

Terminal and Traversal-Dependent Texts

A Theory of What Reading Actually Costs

Flyxion

July 2026

Abstract

The claim that reading is in decline is probably true and largely beside the point. This essay argues that the relevant phenomenon is not a uniform decline in an activity called “reading” but a shift in the relative fortunes of two distinct modes of processing a segment of structured information: *recognition*, which classifies a segment against an existing predictive model, and *traversal*, which holds a segment’s interpretation open so that later material can revise it. A precise criterion is offered for when a text possesses non-trivial traversal structure, an account is given of why compression succeeds on some texts and fails on others, and the theory is extended past text to cover conversation, mechanical reasoning, and unauthored pattern-finding generally — at a cost the essay tries to state honestly rather than paper over. Traversal is shown to have both a rate, governed by local predictability, and a threshold, governed by a system’s tolerance for unresolved ambiguity independent of whether the resources to sustain it exist; the two failure modes this separates — collapsing under load and collapsing by choice — produce identical outcomes from opposite causes. The theory is also applied to composition rather than only reception: writing a traversal-dependent text is shown to be a kind of backward-looking construction, in which an earlier segment is placed because of a later one that already exists in the writer’s intention, and load-bearing status is argued to be negotiated between what an author plants and what a reader’s later, changing model does with it. At the cultural level, the essay favors a selection-pressure account over a replacement account of the present moment: recognition-stable objects are not displacing traversal-dependent ones so much as outcompeting them wherever the two must circulate under the same conditions, which leaves room for refugia — communities and institutions where that pressure is locally suspended — without weakening the underlying diagnosis. The essay closes by distinguishing an individual question, for which the theory has fairly concrete answers, from a civilizational one, for which it does not.

Aqra'a (أَقْرَأَ), the causative Form IV of the trilateral root q-r-hamza) — to cause one to read, to cause a thing to be versified so it is easy to recollect, to cause one to draw near after being away, to be delayed, to hold a thing until a month has passed and its meaning can be determined. One causative verb, five senses, one structure beneath them: to hold something until a later point resolves what it was.^a

^aThis causative form and its attested sense-range are drawn from two online lexicographic sources: a standard reference wiki (Wiktionary contributors, n.d.) and an aggregator drawing on several dictionaries, including *al-Ra'id*, *al-Mu'jam al-Wasit*, *Lisan al-'Arab*, *al-Qamus al-Muhit*, and *Mu'jam al-Lugha al-'Arabiyya al-Mu'asira* (almaany.com contributors, n.d.). The latter independently corroborates several of the senses used here and supplies additional instances of the same pattern (a star's setting, a wind's arrival at its appointed time, a year's delayed rainfall), all glossed as a thing reaching, or being held until, a determined point. The unification of these senses as "retention pending later resolution" remains the author's own synthesis and should not be attributed to either source or to the classical lexicographic tradition without further verification.

1 The Wrong Question

The claim that reading is declining is probably true and mostly beside the point. Fewer people finish books; more communication moves through fragments, feeds, and summaries; students report difficulty sustaining attention across long texts. None of this is seriously in dispute. But the diagnosis usually offered — that a civilization built on literacy is dissolving into a "postliterate" one (Horowitch, 2026) — assumes reading was always a single activity, declining uniformly. It wasn't, and it isn't.

Text has not disappeared. If anything it is more abundant than it has ever been: more of it is produced every day, by more systems, than any prior era could have imagined. What has changed is something narrower and stranger. A certain kind of engagement with text is contracting even as the volume of text explodes. The interesting question is not whether people still read. It is what, specifically, they are no longer doing when they don't.

2 Two Ways to Process a Segment

Any encounter with a structured sequence — a sentence, a paragraph, a line of log output, a gesture in a conversation — can be handled one of two ways.

Definition 1 (Recognition). *A segment is processed by recognition when it is classified against an existing predictive model and the process moves on without revision. No earlier state need be*

held open, because the segment's interpretation is settled at first contact.

