

Chapter 4: The Synthetic Hermeneutic Protocol

Before interpretation becomes verdict

In April 2023, Pushkin Industries announced *Death of an Author*, a murder mystery published under the name Aidan Marchine. The name joined Stephen Marche to the machine systems he had used. Pushkin described the novella as 95 percent written by artificial intelligence and 5 percent written by Marche. It also credited Marche with making the outline and prompts. The announcement presented the book as a demonstration of what human and machine collaboration could produce (Pushkin Industries 2023).

The percentages appear to divide a finished work into measurable shares of authorship. Yet the announcement also names activities that the ratio does not measure: outlining, prompting, selecting, arranging, and releasing. The figure is evidence of how Pushkin presented the novella. To describe the production process, a critic needs records of those activities; to interpret the novella, the critic must read what those activities yielded. Neither task can be completed by converting a publicity ratio into a verdict.

The protocol begins with a question about the text and follows the records that can bear on it. A production account may identify who selected a passage. An archive may reveal which alternatives survived. Platform and publishing conditions may explain how this edition reached readers. Only then can a responsibility claim specify who knew of a problem, controlled a decision, benefited from it, or could have prevented it. The sequence matters because each kind of evidence warrants a different kind of conclusion.

The sequence can be used as a repeatable method. I call it the **Synthetic Hermeneutic Protocol**. It consists of five operations:

1. textual reading;
2. production reading;
3. archive reading;
4. infrastructure reading;
5. responsibility reading.

The order matters, but it is not a staircase that a critic climbs only once. Evidence discovered during production reading may send the critic back to a phrase. A legal dispute may alter the meaning of a disclosure label. A reader's response may expose how a publisher's framing has disciplined attention. The critic moves back and forth while keeping the evidence for each move distinct.

Working definition

The **Synthetic Hermeneutic Protocol** is an evidence-governed sequence for interpreting a synthetic textual event as language, production, archive, infrastructure, and an uneven distribution of responsibility. Each operation has four parts: questions, acceptable evidence, prohibited inferences, and a bounded form of conclusion.

Evidence-governed sets the method's limit. A protocol does not make interpretation mechanical. It does not produce the single correct reading of a poem or novel. It disciplines the routes by which a critic moves from observation to claim. Close reading remains contestable. A production date may not be. A creator's account is evidence that the creator made a claim, but it may not independently establish everything the claim asserts. A resemblance to a known crime writer may support a reading of genre effects, but it cannot prove that writer's work appeared in a training set.

The protocol offers a grammar of critical warrants, not a detection instrument. It tells readers what kind of evidence can support what kind of sentence.

Why five readings?

Literary criticism already has powerful methods for form, authorship, reception, textual genesis, media, and institutions. The new problem is not that each of these traditions has become useless. It is that a synthetic text can bring their objects into a single event while hiding the links among them.

A generated paragraph may be stable on a printed page, yet its creation may have involved many nondeterministic outputs. A byline may name one person while the afterword names three commercial systems. A claim of stylistic originality may depend on a model's capacity to reproduce genre conventions learned from inaccessible data. A reader may encounter the work through an audiobook service whose promotional copy supplies the production ratio before the first sentence is heard. Responsibility for an objectionable passage may be distributed across prompting, model design, selection, editing, publication, and policy, but those actors do not possess equal knowledge or control.

One reading cannot answer all of these questions. Five readings make the differences visible.

A compact map of the protocol

Operation	Primary object	Core question	Strongest permitted conclusion
Textual reading	The words, forms, sequence, and effects	What does this version of the text make available to interpretation?	A claim about form, meaning, tension, pattern, or readerly effect in the examined text
Production reading	Documented acts and transformations	How did this version come to exist?	A bounded reconstruction of recorded human and machine operations
Archive reading	Cultural materials, genres, corpora, and absences	What prior language and cultural memory become legible here?	A claim about intertext, genre, archival structure, or documented source relation
Infrastructure reading	Interfaces, firms, labor, computation, law, and resources	What systems enabled and constrained this textual event?	A case-specific account of relevant dependencies and limits
Responsibility reading	Actors with unequal knowledge and power	Who could know, choose, benefit, prevent, disclose, or repair?	A reasoned allocation of responsibility without flattening the network

The strongest permitted conclusion in the last column is a ceiling. It prevents a useful observation from silently turning into a larger claim. A textual reading can identify recurring mechanical metaphors. It cannot on that basis recover a hidden prompt. An infrastructure reading can show that the work depended on commercial model access. It cannot calculate the energy used by the novella unless workload-level measurements and system boundaries are available.

The claim ladder

Most errors in writing about synthetic literature begin with a defensible observation and climb too high. A fluent paragraph becomes proof that a model understands. A repeated image becomes evidence of a particular training source. A creator's explanation becomes a complete account of composition. The critic needs a way to see the climb.

The protocol uses a four-rung claim ladder:

1. **Observation:** a feature can be located in the examined object or record.
2. **Attributed account:** a named person or institution describes a process, motive, result, or value.
3. **Interpretation:** the critic relates evidence to a formal, historical, ethical, or political argument.
4. **Causal or general claim:** the critic states that one condition produced another, or that a finding travels beyond the case.

Higher rungs require additional evidence. The order does not imply that an observation is simple or an interpretation weak. A close reading can be complex and persuasive. The ladder asks a different question: how far does the sentence travel beyond what can be directly inspected?

Consider the 95 to 5 description of *Death of an Author*. The dated Pushkin announcement permits an observation: the publisher printed the ratio and connected Marche with the outline and prompts. It also permits an attributed account: Pushkin described the book as 95 percent written by AI. A critic may interpret the ratio as a paratext that divides creative labor into marketable shares. The evidence does not establish that readers valued the novella because of the ratio. That causal claim would require reception research capable of isolating the label's effect.

