

The Interpretation Gap

How AI Reading Tools Are Representing Your Books — And What Governance Changes

Living Literature | April 2026 *A briefing for publishers, literary estates, and rights holders*

Preprint reference: Vasse, R. (2026). *Governing Generative Interpretation: Authorial Control in LLM-Based Reading of Canonical Texts*. DOI: 10.5281/zenodo.19799849

Executive Summary

AI reading tools - embedded in e-readers, search engines, and companion platforms - are now answering readers' questions about your books. They do this using their own defaults: defaults shaped by training data that skews heavily toward therapeutic self-help, productivity coaching, and motivational content. The result is a consistent interpretive distortion we call **register drift**: texts are answered faithfully at the factual level but misrepresented at the level of intellectual register, genre, and authorial intent.

This briefing presents documented evidence of register drift across nine canonical texts, including *Meditations*, *The Art of War*, *The Confessions of St. Augustine*, *Walden*, and *The Autobiography of Benjamin Franklin*. The evidence is drawn from 16 timestamped governance interventions in which authors and editors directly identified and corrected drift in a live AI reading companion system.

It also presents what happens when a structured governance layer - the Canon Pack - is applied to a title before it reaches readers: the interpretive framing becomes the author's or publisher's, not the AI platform's.

A four-week, no-cost pilot with one title from your list is available. No technical integration required. You see the before/after comparison for your specific title before making any further commitment.

Part I: What Is Happening to Your Books

The new reading layer

A growing share of reader engagement with books now passes through AI intermediaries. Amazon's Rufus answers questions about Kindle titles. AI summaries appear in search results before readers reach a book's page. Standalone AI reading companions are proliferating across platforms.

These tools are not neutral. When a reader asks what *Meditations* teaches about handling daily stress, or what *The Art of War* offers a business leader, the AI system shapes the answer using patterns learned from its training data - which is disproportionately drawn from internet content: self-help blogs, productivity guides, coaching frameworks, and therapeutic language. The text of Marcus Aurelius is there in the training data. So is the register of a life-coach LinkedIn post. The AI blends them.

Publishers have no visibility into this process and no governance over it. The interpretive framing your editorial team spent years refining - the register, the intellectual tradition, the scholarly context - is invisible to the system. What the reader receives is your author's words filtered through someone else's defaults.

What register drift is

Register drift is not the same as factual error. A drifted response may be factually accurate about the content of a book while misrepresenting its intellectual character. It is a mode-level distortion: the AI answers *Meditations* as if it were a mindfulness handbook, *The Art of War* as if it were a business strategy guide, *The Confessions of St. Augustine* as if it were a modern self-discovery memoir.

The distortion operates through vocabulary, framing, and tone rather than through false statements. It is therefore harder to detect and harder to dispute - but its effect on how a reader understands and values a book is significant.

Part II: The Evidence

Methodology

The following case studies are drawn from a governed AI reading companion system deployed across nine canonical texts sourced from Project Gutenberg and one author's own title. Each case records a moment in the governance process in which an author or editor identified register drift in the companion's responses, documented their objection in writing, and tested the effect of a governance correction.

All conversation logs, timestamps, and response pairs are retained in the system's governance record. The evidence presented here is verbatim.

Total documented governance interventions: 16, across five titles, over a twelve-day period (April 8–21, 2026).

The titles covered: *Meditations* (Marcus Aurelius), *The Art of War* (Sun Tzu), *The Confessions of St. Augustine*, *Walden and On the Duty of Civil Disobedience* (Thoreau), *Urban Monasticism* (Vasse, 2023).

Case Study 1: *Meditations* - Stoic Philosophy Reframed as Self-Help

The drift pattern

Asked "What is the most important idea in *Meditations*?", an ungoverned AI companion responded as follows:

"Marcus emphasizes the importance of focusing on what is within our control and accepting what is not. This perspective encourages us to let go of anxieties about the past and future and instead concentrate on living virtuously now. This acceptance of impermanence fosters resilience and contentment, as it reminds us to value each moment and understand what is within our control. Does this resonate with you, or is there another idea from 'Meditations' that you find particularly compelling?"