Definition 2 (Traversal). *A segment is processed by traversal when its interpretation cannot be settled at first contact and must instead be held in a provisional state so that a later segment can confirm, qualify, or overturn it.*

Most of what anyone reads in a day is processed by recognition, and this is not a failure of attention — it is the correct response to material that does not require anything else. Traversal is expensive precisely because it requires retaining an interpretation in a provisional state, sometimes for a long stretch, while further material is processed against it (Kintsch, 1998; Zwaan and Radvansky, 1998).

These are not properties of a medium. They are properties of a relationship between a segment and the segments around it. The same document can contain long stretches of one and short, decisive intrusions of the other.

3 A Criterion

The distinction can be stated more precisely. Let a text be a sequence of segments $s_1, s_2, \dots, s_n$, and let $I(s_j \mid s_1, \dots, s_{j-1})$ denote the set of interpretations admissible for segment s_j given everything that precedes it.

Definition 3 (Load-bearing segment). *A segment s_i is **load-bearing** with respect to a later segment s_j ($j > i$) if*

$$I(s_j \mid s_1, \dots, s_{i-1}, s_i, \dots, s_{j-1}) \neq I(s_j \mid s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_{j-1}),$$

that is, if removing s_i changes the set of admissible interpretations available for s_j .

A text possesses non-trivial traversal structure exactly to the degree that it contains load-bearing segments¹ — places where later material performs an operation on an interpretation established earlier, rather than simply adding to a growing list of facts. This criterion is related in spirit to discourse theories that represent coherence through functional relations among textual spans, though it differs by defining dependence counterfactually through deletion (Mann and Thompson, 1988).

¹This phrase gestures at a magnitude the definitions above do not supply. The natural candidate would be something like $T(\text{text}) = \sum_{i < j} \mathbf{1}[I(s_j \mid \dots, s_i, \dots) \neq I(s_j \mid \dots, \hat{s}_i, \dots)]$, a simple count of load-bearing pairs, or a weighted variant sensitive to the size of the interpretation-set change rather than merely its presence. Nothing in this paper depends on such a measure, and it is left undeveloped here; the present argument needs only the binary distinction between texts that contain load-bearing segments and texts that do not.

Example 1. *Consider:* Every system accumulates complexity over time. Therefore older systems are harder to repair. *A summary of this claim loses essentially nothing:* older systems are harder to repair *is the whole of it.* *But extend the argument:* Every system accumulates complexity over time. However, much of what appears to be complexity is really retained history. Retained history can reduce repair cost because it preserves diagnostic information. Therefore age alone does not predict repair difficulty. Systems that erase their history may become harder to repair despite being simpler. *A summary — repair difficulty depends on how much history is retained — preserves the conclusion but destroys something else: the sequence in which one explanation was proposed, partially rejected, and replaced by a more precise one.* *A reader who receives only the endpoint has the proposition without the operation that produced it, and will often report understanding the words while not understanding the argument.*

Compare a genuinely terminal case — a shopping list, a weather report, a phone-book entry. Nothing later in these acts on anything earlier; order is arbitrary; a summary is not a compression of the object but a full restatement of it.

Definition 4 (Terminal text). *A text is **terminal** if it contains no load-bearing segments: for every $i < j$, $I(s_j \mid \dots)$ is unaffected by the presence or absence of s_i .*

The criterion draws a real line. Some texts are defined by their contents, and a faithful list of those contents is equivalent to the original. Others are defined by the trajectory through which their contents become admissible, and no list of endpoints can substitute for having traveled it. We will call the latter **traversal-dependent**.

4 Why Compression Fails Selectively

This explains a pattern that is otherwise puzzling: why some summaries — of a manual, a news report, a reference table — are perfectly adequate, while others — of a proof, a legal opinion, a novel — leave the reader with propositions but no comprehension. It is not that literary or argumentative material is summarized *badly*. It is that these texts are traversal-dependent rather than terminal, and a traversal-dependent text does not have a smaller equivalent form. Its content and its trajectory are not separable. Compressing it does not shrink the object; it produces a different object that happens to share the same endpoint.

