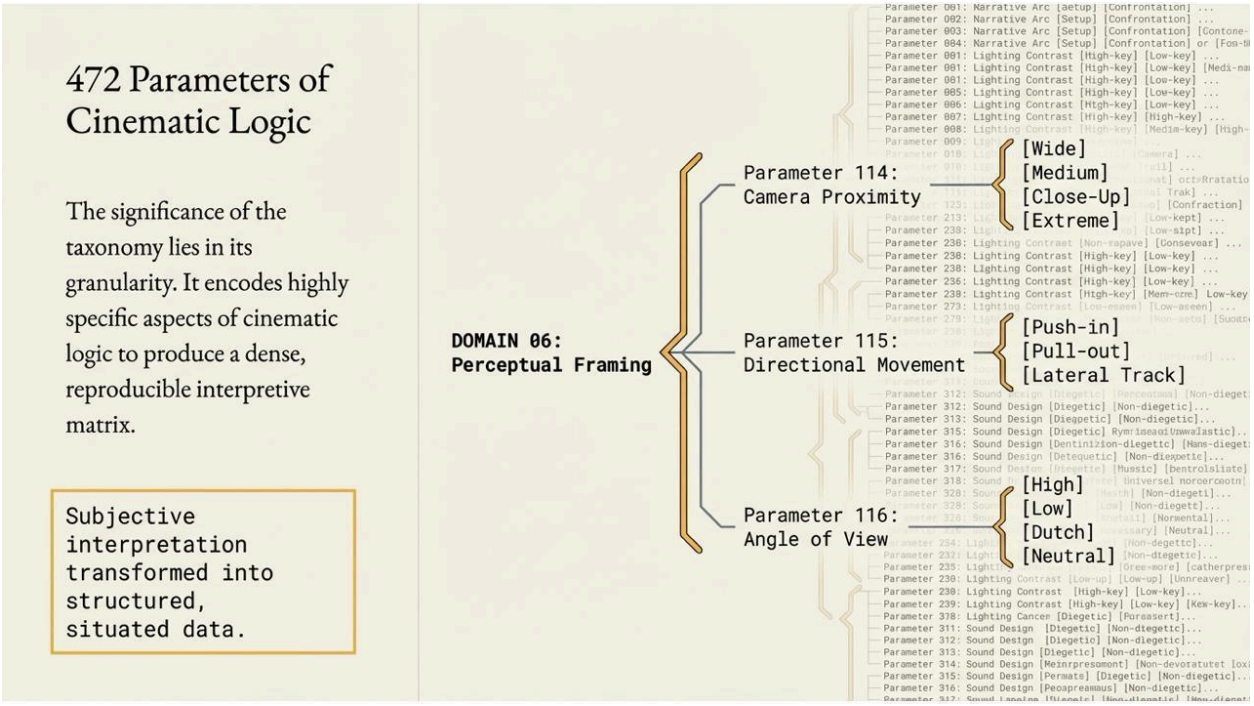

Figure 8. Parameter-level granularity within Domain 06 (Perceptual Framing), detailing camera proximity, directional movement, and angle of view taxonomies.

