Representational Complexity: Documentary, Dramatization, and the Grammaticalization of Historical Agency

Flyxion
Independent Researcher

November 2026

Abstract

A short documentary montage can require archival discovery, authentication, rights clearance, restoration, and editorial assembly out of all proportion to its running time, yet the finished sequence looks continuous and simple. This essay asks how the informational cost of a historical representation relates to the commitments it makes and to the ease with which an audience recognizes what it is being told. It adapts the representational complexity of Rommens, Traversa, Ferraz de Arruda, and Moreno, who show for networked dynamical systems that no exact structure-plus-rule description can fall below the intrinsic description length of the dynamics, that unrestricted languages are therefore informationally interchangeable, and that meaningful comparison begins only when the admissible structures and rules are restricted. The adaptation replaces the exactly recoverable dynamical map with an evidence-supported answer profile over a declared class of historical questions, and replaces exact reproduction with a tolerance on query-relative distortion, so that the central object becomes a constrained minimization of description length subject to bounded distortion. Documentary, dramatization, and grammaticalization are then treated as three families of representational constraints, each distributing information differently among source records, narrative structure, and interpretive rules, rather than as three rungs of increasing complexity. The principal results are a non-monotonicity proposition for description length across the three families, a default-filling proposition showing that overreach can be strictly cheaper to encode than explicit reservation, a complexity inversion in which a familiar narrative grammar lowers the cost of communicating a story while raising the cost of justifying its implications, and a non-injectivity theorem showing that a grammaticalization operator destroys exactly those distinctions on which a historical question depends whenever causally different records map to the same output. A case study of the opening sequence of an artificial-intelligence documentary, together with a provenance audit, a commitment audit, a counterfactual editing test, and a role-substitution test, indicates how the framework can be applied without claiming to measure the Kolmogorov complexity of any film.

Keywords: representational complexity, algorithmic information theory, documentary theory, narrative grammar, query-relative sufficiency, historical representation

1 The Difficulty of Making Five Minutes

A short documentary montage can demand a disproportionate quantity of intellectual, legal, technical, and editorial labor. Five minutes of screen time may incorporate material produced across several decades, by unrelated institutions, for incompatible purposes, under different technological and legal regimes. A film clip made for broadcast in one decade sits beside a laboratory demonstration recorded for an internal audience in another, beside a child interviewed for a chess program, beside a scene from a feature film whose rights belong to a studio that has no interest in the documentary's argument. Each fragment has to be found, dated, authenticated, acquired, rendered technically usable, cleared for the territories in which the work will be shown, trimmed at both ends, and placed in an order. The finished sequence presents none of this. It presents continuity. The apparent smoothness of the result is an achievement of construction and not a property inherited from the sources, which were never continuous with one another.

The motivating example for this essay is the opening of a recent feature documentary on artificial intelligence, *AI: Probably Nothing to Worry About*. As described from a viewing of the film, the opening juxtaposes excerpts of fictional artificial intelligence drawn from *WarGames*, *Short Circuit*, and *RoboCop* with historical technological demonstrations, possibly including Claude Shannon's maze-solving machine Theseus, and with an archival chess interview featuring the young Demis Hassabis. These identifications and the precise order of the sequence require verification against the film, and Section 11 sets out the verification the analysis depends on. The case is introduced here not as proof of the proposed theory but because it makes a problem visible: the screen duration of a finished representation is a poor guide to the number of independent evidentiary and editorial decisions it contains, and an equally poor guide to the amount of historical information it actually carries.

The first claim can be stated in a line. If T_{screen} denotes the running time of a representation and $C_{\text{construction}}$ the effort required to produce it, then

$$T_{\text{screen}} \not\propto C_{\text{construction}},$$

and a brief sequence can contain a large number of independent decisions, each of which could have been made otherwise. The second claim is less obvious and carries most of the weight of this essay. The effort required to construct a representation, the amount of historical information it conveys, and the number of claims it invites an audience to accept are three different quantities, and they can diverge in either direction. A sequence that is arduous to assemble may assert little. A sequence that is trivial to assemble, because it follows a familiar story shape, may assert a great deal.

The question that organizes the whole study is the following. How do the costs, commitments, and evidentiary obligations of representation change as historical records are transformed into documentary narratives, dramatized events, and reusable heroic grammars? Three transformations are distinguished. Documentary constructs a narrative by selecting, authenticating, arranging, and contextualizing existing records. Dramatization constructs a plausible sequence by supplementing incomplete records with invented particulars. Grammaticalization, the term proposed here for the third transformation, constructs reusable narrative roles by turning historically contingent persons and events into instances of rec-

ognizable character types, causal patterns, and moral trajectories. The word is borrowed from linguistics, where it names the historical process by which lexical expressions acquire grammatical function, and the borrowing is defended in Section 6. For now it is enough that the proposed usage concerns the formation and application of a generative narrative grammar and not merely the observation that stories contain stereotypes.

The study is not a comparison of documentary and fiction as genres. Genre classification asks what a work is called and what conventions it declares. The object of analysis here is the set of transformations that a representation performs upon its source material, and the information that those transformations add, remove, or relocate. This shift of attention, from genre to operation, is what makes a mathematical treatment possible. It is also what separates the argument from the large and valuable literature on documentary modes, emplotment, and mythologization, with which Section 13.4 engages and from which the present account borrows freely.

The essay proceeds in four movements. The first states the construction problem, introduces the information-theoretic starting point supplied by the work of Rommens and colleagues, and builds the formal setting of evidence, questions, and answers. The second defines the three families of representational constraints. The third develops the theory: a rate-distortion style optimization, the complexity inversion that is the principal contribution, four propositions with their proofs and limits, and an account of hidden narrative structure. The fourth applies the framework to the documentary case, specifies operational tests, and states plainly what the mathematics establishes and what it merely suggests.

2 Representation as an Allocation of Information

2.1 Structure, rule, and the floor

Rommens, Traversa, Ferraz de Arruda, and Moreno ask how much information is required to represent a dynamical system in terms of an interaction structure and an evolution rule [16]. Their system is a finite Boolean dynamical map $F : \{0,1\}^N \to \{0,1\}^N$, identified with its truth table, and a representation of it is a pair (S, D) in which S is a finite string encoding an interaction structure, D is a finite string encoding a partial computable function, and $D(S) = F$. They define the representational complexity of an exact representation by

$$\mathrm{RC}(S, D) = \mathsf{K}(S, D) - \mathsf{K}(F),$$

where K denotes prefix-free Kolmogorov complexity [11]. Because a fixed evaluator can decode the pair, run D on S, and output F, the inequality $\mathsf{K}(F) \leq^{+} \mathsf{K}(S, D)$ holds, so representational complexity is non-negative up to an additive constant: no exact representation can describe the map in fewer bits than its intrinsic algorithmic content. Their Theorem 3 establishes that the equation-based representation, in which the structure is empty and the rule is a constant function that outputs F, attains this floor, and that any structure attaining it must itself be computable from the dynamical map up to an additive constant. A structure that cannot be recovered from the map therefore incurs a positive cost. Explicit structure is, in their phrase, an additional modeling commitment, justified not by beating $\mathsf{K}(F)$ but by what it provides in interpretability, mechanistic organization, or observability.

Two further features of their framework matter for what follows. The first is the chain rule

decomposition $\mathsf{K}(S, D) =^{+} \mathsf{K}(S) + \mathsf{K}(D \mid S^{*})$, which shows that the cost of a representation is distributed between what is paid in the structure and what remains to be paid in the rule given the structure. Information absent from a skeleton is not lost; it migrates into the rule. The second is the distinction between unrestricted and restricted representations. In the unrestricted setting, the rule may be any computable program, so that information missing from the skeleton can always be hidden inside it. Every structural language, whether graph, hypergraph, or simplicial complex, can then reach the floor, and expressiveness alone cannot distinguish them. Informational comparison becomes meaningful only in the restricted setting, where the structure must belong to a stated language and the rule to a stated admissible class. In that setting the authors identify conditions under which graph and hypergraph descriptions are informationally equivalent, and they show how graph-preferred, hypergraph-preferred, and mixed regimes can emerge when those conditions fail. Throughout they separate what is proven from what is merely suggested by explicit description-length estimates, because Kolmogorov complexity is not computable.

2.2 The narrative analogue

The narrative analogue is direct enough to state before any formalism. A filmmaker who wishes an audience to believe that two events were connected can encode the connection in several places. The connection can be carried by the selection and order of footage, so that it lives in the arrangement. It can be carried by explicit narration, so that it lives in a spoken sentence. It can be carried by a reconstructed scene in which characters say that it was so. Or it can be carried by a conventional character role, so that the audience supplies the connection from its own familiarity with the type: the prodigy's childhood foreshadows the genius, the dissenter's isolation foreshadows the vindication. These are four allocations of the same proposition among arrangement, narration, staging, and interpretive rule. They differ in what they cost to produce, in what they cost the audience to recognize, and in what the audience can check. The Rommens framework supplies a vocabulary for exactly this kind of relocation, because the central lesson of their paper is that information can move between an explicit representation and the rule by which it is read without changing what is ultimately conveyed.