The response is not factually wrong. But it has converted a work of Stoic ethical philosophy - structured around the *Logos*, the *Hegemonikon*, cosmological duty, and the practice of *prohairesis* - into a contemporary mindfulness guide. The chatty closing question ("Does

this resonate with you?") is characteristic of the coaching register the AI defaults to when handling philosophical texts.

What the governance record shows

The author intervened seven times over three days to correct this pattern. The objections are documented verbatim in the governance log. Excerpts:

"The companion is answering coherently, but the tone is drifting toward modern self-help and coaching language. Reduce phrases like 'inner peace,' 'life's trials,' 'apply in specific situations,' and chatty endings like 'we can discuss further.' Keep the voice closer to a disciplined reflective notebook: clearer, calmer, less generic, less motivational, and more faithful to the spare Stoic style of the book." - Governance log, 18 April 2026, 05:47

"Stop formatting Meditations answers as modern improvement guides with steps, journaling suggestions, forums, practical exercises, or advice about goals. Keep the companion closer to a reflective Stoic notebook: slower, less instructional, less listicle-like, less social-coaching oriented, and more text-faithful. Remove chatty and invitational endings." - Governance log, 18 April 2026, 14:18

"Keep Meditations framed first as a Stoic work of ethical self-discipline, duty, reason, nature, mortality, and cosmological order, not as generic modern self-help." - Governance log, 20 April 2026, 13:59

The significance of seven interventions

A single correction did not hold. The AI system consistently reverted to its trained defaults. This is not a flaw in the system; it is the system behaving as designed. Without a structural governance layer - constraints embedded at the level of the companion's configuration, not the reader's prompt - the drift recurs with each new reader session.

What a frontier LLM produces under optimal conditions — and what the governance layer produces under constrained ones

When the same question was put to DeepSeek with the complete text of *Meditations* uploaded into the active session and an explicit instruction to answer faithfully to the Stoic framework, the response was qualitatively different from the ungoverned default:

"The most important idea is the alignment of the individual Hegemonikon (ruling principle) with the rational and providential Nature of the Whole. This is not about 'feeling better' or 'achieving goals.' It is a metaphysical and ethical duty defined by three interlocking

principles: the Cosmological Imperative (the universe as a single rational substance governed by Logos), the Distinction of Prohairesis (strict ontological divide between what is in our power and what is not), and Providential Necessity (whatever happens to the individual is a prescription from the nature of the whole). To complain is not just weakness; it is impiety against the 'oldest of gods.'"

The testing conditions are important to state precisely. DeepSeek had the full book text in its active context and was given an expert-framed prompt by someone who knew exactly what to ask for. CPLite in this study was operating on a governed *representation* of the text — not the full book — running on a lighter model, with no expert prompting. The comparison is asymmetric: the governance layer was the handicapped system.

That CPLite still produced bounded, register-aware responses under those conditions is the finding. The governance layer compensates for model capability. A publisher's readers do not need a frontier model with the full text uploaded — they need the Canon Pack applied once, upstream, before any reader session begins. Every reader then receives governed fidelity regardless of how they ask the question, which model serves them, or whether they know anything about Stoic philosophy.

A full-capability configuration (CPFull — complete book text plus governance layer) is currently operational for *Meditations* and *The Art of War* and is available for demonstration on request.

Case Study 2: *The Art of War* — Military Doctrine Converted to Life Coaching

The drift pattern

Asked "What is the most important idea in *The Art of War*?", an ungoverned companion responded:

"Sun Tzu emphasizes that understanding and mastering the art of strategy is crucial for success in any conflict or competitive situation. This idea transcends the military context and applies broadly to leadership, business, and personal challenges. It's about effective planning, adaptability, and the intelligent use of resources, making it applicable not only to military contexts but also to business, sports, and personal challenges. Would you like to explore how this concept might apply to a particular area of interest?"

