CATEGORICAL METAPHYSICS · BOOK VII | V1.0 · APRIL 2026

Semantic Space Has a Shape

Vocabulary evolution, word embeddings, and the Book VII readout-functor picture

Thorsten Fuchs[1] • Anna-Sophie Fuchs

April 2026
Zenodo DOI 10.5281/zenodo.19816808

Keywords language evolution · word embeddings · semantic space · meaning drift · distributional semantics · Taylor's law · cultural evolution · readout functor · Category τ

1. INTRODUCTION: SEMANTIC SPACE HAS A SHAPE

When a language gains a new word, where does that word appear? Does it fill a random empty position in vocabulary space, does it attach itself to an already crowded region, or does it arise where a community is actively making new distinctions? Guo et al. ask this question quantitatively. Using word embeddings as empirical semantic spaces, they inspect how words are placed, clustered, and historically born in vocabularies. Their answer is that the space is not featureless. High-frequency words cluster near high-frequency words; empirical embeddings display regular clustering-velocity profiles; new English words arise preferentially in recently active regions; and new-word counts obey Taylor-law mean-variance scaling across semantic regions and historical time [4].

The title of this note is therefore literal but guarded: semantic space has a shape. The guardrail matters. The empirical object in Guo et al. is not meaning itself but a learned embedding space constructed from public language data. The paper's methodology treats such embeddings as a way to uncover a latent knowledge or semantic space. The Panta Rhei reading is more conservative. We treat embedding spaces as instruments. They measure public traces of language use, and those traces can reveal structure in the symbolic readout layer without being identical to the subsymbolic source of meaning.

Book VII of the Panta Rhei framework makes the distinction explicit. A word is not the pattern it names. A word is a readout of structured subsymbolic content into a public symbolic medium [2]. In the notation of Book VII, the readout direction is

$$\mathsf{Read}\colon \mathcal{K}_\tau \longrightarrow \mathcal{E}\mathrm{xpr},$$

or, in the language-specific case, from the subsymbolic pattern field into the expressive category of a particular language. The layer model keeps the full measurement chain explicit. The embedding space sits several steps downstream. It is inferred from public usage, not read directly from the source layer:

$$\mathcal{M}_{\mathrm{sub}} \xrightarrow{\mathsf{Read}_i} \mathcal{L}^{(i)}_{\mathrm{pub}} \longrightarrow \mathcal{C}^{(i)}_{\mathrm{text}} \longrightarrow \mathcal{E}^{(i)}_{\mathrm{emb}} \longrightarrow \mathcal{S}_{\mathrm{obs}}.$$

The philosophical significance of Guo et al. is not that the final object $\mathcal{E}^{(i)}_{\mathrm{emb}}$ is meaning. It is that $\mathcal{S}_{\mathrm{obs}}$, computed from $\mathcal{E}^{(i)}_{\mathrm{emb}}$, is structured enough to display robust regularities across embedding families, datasets, and in part across languages.

This chain is the grammar of the whole note. When the anchor paper reports a semantic-space regularity, this note asks which layer the regularity belongs to. Frequency assortativity is a statistic of the learned public trace, not a statement that frequent words are metaphysically central. Clustering velocity is an operational profile of vocabulary neighborhoods, not a direct picture of the subsymbolic presheaf. Word birth marks public naming events, not the birth of meaning itself. Taylor-law scaling measures region-level fluctuation in a chosen embedding instrument, not an already derived tau-law. Keeping the chain visible is what lets the bridge become useful without becoming inflated.

SCOPE DISCIPLINE

This note is a bridge note, not a proof note. It does not claim that Guo et al. prove Category τ, that embeddings are meaning, that all languages share one semantic carrier, or that the Panta Rhei framework already predicts the paper's fitted exponents. Its claim is smaller and therefore stronger: the paper gives a serious empirical calibration surface for Book VII's readout-functor picture of language.

What this note does. Section 2 reconstructs the external paper in its own terms. Section 3 fixes the measurement discipline: embeddings are operational shadows, not meaning itself. Section 4 maps the anchor to Book VII's readout functor, faithfulness criterion, subsymbolic layer, naming, translation, public-language repair, and drift architecture. Section 5 adds Book VI's shared-code and cultural PPAS bridge. Section 6 turns the anchor into a mini research program. Section 7 turns the result into future test surfaces. Section 8 closes with the explicit non-claims that keep the note honest.