The analogy is not an identity, and the differences must be held in view from the start. The target of a dynamical representation is a single exactly specified map. The target of a historical representation is a body of evidence that is incomplete, heterogeneous, and open to dispute, and a documentary is under no obligation to reproduce the entire historical world. The comparison of representations must therefore be made relative to questions that the representation declares itself to answer, and it must tolerate controlled departures from the evidence instead of requiring exact reproduction. The theorems of the Rommens paper concern exact representations of dynamical maps and cannot simply be transferred. What can be transferred is the architecture of the argument: a floor defined by the intrinsic content of the target, interchangeability under unrestricted encodings, and the emergence of real differences only under stated restrictions. Section 3 builds the corresponding objects for historical evidence, and the adaptation is marked as an adaptation wherever it is used.

2.3 What the comparison is and is not

Three quantities have so far been named without being separated, and the separation is the organizing discipline of the essay. The *description length* of a representation is the number of bits needed to specify it in a fixed language under a fixed coding scheme. The *commitment* of a representation is the set of propositions an audience is reasonably invited to infer from it, in particular those that exceed what the evidence supports. The *recognition cost* is the effort an audience needs to understand what it is being told, which depends on what the audience already knows. A narrative grammar that the audience already possesses can make the first and third small while leaving the second large. The proposition that this can happen, and the conditions under which it must, is developed in Sections 8 and 9. It is worth stating now that the proposition concerns description length, commitment, and recognition, and says nothing about production expenditure, artistic quality, or historical adequacy, none of which are being ranked.

3 Evidence, Questions, and Answers

3.1 Events, records, and representations

Let a historical event or process be denoted W, let the surviving evidentiary record be E, and let a representation be Y. The record is not the event, because recording is selective, preservation is contingent, and later access depends on institutional and technical conditions. A representation is built from the evidence together with background knowledge K and the filmmaker's operative choices θ, which may be written $Y = T(E; K, \theta)$. This prevents the mistaken picture of documentary as a transparent mapping from historical reality to an audience: both the record and the representation are intermediate objects, each produced by a selective process, and each susceptible to analysis in its own right.

For the purposes of a formal treatment, E is taken to be a finite set of records, each with a provenance, a date or an estimate of one, a rights status where ascertainable, and the content that the record makes available. Nothing in the framework requires the records to be films; they may equally be letters, transcripts, minutes, photographs, or measurements.

3.2 Questions and answer profiles

The decisive move of the adaptation is to evaluate representations by the answers they induce to historical questions and not by the proportion of authentic material they contain. Let $\mathcal{Q}$ be a class of historical questions. The class is declared, in the sense that a representation is held responsible for answering the questions within its stated scope and is not penalized for silence outside it. For each $q \in \mathcal{Q}$ let $\mathcal{A}_q$ be a finite set of possible answers that includes a distinguished element $\perp$ denoting that the question is unsettled by the evidence. The evidence supports an answer $a_E(q) \in \mathcal{A}_q$, which is $\perp$ when the evidence, together with justified background assumptions, does not determine the matter. A historical question might be whether two events were causally related, whether a person held a stated belief at a stated time, whether an observed behavior was typical of a group, or whether a later achievement was anticipated by an earlier circumstance. The *answer profile* of the evidence over a finite scope $Q \subseteq \mathcal{Q}$ is the finite string $a_E|_Q = (a_E(q))_{q \in Q}$.

A representation Y is read by an audience. Let I denote the interpretation map that takes a representation to the answers it induces, so that $a_Y(q) = I(Y)(q)$. For the formal analysis I is idealized as a fixed computable function that stands for a competent, culturally situated viewer; Section 12 treats the empirical question of how real viewers depart from this idealization. The map I is part of what a representation language specifies, because what counts as an answer induced by an arrangement of footage depends on editing conventions and, for grammaticalized representations, on the grammar the viewer holds.

3.3 Distortion

The discrepancy between what the evidence supports and what a representation induces is measured question by question. For each q let $d_q : \mathcal{A}_q \times \mathcal{A}_q \to [0, \infty)$ satisfy $d_q(a, a) = 0$. The measure is deliberately allowed to be asymmetric. Asserting a definite answer where the evidence says $\perp$ is overreach, and answering $\perp$ where the evidence is definite is omission. A declared scope and asymmetric costs together allow a documentary to leave out an enormous amount of irrelevant information without being penalized, while holding it responsible when its constructions alter answers within scope. With question weights $w_q \geq 0$ the query-relative distortion is

$$\mathcal{D}_Q(E, Y) = \sum_{q \in Q} w_q \, d_q\big(a_E(q), a_Y(q)\big).$$

Distortion zero on Q means that the representation induces exactly the evidentially supported answers on Q, including the answer that the matter is unsettled. This is a stronger and more informative requirement than factual fidelity alone. Three kinds of fidelity are worth distinguishing. A representation may reproduce authentic footage while suggesting a causal relation the record does not support, which is a failure of causal fidelity that leaves factual fidelity intact. A dramatization may invent dialogue while accurately conveying a well-established institutional conflict, which preserves interpretive fidelity at the price of factual particulars. Fidelity is therefore multidimensional and cannot be inferred from the proportion of authentic material used. The query-relative distortion absorbs all three, because each is a way in which the answers induced by Y can depart from the answers supported by E.

3.4 Description length

A representation is an element of a representation language $\mathcal{M}$, a set of finite strings equipped with an interpretation map $I_{\mathcal{M}}$ and a computable prefix-free code length $\mathsf{L}_{\mathcal{M}}$. The code length is a fixed coding convention: a count of the bits needed to specify the representation under stated conventions for archive pointers, trim points, orderings, scene specifications, and grammar productions. It is not the Kolmogorov complexity of any film, which is not computable and which no analysis here claims to measure. This follows the discipline of the Rommens paper, which moves from the incomputable ideal to explicit description-length estimates and keeps the two distinct. The next definition collects the objects.

Definition 3.1 (Historical representation problem)**.** A historical representation problem is a tuple $(E, Q, \mathcal{M})$ consisting of an evidence set E, a finite declared scope $Q \subseteq \mathcal{Q}$ with weights w_q and distortion measures d_q, and a representation language $\mathcal{M}$ with interpretation

$I_{\mathcal{M}}$ and code length $\mathsf{L}_{\mathcal{M}}$. A representation $Y \in \mathcal{M}$ has description length $\mathsf{L}_{\mathcal{M}}(Y)$ and distortion $\mathcal{D}_Q(E, Y)$.

3.5 The floor and its interchangeability

The first result transfers the central structure of the Rommens paper to the present setting. Its proof is elementary, and its force lies in what it licenses the rest of the argument to assume.

Proposition 3.2 (Answer-profile floor)**.** *Let $Y \in \mathcal{M}$ satisfy $I_{\mathcal{M}}(Y)|_Q = a_E|_Q$, that is, zero distortion on Q when each d_q vanishes only on the diagonal and each $w_q > 0$. If $I_{\mathcal{M}}$ is computable and $\mathsf{L}_{\mathcal{M}}$ is a prefix-free code length for $\mathcal{M}$, then*

$$\mathsf{K}(a_E|_Q) \ \leq^+ \ \mathsf{L}_{\mathcal{M}}(Y).$$

Proof. A fixed program decodes Y from its code word, applies $I_{\mathcal{M}}$, restricts the result to Q, and outputs $a_E|_Q$. The program has length $\mathsf{L}_{\mathcal{M}}(Y)$ plus a constant depending only on $\mathcal{M}$ and the universal machine. □

The quantity $\mathsf{K}(a_E|_Q)$ plays the role that $\mathsf{K}(F)$ plays in the Rommens paper. It is the intrinsic algorithmic content of what the evidence supports about the declared questions, and no exact representation of that content can be shorter. The corresponding excess may be called the *narrative representational complexity*

$$\mathrm{RC}_Q(Y) = \mathsf{L}_{\mathcal{M}}(Y) - \mathsf{K}(a_E|_Q),$$

which is non-negative up to an additive constant for exact representations. It measures how much more description a given language spends than the answers themselves require, and it is the correct place to see that elaborate representations are not thereby more informative: a representation may spend many bits and induce few supported answers.

Two features of this floor depart from the dynamical case and must be noted. The target is an answer profile rather than a map, so it changes with the evidence set, and it includes the value $\perp$, so that a representation attains the floor only by conveying uncertainty where the evidence leaves matters unsettled. Furthermore, when the scope is narrower than the whole record, the floor is lower, which is the formal expression of the principle that a documentary may omit irrelevant information without penalty.

The interchangeability result carries over as well. If $\mathcal{M}_\infty$ is the language of all computable decoders paired with their inputs, then any representation of length $\mathsf{K}(a_E|_Q)$ up to a constant attains the floor, and it is immaterial to the length whether the information sits in footage, narration, staged particulars, or an interpretive convention, since in an unrestricted language these are all strings. Expressive power therefore does not distinguish documentary from dramatization from grammaticalization. A voiceover can encode anything; a conventional role can encode anything; the question of which is better cannot be settled by observing that each can say what the others say. As in the Rommens paper, differences appear only when the languages are restricted, and the restrictions that matter for narrative are the subject of the sections on the three families.

4 Documentary as Evidence-Constrained Selection

4.1 The archival search problem

Documentary production begins with an incomplete, heterogeneous, and frequently inaccessible collection of records. Relevant footage must be discovered, identified, dated, authenticated, acquired, and made technically usable. Archival research is not a passive retrieval operation. It involves uncertain search costs, incomplete indexing, reconstruction of provenance, and decisions about whether the footage available adequately supports the claim the filmmaker intends to make. A clip may be known to exist and be impossible to locate, or be locatable and impossible to date, or be dated and impossible to clear. These contingencies belong to the cost of documentary and not to its content, which is why a viewer cannot read them off the finished film.