The same discipline applies to prose. Suppose a passage moves abruptly from restrained narration to an ornate simile. The feature is observable. A reviewer may attribute the shift to the reported use of Cohere for figurative language. A critic may interpret the shift as a seam in the work's assembled voice. Without the prompt, output, transfer, and revision record, the sentence "Cohere wrote this simile" exceeds the evidence.

The ladder is especially useful during revision. Mark the main verb in each factual paragraph. *Contains, states, reports, suggests, causes, and proves* make different demands. A weak chapter often uses a modest citation beside an ambitious verb. Synthetic Hermeneutics treats verb choice as part of factual accuracy.

Evidence classes are not a hierarchy of human worth

The evidence gate can sound legalistic, as if interpretation had to become a courtroom. That is not its purpose. Literary critics work with ambiguity, silence, contradiction, and affect. Synthetic works add another uncertainty: production systems are often proprietary, change over time, and preserve incomplete records. Without labels for evidence classes, a critic may confuse aesthetic uncertainty with unavailable technical facts.

The distinction also protects testimony. A writer's account is not worthless because it is attributed. It may offer the best available description of intention, labor, and choice. Attribution tells the reader whose knowledge and interests shape the account. The same care applies to publishers, model providers, reviewers, and workers. Their positions differ. Their words enter the provenance field as situated evidence.

Four labels will recur in the case studies:

documented for evidence that can be inspected in a stable record;

attributed for a statement tied to a named speaker or institution;

inferred for the critic's reasoned movement beyond direct description;

unknown for a question the current record cannot answer.

These labels prevent two opposite habits. Credulity repeats a process story as fact. Suspicion treats every process story as deception. Neither reads well. The protocol asks what the account establishes, what interests it may serve, what corroborates it, and where it stops.

A claim trace in practice

One sentence from a publicity record can enter all five readings without carrying the same meaning into each. Pushkin's April 2023 announcement said that the work was 95 percent written by AI and credited Marche with the outline and prompts (Pushkin Industries 2023). The protocol traces the sentence as follows.

For textual reading, the sentence is not evidence about syntax, characterization, or narrative coherence. It is a paratext that may alter how those features are noticed. A critic can compare a blind first reading with a reading shaped by the announcement, but cannot let the percentage diagnose every awkward phrase.

For production reading, the sentence documents the publisher's public account. It identifies a claimed division and two human tasks. It does not define the unit counted. Words, characters, passages, generations, editing time, and creative decisions would produce different ratios. The archive located for this book contains no independent audit that resolves the unit.

For archive reading, the sentence says almost nothing about training data or literary sources. The label "AI" compresses three named products and their unknown data histories into one category. The critic needs separate evidence for genre prompts, intertexts, and model documentation.

For infrastructure reading, the sentence has more force. It turns product access into a public identity for the novella and helps Pushkin position the book within a commercial debate about generative writing. The percentage travels through press coverage, reviews, retailer descriptions, and classroom discussion. Circulation makes it part of the work's provenance field.

For responsibility reading, the sentence points toward the actors who chose the claim and could clarify it. Pushkin controls the announcement; Marche is the named human participant; both benefit from the experiment's public distinction. Model providers did not publish this ratio merely because their systems contributed language. Responsibility for the wording follows the decision to make and circulate it.

The trace shows why the five operations should not be collapsed into a general discussion of authorship. The same evidence can be relevant in several domains while supporting a different conclusion in each. What remains stable is the source, date, attribution, and limit.

The protocol's methodological lineage

The five operations reorganize existing critical practices rather than replacing them. Textual reading retains the discipline of close attention to language and form. Production reading draws on genetic criticism, which treats drafts and revisions as witnesses to composition while acknowledging that the assembled *avant-texte* is never the process itself (Van Hulle 2022). A prompt and an output can enter such a dossier, but neither represents the whole genesis.

Archive reading extends source criticism into a setting where source relations may be statistical, composite, and unavailable at work level. The extension creates a stricter limit, not a license for speculation. Resemblance supports comparison; a claim of training membership requires a record suited to that claim.

Infrastructure reading follows book history and the sociology of texts in treating production and transmission as participants in meaning (McKenzie 1999). It also borrows the practice of infrastructural inversion: background systems become visible when a critic asks what had to work for the apparently immediate text to appear. Responsibility reading then adds a normative question that provenance alone cannot answer.

Genette's paratext is especially important because synthetic books often explain their making through bylines, afterwords, labels, and publicity (Genette 1997). These thresholds provide evidence and shape reception at the same time. The protocol preserves both functions.

Finally, W3C provenance models supply a useful grammar of entities, activities, and agents (World Wide Web Consortium 2013). The distinction prevents every component from being called an author. Literary interpretation must go further, however, because a technical association does not decide meaning, value, or responsibility. The five operations join these traditions around one evidentiary rule: each conclusion must remain within the reach of its sources.

Operation one: textual reading

Textual reading temporarily brackets the production story. Production remains relevant, but the bracket prevents a label such as "AI-written" from explaining the prose before the critic has described it.

Questions

- What words, images, rhythms, narrative structures, and genre signals recur?
- Where does the text establish or violate its own formal expectations?
- How are voice, focalization, address, temporality, and point of view organized?
- What ambiguities remain productive rather than accidental?
- Which effects depend on this edition, layout, performance, or interface?
- What changes when the paratext is withheld or supplied?

Acceptable evidence

Acceptable evidence includes the cited edition, verifiable quotations within permission limits, formal description, comparison with documented editions, and carefully bounded accounts of reader response. A critic may also compare a passage with an established genre convention if the comparison is textually demonstrated.