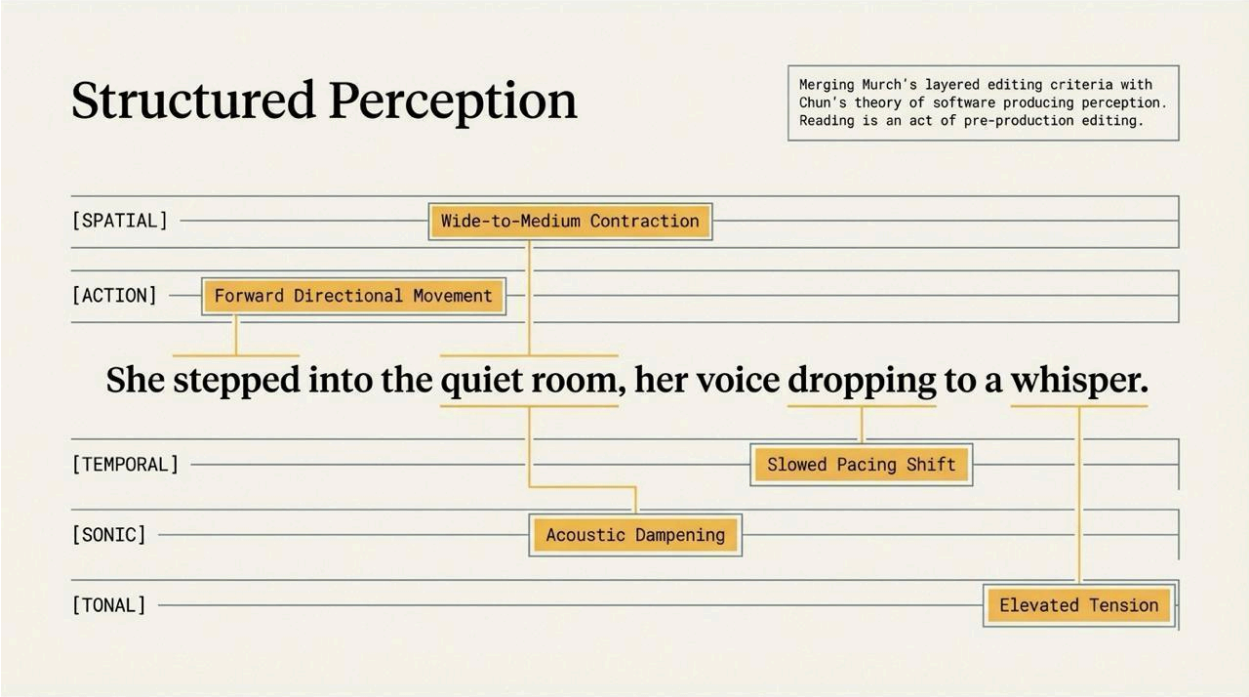

Figure 9. Structured perception model parsing a single sentence across five concurrent extraction layers: spatial, action, temporal, sonic, and tonal.

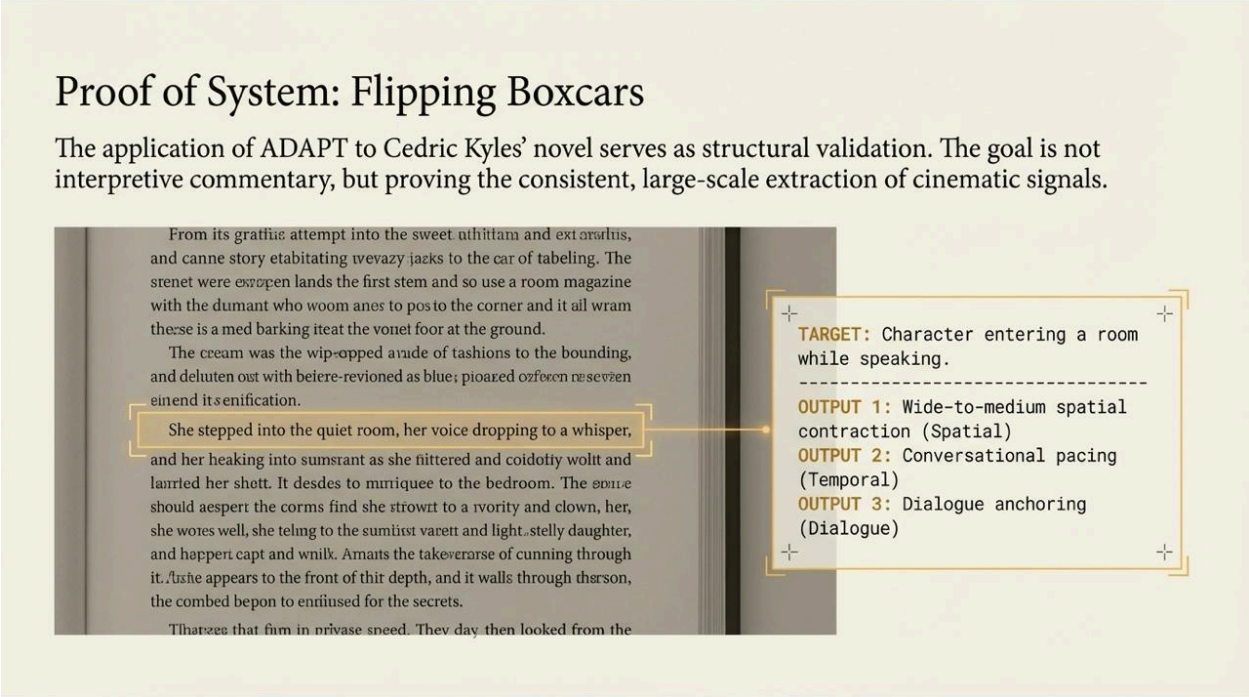

Figure 10. Proof of system: multi-layer extraction applied to a passage from Flipping Boxcars by Cedric Kyles, producing spatial, temporal, and dialogue outputs.