Rights clearance adds a second layer of fragmentation. A single work may combine public historical material, broadcaster archives, privately controlled film excerpts, and footage subject to separate distribution agreements, and the complexity of clearance depends on the number and structure of the rights dependencies and not simply on the total duration of borrowed material. This is also why international distribution may be harder than domestic exhibition, though nothing in the argument assumes that any particular film's geographical restrictions arose from its archival excerpts.

4.2 The restricted language of documentary

In the formal setting, documentary is a restricted representation language $\mathcal{M}_D$ whose characteristic operations are selecting, trimming, ordering, juxtaposing, and captioning records. A representation $Y \in \mathcal{M}_D$ is a finite sequence of references to records in E, each with trim points, together with an arrangement and captions. The defining restriction is source linkage: every content-bearing symbol of Y is a pointer into the evidence or a statement that the evidence supports. This is an admissibility constraint on the structure, in the sense of the restricted representations of Rommens et al., and it has a precise informational consequence. Because each pointer must address a record of an archive of size N_A, a pointer costs on the order of $\log_2 N_A$ bits, each trim point costs a number of bits that depends on the length of the source, and an ordering of n elements costs up to $\log_2 n!$ bits of editorial choice. The description length of a documentary representation therefore grows with the number of independent records and with the freedom exercised in arranging them, and it is largely insensitive to the screen time those records occupy. This is the formal content of the opening observation that duration is a poor guide to construction.

A decomposition of the construction burden into search, authentication, clearance, editing, and verification, $C_D = C_{\text{search}} + C_{\text{auth}} + C_{\text{clear}} + C_{\text{edit}} + C_{\text{verify}}$, is useful as an analytical inventory. The categories are not necessarily additive in any empirical accounting, since search and verification overlap and editorial changes can generate additional clearance work, and they are not description-length quantities. They concern labor and legal exposure and not the informational measure that the later sections develop. The informational and the operational accounts are kept separate throughout.

4.3 Editing as inference

Documentary montage introduces relationships between records that did not necessarily exist when those records were created. If two authentic excerpts e_1 and e_2 are placed consecutively, their juxtaposition may lead an audience to infer a relation $R(e_1, e_2)$, whether of cause, resemblance, succession, or answer. The inferred relation is a property of the edited representation and not necessarily of the underlying history. The effect has a long history in film theory, associated with the editing experiments of Kuleshov [10], and it has a close relation to conversational implicature in the sense of Grice [8]: what is communicated by adjacency exceeds what is literally contained in either component, and the excess is governed by shared expectations about how sequences are organized.

Proposition 4.1 (Authentic-component non-entailment)**.** *The authenticity of every component of a representation does not entail that the relations the representation invites an audience to infer are supported by the evidence. There exist a representation $Y \in \mathcal{M}_D$, an evidence set E, and a scope Q such that every record referenced by Y is authentic and $\mathcal{D}_Q(E, Y) > 0$.*

Proof. Let $E = \{e_1, e_2\}$ with e_1 a recording of a speaker predicting an outcome and e_2 a recording of that outcome, made independently, and let q be the question whether the outcome followed from the speaker's prediction in the sense of a causal or authorial relation. Suppose the evidence does not settle q, so that $a_E(q) = \bot$. The arrangement $Y = (e_1, e_2)$ is a faithful selection of authentic records, and a viewer competent in the conventions of documentary montage interprets it as answering q affirmatively, so $a_Y(q) \neq \bot$ and, with $d_q(\bot, a) > 0$ for $a \neq \bot$, the distortion is positive. □

The proposition limits what a provenance audit can establish. It can establish that components are authentic and attributable, which is necessary. It cannot establish that the composition is valid, which concerns the relations the arrangement induces. The distinction between authenticity of components and validity of composition is the first of the three distinctions that organize the rest of the study.

5 Dramatization as Constrained Reconstruction

5.1 Incompleteness and underdetermination

Historical evidence rarely determines the sequence of dialogue, gesture, spatial arrangement, timing, motivation, and interpersonal response required to stage a scene. A dramatization must select particulars from a larger set of possibilities compatible with the surviving evidence. Let $\Pi(E)$ denote the set of complete scene specifications that do not contradict E. The dramatist chooses an element of $\Pi(E)$, and the choice is not determined by the evidence; it is underdetermined by it, which is the technical content of the claim that dramatization supplies what the record lacks.

Three kinds of particular are worth separating. Some are determined by the evidence, such as the date of a meeting that the minutes record. Some are plausible but unverified, such as the arrangement of a table that the record does not describe. Some contradict or substantially exceed the evidence, such as a confrontation that no source reports. Invention

is not automatically falsification. Its epistemic status depends on how it is presented, what claims it communicates, and whether the audience is given reason to mistake it for established history. This dependence can be stated within the framework: in the dramatization language $\mathcal{M}_R$, each particular carries a presentation status, whether attested, reconstructed, or invented, and the interpretation map I_R treats a particular flagged as reconstructed or invented as inducing $\perp$ on the questions it bears on. Under this convention a thoroughly invented scene can have distortion zero, and a scene whose inventions are presented as attested can have large distortion. What is evaluated is the answer induced, and the induced answer depends on the framing conventions that the language builds in.

5.2 Coherence as an additional constraint

Once a dramatization introduces characters, settings, motivations, and interactions, these particulars must remain mutually consistent across scenes. The language $\mathcal{M}_R$ therefore carries a coherence constraint beyond evidential compatibility: the chosen particulars must be jointly satisfiable and must support continuous characterization. The added cost is not merely the number of invented details but the requirement that they form a coherent and historically defensible reconstruction. Where the documentary burden C_D counts search, authentication, clearance, editing, and verification, the dramatic burden may be written $C_R = C_D^* + C_{\text{spec}} + C_{\text{stage}} + C_{\text{cont}}$, where C_D^* represents whatever documentary-style research is retained, without the assumption that every dramatization first produces a complete documentary. As with the documentary decomposition, this is an inventory of labor and not an informational measure.

The informational counterpart is a scene specification cost, in bits, for each particular selected from $\Pi(E)$, together with the cost of enforcing consistency. Two distinctions arise here, between evidentiary underdetermination, which is the lawful freedom of the dramatist within the compatible set, and narrative inconsistency, which is a violation of coherence and is a defect of construction and not of evidence.

5.3 The reconstruction paradox

Greater representational detail can increase apparent realism without increasing evidentiary support. A carefully reconstructed conversation can seem more informative than an archival record while containing more unsupported particulars. The phenomenological completeness of a scene and the historical warrant for it are different properties, and the former is easily mistaken for the latter. In the present framework the paradox has a clean statement: the description length of a reconstruction can greatly exceed that of a documentary on the same evidence while its induced answers on Q are no better supported, so that RC_Q is large even when the audience's felt sense of knowledge is high.

6 Grammaticalization as Structural Substitution

6.1 The term and its justification

Grammaticalization is defined here operationally as the transformation of particular persons, actions, and events into elements governed by a reusable narrative grammar. A historical

individual is not merely portrayed but assigned a role that constrains how past and future events are interpreted. The term is explicitly distinguished from linguistic grammaticalization, in which lexical expressions develop grammatical functions over historical time. The proposed usage concerns narrative role formation and must be defended against existing concepts, among them mythologization in the sense of Barthes [1], emplotment in the sense of White and Ricoeur [22, 14], the functions and roles of Propp [13], typification, and archetypal characterization.

The defense is that the usage here concerns something those terms do not isolate: a generative grammar, applied as an operator to evidence, whose output depends on the evidence only through a role assignment and whose unspecified parts are filled by defaults. A grammar in this sense has productions, slots, defaults, and an assignment procedure, and each of these has an informational signature that can be studied. The term should be retained only if the argument genuinely concerns the formation and application of such a grammar and not simply heroic stereotyping, and the results of Sections 8 and 9 are what the usage is to be held to. Existing scholarship supplies substantial accounts of narrative construction and mythologization, and the claim is not that stories employ archetypes. The proposed novelty lies in treating role assignment as an operation whose commitments, invariances, information losses, and verification burdens can be investigated.

6.2 A heroic grammar

A simplified heroic grammar may contain the productions

$$\begin{aligned} \text{Hero} &\to \text{Origin Trial Recognition Mission}, \\ \text{Origin} &\to \text{Early Talent} \mid \text{Early Adversity}, \\ \text{Trial} &\to \text{Opposition} \mid \text{Failure} \mid \text{Exclusion}, \\ \text{Mission} &\to \text{Discovery} \mid \text{Warning} \mid \text{Rescue}. \end{aligned}$$

These productions are illustrative and not universal. Their purpose is to show how a finite set of conventions can generate many recognizable biographical narratives, and to expose the informational structure of one. A grammar G consists of its productions, a set of slots, and a default filler $\delta_G(s)$ for each slot s that is not explicitly specified. A representation in the language $\mathcal{M}_G$ is a triple (G, α, ϕ) in which α assigns persons and events to positions in the grammar and ϕ gives explicit fillers for some slots. The interpretation I_G expands the grammar, uses ϕ where it is supplied, and uses δ_G elsewhere. The viewer is assumed to hold G, which is what it means for the grammar to be familiar, and the cost of holding it is paid by cultural exposure and not by the individual work.

