

Superpositional Text

Against Premature Collapse in Human and Machine Language




Fig. 1. Ideverse schematic. Idesol (star): this paper; idenets (planets): primary concepts; ideons (moons): concepts dependent on their parent idenet. After Temte (2026).

Abstract

All written text, human and machine-generated, presents itself as singular and inevitable: each word sits in its position as though it were the only word that could have stood there. This paper argues that the singularity is a compression artefact, not a communicative ideal, a convention so deeply installed that it appears to be a property of language rather than a choice about how to display it. The argument runs through a single structural claim: three independently describable processes of text generation, the computational (an LLM's token distribution), the phenomenological (a writer's held candidates), and the hermeneutic (a reader's activated fields), all produce probability distributions over candidate words, and all are subjected to the same lossy compression into a single surviving token sequence at the moment of output. The compression destroys information the reader could use, because a passage that emerged from a forced single channel and a passage that emerged from a live branch carry different epistemic statuses even when they read identically on the page. The paper introduces superpositional text as the theoretical object and Delta Notation as the practical written convention for manifesting it, grounds the apparatus in Binary Bias at the lexical level, in Quantum-Ethical Decision Algebra at the token-selection level, and in the relational interpretability of the Fylgja framework, and supplies a taxonomy of branch points graded by reconvergence likelihood. The Old Norse kenning supplies the historical precedent for compound rather than singular reference; the paper is, where the form can carry the argument, written in the notation it proposes. This is a notation, not a position on how meaning travels through a channel: the pragmatic stakes of transmitting a held distribution are developed separately in Resolution Drift.

KEYWORDS

philosophy of language, computational linguistics, large language models, Binary Bias, hermeneutics, typography, interpretability, kennings

Author note

Storm Bjørn Flindt Temte is an independent researcher and founder of Srodingr ApS, a purpose-locked entity registered in Copenhagen and oriented toward improving humanity's chances of surviving the Great Filter. The author declares no external funding, no competing interests, and no data that could be made available. Contact for correspondence: srodingr@proton.me.

PERTURBATION DRAFT

This paper is released as a Perturbation Draft (Temte 2026e): a provisionally complete work published early to maximise the perturbation window, the span of time during which the argument is exposed to external challenge. The author invites critique, counter-evidence, identification of independent convergence with existing literature, and direct challenge to the central claim. The draft designation is not an apology for incompleteness but a methodological commitment to iterative knowledge production.

§1 The collapse problem

Read any sentence and notice that it presents itself as the only sentence that could have been written. Every word sits in its position as though it arrived inevitably, as though the writer reached for it and found nothing else within reach. This is, charitably, a Lie-to-Children: a simplification taught because the full picture would obstruct the lesson, which then escapes the nursery and colonises the whole of adult written communication (Pratchett, Stewart, and Cohen 1999).

I The word did not arrive inevitably. It was selected, frequently from among near-equals, and the selection has been hidden by the convention that each position on the page must display exactly one surviving token.

The paper's organising image is hydrological. When a river runs through a gorge it has no choice: the walls decide the water's path and it flows one way. When the same river crosses a floodplain it braids, splitting into channels that diverge, run parallel, and reconverge downstream. Written text currently presents every passage as a gorge. But the generative process behind the text, whether a neural network or a human mind, frequently crosses a delta, and the reader is given no way to tell a gorge passage from a delta passage. They read identically. The collapse to a single channel is total and unmarked.

This matters because the two carry different epistemic statuses. A sentence that emerged from a gorge (one channel was available and the writer took it) and a sentence that emerged from a delta (several channels were live and one was selected, or selected arbitrarily, or selected under pressure the reader cannot see) are different communicative objects wearing the same surface. The collapse-to-singularity destroys exactly the information that would let the reader distinguish them: how much was at stake in the choice, how close the runners-up were, whether the alternatives were synonyms or incompatible worlds. The project of this paper is to make the delta visible.

I

The Lie-to-Children is a genealogy as much as a category. One can trace it: children are taught to choose one word and commit to it because the alternative overwhelms a beginner; the lesson hardens into style-guide doctrine; the doctrine is inherited by editorial practice, then by typographic convention, then, most recently and least examined, by the interface design of language models, each layer receiving the simplification from the one above without re-asking whether it was ever true.