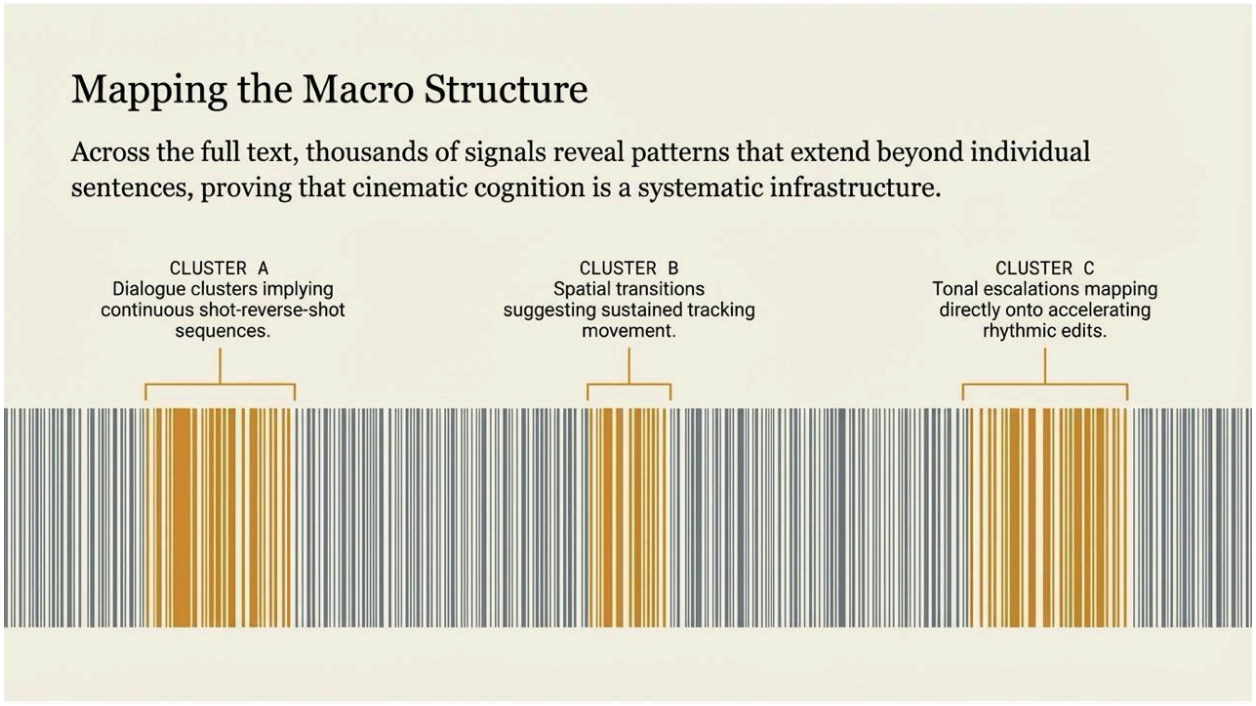

Figure 11. Macro-structural signal patterns — dialogue clusters, spatial transitions, and tonal escalations — identified across the full narrative text.