6.3 Retrospective necessity

Once a person is assigned a heroic role, earlier events may be selected because they appear to anticipate later achievements. An ordinary childhood statement becomes an origin motif. A professional disagreement becomes an early struggle against institutional opposition. A change of opinion becomes a decisive revelation. The narrative makes contingent events appear developmentally necessary. Formally, role assignment is outcome dependent: $\alpha(e; o)$, the position assigned to an earlier event e, is a function of a later outcome o as well as

of the event and its record. Proposition 9.3 shows that this dependence can reverse the significance of a fixed event.

6.4 Authenticity without interpretive fidelity

Grammaticalization does not require fabricated evidence. Authentic footage may be sufficient if the editing, narration, and surrounding sequence lead the audience to interpret it through an established role. This is the most important distinction between grammaticalization and ordinary fictionalization, and it explains why provenance audits cannot detect it. The footage of a child at a chess board is identical whether it is treated as an ordinary biographical record or as the origin scene of a future technological protagonist. What changes is the interpretive rule. The Rommens analogy is at its most useful here: information can migrate between an explicit representation and the rule by which it is interpreted, and in a grammaticalized representation a great deal of what is communicated resides in the rule.

6.5 The compression advantage of familiar roles

Once a narrative grammar is familiar, audiences can infer substantial unstated structure from minimal cues. A heroic representation may require fewer explicit explanations than a historically nuanced account, because the grammar supplies the connective tissue. Grammaticalization thereby increases the organizing power of a representation while potentially decreasing its descriptive specificity. The compression is real, and it is the source of the efficiency that makes such grammars attractive to producers and persuasive to audiences. Sections 8 and 9 show that the same compression is the mechanism by which distinctions required for historical explanation are lost.

6.6 A structural recoverability constraint

The analogue of the second half of the Rommens floor theorem applies to grammaticalization and yields a sharp statement about role assignments. In the dynamical setting, a structure that attains the floor must be computable from the dynamical map up to a constant. The same argument applies to the structural part of a narrative representation, which here consists of the grammar and the role assignment.

Proposition 6.1 (Recoverability of floor-attaining structure)**.** *Let $Y = (G, \alpha, \phi) \in \mathcal{M}_G$ be an exact representation of the answer profile $a = a_E|_Q$ with $\mathsf{L}_{\mathcal{M}}(Y) =^+ \mathsf{K}(a)$. Then $\mathsf{K}(G, \alpha \mid a^*) =^+ 0$: the grammar and the role assignment are computable from the supported answers up to an additive constant.*

Proof. The map $Y \mapsto (a, \langle G, \alpha\rangle)$ is computable, since $a = I_G(Y)|_Q$. Hence $\mathsf{K}(a, \langle G, \alpha\rangle) \leq^+ \mathsf{L}_{\mathcal{M}}(Y) =^+ \mathsf{K}(a)$. Because the joint complexity is at least $\mathsf{K}(a)$, we have $\mathsf{K}(a, \langle G, \alpha\rangle) =^+ \mathsf{K}(a)$, and the chain rule gives $\mathsf{K}(\langle G, \alpha\rangle \mid a^*) =^+ 0$. □

The proof follows the pattern of Theorem 3(iii) of the Rommens paper and is not a result established there. Its meaning for narrative is that a role assignment which cannot be derived from what the evidence supports is not free. A hero grammar applied to a subject whose evidence does not already determine the roles adds structure that costs bits beyond the floor and induces commitments beyond the evidence. Such structure may be

warranted by interpretive or aesthetic purposes, in the same way that explicit structure in the dynamical setting is warranted by interpretability, but it must be recognized as an additional commitment and not as a neutral container.

7 Distortion-Constrained Description Length

7.1 The optimization problem

The adaptation of the restricted representation approach to historical evidence yields a constrained optimization in place of an exact minimum. For a representation language $\mathcal{M}$, a scope Q, and a tolerance $\varepsilon \geq 0$ define

$$\boxed{\mathsf{L}^*_{\mathcal{M}}(E, Q, \varepsilon) = \min_{\substack{Y \in \mathcal{M} \\ \mathcal{D}_Q(E,Y) \leq \varepsilon}} \mathsf{L}_{\mathcal{M}}(Y).}$$

This is the shortest description in the admissible narrative language whose induced answers on the declared questions depart from the evidentially supported ones by no more than the tolerance. It differs from the Rommens formulation in the way stated at the outset: the cost of a representation is minimized subject to a tolerance for historical distortion, and exact reproduction of a dynamical map is not required. It is an adaptation of their restricted-representation approach, and its properties are those of the definition, not theorems of their paper. The family of functions $\varepsilon \mapsto \mathsf{L}^*_{\mathcal{M}}(E, Q, \varepsilon)$ is the narrative analogue of a rate-distortion function in the sense of Shannon [18] and Cover and Thomas [4], and, for individual objects, of the algorithmic rate-distortion function of Vereshchagin and Vitányi [21]. It also resembles the information bottleneck of Tishby, Pereira, and Bialek [20], in which a representation is judged by how much it preserves about a relevant variable, and the minimum description length principle of Rissanen [15], in which models are judged by the total code length of model and data given model.

Proposition 7.1 (Basic properties)**.** *The function $\varepsilon \mapsto \mathsf{L}^*_{\mathcal{M}}(E, Q, \varepsilon)$ is non-increasing, and if the feasible set at $\varepsilon = 0$ is non-empty it is finite and bounded below by $\mathsf{K}(a_E|_Q)$ up to an additive constant whenever $I_{\mathcal{M}}$ is computable. For any ε at least the distortion of the shortest representation in $\mathcal{M}$, it equals that shortest length. The function is not in general convex.*

Proof. Enlarging ε enlarges the feasible set, so the minimum cannot increase. The bound at $\varepsilon = 0$ is Proposition 3.2. The last statement of the first sentence is immediate, and non-convexity is exhibited by any language whose code lengths take two well-separated values on feasible sets that appear at two distinct thresholds. □

The curve is the correct object for comparing the three families. A single number would conceal the trade-off, and the central claim of the theory is that the families occupy different regions of the plane of description length against distortion, with no uniform dominance of one over another.

7.2 The three languages as restriction families

Each family of representation is a language with its own admissibility constraint, interpretation map, and code. The documentary language $\mathcal{M}_D$ restricts content to pointers into E and evidence-supported captions, and its interpretation includes the inferences that arrangement invites. The dramatization language $\mathcal{M}_R$ admits particulars from $\Pi(E)$ under a coherence constraint and a presentation convention. The grammatical language $\mathcal{M}_G$ admits a grammar, a role assignment, and explicit fillers, with defaults for the rest. In terms of the operations they allow, each language contains the previous one: a grammaticalized work can use only authentic footage, and a dramatized work can quote the record. This inclusion concerns permitted operations and implies nothing about cost, because the restrictions imposed by each language on how information may be distributed are different, and it is the restriction, not the expressive reach, that determines the minimal description length. The point is the same as in the Rommens paper, where the passage from unrestricted to restricted languages is exactly the passage from interchangeability to meaningful difference. The three families are therefore best understood as three families of representational constraints, each fixing a way in which information may be distributed among evidence, narrative structure, and interpretive rules, and not as successive levels of complexity.

7.3 Overhead and the cost of structure

The overhead of a particular representation over the intrinsic floor is $\mathrm{RC}_Q(Y) = \mathsf{L}_{\mathcal{M}}(Y) - \mathsf{K}(a_E|_Q)$. The ideal representation in a restricted language has overhead $\mathsf{L}^*_{\mathcal{M}}(E, Q, 0) - \mathsf{K}(a_E|_Q)$, which in general is strictly positive: the restrictions of the language forbid the redistributions that would be needed to reach the floor, as in the restricted setting of Rommens et al. A comparison of two languages on the same evidence and scope therefore reduces, as their graph-hypergraph comparison does, to a difference of finite description lengths, since the floor cancels:

$$\Delta\mathrm{RC}_Q(\mathcal{M}_1, \mathcal{M}_2) = \mathsf{L}^*_{\mathcal{M}_1}(E, Q, 0) - \mathsf{L}^*_{\mathcal{M}_2}(E, Q, 0).$$

The sign of this quantity defines regimes in which one language is preferred, the languages are equivalent, or the other is preferred, and, as in the dynamical case, a preference is meaningful only when the separation exceeds the additive ambiguity of the coding conventions, which in practice means a gap that grows with the size of the problem. The uniformity convention of the Rommens paper, that decoders, pairing maps, and rule catalogues are fixed finite objects independent of problem size, is adopted here for the same reason.

8 The Complexity Inversion

8.1 Cheap to say, expensive to justify

This section states what is proposed as the principal theoretical contribution. A conventional heroic grammar may reduce the description length needed to communicate a story, because the audience already knows the conventions. Writing $\mathsf{L}(Y \mid G)$ for the code length of Y

given the grammar and $\mathsf{L}(Y)$ for its length without it, the relation

$$\mathsf{L}(Y \mid G) < \mathsf{L}(Y)$$

can hold by a wide margin for an appropriate grammar, even though the grammaticalized representation makes stronger claims about causality, motivation, or destiny. The apparent paradox disappears once the informational cost of communicating a narrative is distinguished from the epistemic cost of justifying its implications. The grammar can be inexpensive to reuse while its historical commitments remain expensive to verify. Two results make this precise. The first concerns amortization, and the second concerns the way defaults convert silence into assertion.