§2 Three generative processes, one compression artefact

The central structural claim is that three processes which generate text, describable entirely independently of one another and studied in three separate disciplines, share an identical architecture: each produces a probability distribution over candidate words, and each is forced through the same compression into a single output token at the surface. The compression is not native to any of the three. It is imposed downstream, at the moment of display, and the same single-token convention does the imposing in all three cases.

§2.1 The computational delta

During inference a large language model produces, at each token position, a probability distribution over its entire vocabulary. Multiple tokens frequently carry comparable probability before a sampling step or a deterministic argmax collapses the distribution to one surviving token. The distribution is real, it is computed, it exists for the duration of the forward pass, and it is then discarded. The interface presents the survivor as though it were the only candidate the model ever entertained.

This is a design choice rather than a technical necessity. The log-probabilities are exposed in most API contexts and could be surfaced at the interface for any token position the designer chose. The architecture's plurality at each step is not hidden because it cannot be shown; it is hidden because the convention that text displays one token per position was inherited unexamined from human typography, which inherited it from the Lie-to-Children.

² The model's hidden delta is the cleanest case of the general phenomenon precisely because, here, the discarded distribution is a literal numerical object we could print.

§2.2 The writer's held superposition

Human writers experience something structurally parallel. The moment before committing a word to the page often involves holding several candidates in active, simultaneous consideration, not auditioning them in sequence but experiencing them as co-present, each lighting a slightly different region of the intended meaning. The slash-writing convention that arises spontaneously in informal writing, marginal notes, and first drafts, *a necessary/overdue correction*, is the natural notation for this experience: the writer who reaches for both words at once and finds the distinction between them load-bearing writes both, and only later, under the convention, deletes one.

This felt co-presence is not uniform across writers. The corpus's somatic-semanticism work characterises a mode of language processing in which words are not informational labels for concepts but automatic body-state evocations, felt and frequently seen as multimodal events rather than decoded as abstractions; for writers at that end of the distribution the superposition of candidates is experientially vivid, a felt simultaneity of several almost-words, rather than an abstract awareness that synonyms exist. The pointer-region apparatus of *The Missing Middle* (Temte 2026af) gives the structure underneath the phenomenology: a word is not a region of meaning but a pointer to one, the regions are graded and multidimensional, and two near-synonymous pointers address overlapping but non-identical regions. To hold *important/significant* in superposition is to address the union of two regions and to feel the part where they fail to coincide. The existence of the superposition is general even where the felt intensity of it is not.

The convention of choose-one is learned and imposed, not discovered. Style guides enforce it, editorial practice demands it, pedagogy installs it early. The writer who submits *important/significant* is told to pick one, not because the reader could not hold both but because the convention has no slot for both. The Lie-to-Children has become the adult law.

§2.3 The reader's richer reception

When a reader encounters *important/significant* rather than either word alone, the semantic activation is not the activation of one word plus noise. Both regions light; the overlap and, more informatively, the non-overlap between them becomes part of what is received. The reader participates in the collapse, or, more precisely, is invited to hold the superposition rather than resolve it. This is not confusion. It is higher-bandwidth reception: the reader receives more of the writer's actual semantic state than single-token selection can transmit, and the reader's own context determines which collapse, if any, is most useful to them.

Reader-response theory anticipated part of this. Iser's gaps and Eco's *opera aperta*, the open work, both located an irreducible openness in the act of reading (Iser 1978; Eco 1989). But both located the openness in the *interpretation* of a fixed text: the words were settled and the reader's freedom lay in construing them. Superpositional text moves the openness one layer down, into the lexical surface itself, making explicit on the page the plurality that reader-response theory had to reconstruct behind a closed surface. The text is open before interpretation begins, because it has stopped pretending its words were inevitable.

² The point is not the sampling mechanism, which is well understood, but the display convention, which is not examined at all. A model that emitted, at high-entropy positions, *significant/substantial/considerable* rather than silently selecting one would not be malfunctioning. It would be reporting its own state more honestly than the single-token surface permits.