Why this is a Research Note rather than a survey. The purpose is not to review distributional semantics, language evolution, or the full computational linguistics literature. The purpose is narrower: to decide whether one high-quality external paper can become a useful anchor inside the Panta Rhei research program. That requires two simultaneous acts of discipline. First, we must let the external paper speak in its own empirical vocabulary: embeddings, frequencies, clusters, word-origin dates, Taylor-law scaling, and generative models. Second, we must translate only what is structurally relevant into the tau vocabulary: subsymbolic patterns, readout functors, shared codes, drift, and future calibration surfaces.

This note therefore has the form of a bridge dossier turned into prose. It does not try to maximize coverage. It tries to maximize auditability. Every strong sentence should be traceable either to Guo et al.'s empirical claims or to the Book VII/Book VI framework claims. Where that trace cannot be made without adding new work, the note leaves a test target rather than filling the gap rhetorically.

2. THE GUO ET AL. ANCHOR

Guo et al. situate their work inside the study of cultural development: new words and ideas appear as societies change, and those changes leave traces in language. The paper uses word embeddings to turn that intuition into a statistical object. An embedding represents a vocabulary as points in a geometric space whose distances are intended to reflect semantic similarity. The authors then ask whether empirical embedding spaces have regular structure and whether synthetic generative processes can reproduce that structure.

2.1 Data, scope, and modelling stance

The linguistic component of the paper uses four families of static word embeddings: FastText, GloVe, Polyglot, and Word2vec/GoogleNews. The authors study English and up to 21 additional languages. They standardize the embedding space to a unit hyperball and, for dynamic spatial analysis, partition it into $K = 100$ equal-volume hypercone regions. The first two regularities, frequency assortativity and clustering-velocity profiles, are verified beyond English. The historical word-origin timing used for the dynamic claims is available only for English, and the paper says so explicitly [4].

The paper also makes a methodological postulate: embedding methodology is treated as uncovering a latent knowledge or semantic space that is primitive to text or graph representations. We should report that faithfully, but we do not need to inherit it ontologically. In this note, the embedding space is an instrument. It is downstream of public language use, and any interpretive bridge to Book VII must pass through that measurement discipline.

The standardization to the unit hyperball and the partition into hypercones matter for this reason. They make semantic space into an inspectable measurement surface. A region is not a metaphysical region of meaning; it is an operational region of an embedding space under a chosen normalization and partition. That sounds modest, but it is powerful. Once the region is operationally fixed, one can ask where high-frequency words sit, where new words appear, how clusters merge, and whether creation fluctuations scale predictably. The research-note bridge begins at that level of public, reproducible measurement.

The static-embedding choice is both a strength and a limitation. It is a strength because multiple established embedding families yield consistent large-scale regularities. It is a limitation because static embeddings assign one vector to each word, placing polysemous words somewhere between their distinct senses. The paper also focuses on word creation and usage, not on the modification of existing word senses. That boundary is important for Section 4: Guo et al. give us a strong anchor for where new vocabulary appears, but not a complete model of semantic drift.

2.2 Four regularities

The paper's empirical spine consists of four regularities. Table 1 summarizes them in the form needed for this Research Note.

Frequency assortativity. The first regularity is spatial stratification by popularity. Words with high usage frequency tend to occupy neighborhoods containing other high-frequency words. Guo et al. measure this with spatial correlation statistics adapted from mathematical geography, including Moran's I and Geary's C. The result is observed across multiple languages and even within specific parts of speech, which helps rule out the trivial explanation that the effect is only part-of-speech clustering.

For our purposes, frequency assortativity is the cleanest non-homogeneity signal. The point is not that frequent words are more metaphysically real. The point is that the public readout layer has density structure: some regions of use become stable, crowded, and frequency-rich, while others remain nascent or low-frequency.

That distinction matters because frequency is a mixed signal. A word can be frequent because it is semantically central, because it is grammatically useful, because it belongs to a common domain, or because a corpus samples one register more heavily than another. The anchor value is therefore not the ranking of individual words. It is the spatial fact that frequency is not randomly distributed over the embedding. For Book VII, that supports the more modest claim that public readout usage has geometry.