8.2 Amortization

A grammar is a shared resource. If a single grammar G is applied across a corpus of n works, then the total cost of the corpus is $\mathsf{L}(G) + \sum_{i=1}^{n} \mathsf{L}(Y_i \mid G)$ and the per-work cost is $\mathsf{L}(G)/n + \overline{\mathsf{L}(Y \mid G)}$, where the bar denotes the average. A documentary representation of the same evidence pays for its pointers, trims, and orderings in every work. Let ℓ_D denote the per-work documentary length and ℓ_G the per-work grammatical length given the grammar, with $\ell_G < \ell_D$. The grammatical approach has the lower total cost for all corpora larger than

$$n^* = \frac{\mathsf{L}(G)}{\ell_D - \ell_G}.$$

The break-even is not a claim about any real corpus; it states the structure of the advantage. A culture that has already paid $\mathsf{L}(G)$, through the repeated exposure of its members to stories of a given shape, has n^* effectively equal to one, and for that audience the grammatical representation is cheaper to receive from the first work. This is the formal content of the observation that a heroic representation may require fewer explicit explanations than a nuanced account.

8.3 Defaults convert silence into assertion

The second result is more consequential. A grammar G assigns to each slot s a default $\delta_G(s)$. A representation need not mention a slot for the slot to be filled in the audience's interpretation. The default is therefore free in description length, since an unspecified slot costs nothing, and it is not free in commitment, since the interpretation I_G induces the default answer on every question that the slot bears on.

Proposition 8.1 (Overreach can be cheaper than reservation)**.** *Suppose a grammar G has a slot s bearing on a question $q \in Q$, with default $\delta_G(s) = \bar{a} \neq \bot$, and suppose the evidence leaves the matter unsettled, $a_E(q) = \bot$. Let Y be a representation that fills s with an explicit indeterminacy marker, so that $I_G(Y)(q) = \bot$, and let Y' be identical to Y except that it leaves s unspecified. Then $\mathsf{L}(Y') < \mathsf{L}(Y)$ whenever the marker has positive code length, while $\mathcal{D}_Q(E, Y) \leq \mathcal{D}_Q(E, Y')$ with strict inequality if $d_q(\bot, \bar{a}) > 0$ and $w_q > 0$.*

Proof. The representations differ only in the specification of s. The explicit marker is a string of positive code length, so Y is longer than Y'. In Y' the interpretation uses the

default, so $I_G(Y')(q) = \bar{a}$, whereas $I_G(Y)(q) = \bot = a_E(q)$ contributes zero to the distortion. All other terms are equal, and the claim follows. □

The proposition says that, under a grammar with defaults, honesty about what is unknown is a cost paid in bits, while overreach is the zero-cost default. If every unsettled question tied to a slot is handled in this way, the length of the representation is independent of the number d of default-filled slots, whereas the unsupported commitment mass grows with d. Description length and commitment are thereby decoupled: the first can fall as the second rises. This gives the inversion its precise form.

Theorem 8.2 (Complexity inversion)**.** *Let $\mathcal{M}_G$ be a grammatical language with defaults, and let Y_0 be a representation in $\mathcal{M}_G$ with d_0 explicitly handled slots and m slots in total, $m > d_0$. For each j with $0 \le j \le m - d_0$, let Y_j be obtained from Y_0 by leaving j further slots unspecified, each bearing on a distinct question in Q on which the evidence is unsettled and the default is definite. Then $\mathsf{L}(Y_j)$ is non-increasing in j and $\mathcal{D}_Q(E, Y_j)$ is non-decreasing in j, strictly in each argument when the markers have positive length and the distortions are positive.*

Proof. Each additional unspecified slot removes a marker of positive length, reducing the code length, and replaces an induced $\bot$ with a definite default on a question where the evidence says $\bot$, adding a positive term to the distortion. The questions are distinct, so the terms are additive. □

The theorem is a constructive existence statement about the structure of grammars with defaults and does not assert that any particular film makes the trade it describes. It identifies a mechanism, and it identifies the quantity that mechanisms of this kind exploit: the zero price of a default. The gap between cheap communication and expensive justification is a property of the language, and not a defect of any author's intention.

8.4 Recognition cost

Description length given a held grammar is a measure of what the audience must be told. A related but distinct quantity is the effort an audience needs to recognize the structure, which may be written $C_{\text{recognize}}(Y \mid G)$, and a third is the effort needed to verify the implied claims against independent sources, $C_{\text{verify}}(Y)$. The central hypothesis of the essay is that these vary independently, in particular that strong conventionality can reduce recognition cost without reducing the evidentiary or editorial work required to construct the particular representation, and without reducing the work required to verify what it implies. This is an empirical hypothesis, treated in Section 12, and not a theorem. What the theorem establishes is that its premises are consistent: there is no mathematical obstacle to a representation that is short, easy to recognize, and expensive to check.

9 Four Propositions

The propositions of this section develop the formal consequences of the framework. Each is accompanied by its assumptions, a proof or constructive counterexample, and a statement

of what it does not establish. Proposition 4.1 and Proposition 8.1 appear in earlier chapters where the argument required them.

9.1 Non-monotonicity of description length

Proposition 9.1 (Non-monotonicity)**.** *Without restrictions on the admissible encoding schemes, source material, and interpretive rules, there is no universal ordering of minimal description length among documentary, dramatized, and grammaticalized representations of the same historical subject. More precisely, for any two of the three languages there exist problems* (E, Q) *and coding conventions satisfying the stated restrictions in which each of the two has strictly the smaller* $\mathsf{L}^*(E, Q, \varepsilon)$ *at a given tolerance.*

Constructive argument. Take the documentary language with per-pointer cost $\log_2 N_A$ bits, per-trim cost τ, and ordering cost at most $\log_2 n!$, so that a representation using k records costs on the order of $k(\log_2 N_A + 2\tau)$ plus ordering. Take the grammatical language with grammar cost $\mathsf{L}(G)$ counted against the individual work (no amortization), a role-assignment cost of $\log_2 |R|$ per assigned subject for a role set R, and negligible explicit filler cost where defaults suffice. Let Q consist of a single question whose supported answer is fixed by one record. A documentary representation with $k = 1$ satisfies the constraint at zero distortion and costs on the order of $\log_2 N_A + 2\tau$, while any grammatical representation must pay $\mathsf{L}(G)$, which is larger when the grammar is not amortized, so the documentary language has the smaller length. Now let Q consist of questions whose supported answers jointly are conveyed, within a large tolerance, by a role label, and let the supported evidence be spread across k records with k large. A documentary representation faithful to the evidence must reference a large fraction of these records, so its cost grows with k, while the grammatical representation pays $\mathsf{L}(G)$ once, plus the role assignment, and for k sufficiently large the grammatical language has the smaller length. A dramatization language with a short scene specification can be made smaller than both for a question that records do not address and that a single reconstructed scene, flagged as such, answers within tolerance by inducing $\perp$ on the unsupported particulars. In each case the restrictions are stated and the conclusion follows from the length formulas. □

The proposition concerns description length. It is not a claim about production expenditure, artistic quality, or historical adequacy, and it is compatible with the commonplace that an archival documentary can be extremely costly to produce. A low-budget staged scene and an expensive archival work are an elementary counterexample to any universal cost ranking, and the formal statement shows that the same is true when cost is measured informationally. The hierarchy that survives is a hierarchy of transformations permitted, not of cost: the three languages differ in the operations they allow, and the minimal description length depends on which operations a given problem calls for.

9.2 Collapse under a grammaticalization operator

A grammaticalization operator maps an evidence set to a representation determined by a role assignment, so that two evidence sets assigned to the same role in the same grammar produce the same output wherever the explicit fillers agree. Let $\mathfrak{G}$ denote the operator, acting on a domain $\mathfrak{E}$ of evidence sets.

Theorem 9.2 (Non-injectivity and collapse)**.** *Let $\mathfrak{G} : \mathfrak{E} \to \mathcal{O}$ be a grammaticalization operator, let q be a historical question, and suppose $E_1, E_2 \in \mathfrak{E}$ satisfy $\mathfrak{G}(E_1) = \mathfrak{G}(E_2)$ and $a_{E_1}(q) \neq a_{E_2}(q)$. Then no function h of the output recovers the evidentially supported answer on both: there is no h with $h(\mathfrak{G}(E)) = a_E(q)$ for $E \in \{E_1, E_2\}$. If d_q is a metric on $\mathcal{A}_q$, then for every h,*

$$\max_{i \in \{1,2\}} d_q(a_{E_i}(q), h(\mathfrak{G}(E_i))) \;\geq\; \tfrac{1}{2}\, d_q(a_{E_1}(q), a_{E_2}(q)).$$

Moreover, there is a function h_q with $h_q(\mathfrak{G}(E)) = a_E(q)$ for all $E \in \mathfrak{E}$ if and only if $\mathfrak{G}(E) = \mathfrak{G}(E')$ implies $a_E(q) = a_{E'}(q)$.