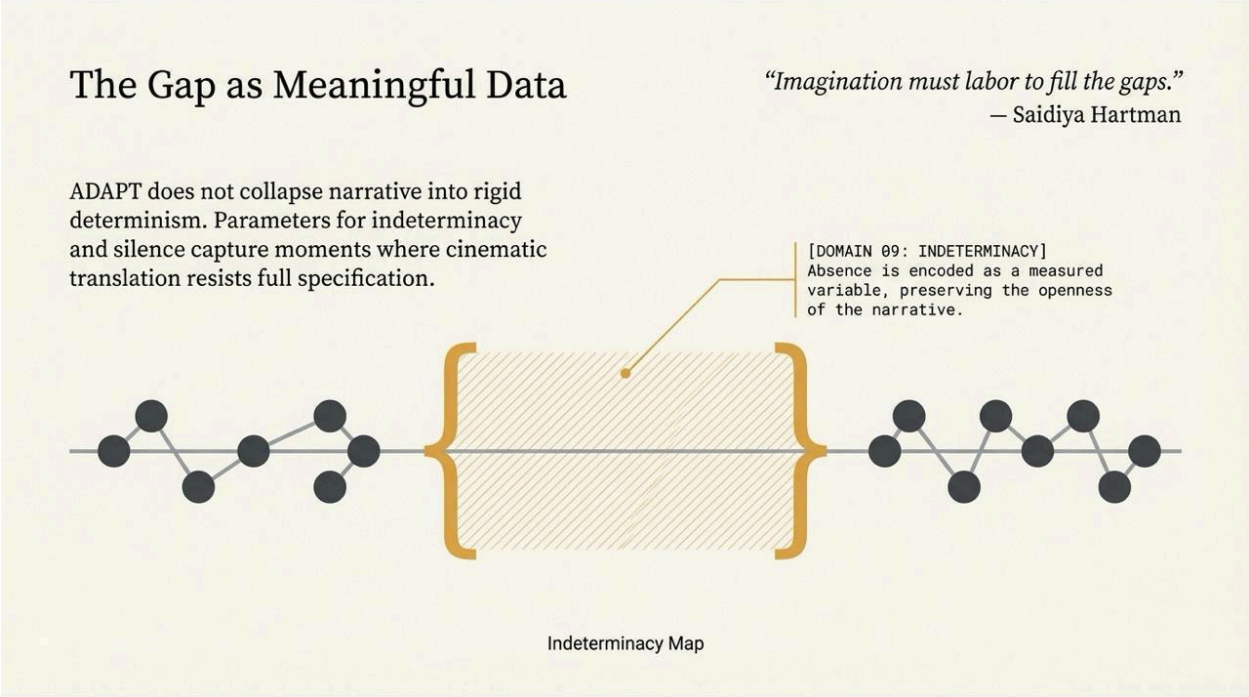

Figure 12. Indeterminacy map encoding narrative gaps as measured variables within Domain 09, after Hartman's critical fabulation framework.

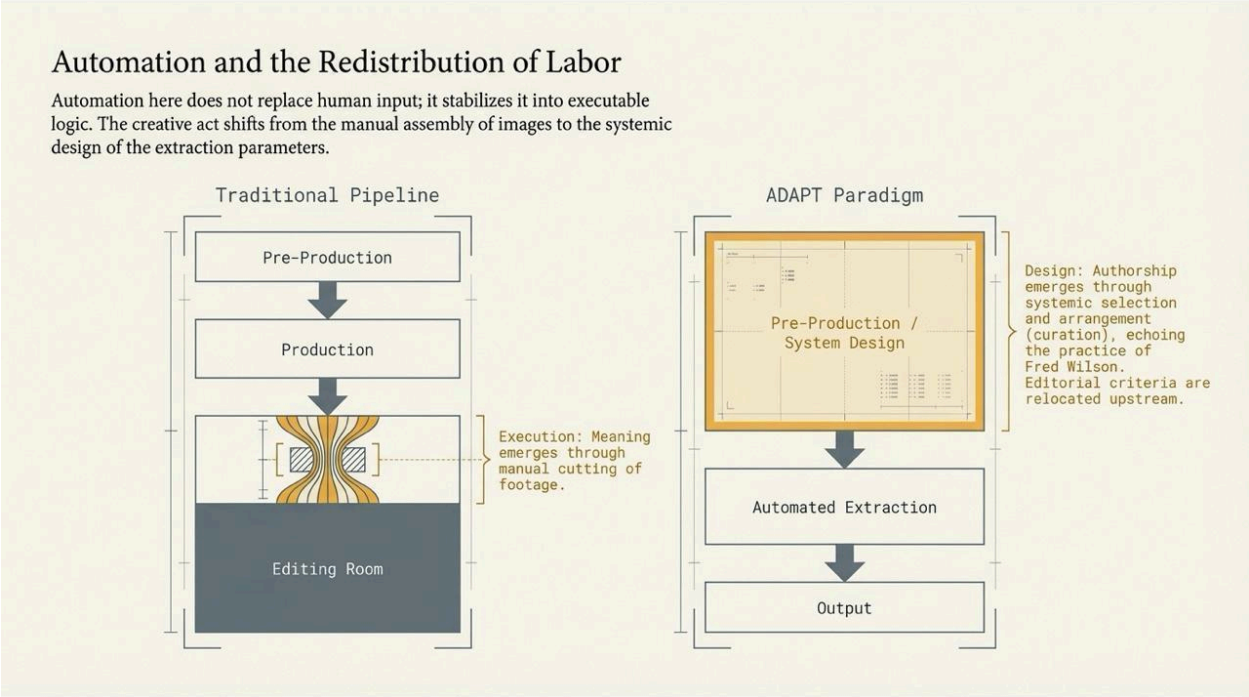

Figure 13. Redistribution of creative labor: traditional editing-room pipeline compared with ADAPT's upstream system-design paradigm.