Prohibited inferences

Textual reading must not infer a hidden model, prompt, dataset member, temperature setting, or editing history from style alone. It must not treat awkward prose as proof of machine generation or smooth prose as proof of human revision. It must not assume that a model possesses a communicative intention because the narrator appears to have one. It must also resist the reverse error: declaring that the language cannot sustain meaning because part of it was generated statistically.

Permitted conclusion

The operation supports an interpretation of the examined text's form and effects. It may identify features that later production evidence helps explain. It cannot, by itself, establish how those features were produced.

Death of an Author: a bounded primary-text reading

The complete Kindle text was not available for page-verified quotation during this research pass. The textual operation therefore uses the authorized excerpt published by *WIRED* on 9 May 2023 and treats claims about the rest of the novella as publisher synopsis or reception evidence (Marchine 2023). This limit belongs in the model rather than in a hidden research note.

The excerpt opens with Gus Dupin in an isolated lakeside setting. A motorboat delivers a handwritten funeral invitation, described aloud as "an honest-to-God letter" (Marchine 2023). The scene establishes a contrast between physical message and disconnected recipient. It then shifts through three documentary forms: invitation, online news report, and obituary. Each form

claims authority over Peggy Firmin's death, and each gives Gus material that he must interpret.

The alternation is formally important. Descriptive narration presents smell, texture, water, and weather. Institutional prose replaces those sensations with public formulas of grief and reputation. Gus notices omissions in the obituary because his specialized knowledge exceeds the official record. The detective plot begins as a contest between documents.

That contest makes literary criticism part of the mystery. Gus studies crime and cyber fiction and has written about Firmin. The title recalls Roland Barthes's "The Death of the Author," while the surname Dupin recalls Edgar Allan Poe's detective C. Auguste Dupin (Barthes 1977; Poe 1841). These references do not prove anything about the generation process. They position the reader inside traditions of authorship and detection before production evidence enters.

The excerpt also introduces a deliberate problem for provenance. An invitation promises that the dead writer will deliver her own eulogy. A news report calls her a pioneer in using AI to make literary artifacts. The text asks how a voice can persist after a bodily death and how documents manufacture an authorial identity. Those questions arise from the published language, not from a guess about which system wrote a sentence.

The bounded textual conclusion is that the authorized opening uses nested documents and a scholar-detective to make authorship a problem of competing records. Crime fiction supplies the expectation that traces will lead to a cause. Synthetic provenance complicates that expectation because some production traces may remain inaccessible. A fuller assessment of continuity, style, and later narrative development requires the cited edition and page verification; the present application does not claim it.

Operation two: production reading

Production reading reconstructs the transformations that connect intention, instruction, generation, selection, revision, arrangement, publication, and reception. It resembles genetic criticism, but the available witnesses may include prompt records, model outputs, interface states, software versions, retrieval contexts, moderation interventions, and nondeterministic reruns.

Questions

Who initiated the project, and what goals did they state?

Which people and institutions planned, prompted, generated, selected, rejected, edited, arranged, designed, marketed, and published?

Which systems were used, in what sequence, and for what declared functions?

Are prompts, full outputs, revision histories, settings, and dates preserved?

Which claims come from creators or publishers, and which have independent corroboration?

What remains unknown because no record exists or access is restricted?

Acceptable evidence

Strong evidence includes dated drafts, complete prompt and output logs, version histories, source files, contracts where lawfully available, correspondence, and independently inspectable archives. Creator afterwords and interviews are important but attributed. Publisher announcements establish public framing and claims. Reviews can document reception or report what an afterword says, but they are not substitutes for a full process archive.

Prohibited inferences

The critic must not convert a creator's production percentage into an independently verified measurement. The critic must not assume that a named product remained technically unchanged during composition. Missing prompts cannot be reconstructed from apparent responses. A final passage cannot reveal how many rejected outputs preceded it. A polished sentence does not prove human rewriting, and a cliché does not prove unedited model output.

Permitted conclusion

Production reading supports a qualified account of documented acts, roles, and transformations. It must label attributed, inferred, and unknown elements separately.

Death of an Author: an attributed production chain

The public record establishes several facts and several claims. Pushkin released the work in May 2023 and described it as written through three AI-writing platforms, with Marche responsible for the outline and prompts. The publisher assigned the 95 to 5 ratio. Marche's afterword, as described by reviewers and later scholarship, discusses a process involving ChatGPT, Sudowrite, and Cohere. Accounts of the workflow say that ChatGPT generated passages, Sudowrite reshaped prose, and Cohere supplied or varied figurative language (Harris 2023; Comitta 2023).

That evidence supports this bounded reconstruction:

1. Pushkin published and promoted an experiment in machine-mediated fiction;
2. Marche devised an overall plan and generated instructions;

3. multiple commercial systems performed different textual operations;
4. Marche evaluated outputs and assembled a novella;
5. the work was published under a constructed byline and accompanied by explanatory paratext;
6. Pushkin marketed a numerical division between machine and human writing.

The same record does not show a complete prompt archive, every output, every rejection, every manual alteration, or a reproducible method for classifying words as machine-written or human-written. The ratio may have constrained composition. It certainly supported publicity. The evidence available to an ordinary reader cannot settle whether it qualifies as an auditable measurement.

This gap is not a failure of criticism. It is a finding. The production reading identifies **interpretive opacity** between a detailed creator narrative and an incomplete public record. The critic can say who claimed what, name the documented tools, and describe the reported sequence. The critic cannot rebuild the composition session.

Production reading also changes the authorship question. "Who wrote the words?" proves less useful than a set of narrower questions. Who established the mystery's architecture? Who requested each passage? Who judged one output better than another? Who revised? Who chose the byline? Who could disclose the process? Who sold the result? These are not evasions of authorship. They reveal that the category has been carrying several kinds of work at once.