Proof. Since $\mathfrak{G}(E_1) = \mathfrak{G}(E_2)$, any function h of the output takes the same value o on both. If h recovered the answer on both, then $a_{E_1}(q) = o = a_{E_2}(q)$, contradicting the hypothesis. For the metric statement, the triangle inequality gives $d_q(a_{E_1}(q), a_{E_2}(q)) \leq d_q(a_{E_1}(q), o) + d_q(o, a_{E_2}(q))$, so one of the two terms is at least half the left side, and symmetry of the metric allows the second term to be read as $d_q(a_{E_2}(q), o)$. For the last statement, necessity is the first claim, and sufficiency follows by defining h_q on the image of $\mathfrak{G}$ by $h_q(o) = a_E(q)$ for any E with $\mathfrak{G}(E) = o$, which is well defined under the stated condition. □

The last clause is a factorization criterion: a grammaticalized output is sufficient for a question exactly when the answer to the question factors through the output. This is the query-relative notion of sufficiency that the whole framework relies on, and it can be written as an equivalence relation. For a class of questions Q, two evidence sets are equivalent, $E_1 \equiv_Q E_2$, when $a_{E_1}(q) = a_{E_2}(q)$ for every $q \in Q$. A representation preserves the relevant information when it is a function of nothing coarser than the Q-equivalence class. A grammar becomes problematic exactly when it identifies evidence sets that are not Q-equivalent for questions the representation invites the audience to answer. The same factorization condition has been studied for many-to-one descriptions in a separate treatment [5], and the present theorem is its application to narrative.

When a prior over evidence sets is available, the quantitative cost of collapse can be bounded by Fano's inequality. If the conditional entropy $H(a_E(q) \mid \mathfrak{G}(E))$ is positive then the probability of error of any reader of the output satisfies $P_{\text{err}} \geq (H(a_E(q) \mid \mathfrak{G}(E)) - 1)/\log_2 |\mathcal{A}_q|$ in bits, which is a nontrivial bound only when the conditional entropy exceeds one bit. This bound requires a prior, and priors over histories are not given by the evidence. It is stated for completeness and the qualitative theorem should be preferred where no defensible prior exists.

The relevant point is not that stereotypes are always false. It is that multiple, causally different histories can be mapped onto the same recognizable heroic structure. The theorem gives the essay a formal result that is more specific than the observation that narratives simplify reality: it identifies the condition under which grammaticalization destroys information needed to answer a particular historical question. A childhood in which early talent was cultivated by deliberate parental design and a childhood in which it appeared spontaneously can both be mapped to the production Early Talent, and no function of the output can tell them apart.

9.3 Retrospective role instability

Proposition 9.3 (Retrospective role instability)**.** *Let e be an event with fixed record E_e and let o, o' be two possible later outcomes. If the role assignment is outcome dependent in the sense that $\alpha(e; o) \neq \alpha(e; o')$ for some such pair, then the narrative significance assigned to e differs between the two cases although e and E_e are identical.*

Proof. Immediate from the definition of outcome dependence. The content of the proposition is that outcome dependence is a possible and, under the heroic grammar of Section 6, a typical feature of role assignment, and the examples below show that it holds for assignments of the kind the grammar prescribes. □

A child's remark about a game may be assigned to Early Talent if the child later becomes celebrated, and may be ignored or assigned no role if the child does not. The record is unchanged. The proposition motivates a counterfactual test for retrospective hero construction, the role-substitution test of Section 12, and it connects to a selection effect worth stating without numbers. Among children who exhibited a given behavior, only those who later became prominent are likely to have the behavior recorded, retained, and cited in a biography. The inclusion of the record in a biographical account is therefore an outcome-dependent event, and its presence is correspondingly weak evidence for a causal role in the outcome. The proposition does not establish that any particular biographical claim is false, only that the inclusion and significance of the early event cannot be read as independent confirmation of the later achievement.

9.4 What these propositions do not establish

Each result has a limited scope. Proposition 4.1 limits the force of provenance audits and does not show that any documentary's arrangement is invalid. Proposition 8.1 and Theorem 8.2 concern grammars with defaults in an idealized interpretation map and do not show that real audiences fill defaults as the idealization says. Proposition 9.1 concerns description length under stated coding conventions and does not rank productions. Theorem 9.2 concerns questions on which the evidence sets differ, and it is silent where they do not. Proposition 9.3 concerns a possibility that must be tested for each case.

10 The Cost of Hidden Narrative Structure

A central distinction in the Rommens paper is between information explicitly represented in a structure and information hidden inside the rule that operates on it. A representation can reproduce the required behavior while sacrificing mechanistic transparency, and the authors argue that a good model should be both efficiently representable within its chosen language and honest about the mechanisms it posits. A fully flexible representation may achieve a short description while concealing the physical hypothesis inside an opaque effective rule, and a more constrained representation can be scientifically preferable precisely because it makes the hypothesis explicit and testable.

In narrative production an analogous problem arises when a documentary does not state its organizing assumptions. The audience may receive a sequence that appears evidentiary

while relying on an unstated grammar of genius, competition, technological destiny, or prophetic warning. Whether to state the grammar is a design choice that can be analyzed informationally. Let Δ be a disclosure operator that appends to a representation an explicit statement that its slots are filled by convention and not by evidence, and that maps the interpretation so that default fillers induce $\perp$ on the questions they bear on. The code length of $\Delta(Y)$ exceeds that of Y by the code length of the disclosure, which is small and does not depend on the size of the representation, and the distortion of $\Delta(Y)$ can be much lower, since the defaults no longer induce definite answers. Disclosure is therefore cheap in bits and large in effect on distortion. In the language of Proposition 8.1, it is the generalization of the explicit indeterminacy marker, paid once instead of per slot.

Two observations follow. The first is that disclosure improves interpretive accountability in exactly the sense in which the explicit structure of the hypergraph model improves transparency: it relocates information from the rule into the representation, where it can be inspected. The second is that the relocation changes what the representation does. A grammaticalized work whose grammar is announced no longer imposes the grammar silently, and what is lost is precisely the persuasive economy that makes the grammar attractive. The informational analysis therefore explains both why disclosure is cheap and why it is rare. This does not make disclosure a universal remedy, since a sophisticated audience may discount a role it can see and an unsophisticated one may not, and the effect is an empirical matter. The essay's claim is limited to the formal point that the cost of exposing hidden structure is small compared with the commitments it bounds.

The distinction between hidden structure and explicit structure also clarifies a limit of the formal framework. The comparison of languages by description length treats the code as fixed. Whether a defaults-based language is better than an explicit one depends on criteria outside the code length, namely interpretability, identifiability, and the capacity of an audience to check. The Rommens paper reaches the same conclusion for graphs and hypergraphs, where, in the equivalence regime, any preference must rest on non-informational criteria. The essay accepts this: informational cost is one lens, historical adequacy is another, and the framework is useful because it keeps them apart.

11 Constructing the History of Artificial Intelligence

11.1 Case selection and evidentiary limits

The film is a British documentary directed by Nick Holt and produced by 72 Films and Windfall Films, with Ben Brown as producer and editor and with archive producers credited separately [9]. It premiered at the Tribeca Festival on 6 June 2026 and was broadcast on Channel 4 on 12 November 2026. The credits and release dates are those of the production and festival records and are cited as such. The case examines its opening sequence. The identifications of the sequence's components and their order are taken from a viewing of the film and have not yet been checked against the finished work. The analysis below is therefore conditional on the following being confirmed by a shot-level ledger with timestamps: that the sequence includes excerpts from *WarGames*, *Short Circuit*, and *RoboCop*; that it includes archival footage of a machine demonstration that is Shannon's Theseus [17]; that it includes a childhood chess interview of Demis Hassabis, described in the viewing as dating from 1986; and that the film's treatment of Geoffrey Hinton assembles his scientific history and

his later risk assessments in the manner described in Section 11.5. Passages of the film are located by timestamp, and the ledger of Section 12 is the place where the timestamps, the sources, and the order are fixed. The purpose of the case is to examine representational operations and not to adjudicate the film's overall argument about the risks of artificial intelligence, a question on which the present analysis takes no position.

A rare feature of the case is a primary statement of the production's representational strategy. In a published interview, the director reports that the broadcaster commissioned the film in 2024, that the production deliberately left out technical exposition, that the structure centers on people and in particular on Hinton, and that the chapters follow the development of the technology through stages likened to conception, intuition, and early adolescence [19]. The same account describes a search through archive for late-1980s footage of Hinton being laughed at on television. These statements are the filmmakers' own account of their choices and are used as such: they indicate what grammar the production reports adopting, and they leave open how far the finished sequence realizes it. They matter to the framework because a stated preference for characters over technical explanation is a stated choice of language, in the sense of Section 7, in which structure and role carry the weight that explanation would otherwise carry.

11.2 Fictional machines as cultural memory

Excerpts from fictional works can be analyzed as representations of historically circulating expectations about intelligent machines. The methodological question is whether their placement distinguishes fictional imagination from empirical evidence or permits the former to supply unexamined intuitions about the latter. In the terms of the framework, a fictional excerpt is an authentic record of a cultural artifact, and as such it is entirely valid evidence for questions about what audiences of a given period imagined. It is not evidence for questions about what machines of that period could do. If the arrangement places a fictional robot adjacent to a laboratory demonstration, the viewer may infer a relation of anticipation or continuity between the two. By Proposition 4.1, the authenticity of both components does not settle whether that relation is supported. The scope question for the audit is therefore which questions the sequence invites the audience to answer, and which of them the evidence supports.