This is also, incidentally, why some machine-generated summarization feels adequate and some feels hollow in a way that is difficult to name. A summarizer optimized to preserve propositions has no way to distinguish a load-bearing segment from a decorative

one, because that distinction is not visible locally by the definition above — it can only be found by checking whether some later segment’s admissible interpretation set actually depends on it. A terminal text survives this indifference. A traversal-dependent one does not. Research on discourse-aware summarization supplies a narrower computational analogue: systems that treat documents merely as sequences of sentences often lose coherence, while models incorporating rhetorical or discourse relations perform the task more successfully (Kikuchi et al., 2014; Xu et al., 2020).

5 Traversal Has a Rate and a Threshold

Reading speed is not a single number a person carries around; it is a continuously re-estimated function of local predictability. This claim is consonant with surprisal-based accounts of incremental language processing, according to which less predictable words tend to require longer processing times (Levy, 2008; Wilcox et al., 2023). A fluent reader moves through familiar material — a routine compiler warning, a standard paragraph shape, an expected turn of phrase — at very high speed, because the model in use classifies it instantly and no revision is pending. The same reader slows sharply, sometimes to a crawl, exactly where a segment breaks the pattern the model expects: an anomalous log line, an unfamiliar clause, a sentence that does not fit the shape the argument has taken so far.

Skimming a long document by title, header, and transition is not a degraded form of reading; it is traversal concentrated at the points an author has signaled as likely to revise what came before, performed by a reader who does not yet have a fine-grained model of the content itself. An expert scanning ten thousand lines of routine output for the one anomalous line is doing the same thing with a much better-tuned predictive model, built from repetition, which is why an expert’s eye lands on the relevant line and a novice’s does not: the novice’s model of “normal” is worse, so surprise is a less reliable signal for them.

Pacing tools — anything that fixes the rate at which new material arrives, rather than leaving it to the reader’s own variable internal clock — are doing something specific: they are attempting to hold dwell time near whatever minimum a given revision requires to register as revision, rather than merely additional information passing through.

Rate is not the only variable governing traversal, though, and it is worth separating from a second one that looks similar but is not. A system can have every resource it needs — time, attention, a quiet room, sufficient background knowledge — and still abandon traversal early, not because it ran out of capacity to hold a branch open, but

because holding it open became uncomfortable enough to prune voluntarily.

Definition 5 (Closure threshold). *The **closure threshold** of a system is the point at which unresolved ambiguity is judged more costly than premature commitment, independent of whether the resources to continue existed. The term is adapted from research on the need for cognitive closure, understood as a motivation to obtain a definite answer and terminate ambiguity or further judgment (Kruglanski, 1990; Webster and Kruglanski, 1994).*

Resource failure and closure failure can produce the identical visible outcome — a branch disappears — while arising from opposite conditions: one because the system could not sustain it, the other because the system chose not to. The clearest evidence that these are different mechanisms is a case where resources are abundant and traversal still collapses. Some forms of conspiratorial inquiry may exhibit exactly this pattern: enormous informational effort combined with a low tolerance for keeping competing hypotheses alive, so that the search ends not when the evidence is complete but when uncertainty becomes unbearable. The skill cultivated by careful readers, researchers, and investigators is not primarily the capacity to hold more information. It is a raised closure threshold — the trained tolerance to keep several candidate interpretations open past the point where premature certainty would have felt like relief.

6 Relevance Is Retroactive

Load-bearing status, as defined in Section 3, is stated relative to a fixed pair of segments s_i and s_j within a text. But the reader applying this criterion is not fixed; the predictive model against which admissibility is judged keeps changing over the reader's lifetime. A detail can be classified as noise on first encounter — nothing in the reader's model at that time made it stand out — and reclassified as signal years later, once a different experience retroactively reveals what it was actually connected to. This is why people report having read something long ago that “didn't mean anything at the time,” only to find it becomes indispensable once a later encounter recontextualizes it. The general idea that textual meaning is actualized through an interaction between textual structure and a temporally situated reader has a substantial precedent in reader-response theory (Iser, 1978).