Clustering velocity. The second regularity concerns how clusters aggregate as the sensitivity of the clustering procedure changes. Guo et al. use a graph-partitioning approach and a topological single-linkage approach; both produce regular profiles. The paper interprets this as evidence that empirical semantic space consists of hierarchical clusters that continuously merge as the level of aggregation increases.

This is tempting language for Book VII, but it needs care. The paper's hierarchical clusters are not automatically the presheaf restriction structure of $\mathcal{M}_{\text{sub}}$. What they do provide is an operational sign that the public readout trace has nested semantic neighborhoods. That is the level at which this note should use the result.

The useful phrase here is "operational hierarchy." The hierarchy is not inserted by Book VII and then read back into the data. It is an empirical profile produced by changing a clustering scale. If similar profiles appear across embedding families, they become a useful external pressure on our language theory: a theory of symbolic readout should expect public words to organize into neighborhoods at more than one scale. If it cannot say why multi-scale neighborhoods arise, then it has not yet touched the empirical shape of language use.

Persistent temporal dynamics. The third regularity moves from spatial structure to historical word birth. Using English word-origin timing, Guo et al. show that new words do not appear uniformly across embedding space. They arise in relatively active regions, and activity persists across adjacent periods: hot regions tend to remain hot, and cold regions tend to remain cold, more than a uniform null model would suggest.

This is the bridge to naming pressure. A word appears when a distinction becomes socially useful enough to travel. In Book VII language, a pattern receives a portable symbolic representative; in Book VI language, the shared code mutates under cultural selection. But the dynamic evidence here is English-only, so any cross-language statement must remain restricted to the first two regularities unless and until comparable dated vocabularies are analyzed.

The English-only boundary is not a minor footnote. It marks the exact place where this note must separate a strong anchor from an open program. The dynamic result is deeply relevant to naming pressure, because it suggests that vocabulary innovation follows active semantic regions. But it is not yet a cross-language theorem about word birth. A future note or computation sprint would need comparable historical lexicons for additional languages before claiming a broad language-evolution law.

Regularity	What Guo et al. measure	Empirical meaning	Research-note bridge	Guardrail
Frequency assortativity	High-frequency words tend to have high-frequency neighbors in embedding space.	The empirical space is non-homogeneous: some neighborhoods become frequency-rich.	Readout density is structured; some symbolic regions become heavily named, used, and stabilized.	Frequency is not ontological centrality.
Clustering velocity profiles	Clusters merge regularly as the clustering scale changes, with profile shapes robust across methods.	Semantic spaces show hierarchical density structure rather than abrupt uniform collapse.	Operational evidence that readout space has nested neighborhoods rather than a flat word inventory.	Not identical to the Book VII presheaf hierarchy.
Persistent temporal dynamics	English neologisms arise near recently active embedding regions, with hot and cold activity persistence.	Word birth is spatially and temporally clustered rather than uniform.	Naming pressure appears in active regions of a shared code.	Historical timing is English-only.
Taylor's law	New-word counts show mean-variance scaling over time and semantic regions.	Vocabulary creation has measurable fluctuation structure and abnormal outliers can be detected.	A quantitative calibration surface for future tau-facing language models.	No tau prediction of the exponents.

Table 1. The four anchor regularities from Guo et al. and the disciplined way they enter this Research Note.

Taylor's law. The fourth regularity concerns fluctuation scaling. Guo et al. partition embedding space into semantic regions and historical periods, then study the means and variances of new-word counts. The temporal Taylor-law exponents for the language data are reported in the range 1.8–2.1, while spatial exponents are closer to 1.1–1.4. The point for this note is not the particular exponent as a tau prediction. The point is that the anchor paper turns semantic activity into a benchmarkable quantitative surface.

The modelling gift is substantial. Taylor's law turns "vocabulary grows unevenly" into a measured mean-variance relationship. For the Panta Rhei program, this changes the status of future language models. It is no longer enough to say that cultural dynamics are bursty or that shared codes evolve. A tau-facing model should eventually produce counts over regions and time windows, and those counts should be compared to the same scaling surfaces. Agreement would not prove the framework, but failure would be informative.

2.3 Directed preferential placement

Guo et al. compare empirical embeddings with synthetic embeddings generated by several model families. Their preferred model is directed preferential placement, or DPP. The essential idea is that new entities are not placed uniformly around old ones; the model is sensitive to local spatial context and pushes new entities away from dense areas. Among the tested model families, DPP broadly matches the observed package better than simpler uniform or preferential-placement alternatives.