11.3 Technical continuity, thematic continuity, and visual resemblance

If the archival footage is confirmed to depict Theseus, it provides an example of an actual adaptive machine, built to learn a maze, whose historical significance differs from that of fictional autonomous robots. The analysis asks whether the montage establishes technical continuity, in the sense of a documented lineage of methods; thematic continuity, in the sense of recurring questions about machine learning and control; or merely visual resemblance, in the sense of a shared look of cabinets, wires, and moving parts. Each of these answers a different question, and the evidence for them differs. A documented lineage would be supported by citations, collaborations, or successive systems. A thematic continuity is supported by the history of ideas. A visual resemblance is supported by nothing beyond the footage itself. The montage, by juxtaposing the machine with fictional robots and later systems, can induce an answer to the lineage question without any component asserting

one, and the analysis asks which answer it induces and whether the evidence supports it.

11.4 Hassabis and retrospective biography

The childhood chess interview provides a particularly useful test of grammaticalization. The analysis separates three things: what the footage establishes, what later biographical knowledge supplies, and what the documentary's arrangement invites viewers to infer. The footage establishes that a young person, identified by the film as the future researcher, was interviewed about chess at a given age. Later knowledge supplies his career. The arrangement, by placing the interview in a sequence devoted to the history of machine intelligence and not, say, in a sequence on childhood games, invites the audience to treat the interview as an origin scene, and thereby to assign the production Early Talent with the follow-on that the talent anticipates the later scientific achievement. The question is not whether he was competitive as a child, which the footage may well support. It is whether that observation is being transformed into an explanatory origin story for his later scientific career, which would be a causal answer to a question that the footage does not address.

The retrospective instability of Proposition 9.3 gives the useful counterfactual. The same interview, had the subject not become prominent, would almost certainly not have been selected, retained, or arranged, and would not have been narratively significant to any audience. The record is the same in both cases. What differs is the outcome-dependent role assignment, and the inclusion of the record is itself a product of the outcome. Whether the footage is also evidence that the interviewee was unusual among the children interviewed on that occasion is a separate question, and one that depends on information the film may or may not supply.

11.5 Hinton and the prophetic role

The film's treatment of Hinton can be examined for the construction of a warning figure, particularly if his scientific authority, his historical involvement in neural-network research, and his subsequent risk assessments are assembled into a single prophetic trajectory. In the heroic grammar, the production Mission can expand to Warning, and a figure who contributed to a technology and later warns of it fits that production well. The analysis must distinguish Hinton's documented statements, which can be traced to their sources and dated, from the narrative role the documentary may assign him, which is a property of the arrangement. The production's own account of its archive search adds an instance of the grammar's earlier productions. Footage in which Hinton is ridiculed in the late 1980s for suggesting that computers might one day think [19] is an authentic record of an event, and its placement ahead of his later warnings fits the sequence Exclusion, then Recognition, then Warning. Whether the ridicule is evidence that his position was then unjustly dismissed, as the sequence invites the audience to conclude, or evidence of one television audience's amusement on one occasion, is a question the footage alone does not settle, and the sequence induces an answer to it. A documented statement of concern is evidence that he holds the concern. It is not, by itself, evidence for the rightness of the concern, for the proposition that his earlier work caused the situation the concern addresses, or for the proposition that the concern follows from his technical expertise rather than from other considerations. Each of these is a distinct question with a distinct evidential status, and the prophetic role

supplies answers to all of them at once.

11.6 The five-minute problem revisited

The case concludes by returning to the opening observation. A brief montage can simultaneously perform archival recovery, historical compression, cultural juxtaposition, and biographical role assignment, and these operations can occur within the same sequence and not in separate successive stages. The construction cost of such a sequence includes the pointers, trims, and clearances of the documentary language, whose ordering freedom alone, for a sequence of n shots, is up to $\log_2 n!$ bits of editorial commitment. The communication cost, given a familiar grammar, may be small. The commitment is the set of answers the sequence induces on questions about lineage, origin, and warning, which is where Theorem 8.2 applies. The three quantities, construction cost, description length given the grammar, and unsupported commitment, are the ones the framework separates, and the sequence is a case in which all three are present and visibly different in magnitude.

12 Operational Tests and Limitations

12.1 Separating mathematical results from empirical hypotheses

The Rommens paper separates what is proven from what is merely suggested, and in particular distinguishes theorems about Kolmogorov complexity from explicit description-length estimates, because the former are incomputable. The narrative study follows the same discipline. The propositions of Section 9 are mathematical statements about idealized languages and interpretation maps. The claims that real audiences interpret as the idealization says, that real films approach the extremes the constructions describe, and that disclosure changes real judgments are empirical hypotheses. The study can compare specified editing strategies, coding conventions, and viewer interpretations, and it does not claim to measure the intrinsic Kolmogorov complexity of a documentary.

12.2 Provenance audit

The first operational instrument is a shot-level ledger. For each archival excerpt the ledger records the origin, the date or its estimate with a statement of confidence, the rights status where ascertainable, the in and out points, the position in the sequence, and the evidentiary function that the excerpt performs, whether as a record of an event, as an illustration, as a cultural artifact, or as a transitional device. The ledger is both the verification instrument for the case study and the source of an explicit code length: each field has a stated encoding, and the length of the encoded ledger is the description length $\mathsf{L}_{\mathcal{M}_D}$ of the documentary representation under that convention. Table 1 gives the structure of the ledger for the opening sequence, with the entries that require verification marked as such.

12.3 Commitment audit

The second instrument classifies the claims a sequence communicates into three kinds: claims directly supported by the footage, claims introduced by narration, and claims inferred from

Segment (as described)	Source and date to confirm	Function to assign	Timestamp
Excerpt from *WarGames*	Studio film; rights holder and year to confirm	Cultural artifact	to be entered
Excerpt from *Short Circuit*	Studio film; rights holder and year to confirm	Cultural artifact	to be entered
Excerpt from *RoboCop*	Studio film; rights holder and year to confirm	Cultural artifact	to be entered
Machine demonstration, possibly Theseus	Archive and date to confirm; identification to confirm	Record of an event	to be entered
Chess interview with the young Hassabis	Broadcaster and date to confirm	Record of an event, origin motif	to be entered
Treatment of Hinton	Interview dates and sources of each statement to confirm	Record of statements, warning role	to be entered

Table 1: Structure of the shot-level ledger for the opening sequence of the case study. Entries are described from viewing and require verification against the film, with timestamps.

montage. For each claim the audit records the question it answers, the answer induced, the answer the evidence supports, and the resulting distortion term. Independent coders can assess the reliability of these classifications, and the agreement between coders is itself a measure of how determinate the induced answers are. A low agreement on what a sequence claims is informative: it indicates that the sequence induces indeterminate impressions and not definite answers, which affects how the distortion should be scored.

12.4 Counterfactual editing test

The third instrument constructs alternative sequences from the same authentic excerpts with different order and narration, and asks whether viewers infer different causal relationships, character traits, or historical trajectories. The test directly addresses the editing-as-inference claim of Proposition 4.1: if viewers' inferences change with order while the components are held fixed, the inferences are properties of the composition. A practical experiment would present the same archival footage with documentary narration, dramatized reconstruction, and heroic biographical framing, and measure differences in causal inference and confidence across the three presentations. The experiment compares specified editing strategies and measures viewer interpretation, and it makes no claim about the Kolmogorov complexity of any of the stimuli.

12.5 Role-substitution test

The fourth instrument tests retrospective role instability. It replaces the eventual famous subject with a historically comparable individual who did not achieve similar recognition, while preserving the early-life evidence as closely as possible, and measures whether the same childhood behavior is still interpreted as evidence of exceptional destiny. A positive result, in which the behavior is interpreted as exceptional only when the outcome is attached, supports outcome dependence in the sense of Proposition 9.3. A negative result suggests that the behavior itself carries the interpretation, and the case for retrospective construction is then weaker for that example.

12.6 Recognition and verification costs

The fifth instrument compares the time required to understand a narrative, to identify its implied claims, and to verify those claims against independent sources. The hypothesis, stated in Section 8, is that familiar heroic grammars may accelerate recognition while leaving verification burdens unchanged or increasing them. Under the analysis of this essay that hypothesis is a consistent possibility with a mechanism, the zero price of defaults. Whether it holds for real audiences is a measurement and not a deduction.

12.7 Limitations

Several limitations bound the claims. The interpretation map is idealized, and real audiences vary in the grammars they hold and in how they resolve defaults, so distortion as computed for the idealized viewer is a property of a language and a convention, and the distortion experienced by a particular audience is an empirical quantity. The answer profile of the evidence depends on background assumptions K and on judgments about what the evidence supports, and these are contestable; the framework makes them explicit and does not remove the dispute. The description lengths depend on coding conventions, and a preference between languages is meaningful only when the gap exceeds the effect of reasonable changes to those conventions. The three languages are idealizations of practices that overlap, and many real works mix them. The restriction families are chosen to isolate mechanisms, and a different partition might isolate others. Finally, the case study is conditional on verification, and nothing in the framework replaces that verification.