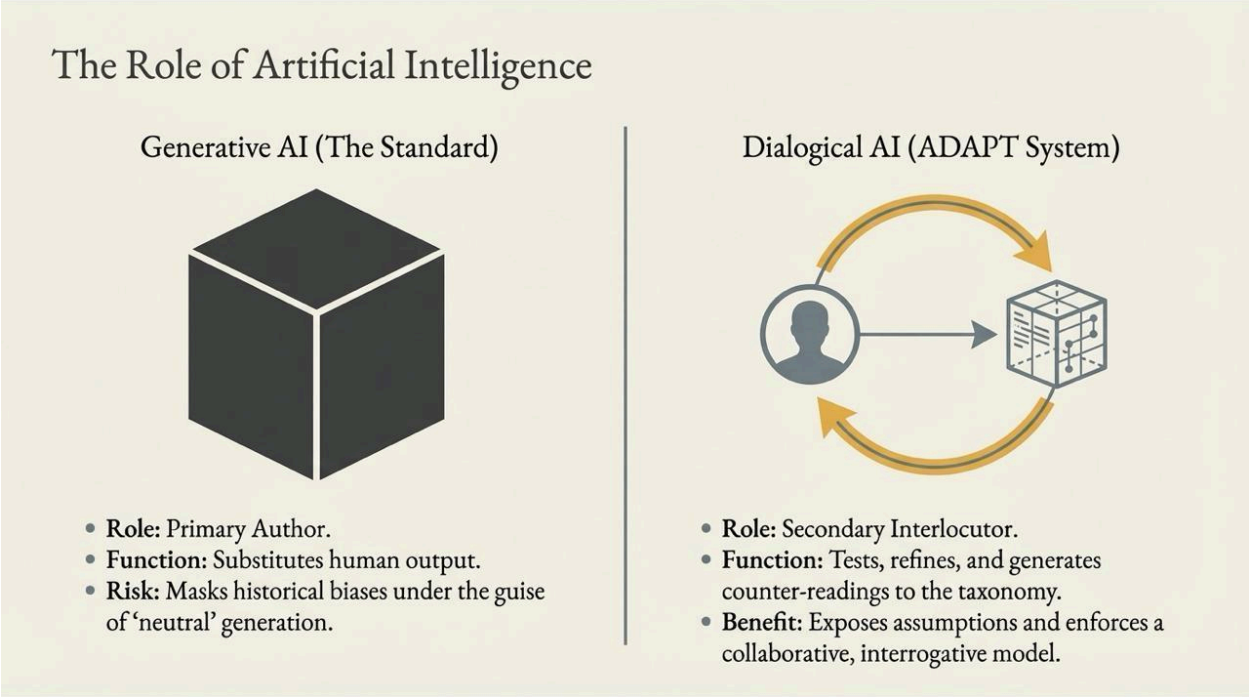

Figure 14. AI repositioned from generative author (primary) to dialogical interlocutor (secondary) within the ADAPT extraction model.

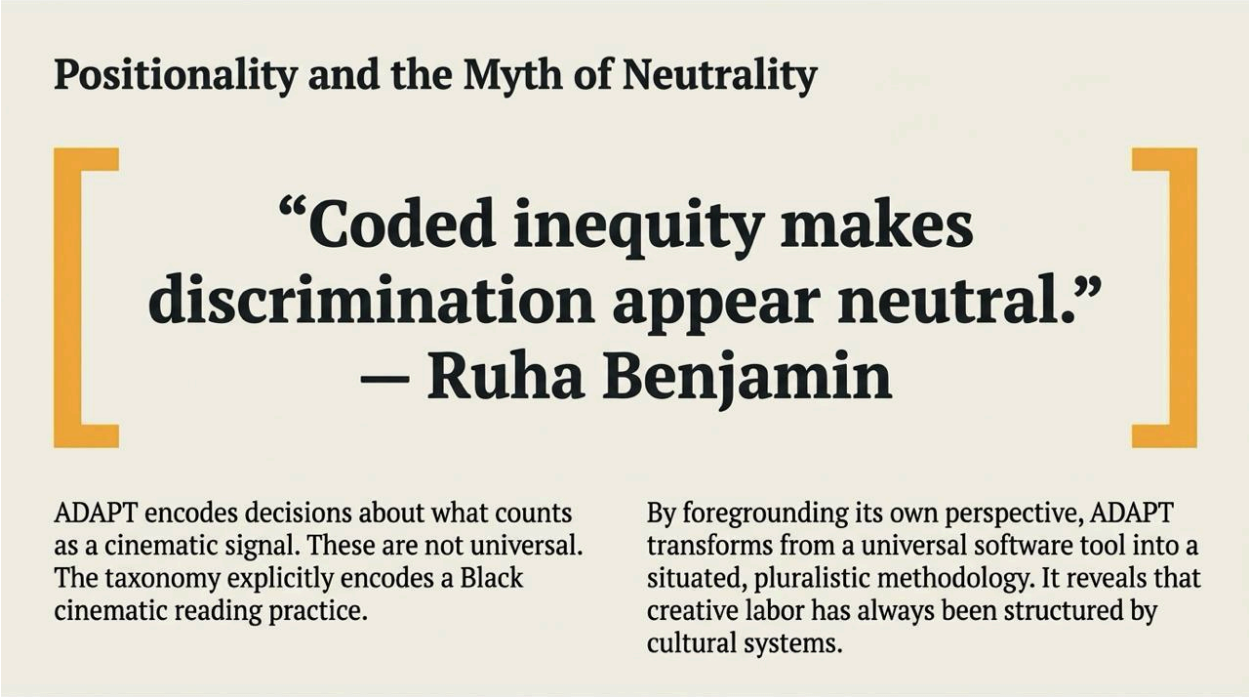

Figure 15. Positionality and the myth of neutrality: ADAPT as a situated, Black cinematic reading practice, after Benjamin's coded inequity framework.