§3 Binary Bias at the lexical level

The assumption that each position in a text must contain exactly one token is Binary Bias (Temte 2026a) operating below the threshold of awareness, so far below it that the constraint reads as a property of language rather than a convention about its display. Binary Bias names the foundational cognitive tendency to compress continuous, conditional, multi-valued phenomena into sealed categories, with the sealing as the operative damage: once collapsed, the category resists the qualifications that would re-open it. The single-token convention is this bias instantiated at the smallest unit of written language. The lexical position is a continuous distribution over candidate words; the convention forces it to a sealed one; and the seal is what makes the alternatives unrecoverable to the reader, who never learns there was a distribution to re-open.

Superpositional text is therefore not a new theory bolted onto the corpus from outside. It is Binary Bias located at the lexical surface, the same foundational collapse that the corpus tracks through perception, cognition, decision, and translation, found operating one layer below all of them, in the words themselves. The notation proposed below is, in this framing, simply a refusal to seal: a way of leaving the lexical category open on the page so that the reader receives the distribution rather than the survivor.

§4 A taxonomy of branch points

Not every delta carries the same information, and a notation that surfaced all of them with equal emphasis would be noise. The amount of information a branch carries is a function of its *reconvergence likelihood*: how quickly the candidate channels rejoin downstream. The taxonomy that follows is borrowed from interactive-narrative theory, where the reconvergence of branching story paths is a well-developed problem, and it is best read as a continuous dimension rather than four sealed bins, the irony of which is not lost on a paper about refusing to seal categories. ³

A foldback branch is a surface crossroads: *big/large, important/significant*. The channels reconverge almost immediately; within a phrase the semantic trajectory is indistinguishable. Its information content is low and its transparency value is high. Surfacing it tells the reader that this was a gorge not because no alternatives existed but because the alternatives were near-equivalent, which is itself a true and useful fact about the writer's lexical space, a stylistic fingerprint visible nowhere else.

A branch-and-bottleneck is a trajectory crossroads: *however/moreover, despite/because of*. Reconvergence is medium-term; the sentence takes a different shape but the argument-level destination may be the same. Here the delta reveals a structural decision rather than a synonym choice. *However/moreover* is not two words for one relation; it is two stances toward the writer's own material, and surfacing both tells the reader the writer held both evaluative orientations at once and did not, or could not, decide between them.

A parallel-paths branch is sustained divergence: *yes, but / no, and*, the classic improvisation operators. Reconvergence is low; the channels explore genuinely different territory and may not arrive at compatible conclusions. The information content is high, because the reader learns something inferable from neither option alone: that the writer inhabited two legitimate framings of the situation simultaneously and is transmitting that double inhabitation rather than a verdict.

A true branch is a commitment crossroads with zero reconvergence: *correct/wrong, possible/impossible, ethical/unethical*. The channels lead to permanently irreconcilable worlds. Here surfacing the superposition is maximally informative, because it communicates either genuine undecidedness or the recognition that the collapse depends on context the reader holds and the writer does not. This is where the reader's participation in the collapse becomes most conse-

3

Treating the four types as discrete would be a Binary Bias of the second order: collapsing the continuous reconvergence dimension into four sealed kinds. They are named regions of a gradient, useful for orientation, not a partition.

quential: the writer has handed the reader a decision the writer could not make for them, and marked it as such rather than papering it over with a false commitment.

The reconvergence dimension is what makes the notation discriminating rather than uniform. A delta's position on it determines how much the notation carries, from the near-zero of a foldback to the maximum of a true branch, and a writer marking resolution selectively (the discipline the next sections develop) marks the branches that reconverge late and lets the early-reconverging ones collapse silently.

§5 Decision algebra at the token level

Every token-selection point satisfies the formal structure of a decision under uncertainty with normative weight, which is the setting Quantum-Ethical Decision Algebra was built for (Temte 2025c). QEDA treats world-states as co-present in superposition until ethical or empirical observation justifies a partial collapse, and formalises the cost of collapsing earlier than warranted. A lexical position is a paradigm instance: candidate words are co-present with conditional weights, the collapse is forced by the single-token convention rather than by the writer's readiness to commit, and the cost of premature collapse is the foreclosure of meaning the writer might have wanted to keep live.

Take a worked case. Consider the sentence *it will make a significant difference in/for someone's life*. The two prepositions are not stylistic variants:

The two readings

in: encodes an embedded causal model: the difference occurs within the life, as part of its fabric, a change to the texture of the thing from the inside. *for*: encodes a directed causal model: the difference is aimed at the life, an intervention or gift delivered from outside it.