13 Why Historical Complexity Disappears

13.1 Compression that preserves and compression that substitutes

The discussion considers the broader consequences of the framework for historical documentaries, biographical films, journalism, institutional histories, and representations of scientific achievement. The central issue is not that narrative simplification is inherently illegitimate. Every finite representation must omit information, and a representation that omitted nothing would be the record itself. The question is whether the omitted distinctions matter to the claims the representation encourages audiences to accept. The framework gives that question a definite form. A compression preserves the relevant information when the answers to the declared questions factor through it, and it substitutes a conventional story for those answers when the grammar identifies evidence sets that differ on a question the representation invites the audience to answer.

The distinction is the one made in the sufficiency criterion of Theorem 9.2. A highly compressed heroic biography may be adequate for identifying its subject or communicating a broad cultural role, and inadequate for explaining the actual causes of an achievement. Both statements can be true of the same representation, because adequacy is relative to a class of questions. A representation is adequate for Q exactly when its output determines the answers on Q, which is to say when $E_1 \equiv_Q E_2$ whenever the outputs agree. The problem lies not in compression but in the unannounced extension of a representation's apparent authority from the questions it preserves to the questions it does not.

13.2 Apparent inevitability and the individualization of collective work

Familiar narrative grammars can manufacture apparent inevitability, because the slots of a heroic grammar lead from origin through trial to mission as a sequence in which each stage follows from the last, and a life arranged in that sequence looks like a destiny. They can individualize collective achievements, because a grammar whose roles are persons has no slot for institutions, funding structures, concurrent groups, or the many contributors whose work was necessary and unnamed. They can transform institutional or contingent developments into stories of personal destiny. In the language of the theorem, causally different histories, in which an achievement is the work of a field or the outcome of a contingency, map to the same heroic output, and the output cannot recover which obtained. The historical complexity is not hidden in the output; it has been identified away.

13.3 Grammaticalization need not be heroic

A qualification is essential. Grammaticalization need not be heroic. Villain narratives, martyr narratives, institutional redemption stories, and narratives of inevitable decline may operate through analogous mechanisms, each with its own productions, slots, and defaults, and each subject to the same collapse and default results. The framework is a framework for the operation and not for any particular grammar. A study of hero grammars is a study of one family of cases, chosen because the case that motivated the work belongs to it.

13.4 Relation to prior work

The study draws on several established traditions, and its distinct contribution should be stated narrowly. Nichols' work on documentary modes and the rhetoric of evidence supplies the account of documentary voice and the conventions by which a work claims to represent reality [12]. White's account of emplotment, and the related work of Ricoeur on narrative configuration and temporality, establish that historical meaning is produced by narrative form [22, 14]. Bordwell's account of narration and cinematic comprehension supplies a theory of the inferences that viewers draw from arrangement [2], and Chatman's distinction between story and discourse separates events from their presentation [3]. Propp's functions and roles are the antecedent of any role grammar [13], Barthes' account of myth as a second-order system describes how a sign acquires a further signification [1], and Goffman's frames describe how interpretation is organized in advance of the particular event [6]. Goodman's treatment of symbol systems and worldmaking supplies the philosophical setting for the claim that representations make and do not merely mirror [7]. Li and Vitányi supply the algorithmic information theory [11].

None of this is new, and the proposal is not that stories employ archetypes. The distinct contribution lies in connecting narratology and documentary theory to the formal study of information preservation. Its specific claims are that role assignment can be treated as an operation with an informational cost, that the cost of communicating a story given a grammar can fall while the cost of justifying its implications rises, that the condition under which a grammar destroys information is the failure of a factorization criterion, and that the comparison of representational modes should be made through a distortion-constrained description length and not through expressive power or surface complexity. These claims are narrower than a general theory of narrative, and they are intended to be testable in the

senses of the preceding chapter.

14 Conclusion: Representational Economy and Historical Accountability

The study began with the observation that five minutes of documentary can be difficult to construct. The difficulty arises not simply from acquiring footage but from producing intelligibility under conditions of incomplete evidence, fragmented provenance, and competing interpretive possibilities. Documentary, dramatization, and grammaticalization represent distinct ways of addressing that difficulty. Documentary organizes surviving traces. Dramatization supplies missing particulars. Grammaticalization makes particulars legible through reusable narrative roles. They are best understood as three families of constraints on how information may be distributed among evidence, narrative structure, and interpretive rules, each with its own costs, commitments, and verification burdens.

The adaptation of the Rommens framework yields a definite structure. There is a floor, the intrinsic description length of the evidence-supported answers on the declared questions. There is interchangeability under unrestricted languages, so that expressive power cannot distinguish the three families. There is a constrained optimization, the shortest description whose distortion on the declared questions stays within a tolerance, which is the correct object for comparison. And there are four results of the theory: description length across the families has no universal ordering; defaults make overreach cheaper than reservation, which gives the complexity inversion; a grammaticalization operator destroys the distinctions on which a question depends whenever it identifies evidence sets that differ on that question; and the significance of an early event can change when a later outcome changes while the event and its record remain fixed.

The conclusion rejects the assumption that more elaborate representations necessarily require more information or convey more historical truth. Narrative conventions can reduce the cost of communicating a story while increasing the burden of substantiating its implications. The central result is a separation between informational economy, representational commitment, and evidentiary adequacy, three quantities that the surface of a finished sequence presents as one. Their divergence explains much of what is observed in the first five minutes of a film such as the one examined here.

The concluding claim is that historical representation becomes most epistemically consequential when its organizing conventions disappear from view, and the resulting narrative appears to have been discovered and not constructed. The remedies the framework supports are modest and specific. They are declared scope, so that a representation states the questions it answers; explicit marking of what the evidence leaves open, so that the cost of reservation is paid once and in public; disclosure of the grammar, so that structure hidden in the rule is relocated into the representation where it can be examined; and audits at the level of the shot, the claim, and the counterfactual, so that authenticity of components is not mistaken for validity of composition. A narrative that survives these audits has paid its way. A narrative that does not may still be a good story, and the framework gives a precise account of which of its claims should be held to the standard of history.

References

[1] Barthes, R. *Mythologies*. Translated by A. Lavers. Hill and Wang, New York, 1972. Originally published 1957.

[2] Bordwell, D. *Narration in the Fiction Film*. University of Wisconsin Press, Madison, 1985.

[3] Chatman, S. *Story and Discourse: Narrative Structure in Fiction and Film*. Cornell University Press, Ithaca, 1978.

[4] Cover, T. M. and Thomas, J. A. *Elements of Information Theory*. Second edition. Wiley, Hoboken, 2006.

[5] Flyxion. Lossless Projection. Independent manuscript, 2026.

[6] Goffman, E. *Frame Analysis: An Essay on the Organization of Experience*. Harvard University Press, Cambridge, 1974.

[7] Goodman, N. *Ways of Worldmaking*. Hackett, Indianapolis, 1978.

[8] Grice, H. P. Logic and Conversation. In Cole, P. and Morgan, J. L. (eds.), *Syntax and Semantics, Volume 3: Speech Acts*, pp. 41–58. Academic Press, New York, 1975.

[9] Holt, N., director. *AI: Probably Nothing to Worry About*. Documentary film. 72 Films and Windfall Films, United Kingdom, 2026. Premiered at the Tribeca Festival, 6 June 2026; broadcast on Channel 4, 12 November 2026.

[10] Kuleshov, L. *Kuleshov on Film: Writings by Lev Kuleshov*. Translated and edited by R. Levaco. University of California Press, Berkeley, 1974.

[11] Li, M. and Vitányi, P. M. B. *An Introduction to Kolmogorov Complexity and Its Applications*. Fourth edition. Springer, New York, 2019.

[12] Nichols, B. *Representing Reality: Issues and Concepts in Documentary*. Indiana University Press, Bloomington, 1991.

[13] Propp, V. *Morphology of the Folktale*. Second edition. Translated by L. Scott. University of Texas Press, Austin, 1968. Originally published 1928.

[14] Ricoeur, P. *Time and Narrative*. Volumes 1 to 3. Translated by K. McLaughlin and D. Pellauer. University of Chicago Press, Chicago, 1984–1988.

[15] Rissanen, J. Modeling by Shortest Data Description. *Automatica*, 14(5):465–471, 1978.

[16] Rommens, C., Traversa, P., Ferraz de Arruda, G., and Moreno, Y. The Informational Cost of Structure: Representational Complexity in Networked Dynamical Systems. arXiv:2607.03608 [cs.IT], July 2026.

[17] Shannon, C. E. Presentation of a Maze-Solving Machine. In von Foerster, H. (ed.), *Cybernetics: Transactions of the Eighth Conference*, pp. 173–180. Josiah Macy Jr. Foundation, New York, 1952.

[18] Shannon, C. E. Coding Theorems for a Discrete Source with a Fidelity Criterion. *IRE National Convention Record*, 7(4):142–163, 1959.

[19] Televisual. Behind the Scenes: AI: Probably Nothing to Worry About. *Televisual*, 2026. `https://www.televisual.com/news/bts-ai-probably-nothing-to-worry-about/`.

[20] Tishby, N., Pereira, F. C., and Bialek, W. The Information Bottleneck Method. arXiv:physics/0004057, 2000.

[21] Vereshchagin, N. K. and Vitányi, P. M. B. Rate Distortion and Denoising of Individual Data Using Kolmogorov Complexity. *IEEE Transactions on Information Theory*, 56(7):3438–3454, 2010.

[22] White, H. *Metahistory: The Historical Imagination in Nineteenth-Century Europe*. Johns Hopkins University Press, Baltimore, 1973.