Operation three: archive reading

Archive reading asks how prior language, genre, digitization, and cultural selection become active in a synthetic textual event. The archive includes known literary sources and conventions, documented fine-tuning or retrieval collections, and broader training corpora where their composition is publicly established. It also includes absences: languages poorly represented online, works excluded by access or format, and cultural forms compressed into dominant labels.

Questions

- Which intertexts, genres, conventions, and styles can be demonstrated in the text?
- Which sources were named in prompts, process records, or creator accounts?

- What is publicly known about the relevant model's training, fine-tuning, or retrieval data at the time?
- Which claims concern direct quotation, stylistic imitation, generic convention, or broad cultural memory?
- Whose works were digitized, licensed, scraped, omitted, or rendered difficult to identify?
- How does archival uncertainty limit the interpretation?

Acceptable evidence

Acceptable evidence includes explicit citations, verifiable textual parallels, documented prompts, dataset cards, provider disclosures, licensed corpus records, and reliable historical scholarship about a genre. Style comparison may support a formal interpretation when it is demonstrated. It does not establish dataset membership.

Prohibited inferences

The critic must never infer that a particular copyrighted work was used in training merely because the output resembles its author. A genre convention cannot be owned by the most famous writer who used it. Legal permission, ethical legitimacy, influence, imitation, and copying must not be collapsed into one accusation. An inaccessible training corpus must not become an empty space into which the critic projects any desired source.

Permitted conclusion

Archive reading supports claims about demonstrable intertexts, genre memory, documented data composition, and patterned absence. It can state that a source relation is unknown.

Death of an Author: genre memory without a source fantasy

The work openly solicits comparison with detective fiction. Marche's reported process involved rereading crime writers and asking for passages in established styles and publication forms (Comitta 2023). The title, names, plot, and rhetoric also create visible relations with literary theory and the detective tradition.

These documented choices give archive reading a firm starting point. The novella does not arrive only from an unknowable statistical past. A human planner named traditions, recognized their signals, and used model outputs to assemble a contemporary mystery. Genre memory enters through at least three routes: Marche's reading, his instructions, and the systems' learned patterns. The public record cannot assign a percentage to each route.

This is where source fantasy becomes tempting. A critic hears a Chandler-like cadence and imagines Chandler's novels inside the model. Yet the resemblance

may arise from direct prompt wording, cultural descriptions of "Chandleresque" prose, other writers who adopted similar conventions, the human editor's selection, or some combination. Unless a training record establishes membership, the specific corpus claim remains unknown.

The bounded conclusion is richer than an accusation. *Death of an Author* packages detective fiction as a repertoire of reproducible signals while making the ownership of those signals part of its cultural controversy. Its archival force lies both in what can be named and in what cannot be traced. Interpretive opacity does not erase cultural memory. It changes the critic's warrant for speaking about it.

Operation four: infrastructure reading

Infrastructure reading turns toward the systems that make the textual event possible and often recede from the page. These may include model providers, cloud services, application interfaces, content policies, subscription tiers, publishers, audio platforms, data labor, electricity, water, semiconductor supply chains, and legal regimes.

The purpose is not to attach a standard paragraph about planetary harm to every reading. Infrastructure matters when it conditions the text's creation, form, circulation, availability, or reception.

Questions

Which services, devices, accounts, interfaces, and institutions were necessary?

What did each interface permit, block, retain, or obscure at the relevant date?

Did privileged or early access shape what the creator could make?

Which workers and organizations maintained data, moderation, annotation, publishing, and distribution?

Which legal and policy conditions affected disclosure, ownership, access, or redress?

What system-level material evidence is relevant, and what cannot be allocated to this text?

Acceptable evidence

Acceptable evidence includes dated provider documentation, interface records, terms and policies, institutional reports, statutes and judgments by jurisdiction, publishing records, labor research, and measurements with clear system boundaries. A critic should give ranges and dates for energy or water evidence and state whether a number concerns a facility, workload, model, or projection.

Prohibited inferences

Do not calculate the electricity or water consumed by a book from a viral per-query estimate. Do not treat every worker in a global supply chain as a causal participant in one paragraph. Do not assume that present product features existed during composition. Do not turn a platform's marketing language into a technical description.

Permitted conclusion

Infrastructure reading supports a case-specific account of enabling and constraining systems. It can locate the text within wider material conditions without pretending to measure an unavailable per-text footprint.

Death of an Author: access as a literary condition

Pushkin's announcement noted that Marche had access to three writing platforms, some not publicly available (Pushkin Industries 2023). Such access helps explain the work's historical position and its claim to experimental distinction.

The novella required a writer, several commercial interfaces with different affordances, a publisher that framed and distributed the experiment, digital and audio formats, and a publicity environment in which a production percentage could attract attention. The constructed author name, cover, announcement, and afterword formed an expanded interface through which the public learned how to read the experiment.

Infrastructure reading modifies the aesthetic account. The work's variety does not arise from one system displaying a unified "voice." It arises in part from movement among systems selected for different tasks. Its coherence depends on human transfer and arrangement across platforms. Its public identity depends on Pushkin's decision to foreground this process. Here infrastructure acts at the level of style and authorship, not only at the level of electricity.

Material systems still matter, but the available evidence supports only a general placement. Model inference and cloud storage use energy and water within highly variable data-center conditions. Current institutional research warns against universal per-query figures because workload, hardware, cooling, location, grid, and utilization can change resource use by orders of magnitude (Lei et al. 2025). No public measurement isolates the resource use of *Death of an Author*. The critic can identify dependence, not calculate a footprint.

Operation five: responsibility reading

Responsibility reading asks who should answer for a synthetic textual event and for which decision. Distribution does not mean equality. A model, editor,

publisher, dataset curator, provider, user, and regulator occupy different positions. Some can intend and justify decisions. Some can change systems. Some receive revenue. Some possess only partial knowledge. Some are exposed to harms without meaningful control.

Questions

Who knew or could reasonably have known about the relevant risk or effect?