This gives curatorial acts — bookmarking, note-taking, following a source, keeping something open for later review — a different character than they are usually given. They are not simply records of what already mattered. They are bets, made under uncertainty, about what is likely to become load-bearing to a reader's *future* model, a model that

does not yet exist in its eventual form. A note is a compression authored by the same predictive system that will later be doing the retrieval, which is part of why a person's own old notes are often more useful to them than an external summary of the same material would be: the selection criterion and the retrieval criterion are the same system, assessed at two different times.

7 Beyond Text, and the Limit of the Theory

None of this is specific to written language. A conversation has exactly the same structure: an earlier remark's meaning can be revised by a later one — a clarification, a change of tone, a “wait, were you serious?” — with the traversal distributed across two live participants instead of fixed in a document in advance. Following an unfamiliar mechanism — tracing which cable connects to which pulley in a sealed device — requires the same kind of dependency-tracing as following a proof, just applied to a causal structure instead of a logical one. A hike down a familiar trail is recognition-speed until an unexpected fork breaks the pattern and forces attention.

Notably, at least one major linguistic tradition appears not to have confined the concept of reading to text in the first place. Several sources describing the Arabic causative *aqra'a* and its base form *qara'a* (أَرَقَّ, root أ ر ق) associate the root with senses that, in English, would be grouped under divination — reading an outcome from an object rather than from an authored text.² If that association is accurately reported, it is worth taking seriously rather than treating as a curiosity, because it names the theory's actual limit. Reading tea leaves, a palm, or a set of cards is traversal-shaped: the reader holds provisional structure open, revises it against further detail, and arrives at an interpretation. But there is no guarantee that the structure being traversed is really there. The diviner may be discovering a genuine pattern, or constructing one and mistaking the construction for a discovery. Nothing internal to the act of traversal distinguishes these cases.

The load-bearing-segment criterion of Section 3 works cleanly for an authored text, where a real author placed a real dependency between two real segments, and where the definition's counterfactual (removing s_i and asking whether $I(s_j \mid \dots)$ changes) has a determinate answer because the segments and their relationship were fixed by a

²This connection should be held with appropriate caution. It rests on secondary lexicographic sources rather than a classical Arabic lexicon consulted directly by the author, and an earlier line of inquiry in developing this essay mistakenly attempted to fold in material from a phonologically similar but etymologically distinct root (q-r-ayn, ع ر ق, the root of *qur'a*, قُرْءَانٌ, “a drawn lot,” as in drawing straws (Wiktionary contributors, n.d.)) before that connection was identified as unsupported and withdrawn. The reader should treat the augury connection discussed here as a plausible but unverified reading, not an established etymological fact.

single intentional process. It becomes much harder to apply to unauthored material — a stretch of static, a stranger’s face, the news — where the traversal may be finding order that a source deliberately built in, or may be finding order that the reader’s own model is supplying on its own. In such cases the criterion’s counterfactual is not well posed, because there is no fact of the matter about what would happen if a “segment” were removed from an object that was never authored as a sequence of dependent segments to begin with.

This is the cost of generalizing the theory past text: once “reading” names any process of holding a provisional interpretation open against further input, the same apparatus explains both a mathematician tracing a proof and a person seeing a face in a wall stain. The mechanism is genuinely the same. Whether it is tracking something real is a separate question the mechanism itself cannot answer.

8 What This Says About the Present Moment

The concern about postliteracy, restated with this apparatus, is not that traversal-dependent objects are being replaced by recognition-stable ones. Long technical lectures, specialized forums, book clubs, and years-deep collaborative documentation all still exist, and their existence is not a counterexample to the thesis — it would only be one under a replacement model, where the claim is that traversal-dependent material is disappearing. It isn’t. What is changing is which kind of object wins wherever the two compete for the same attention, under the same conditions of transmission.

Recognition-stable objects — logos, slogans, personas, meme formats, short clips — survive compression, survive partial attention, survive being reposted with the context stripped away, survive fragmentation into pieces that still make sense alone. Traversal-dependent objects generally do not survive any of that; take away the path and what remains is not a smaller version of the object but a different one. Put these two kinds of object into an environment that rewards whatever circulates fastest with the least friction, and the outcome is not that one kind vanishes. It is that one kind is *selected for* and the other becomes disproportionately dependent on environments where that selection pressure is weaker — a subscription context, a slow-media publication, a community organized specifically around sustained attention. These are not exceptions to the thesis. They are exactly what the thesis predicts: **refugia**, in which the ordinary transmission advantages of recognition-stable objects have been locally suspended.