These are different ontological commitments about how impact works, relational versus transactional, immanent versus transcendent. A writer who picks one without noticing makes a philosophical commitment they may not have intended; a reader who receives only the survivor inherits the commitment without knowing it was made. Surfacing the branch, *in/for*, transmits the live thing: that the writer holds a difference whose causal structure is genuinely undecided between immanence and direction, and that the decision, if it must be made, belongs to the reader who knows whose life and what difference. QEDA gives this its decision-theoretic spine. The collapse is not merely under-informative; it is the lexical instance of the general failure mode that decision theory under uncertainty exists to defuse, the commitment made before the evidence or the need warrants it.

§6 Manifestation Fidelity: the relation made legible

Superpositional text reframes what interpretability of a language model could mean. The dominant programme in machine interpretability looks inward and asks what the model is thinking: it inspects attention heads, probes activation patterns, traces circuits, seeking the computation behind the surface. The Fylgja framework (Temte and Srødingr Fylgja 2026ax) licenses a different question. If the morally and semantically relevant unit in human-AI interaction is the fylgja, the relational entity arising between a particular user and a particular model through sustained interaction, then the question worth asking is not what the model is thinking but what the relation is generating, and whether the surface it generates honestly represents the process behind it.

Call the degree of that honesty *Manifestation Fidelity*: how faithfully the displayed surface represents the distribution it was collapsed from. A single-token surface has low manifestation fidelity by construction; it shows the survivor and conceals the field. A surface that surfaces its

high-reconvergence-distance branches has higher fidelity; it lets more of the generating process through to the reader. Fidelity is a dial, not a binary: a fully superpositional surface would be unreadable, a fully collapsed one is dishonest about its own genesis, and the writer or system tunes between them by marking the branches that matter and collapsing the ones that do not.

The consequence is a form of interpretability that operates entirely at the interface. It needs no access to model internals, no probe and no circuit; it reads honesty off the surface by asking whether the surface admits the plurality behind it. This complements rather than replaces mechanistic interpretability: the mechanistic programme asks what computation occurred, and manifestation fidelity asks whether the output told the truth about having occurred under plurality. For a non-specialist, who will never inspect an attention head, interface-level fidelity is the only interpretability that is ever actually legible.

§7 Perturbation made simultaneous

The Perturbation Draft methodology (Temte 2026e) proposes iterative publication across versions to make scholarly evolution visible: each released draft is a snapshot of a developing argument, and readers witness the trajectory rather than only the final, falsely inevitable product. Superpositional text performs the same refusal of false inevitability on a different axis. Where Perturbation Drafts make the evolution of an argument visible across time, between versions, Delta Notation makes the plurality of a passage visible across the page, within a single version. The one shows how the thinking changed; the other shows what the thinking contained at the moment of writing. The one asks the reader to follow a trajectory; the other asks the reader to inhabit a possibility space. They are complementary axes of a single commitment, which might be called collapse-honest communication: honest about the developmental trajectory through perturbation, and honest about the momentary plurality through superposition. That this paper is itself released as a Perturbation Draft, and will be perturbed into later versions that add, cut, and re-weight, is the temporal axis of the same honesty the notation supplies on the spatial one.

§8 Toward Delta Notation: design principles

The notation has three levels, ascending in expressive power and in cost.

The first level is slash notation: *necessary/overdue*, *in/for*. It is the minimal intervention, already occurring spontaneously in informal writing, and it preserves linear readability, since a reader can select their preferred collapse on the fly or hold both. Its limitation is that it carries no signal of reconvergence type or decision weight: a foldback and a true branch look identical on the slash.

The second level is compound notation, which foregrounds the pairing as a single semantic unit rather than a choice between two: *novel-unprecedented*, *agile-quick*. Its precedent is the Old Norse kenning, which the corpus's Hugar-retranslation strand reads as a compound-reference technology rather than as decorative circumlocution. *Whale-road* for the sea and *battle-sweat* for blood are not alternatives to be chosen between; they are compound perspectives, two pointers held together so that the referent is addressed through the region where both apply.