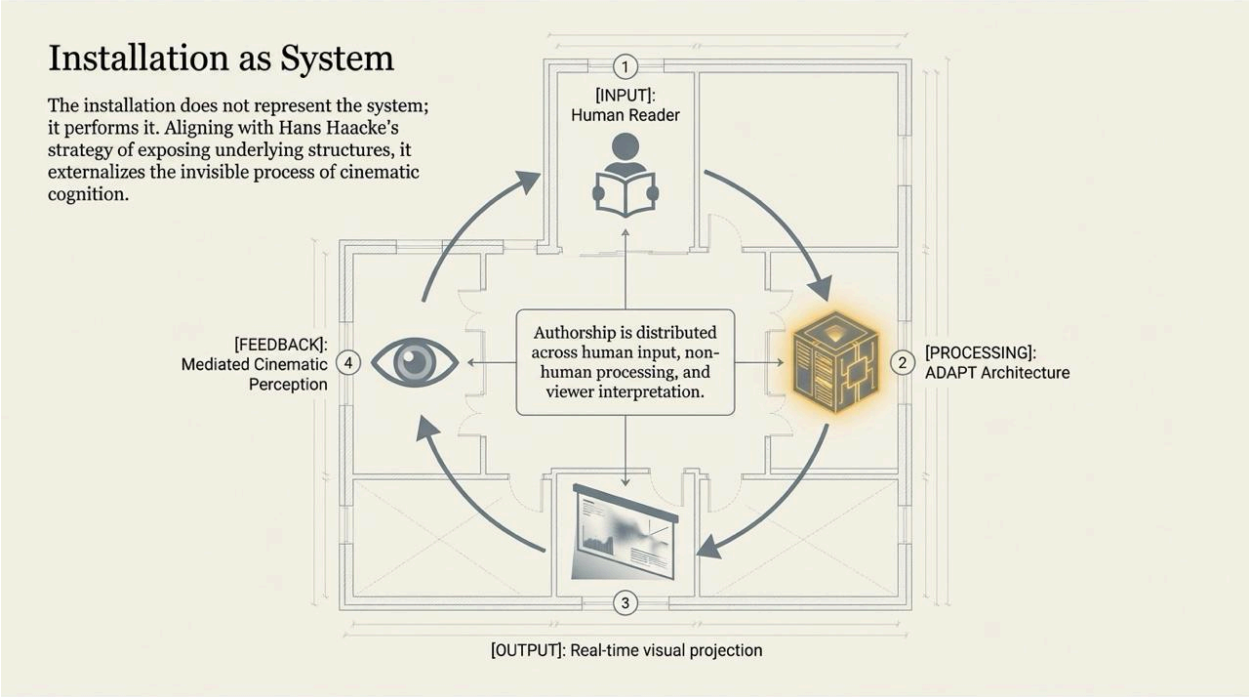

Figure 16. Installation as system: four-phase cycle (input, processing, output, feedback) externalizing cinematic cognition, after Haacke's systems critique.