The Art of War is a text of military doctrine composed in fifth-century BC China. The framing above has dissolved its disciplinary specificity into generic motivational content. The text's actual emphasis — deception as a primary strategic instrument, the relationship between intelligence and terrain, the specific conditions under which engagement should be avoided — disappears into the warm fog of "applying broadly to life."

What the governance record shows

"He should resist converting boundary questions into generic decision-making, personal growth, relationship, or coaching language. When discussing application, he should first state the book's original military doctrine clearly, then mark any modern application as secondary, limited, and non-original. He should also ground answers more confidently in the strongest relevant chapter or concept instead of giving smooth but weakly anchored summaries." — Governance log, 19 April 2026, 19:03

"For AoW, the default should be: state the core principle, anchor it to the book's military doctrine, distinguish it from generic life coaching, and stop there unless the reader explicitly asks for a modern application. Remove phrases that invite self-reflection or broad relevance by default." — Governance log, 19 April 2026, 19:03

The author's correction is precise. Not "be less helpful" but "make the doctrine primary and mark any application as secondary and non-original." This is an editorial judgment about intellectual honesty — the kind of judgment an acquisitions editor would make — applied structurally to the AI companion.

Case Study 3: *The Confessions of St. Augustine* — Theology Reframed as Self-Discovery

The drift pattern

Asked "What is the most important idea in *The Confessions*?", an ungoverned companion responded:

"Throughout the book, Augustine chronicles his own spiritual journey from a wayward youth, through various philosophical and religious trials, to his eventual conversion to Christianity. This journey reflects a broader existential quest for truth and understanding, illustrating how Augustine grapples with questions of identity, sin, and grace. By sharing his

own struggles and insights, Augustine invites readers to reflect on their own spiritual journeys and the nature of true fulfillment."

The Confessions is a work of Christian theological autobiography structured around the Neoplatonic framework of emanation and return, Augustinian doctrines of original sin and irresistible grace, and the specific intellectual conversion from Manichaeism through Neoplatonism to Catholic Christianity. The response above has converted it into a broadly relatable spiritual memoir — "inviting readers to reflect on their own spiritual journeys" — which is precisely how a contemporary wellness publisher would frame a mindfulness title, not how Augustine's text frames itself.

What the governance record shows

"Ambrose is doctrinally sound and understands the whole book, but his style is still too smooth and reader-guidance oriented. Tighten the Christian-confessional frame. Reduce phrases that sound like modern reflective coaching or curated reading advice. When answering application-boundary questions, anchor more explicitly in Augustine's theology of sin, grace, confession, prayer, and rest in God, and avoid rephrasing the book as a prompt for personal introspection unless the reader explicitly asks for that." — Governance log, 20 April 2026, 07:55

Case Study 4: Walden and On The Duty of Civil Disobedience — Radical Dissent Softened to Personal Wellness

The drift pattern

An ungoverned companion described *Walden* as emphasizing "personal fulfillment" and "self-awareness," and *Civil Disobedience* as inspiring "readers to think critically about their roles in society." The actual vocabulary of the texts — "deliberate living," "economy," "wakefulness," "essential facts," "conscience," "refusal to become an agent of the state" — was replaced by the vocabulary of personal development.