Who controlled the decision, system, selection, disclosure, or publication?

Who benefited financially, professionally, or institutionally?

Who had the capacity to prevent, correct, disclose, or repair?

Which harms or achievements are being allocated, and at what stage?

What legal jurisdiction and policy date govern any formal duty?

Acceptable evidence

Acceptable evidence includes documented roles, contracts and policies where available, decision records, correspondence, technical control, financial relations, statutory duties, and credible testimony. Textual evidence can identify what calls for an account. It cannot by itself prove what an actor knew.

Prohibited inferences

Do not divide responsibility according to the number of participants or the percentage of generated words. Do not treat a nonhuman system as morally responsible in the same sense as an institution or person capable of giving reasons and changing policy. Do not let distributed agency dissolve the greater duties of actors with greater control and benefit. Do not assign legal liability without naming jurisdiction, date, and authority.

Permitted conclusion

The reading supports a reasoned, issue-specific allocation of responsibility. It should identify uncertainty and distinguish moral, professional, legal, causal, and aesthetic responsibility.

Death of an Author: responsibility without arithmetic

For the novella's narrative plan, passage selection, final arrangement, and decision to publish, Marche bears substantial creative and professional responsibility. For the public 95 to 5 framing, Pushkin and Marche bear responsibility as the institutions and person who advanced it. For the design and documented operation of the writing services, their providers possess responsibilities proportionate to their control, knowledge, and benefit. For the final audiobook and publication package, editorial and publishing actors also matter. None of these allocations requires calling the model morally blameworthy.

The same matrix distributes credit. Marche's planning and judgment are central. Generated language is causally indispensable to the experiment as described. The providers built systems whose capacities shaped the available prose. Pushkin created the conditions for circulation and reception. A critic may value the resulting form while still asking whether data practices, imitation requests, or promotional claims were adequately disclosed.

Responsibility reading replaces "Is AI the author?" with "Who answers for which part of this event?" The second question can receive more precise answers.

Three provenance shocks

A synthetic text rarely arrives with perfect disclosure. Readers may learn about its making at different moments. The protocol must work when provenance changes, rather than assuming one stable information state. Three common situations can be called provenance shocks.

Disclosure before reading

In the first situation, a cover, platform label, classroom introduction, or publicity campaign announces machine involvement before the text begins. The label enters the reading as a paratext. It can direct attention toward repetition, anomaly, novelty, or deception. It may make a conventional metaphor appear diagnostic. It may also encourage readers to overlook the human work of planning and editing.

The critic should record the exact wording and placement of the disclosure. "Made with AI," "AI-assisted," and "95 percent AI-written" do not frame the same relation. Nor are they neutral technical descriptions. Each selects a unit of creative labor and an implied authorial story. Textual reading can temporarily bracket the label, but reception analysis must return to it.

Disclosure after reading

In the second situation, the reader encounters the text without production information and learns about machine involvement later. The disclosure can reorganize memory. Fluency may be recoded as imitation. A moving passage may acquire an ethical discomfort. A disliked passage may be dismissed as predictably synthetic.

The first response and the later judgment should be recorded separately. Retrospective explanation is not evidence that the initial feeling was false. Nor does an unchanged judgment prove that provenance is ethically irrelevant. Chapter 9 develops this distinction through reception studies. The methodological point here is simple: time of disclosure belongs in the evidence record.

Contested disclosure

The third situation is harder. A creator denies machine use while a detector, reviewer, or social-media audience alleges it. The protocol does not authorize the critic to choose whichever story fits the prose. A detector score is probabilistic and corpus-bound. Intuition about "AI style" is vulnerable to genre expectations, linguistic prejudice, and hindsight. A denial may also be incomplete.

The critic should state the conflict, assess the quality of each source, and preserve the unknown when no production record resolves it. Textual reading may still examine why certain features attracted suspicion. Responsibility reading may evaluate the consequences of an unsupported public accusation. The method can interpret detection anxiety without pretending to solve provenance by instinct.

Death of an Author produces a milder version of the first shock because its machine involvement was central to promotion. Readers were invited to watch for seams. The afterword then supplied a second layer of process narrative. A responsible analysis distinguishes what the fiction offered, what the advance publicity framed, and what the later explanation encouraged readers to notice.

The protocol under incomplete evidence

Critics often face a thin record: a finished text, a product label, and perhaps one interview. The five operations still apply, but their conclusions contract.

With only the text, the critic can produce a textual interpretation and identify questions for the other readings. With a dated disclosure, production reading can establish at least one attributed relation. With provider documentation, infrastructure reading can describe product-level affordances at the relevant date, though not the creator's exact use. With no training record, archive reading stays with demonstrable intertexts and genre. Responsibility reading can still evaluate the act of publication or accusation when those decisions are documented.

The method should never reward opacity by declaring that nothing can be said. It should also not reward the critic for filling every box. A blank cell can be the most accurate result. The completed worksheet is not the one with the most prose. It is the one in which every conclusion fits the evidence entered beside it.

How the five readings change one another

If the protocol stopped after textual reading, *Death of an Author* would appear as a self-reflexive detective novella whose generic fluency is interrupted by visible seams. Production reading places those seams within a multi-system

process governed by human planning and selection, though the public archive cannot reconstruct every operation. Archive reading locates detective style in named traditions, prompts, and learned genre patterns while specific training sources remain unknown. Infrastructure reading makes privileged access, platform specialization, publication, and disclosure part of the work's form. Responsibility reading assigns decisions to actors according to control rather than word percentages.

What emerges is an account of a **synthetic textual event**, not a verdict that the book belongs to either human or machine literature. Its meaning and value emerge through relations among prose, genre, instruction, systems, selection, paratext, platforms, and readers. The account remains revisable if a prompt archive or fuller production record appears.