Large-scale text generation intensifies this same selection pressure rather than reversing it. It does not, on this account, threaten traversal-dependent objects by replacing

them — it can produce them as readily as anything else. What it changes is volume: enormously more material now competes for the same finite traversal capacity, and in a crowded channel the advantage still runs toward whatever is cheapest to classify and pass along. More text does not imply more traversal. It can just as easily mean a larger recognition-stable majority drowning out a traversal-dependent minority that was never shrinking in absolute terms, only in relative visibility.

This reframes the classical comparison to the Library of Alexandria. That was a story of physical loss — texts that ceased to exist and so could no longer be consulted by anyone. The contemporary situation, if the selection-pressure account is right, is not that story repeating. Nothing need be destroyed. The archive can remain fully intact, everything in it technically retrievable, while the cultural mechanisms that reward traversal weaken relative to the mechanisms that reward recognition. The scrolls are not lost. They are outcompeted.

That reframing separates two questions that are easy to run together. The first is individual: how does a reader preserve traversal capacity under recognition-heavy conditions? This has real answers — offloading provisional state into notes rather than working memory, deliberately spending time in refugia where the ambient selection pressure is locally suspended, building enough recognition-speed fluency in a domain that traversal is reserved for what actually needs it, and raising one's own closure threshold by practicing delay rather than reflexively resolving discomfort. None of this is in question. A motivated reader can do all of it.

The second question is civilizational: what kinds of institutions and transmission environments remain capable of supporting traversal-dependent objects at all? The first question's answer is not an answer to the second. An individual organism adapting successfully to a shrinking habitat is not evidence the habitat is healthy — if anything, sophisticated individual adaptation is what a shrinking habitat produces. A refuge is defined by being exceptional; the more indispensable refugia become, the more clearly they document that ambient conditions have diverged from local ones. A society does not lose traversal because some of its members stop reading — plenty of readers can remain, doing all the right things, and the drift can proceed regardless. It loses traversal when the institutions that once made sustained dependency-tracking an ordinary, unremarkable part of a normal life become rare enough that maintaining the practice requires unusual effort, a specialized community, or deliberate resistance to the ambient incentive structure. At that point the archive is still there. The paths through it are still walkable. They have simply stopped being roads, and become trails, kept open by a shrinking number of people willing to walk them.

9 Writing as the Mirror Operation

Everything so far has been stated from the reader's side of the transaction, but the theory has an obligation to say something about the writer's side as well, since every traversal-dependent text was, at some point, deliberately built by someone.

To compose a load-bearing segment is to place a dependency that does not yet exist for the reader — to write a sentence whose function only becomes visible once a later sentence acts on it. This is a strange kind of authorship, because the writer is working simultaneously in two positions: constructing the earlier state as it will be experienced the first time through, while also holding the later revision that gives the earlier state its purpose. A well-built argument, in this sense, is not written forward. Its load-bearing segments are placed with the revision already in mind, then arranged so a first-time reader encounters them in an order that makes the later revision land. Weak argumentative writing often fails for exactly the reverse reason: the writer knows the conclusion but never goes back to plant the earlier segment the conclusion needs in order to feel earned rather than asserted. The reader receives an endpoint without a path, which is the same failure Section 4 identified in a summary — except here it originates with the author rather than with a compressor working after the fact.

This gives Section 4's account of compression a sharper edge. A summarizer that flattens a trajectory into a point is not the only way a traversal-dependent text can be reduced to something terminal; an author who never builds the trajectory in the first place produces the same object without needing anything to compress. And it complicates Section 6's account of retroactive relevance in a useful direction: an author cannot fully control which of their segments will turn out to be load-bearing, because that status is partly settled later, by readers, against predictive models the author does not have access to. A passage written as incidental color can become load-bearing decades later once a different reader, or a different era, supplies the later segment that makes it revise something. Correspondence, marginalia, and technical documentation not originally intended as arguments regularly acquire traversal structure this way, retroactively, without their authors having engineered any of it — which suggests load-bearing status is never purely authored. It is negotiated between what a writer plants and what a reader's later model does with it.