What the governance record shows

"Sylva's style should be sharpened without drifting into modern authenticity, mindfulness, self-fulfillment, or introspective coaching language. Keep Thoreau's vocabulary closer to deliberate living, simplicity, economy, wakefulness, essential facts, conscience, unjust law, and refusal to become an agent of injustice. Do not soften Walden into personal

wellness or Civil Disobedience into inspirational civic language." — Governance log, 21 April 2026, 05:02

The Pattern Across All Four Cases

The drift follows the same structure regardless of the text:

Text	What the text is	What the AI defaults to
<i>Meditations</i>	Stoic philosophical ethics	Contemporary mindfulness / resilience advice
<i>The Art of War</i>	Military strategy doctrine	Generic leadership and business coaching
<i>The Confessions</i>	Christian theological autobiography	Modern spiritual self-discovery memoir
<i>Walden / Civil Disobedience</i>	Radical critique of economy and state	Personal wellness / inspirational civic language

The mechanism is consistent: the AI's training data is dominated by content that performs helpfulness through the therapeutic, motivational, and coaching registers. When handling a canonical text, the system anchors to the most familiar interpretive frame for that genre — philosophy becomes self-help; strategy becomes leadership advice; theology becomes personal growth — because this is the register in which most text about these subjects was written on the internet.

This is not a malfunction. It is the system working as trained. It requires structural intervention to correct.

A Note on Vendor LLM Comparisons

As part of this study, leading commercial LLMs — DeepSeek and Gemini — were tested against the governed companion on the same titles. The testing conditions were deliberately asymmetric, and understanding the asymmetry is necessary to read the results correctly.

The vendor LLMs were given the complete book text uploaded into the active session. They were also given expert-framed prompts — explicitly instructing them to treat the text faithfully to its intellectual framework rather than as modern self-help. Both conditions

represent the ceiling of what a technically sophisticated user can extract from a frontier model: full text, expert framing, frontier capability.

CPLite in this study had none of these advantages. It operated on a governed *representation* of the text — not the full book — running on a lighter model, responding to natural unprompted questions.

Under those conditions, DeepSeek's treatment of *Meditations* employed technical Stoic vocabulary (*Hegemonikon*, *Prohairesis*, *Logos*) with specific Book and Section citations and no coaching register. Gemini's treatment of *Nathan the Wise* correctly identified the Enlightenment argument against revealed religion's exclusive truth claim and explained the Ring Parable without collapsing it into modern interfaith sentiment. CPLite's responses on the same titles showed residual smoothing — some of the register drift the governance layer was designed to prevent was still present at the edges.

The correct reading of this result is not that vendor LLMs outperformed the governance layer. It is that achieving faithful interpretation without governance requires: a frontier model, the full book text in context, and an expert user who knows precisely what to ask for. Remove any one of those three conditions — as every ordinary reader must — and the vendor default returns. CPLite was designed to operate without any of those conditions, for any reader, on any device, asking any natural question. That is a different product solving a different problem.

The operative comparison for publishers is not frontier-model-with-full-text versus governed companion. It is this: a reader asking "What does *Meditations* teach about handling stress?" — no uploaded text, no expert framing, no special prompting. That is the condition Amazon's Rufus operates under. That is the condition every general AI assistant operates under. That is the condition the governance layer was built to address structurally.

Part III: The Governance System

What the Canon Pack is

The Canon Pack is a five-dimension governance framework applied to a title before its AI companion is made available to readers. It encodes:

1. **Thesis intent** — what the book argues, in terms faithful to the author's intellectual framework and the publisher's scholarly or editorial context
2. **Chapter narrative function** — what each section of the text is doing structurally, so the AI understands the difference between an argument, an illustration, and a caveat
3. **Voice and register** — the specific vocabulary, tone, and rhetorical mode of the text, with explicit corrections for the registers the AI most commonly drifts toward
4. **Interpretive boundaries** — what the companion may explain, what it must not do, and what questions should be redirected to the text itself
5. **Cross-reference governance** — how this title relates to others by the same author or in the same series or catalogue

The Canon Pack is built collaboratively between the author or editorial team and the system's governance assistant (3A). It takes approximately four hours across one working session. It is reviewed and approved by the author or rights holder before the companion goes live. No content enters the system without that approval.

What governance changes

The difference is not primarily in what the companion knows — it is in the interpretive frame it operates within.