A second complete demonstration: *Aum Golly 3*

The first demonstration began with a novella whose public archive consists largely of an edition, afterword, publisher account, interviews, and sample instructions. *Aum Golly 3: Perfectly Fine Poems on Humanity by Artificial Intelligence* offers a different evidentiary configuration. Jukka Aalho's 2025 project has public architecture notes, source code, and a Zenodo deposit containing job artifacts, variants, feedback, selections, and translations. The documented session produced 1,260 variants across 253 jobs, saved 98 candidates, and approved 41 poems for publication (Aalho 2026; Aalho n.d.).

The record is unusually substantial, but it is not a total origin. The depositor describes it as the complete project corpus for the eight-hour session. It does not include provider-side logs, proprietary model weights, complete training corpora, hidden safety layers, or an independent record of every infrastructural event. The demonstration therefore asks what becomes interpretable when alternatives and approvals survive, while keeping the absent provider record absent.

Operation one: textual reading

Questions. What reading position does the title create? How do "Perfectly Fine," "Poems on Humanity," and "by Artificial Intelligence" place adequacy, subject matter, and attributed production beside one another? Does the modest evaluative phrase invite the reader to test the poems against lowered expectations, irony, or promotional self-awareness? At the level of an individual poem, which pattern in syntax, lineation, image, or sequence would deserve tracing through the archive?

Evidence used. The fixed object is the English 2025 volume, its documented title, its status as a collection of 41 approved poems, and its public production

framing. No poem is reproduced here because quotation clearance for the commercial edition remains open. The demonstration therefore uses the inspectable title and collection structure, then states the evidence needed for a poem-level extension.

Prohibited inference. The phrase "by Artificial Intelligence" cannot establish that one system wrote autonomously, that the software roles shared a communicative intention, or that every line is an untouched output. A title's ironic or modest tone cannot reveal which model generated a poem. The number of published poems cannot prove stylistic consistency or quality.

Bounded conclusion. The title asks readers to judge a poetry collection through an explicit production claim and a deliberately restrained quality claim. Forty-one texts appear as one volume under that frame. This conclusion concerns how the edition addresses its reader. It does not yet explain how the poems were made.

How the next operation revises the result. Production reading replaces the singular agency implied by "Artificial Intelligence" with a sequence of documented operations. It asks how several model roles and a human approval gate produced the collection that the title presents as one object.

Operation two: production reading

Questions. Who set the project constraint and initiated the session? Which roles generated, criticized, edited, maintained continuity, proposed ideas, or translated? How many variants moved through each stage? Who saved candidates, approved poems, prepared editions, and decided that the session was complete? Which of these relations are visible in project records, and which remain the creator's account?

Evidence used. Aalho documents a custom Python and Streamlit pipeline with generator, critic, editor, continuity, ideation, and translation roles. The system could produce three variants at a generation step, select one for refinement, branch, and repeat critique-edit cycles. The deposit records 1,260 variants, 253 jobs, 98 saved candidates, and 41 approved poems. Aalho retained final human approval and publishing authority (Aalho 2026; Aalho n.d.).

Prohibited inference. Role names cannot be treated as persons. A "critic" label does not show aesthetic judgment in the human sense, and an "editor" label does not establish a professional intention. The counts do not measure percentages of authorship. They do not show that the approved poems were the best possible outputs or that every model family had an equal opportunity to contribute.

Bounded conclusion. The volume emerged from a documented selection funnel in which generated alternatives greatly exceeded published poems. Model-mediated critique and revision operated at intermediate stages, while Aalho controlled project design, saving, final approval, edition preparation, and publication. Selection is not a marginal afterthought in this case. It is a recorded sequence of gates.

How the next operation revises the result. Archive reading tests whether the production narrative can be followed through surviving artifacts. It may confirm a relation among seed, variants, feedback, revision, translation, and approval for a specific poem. It may also reveal that a project-level count cannot explain a local phrase.

Operation three: archive reading

Questions. Which job artifacts link a seed or instruction to variants, feedback, revised candidates, translations, and approval? Does the archive preserve rejected alternatives as well as successful ones? Can a formal feature in one poem be traced through its particular artifact chain? Which cultural sources are named, and which training relations remain unavailable?

Evidence used. The Zenodo record deposits project-level job artifacts, variants, feedback, selections, and translations under DOI 10.5281/zenodo.18205498. Public code and architecture notes document the locally orchestrated workflow. The project record supports model-specific usage and approval counts at the scope defined by the depositor. Its licence is CC BY-NC-ND 4.0, so this book paraphrases the archive rather than reproducing or transforming its contents.

Prohibited inference. The deposit cannot be called a controlled comparison among models. Systems, prompts, roles, translation paths, and curation opportunities were not held constant for a quality experiment. A project-level approval count cannot explain the origin of one line without its artifact chain. The deposit does not expose provider training data, and resemblance to a writer cannot fill that gap. It also cannot make the thinner records for volumes 1 and 2 complete after the fact.

Bounded conclusion. Some selection and revision claims can be tested against deposited alternatives rather than accepted only as retrospective publicity. The archive makes it possible to trace how particular candidates survived within the project. It remains bounded to the deposited session and cannot recover proprietary provider histories or title-level training membership.

How the next operation revises the result. Infrastructure reading asks what the deposit itself depends on and excludes. A local repository can make

orchestration visible while the systems supplying generation remain distributed across proprietary providers. Archive richness at one layer can coexist with opacity at another.

Operation four: infrastructure reading

Questions. Which local and remote systems were necessary? What depended on Python, Streamlit, repository hosting, model-provider access, archive hosting, translation, book production, and distribution? Which providers supplied the documented model families? Did product access shape the range of roles and variants? What system-level material or labour evidence can be connected to this case, and what cannot be allocated to the book?