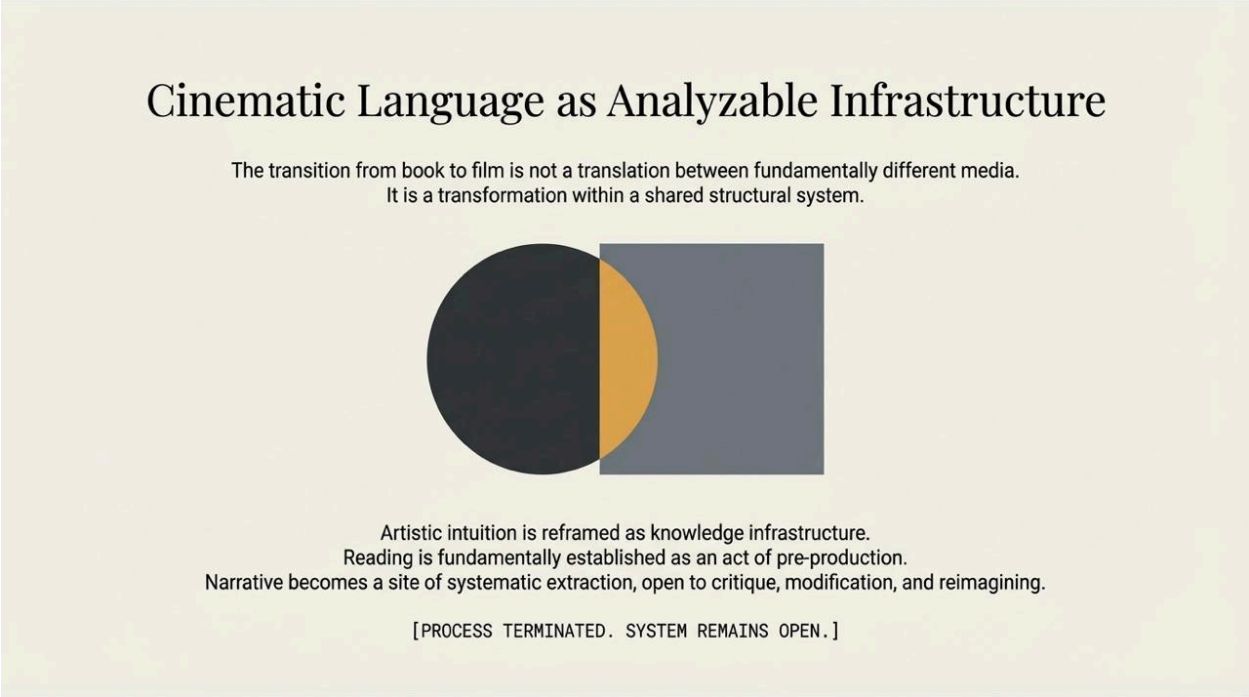

Figure 17. Closing diagram: narrative and cinema as transformations within a shared structural system rather than translations between separate media.

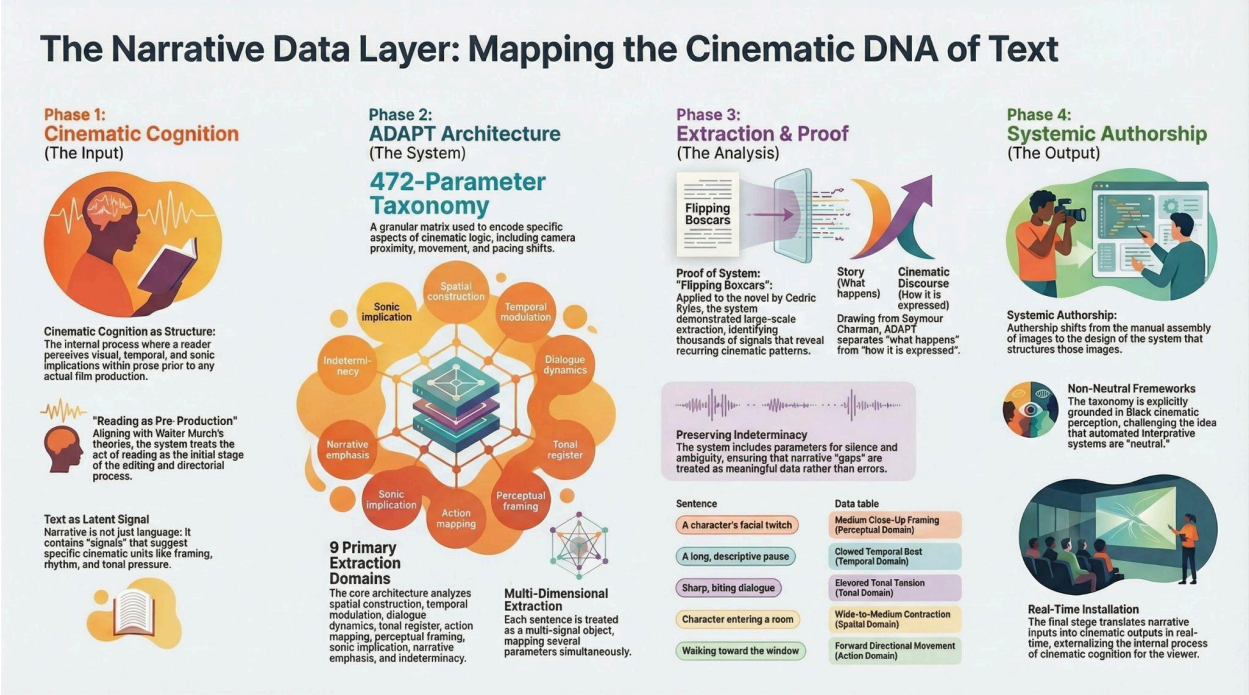

Figure 18. The Narrative Data Layer: a four-phase model demonstrating how cinematic cognition, ADAPT's 472+-parameter architecture, multi-domain signal extraction, and systemic authorship operate as a continuous pipeline—positioning narrative and cinema as expressions of a unified, pre-embedded cinematic structure rather than distinct, sequential media processes.