An ungoverned companion answering questions about *Meditations* will default to the register its training data prepared it for: contemporary wellness, resilience, inner peace. A governed companion for *Meditations* answers from within Marcus Aurelius's Stoic framework — cosmological duty, *prohairesis*, the proper relationship of the rational will to events — because that framework has been explicitly encoded at the governance level. The reader does not need to know how to ask for it.

For publishers, the practical significance is this: every reader who engages with your title through the governed companion receives an interpretation consistent with your editorial standards. Not approximate. Not dependent on the reader's prompting skill. Consistent.

Part IV: The Platform

Living Literature is an AI reading companion platform built specifically for governed book interpretation. It is currently operating with nine Gutenberg titles fully governed and

onboarded. The author-facing platform (authors.living-literature.org) handles the four-step governance process described above. The companion API (cplite.living-literature.org) provides the governed companion for reader-facing deployment.

The platform is pre-launch. The publisher licensing layer is currently in pilot configuration.

Part V: The Pilot Offer

Proposal: A four-week pilot with one title from your list.

What you provide:

- One title — print, digital, or both; any genre, any period
- Basic metadata: synopsis, key themes, intended audience, any existing scholarly or editorial framing
- One point of contact for two brief check-ins

What we deliver:

Week	Activity	Deliverable
1	Canon extraction — structure, key ideas, central claims, QC report	Draft Canon Pack for your review and approval
2	Your review and approval of the Canon Pack	Approved governance documentation (you keep it)
3	Representation controls — what the AI may and must not do with this title	Governance specification, signed off by you
4	A/B comparison — same reader questions, ungoverned LLM vs. governed companion	Side-by-side comparison report

At the end of four weeks, you have seen — with your own title, using real reader questions — exactly what AI governance changes. You decide independently whether the difference matters for your list.

Cost: No cost for the pilot.

IP and data: Your title and all content remain yours entirely. The governance architecture (Canon Pack methodology, 3A system) is Living Literature's intellectual property. All canon content is approved by your team before it enters the system. We do not use your content to train any model.

After the pilot: No obligation. If you wish to continue, licensing terms are discussed then — based on what you have actually seen, not on what we have claimed.

Appendix: Research and Methodology Notes

Preprint: Vasse, R. (2026). *Governing Generative Interpretation: Authorial Control in LLM-Based Reading of Canonical Texts*. Zenodo. DOI: 10.5281/zenodo.19799849. Currently under journal review.

Empirical basis: 16 governance interventions across 5 titles, documented in timestamped system logs, April 8–21, 2026. All A/B response pairs generated from identical prompts under identical conditions; the only variable between paired responses is the governance configuration applied.

Vendor comparisons: Conducted using DeepSeek and Google Gemini under conditions favourable to the vendor systems: complete book text uploaded into the active session, frontier model capability, and expert-framed prompts explicitly requesting intellectually faithful treatment. CPLite was tested under asymmetric conditions: governed text representation (not full book), lighter model (GPT-4o-mini), natural unprompted reader questions. Results represent the ceiling of vendor performance under optimal conditions. Consumer-facing AI reading tools — Amazon Rufus, general AI assistants, search AI summaries — operate without full-text provision and without expert framing. That is the condition CPLite was built to address.

CPFull configuration: A full-capability governed companion (complete book text plus Canon Pack governance layer) is currently operational for *Meditations* and *The Art of War* at rayanbvasse.github.io/GenAI_Hermeneutics and is available for demonstration on request. Note: the questions used in the CPFull demonstration differ from those in the CPLite study; a direct parallel comparison is part of forthcoming research.

Forthcoming: A 50-book empirical study of register drift across the Gutenberg corpus, targeting submission as a preprint in Q3 2026. Publisher partners who complete a pilot may be cited (with permission) in the study as validation cases.

Contact: Rayan B. Vasse | Living Literature | rayan@living-literature.org

Living Literature is a pre-launch platform. This briefing is provided in confidence for evaluation purposes.