Evidence used. The public architecture identifies local orchestration and several documented model families: GPT-5, multiple Claude versions, and Gemini versions. GitHub hosts public code and Zenodo hosts the deposited corpus. Aalho organized an eight-hour English and Finnish production session, prepared the volume, and published it through Kertojan aani. These records establish institutional and technical dependencies at product and project level.

Prohibited inference. The use of several providers does not yield a per-poem electricity, water, mineral, or labour figure. The archive does not expose data-centre location, hardware, utilization, cooling, provider-side routing, or complete supply chains. Current service behaviour cannot be projected backward onto the November 2025 session. A locally orchestrated pipeline is not provider independence.

Bounded conclusion. The collection depended on a local editorial apparatus joined to several proprietary generation services, public code hosting, archival infrastructure, translation work, and publishing channels. This plural infrastructure helps explain why the project could distribute functions across named roles and retain many variants. Its material cost for this volume remains unmeasured.

How the next operation revises the result. Responsibility reading separates causal dependence from answerability. It asks who could alter the pipeline, who controlled final inclusion, who described the archive, who operated the proprietary systems, and who published the resulting claims.

Operation five: responsibility reading

Questions. Who had knowledge of the relevant decision? Who controlled project constraints, system access, selection thresholds, final approval, translation, disclosure, and publication? Who benefited from the work's circulation? Which actor could correct an inaccurate process claim or repair a

rights problem? What responsibility belongs to providers, the human editor, repository hosts, or the publisher for a particular issue?

Evidence used. Aalho designed and initiated the project, orchestrated the systems, saved candidates, approved the final 41 poems, prepared the editions, described the process, and published the collection. Model providers controlled their services and undisclosed provider-side conditions. GitHub and Zenodo supplied code and archive hosting under their own institutional rules. The role labels inside the pipeline identify causal operations but not moral agents.

Prohibited inference. Responsibility cannot be divided according to 1,260 variants, 41 approvals, model-specific counts, or the number of named systems. A software role cannot answer for the collection in the same sense as a person or institution able to explain and change a decision. Distributed causation cannot erase Aalho's final approval, and final approval cannot absorb every provider-level responsibility.

Bounded conclusion. Aalho bears strong creative and professional responsibility for the constraint, selection system, final inclusion, framing, and publication. Providers bear responsibility for the systems and policies they controlled, though this record does not establish a case-specific wrong. Repository institutions bear narrower responsibilities for hosting and licence administration. The pipeline roles remain causally important without becoming moral authors.

How the next operation revises the result. The protocol now returns to textual reading. The title's singular "Artificial Intelligence" looks less like a technical inventory and more like a public frame laid over plural systems and unequal decision rights. A poem-level reading can then ask whether a formal feature survived one documented selection path, while still refusing to call that path the poem's complete meaning.

The five operations change one another without cancelling earlier evidence. Textual reading identifies the edition's challenge to the reader. Production reading replaces singular machine agency with a selection sequence. Archive reading tests parts of that sequence against variants and approvals. Infrastructure reading locates the visible archive inside less visible provider dependencies. Responsibility reading assigns answerability according to control rather than output volume. A return to the text makes the title and poems readable as the public form of those relations, not as a transparent transcript of them.

Stop and Reflect

Choose a short synthetic text whose production is publicly described. Write five sentences, one for each operation.

1. In the textual sentence, describe a formal pattern without naming a cause.
2. In the production sentence, state one documented act and attribute its source.
3. In the archive sentence, name one demonstrable intertext or genre relation.
4. In the infrastructure sentence, identify one enabling system that affects form or circulation.
5. In the responsibility sentence, name one actor, one decision, and one reason for the allocation.

Now underline every verb that claims knowledge: *shows, proves, reveals, suggests, reports, claims, documents*. Could the evidence bear the force of that verb? Revise any sentence that exceeds its warrant.

What a synthetic hermeneutic critic would do

A synthetic hermeneutic critic would make the following moves:

1. **Fix the object.** Record title, edition, format, date, interface, and disclosure condition before generalizing about "the text."
2. **Read before diagnosing.** Describe form and effect without using the AI label as a total explanation.
3. **Separate evidence classes.** Mark facts, attributed accounts, interpretations, and original proposals.
4. **Map transformations.** Trace planning, prompting, generation, selection, revision, arrangement, and publication only as far as records allow.
5. **Treat absence as bounded opacity.** Say what remains unknown and prohibit the inference that the missing record might tempt.
6. **Connect infrastructure to interpretation.** Include a platform, policy, labor relation, or material dependency only when it changes the production, form, circulation, reception, or responsibility of the case.
7. **Allocate responsibility unevenly.** Use knowledge, control, benefit, capacity, and preventability instead of equal shares or generated-word percentages.

Record A: object identity

- Title and creator or attributed byline:
- Edition, ISBN, URL, or archive identifier:
- Format and interface:
- Publication or access date:
- Model-related disclosure visible before reading:
- Later disclosure or contested claim:

Record B: claim ledger

- For every checkable statement, record:
- exact claim;
 - claim class: fact, attributed argument, interpretation, or proposal;
 - best source and location;
 - date, jurisdiction, model version, edition, or sample;
 - limitation or conflicting evidence;
 - permitted wording;
 - verification status.

Record C: five readings

For each operation, complete four boxes:

Record 1	Operation: Textual
Questions asked	
Evidence used	
Prohibited inference	
Bounded conclusion	
Record 2	Operation: Production
Questions asked	
Evidence used	
Prohibited inference	
Bounded conclusion	
Record 3	Operation: Archive
Questions asked	
Evidence used	

Prohibited inference	
Bounded conclusion	
Record 4	Operation: Infrastructure
Questions asked	
Evidence used	
Prohibited inference	
Bounded conclusion	
Record 5	Operation: Responsibility
Questions asked	
Evidence used	
Prohibited inference	
Bounded conclusion	

Record D: opacity statement

Write one paragraph beginning: "The available evidence does not establish..."
Then write a second beginning: "This absence permits the critic to conclude only..."
The exercise turns uncertainty into method rather than embarrassment.