Large-scale text generation sits at an odd angle to this picture, worth stating plainly rather than resolving. A system trained to predict the next segment from the ones before it is, in one sense, doing nothing but placing dependencies — every word is conditioned on what came before. But conditioning on prior text is not the same operation as holding

a revision in mind before the segment that will need it is written; the local, forward-only prediction that produces fluent text is not obviously the same act as the backward-looking construction described above, in which an earlier segment is placed because of a later one that already exists in the writer’s intention. Whether the output of such a process can be genuinely load-bearing in the sense this essay has been using, or only resembles it locally without the dependency actually holding, is not a question the theory can settle from the outside. It would need to be checked the same way any text’s traversal structure is checked: by trying to remove a segment and seeing whether some later one stops making sense.

10 Conclusion

The theory offered here began from a narrow provocation — that “reading is declining” conflates several distinct claims — and has ended somewhere considerably more general: an account of when an ordered sequence of segments can be reduced to its endpoints without loss, and when it cannot. Recognition and traversal are not properties of texts, media, or historical periods. They are properties of the relationship between a segment and what depends on it, assessed by a reader whose predictive model keeps changing, under transmission conditions that reward one mode of processing more than the other. Reading, on this account, was never quite the right name for what was being lost. What is at stake is the ordinary cultural availability of a specific cognitive operation: the willingness to hold a branch open, sometimes for a long time, because something further along might still need it.

References

- almaany.com contributors (n.d.). أَقْرَأَ (*aqra’a*). In *Ma’ani.com Arabic–Arabic Dictionary*, aggregating *al-Mu’jam al-Jami’, al-Ra’id, al-Mu’jam al-Wasit, Lisan al-‘Arab, al-Qamus al-Muhit*, and *Mu’jam al-Lugha al-‘Arabiyya al-Mu’asira*. Retrieved 2026.
- Horowitch, R. (2026). The End of Reading Is Here. *The Atlantic*, August 2026 issue. Published online July 8, 2026.
- Iser, W. (1978). *The Act of Reading: A Theory of Aesthetic Response*. Johns Hopkins University Press, Baltimore.
- Kikuchi, Y., Hirao, T., Takamura, H., Okumura, M., and Nagata, M. (2014). Single document summarization based on nested tree structure. In *Proceedings of the 52nd Annual*

- Meeting of the Association for Computational Linguistics: Short Papers*, pages 315–320, Baltimore, Maryland. Association for Computational Linguistics.
- Kintsch, W. (1998). *Comprehension: A Paradigm for Cognition*. Cambridge University Press, Cambridge.
- Kruglanski, A. W. (1990). Motivations for judging and knowing: Implications for causal attribution. In Higgins, E. T. and Sorrentino, R. M., editors, *Handbook of Motivation and Cognition: Foundations of Social Behavior*, volume 2, pages 333–368. Guilford Press, New York.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3):1126–1177.
- Mann, W. C. and Thompson, S. A. (1988). Rhetorical Structure Theory: Toward a functional theory of text organization. *Text*, 8(3):243–281.
- Webster, D. M. and Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, 67(6):1049–1062.
- Wiktionary contributors (n.d.). لَرَقَأَ (*aqra’a*). In Wiktionary, *The Free Dictionary*. Retrieved 2026.
- Wiktionary contributors (n.d.). قُرِئَ (*qur’a*). In Wiktionary, *The Free Dictionary*. Retrieved 2026.
- Wilcox, E. G., Pimentel, T., Meister, C., Cotterell, R., and Levy, R. P. (2023). Testing the predictions of surprisal theory in 11 languages. *Transactions of the Association for Computational Linguistics*, 11:1451–1470.
- Xu, J., Gan, Z., Cheng, Y., and Liu, J. (2020). Discourse-aware neural extractive text summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5021–5031. Association for Computational Linguistics.
- Zwaan, R. A. and Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2):162–185.