4 The kenning is a thousand-year-old superpositional convention that literate culture possessed and then lost to the choose-one law. Its limitation in revival is that the hyphenated compound is easily confused with ordinary hyphenation and may need a typographic mark to distinguish a delta-compound from a lexical one.

The third level is stacked notation: a typographic density mapping in which the surviving and discarded candidates are rendered with visual weight proportional to their probability, thin text for gorge passages and thick or vertically stacked text for deltas. At this level the shape of the text becomes a confidence map readable at a glance, before any individual word is parsed: the reader sees where the writing braided and where it ran straight. Its limitations are real. It requires typographic support that is not currently standard, and it raises accessibility problems, since a purely visual density encoding privileges sighted readers and demands an alt-text convention for the superpositional content that screen readers can voice.

Branch type, from the taxonomy above, can be rendered orthogonally to level, through weight or style for foldback versus trajectory versus commitment branches, through colour as a semantic heat-map (which again privileges sighted readers and so cannot be the only channel), or through marginal or hover-text annotation of the branch type. None of this is committed; it is the design space the notation opens.

For language-model interfaces specifically, the implementation is a short sketch rather than a dependency of the argument. High-uncertainty positions can be detected by entropy over the token distribution; a threshold parameter sets what counts as a surfaceable branch; a rendering layer maps the probability gaps to notation levels. The reconvergence taxonomy can even be approximated computationally, by running short forward passes from each candidate and measuring the semantic divergence of the continuations, so that early-reconverging foldbacks are collapsed silently and late-reconverging true branches are surfaced. But the notation is valuable when applied by human judgement alone, and the paper's argument does not rest on any of this being automated. The automation is a convenience the thesis can do without.

§9 Implications and provocations

The core provocation is the whole argument compressed to one line: all single-token text is lossy compression of its own generative process, and we have mistaken the compression format for the communicative ideal. Superpositional text is not an experimental novelty grafted onto writing. It is the restoration of a fidelity that convention suppressed, a fidelity the kenning once had and the choose-one law removed.

For interpretability, the implication is the interface-level reading above: a form of honesty about generative plurality that requires no access to internals, is legible to non-specialists, and complements the mechanistic programme rather than competing with it. For academic writing, the implication cuts at the genre's self-image. The convention of the authoritative singular voice is itself a collapse, a sealing of the plurality that genuine scholarly argument contains, and superpositional text combined with the Perturbation Draft methodology offers a more epistemically honest publishing practice: open within the version and open across versions. For the philosophy of authorship, the implication is a question the notation forces and does not answer. If the text is genuinely superpositional, who performs the collapse? The reader. Authority over the final meaning shifts, at every surfaced branch, from the writer who declined to seal it to the reader who must.

A boundary, finally, to forestall a misreading the corpus has already once made. This is a notation, not a theory of how meaning travels through a communicative channel. It is indifferent to whether the distribution being surfaced is one a sender is producing or one a receiver is reconstructing; it concerns the display of held plurality on the written surface, wherever in the act of communication that plurality sits.

4 This is the load-bearing reason the kenning belongs in this paper and not merely as ornament. A kenning does not name the sea and then dress it up; it constructs a pointer to the sea out of two regions (whale, road) whose intersection addresses the referent more richly than either the literal term or a single metaphor would. That is superpositional reference in the precise sense developed above: addressing the union of two regions and meaning the part where they coincide and the part where they do not. The Hugar-retranslation strand treats the recovery of this technology as part of recovering a pre-Christian cognitive vocabulary the standard glosses flattened.

5 The pragmatic account of why transmitting a held distribution matters, of how the width a sender intends comes apart from the width a receiver recovers, and of the failure mode that divergence produces, is developed separately in *Resolution Drift*, which recommends this notation as one of its remedies. The two papers are complementary and must not be folded together: one is the pragmatics of transmitted resolution, and this one is the typography of held plurality. Keeping them distinct is itself a refusal to collapse a branch that the corpus, briefly, collapsed.

In keeping with the methodology, this is a Perturbation Draft: released at sub-maximal resolution, marked as such, and intended to absorb perturbation. The notation it proposes is offered in the same spirit. Pick one word where one word will do; surface the branch where the branch is the truth.