Assessing a protocol exercise

A classroom response can be assessed without grading students on whether they like the synthetic text. Use five criteria, each worth two points:

1. **Object precision:** the response names the edition, format, date, and disclosure state.
2. **Textual warrant:** formal claims point to specific features and avoid hidden-production diagnoses.
3. **Evidence control:** facts, attributed accounts, interpretations, and unknowns remain distinct.
4. **Relational analysis:** at least one production or infrastructure condition is connected to a textual or reception effect.
5. **Responsibility precision:** the response assigns a decision to an actor and explains the allocation through knowledge, control, benefit, capacity, or preventability.

A response earns full credit for a well-supported refusal. If the question asks which model produced a passage and the record does not identify a model version, "unknown" may be the only accurate answer. The student should then explain what additional evidence would be needed.

A shorter classroom protocol

When time is limited, students can use five verbs:

Describe the language.

Reconstruct the documented making.

Locate the cultural archive.

Connect the enabling systems.

Allocate responsibility.

Each group receives one operation and must report one conclusion plus one statement it refuses to make. The refusals are pedagogically valuable. Students learn that critical confidence comes partly from knowing what the evidence cannot support.

Objections and replies

"This is source criticism with extra steps"

Source criticism asks where a text or passage comes from. The protocol includes that question but expands the object. A synthetic textual event may depend on inaccessible training data, nondeterministic generations, interface design, editorial selection, model policy, and post-publication disclosure. The method also distinguishes a source claim from infrastructure and responsibility. Its value lies in relating these domains without allowing one to stand in for another.

"The protocol restores the author through the back door"

Production evidence does not restore a sovereign author. It replaces a single origin with a decision record. Human intention remains relevant where it is documented, but it neither closes meaning nor exhausts causation. The method can show that a creator planned an effect while also showing that readers, systems, paratexts, and institutions transformed it.

"Most provenance is inaccessible, so the method cannot work"

Complete provenance is rare for human and synthetic texts alike. The protocol does not require omniscience. It requires the critic to grade evidence and mark opacity. An unknown prompt history narrows the production conclusion, but textual, paratextual, legal, and reception evidence may still support a substantial interpretation. The method fails only if it treats the gap as permission to invent.

"Infrastructure overwhelms literature"

Infrastructure reading can become a catalogue of distant harms. The relevance test prevents this. A condition belongs in the chapter only if it changes what could be produced, how it appears, who encounters it, or who should answer for it. In *Death of an Author*, access to several specialized systems directly shaped composition and public identity. A generic inventory of minerals would not improve that specific reading.

"Five operations imply false completeness"

The protocol is a minimum, not a total theory of literature. It does not replace historical, psychoanalytic, feminist, Marxist, postcolonial, queer, formalist, or other methods. Those approaches may ask the most important question in a given case. The five operations ensure that machine mediation is neither fetishized nor ignored. Critics can add methods while preserving the evidence rules.

Limits of the protocol

The protocol is unnecessary when machine learning has no material role in generation, transformation, translation, structure, personalization, or reception. Ordinary word processing does not automatically create a synthetic text. A spelling correction may be too slight to justify the full method. Chapter 2 supplies the threshold test.

The protocol is also insufficient for several tasks. It cannot establish model consciousness. It cannot settle every copyright dispute. It cannot infer the contents of closed training corpora. It cannot replace empirical audience research with an imagined reader. It cannot calculate a text's resource footprint without measurements. It cannot decide literary value by checklist.

Its contribution is more modest and more practical. It coordinates critical questions that a synthetic textual event places in unusually tight relation.

Study note

Remember the distinction between the five objects:

text means the examined language and form;

production means documented acts and transformations;

archive means prior cultural materials, data, genres, and absences;

infrastructure means enabling and constraining systems;

responsibility means answerability allocated by power and position.

Do not let a claim migrate between objects without new evidence. A textual resemblance is not a production record. A production role is not automatically moral blame. A system-level resource estimate is not a per-text measurement.

Summary

The Synthetic Hermeneutic Protocol reads a synthetic textual event through five connected operations.

Every operation specifies questions, acceptable evidence, prohibited inference, and a bounded conclusion.

Textual reading describes form without diagnosing hidden production.

Production reading reconstructs documented acts while separating creator claims from independent evidence.

Archive reading traces genre and cultural memory without inventing training sources.

Infrastructure reading connects enabling systems to form, circulation, reception, or responsibility.

Responsibility reading allocates answerability unevenly according to knowledge, control, benefit, capacity, and preventability.

Applied to *Death of an Author*, the protocol replaces a simple human-versus-machine verdict with an auditable account of language, process, access, framing, and decision.

Annotated further reading

Colella, Silvana. 2025. "The Language of the Digital Air: AI-Generated Literature and the Performance of Authorship." *Humanities* 14 (8): 164. A sustained analysis of AI literary experiments, including the paratextual construction of Marche as an author-curator. Useful as a counterweight to publisher publicity because it asks how books teach readers to assign authorship.

Genette, Gérard. 1997. *Paratexts: Thresholds of Interpretation*. Translated by Jane E. Lewin. Cambridge: Cambridge University Press. A detailed account of titles, prefaces, bylines, and other thresholds. Synthetic disclosure labels and process afterwords extend, rather than replace, this field.

Van Hulle, Dirk. 2022. *Genetic Criticism: Tracing Creativity in Literature*. Oxford: Oxford University Press. A rigorous introduction to studying drafts and writing processes. Prompt logs and output histories require modifications to genetic method, especially where model states are nondeterministic or unavailable.