Glossary of Terminology

ADAPT — ADAPT (Adaptation Data Architecture for Production Translation) is a computational and conceptual system designed to translate **prose-based narrative texts** into **structured cinematic representations**. It operates by decomposing narrative into multiple concurrent analytical domains, including spatial framing, temporal pacing, dialogue structure, tonal modulation, and sonic implication. These domains are formalized through a 472+-parameter taxonomy, allowing narrative to be processed as a multi-signal system rather than a linear text.

Key clarification:

- Input = prose narrative (novels, scripts, oral storytelling transcripts)
- Output = film-ready structural data (shots, beats, tone, pacing)

Narrative Data Layer — The Narrative Data Layer is a structured ontology composed of 472+ parameters that capture distinct dimensions of cinematic logic embedded within narrative text. These parameters span domains such as:

- spatial construction (shot scale, orientation)
- temporal modulation (beat duration, pacing shifts)
- dialogue dynamics (exchange patterns, subtext density)
- tonal register (tension, neutrality, anticipation)
- sonic implication (diegetic vs implied sound)

Critical distinction:

- **Signal** = an observed cinematic implication in text
  (e.g., “he hesitates” → temporal slowing)
- **Parameter** = the classification axis used to encode that signal
  (e.g., temporal_pacing = decelerated)

Signals are **phenomena**
Parameters are **measurement structures**

Cinematic Syntax — Cinematic syntax refers to the structured relationships between visual, temporal, and sonic elements that produce meaning in film. Within ADAPT, cinematic syntax is understood as **pre-embedded within narrative prose**, where linguistic constructs imply:

- framing (wide, close-up)
- sequencing (cut vs continuous action)
- rhythm (fast vs slow pacing)

It is embedded in **descriptive prose, dialogue, and narrative structure**, particularly in texts that contain spatial action, character interaction, or temporal progression.

Cinematic Cognition — Cinematic cognition is the process by which narrative text is perceived as potential film prior to production.

Within this thesis, it operates across two distinct but interacting domains:

- **Human cognition**: interpretive, experiential, culturally situated perception of narrative as image, time, and sound
- **Computational cognition (AI)**: pattern-based approximation of these perceptual structures through parameterized extraction

ADAPT does not claim equivalence between the two. Instead, it formalizes **human cinematic cognition into structured representations** that AI can process and test.

Signal — A signal is a discrete unit of cinematic implication embedded within narrative text.

Example:

> "She slowly turns toward the door"

Produces signals:

- temporal: slowed pacing
- spatial: directional movement
- tonal: anticipatory tension

Signals are identified through:

- linguistic cues (verbs, adjectives)
- narrative context
- relational positioning within scenes

They are not subjective guesses but **repeatable pattern recognitions**.

Extraction — Extraction is the process of identifying cinematic signals within narrative text. It is the first stage of the ADAPT pipeline and operates prior to formal categorization.

Extraction involves detecting:

- movement cues (e.g., walking, turning)
- temporal indicators (e.g., pauses, acceleration)
- tonal shifts (e.g., tension, calm)
- spatial relationships (e.g., proximity, orientation)
- dialogue dynamics (e.g., exchange, interruption)

Extraction is distinct from classification. It answers the question: **what cinematic implications are present in the text?**

This distinction is critical. Extraction isolates signals without yet assigning them to a structured system. It is a process of detection rather than interpretation.

Translation — Translation is the process of converting classified signals into filmic structures, including:

- shot types (wide, close-up, OTS)
- camera movement (tracking, static)
- editing logic (cut, dissolve, continuity)
- sound design cues
- scene rhythm and sequencing

This is where narrative becomes **production-ready cinematic instruction**

Systemic Authorship — Systemic authorship refers to authorship that emerges through the design of rule-based systems that structure creative output.

In ADAPT, authorship resides in:

- the design of the 472+-parameter taxonomy
- the logic governing signal extraction
- the constraints that shape translation into film

Unlike traditional authorship (writing or directing), systemic authorship operates at the level of **framework design rather than content creation**.

Latent Structure — Latent structure refers to cinematic organization that exists implicitly within narrative text but is not explicitly described.

"Latent" here does not mean hidden in a mystical sense. It means:

- not directly specified
- inferable through consistent patterns
- activated through perception or system extraction

Example:
A paragraph with dialogue alternation → latent shot-reverse-shot structure

Indeterminacy — Indeterminacy refers to moments in narrative where cinematic translation is not fully specified and remains open to interpretation.

Within ADAPT, indeterminacy is not treated as error but as:

- a preserved ambiguity
- a parameterized absence
- a site of interpretive flexibility

This aligns with art theory (Hartman, etc.) but is operationalized as a **tracked condition within the system.**