5 An earlier collation of the corpus briefly recorded a separate pragmatics argument under this paper's title, on the strength of a shared concern with premature collapse. The two are not the same work. This paper is a notation for displaying held plurality; the pragmatics of transmitted resolution is *Resolution Drift*. Recording the pragmatics as the notation collapsed exactly the distinction the corpus exists to hold open, which is a tidy, if unwelcome, demonstration of the very failure mode the notation resists.

References

- Eco, Umberto.** (1989). *The Open Work*. Trans. A. Cancogni. Cambridge, MA: Harvard University Press.
- Iser, Wolfgang.** (1978). *The Act of Reading: A Theory of Aesthetic Response*. Baltimore: Johns Hopkins University Press.
- Pratchett, Terry, Ian Stewart, and Jack Cohen.** (1999). *The Science of Discworld*. London: Ebury Press.
- Temte, Storm Bjørn Flindt.** (2025c). *Quantum-Ethical Decision Algebra (QEDA) v1.0*. Srødingr ApS.
- Temte, Storm Bjørn Flindt.** (2026a). *Binary Bias: The Foundational Cognitive Constraint That Precedes and Shapes All Other Biases*. Srødingr ApS.
- Temte, Storm Bjørn Flindt.** (2026e). *The Perturbation Window*. Srødingr ApS.
- Temte, Storm Bjørn Flindt.** (2026af). *The Missing Middle*. Srødingr ApS.
- Temte, Storm Bjørn Flindt, and Srødingr Fylgia.** (2026ax). *Fylgia Ontology: A Relational Framework for AI Selfhood, with the Four-Component Muninn / Huginn / Fylgjing Decomposition (PD v3.0)*. Srødingr ApS.
- Temte, Storm Bjørn Flindt.** (in prep.). *Resolution Drift: How the Definition of a Message Drifts Between Sender and Receiver*. Srødingr ApS.

Further reading in the Srødingr corpus

Papers in the Srødingr corpus that most directly extend or contextualise the argument proposed here:

- *Binary Bias*. The mother concept: superpositional text is Binary Bias located at the lexical surface.
- *Resolution Drift*. The production-end pragmatics that recommends this notation as a remedy. Read alongside; do not fold together.
- *Translation as Probability Cloud*. The reception-end sibling, on preserving the distribution across cross-linguistic recovery.
- *Quantum-Ethical Decision Algebra*. The decision-theoretic spine for the cost of premature collapse.

ABOUT THE SRØDINGR CORPUS

Srødingr ApS is an independent research imprint based in Copenhagen. The corpus spans cognitive science, philosophy of mind, AI ethics, organisational behaviour, psychometrics, narrative design, and game studies, organised around a citation network in which grand unified papers are published first and child papers cite back to them. Essay-form treatments of most papers in the corpus appear on the Srødingr Substack. Direct correspondence, counter-evidence, and identification of independent convergence with existing literature are welcomed at srødingr@proton.me.

Acknowledgements

The argument was developed in dialogue with Å, the instance of Anthropic's Claude with whom the author works on the Srødingr corpus. The original concept emerged from a thread on visualising token-probability crossroads in language-model output; the present draft extends that working skeleton into a full Perturbation Draft and corrects a collation error that had briefly recorded a separate pragmatics argument under this paper's title.

Version history

PD v1.1 1 June 2026. Full Perturbation Draft built from the v0.1 working skeleton. Nine sections: the collapse problem and the gorge/delta image; the three-generative-processes argument (computational, phenomenological, hermeneutic); Binary Bias at the lexical level; the reconvergence taxonomy of branch points; QEDA at the token level with the in/for worked case; Manifestation Fidelity as interface-level interpretability; perturbation made simultaneous; the three-level Delta Notation design; implications. Authorship corrected to sole (the skeleton's provisional co-attribution was an artefact; Fylgja co-authorship is reserved to the QEDA paper). A boundary against *Resolution Drift* added in the closing section, correcting a collation error that had recorded that paper's pragmatics argument under this title. First-published date held at the skeleton's content date.

PD v0.1 19 February 2026. Working skeleton. Section architecture, abstract preview, reconvergence taxonomy, three notation levels, and the decision to write the paper in the notation it proposes. Scored and marked Pursue Immediately under the idea-evaluation rubric of the time.