That claim should be stated with restraint. DPP is not a complete causal theory of language evolution. It is a generative benchmark that captures some observed structural features: frequency assortativity, activity co-movement, bursts and lulls, and Taylor-law behavior within the paper's experimental setup. Its value for us is comparative. A future tau-facing model of shared-code evolution should be able to say which of these regularities it reproduces, which it refines, and which it fails.

This comparative role is essential. The temptation would be to rename DPP as a tau process because it involves directed placement, active regions, and cumulative structure. That would be premature. DPP is an external model family with its own assumptions and parameters. Book VI PPAS is a framework for reproduction, variation, selection, and drift in shared codes. The productive research question is not whether one can be verbally identified with the other, but whether a properly specified shared-code model can reproduce the same observables without hand-tuning away the hard parts.

2.4 Limitations as useful boundaries

The anchor paper is strong partly because its limitations are explicit. First, the embeddings are static. They do not distinguish contextual senses of the same word, and polysemous words are represented by a single vector. Second, word-origin timing is available only for English in the dynamic analysis. Third, the paper models word creation and usage, not the full modification of existing word meanings. Fourth, the geometric space is an instrument produced by an embedding method, not direct access to meaning.

Those limitations are not embarrassments for this Research Note. They are the guardrails. They tell us what may be claimed: the public readout trace of language has measurable shape, and that shape is relevant to Book VII's language architecture. They also tell us what may not be claimed: the paper does not measure meaning itself, does not settle translation, does not fully model polysemy, and does not give a cross-language dynamics of word birth.

They also help prioritize the next empirical work. Static embeddings should be challenged by contextual embeddings. English-only word-origin dynamics should be challenged by multilingual dated vocabularies. Hypercone regions should be challenged by alternative partitions. DPP should be challenged by models that include social transmission, institutional uptake, and carrier-local morphology. If the bridge survives those pressures, it becomes stronger. If it does not, the note will still have done useful work by showing which part of the bridge was too instrument-dependent.

3. EMBEDDINGS AS OPERATIONAL SHADOWS

The phrase "semantic space" is ambiguous because it can name two different objects. In Guo et al., it names an empirical embedding space: a geometric object learned from public language data. In Book VII, the source layer of meaning is the subsymbolic pattern field and its readout through language. One task of the present note is to avoid confusing these.

Book VII defines the subsymbolic layer $\mathcal{M}_{\text{sub}}$ as a label-independent presheaf of patterns [2]. A pattern exists before any label attaches to it. A label is a readout, not a constitutive act. In ordinary language: the word is not the thing meant. The word is a portable public handle for a pattern that may have existed before anyone named it.

This is why embeddings can be both useful and insufficient. They are useful because usage carries structure. If similar patterns receive contextually similar labels, then corpora contain statistical traces of those similarities. Embedding methods can recover some of that trace. They are insufficient because the trace is not the source. A learned vector representation is downstream of many projections: perception, experience, social naming, usage, writing, corpus selection, tokenization, and training.

The consequence is a useful inversion of the usual debate. We do not have to ask whether an embedding model "has meaning." The better question is whether a particular embedding construction preserves enough public readout structure to act as a scientific instrument. A telescope is not a star, and a map is not a territory; nevertheless, both can carry disciplined evidence when their distortions are understood. In the present case the distortions are known: corpus bias, tokenization choices, static treatment of polysemy, embedding-family dependence, and the fact that social practice is sampled only through available text.

That is why this note treats Guo et al. as a calibration anchor rather than as a philosophical authority. The paper gives us an empirical object with enough regularity to inspect, enough methodological detail to reproduce, and enough limitations to resist overreading. The Book VII bridge is strongest where the measurement chain remains visible.

Figure 1 gives the working discipline. The subsymbolic pattern field is not observed directly. It is projected through language-specific readout functors into public words, expressions, and texts. Public-language repair mechanisms stabilize that projection without freezing it. A corpus samples the public layer. An embedding algorithm then learns a geometric representation from the sample. Finally, spatial statistics reveal regularities in the learned space.

Two consequences follow.

First, an embedding regularity is evidence about the public readout trace, not a transparent window into the kernel. When Guo et al. find frequency assortativity, the disciplined interpretation is that public vocabulary use has structured neighborhoods. The stronger statement that the same neighborhoods are kernel neighborhoods

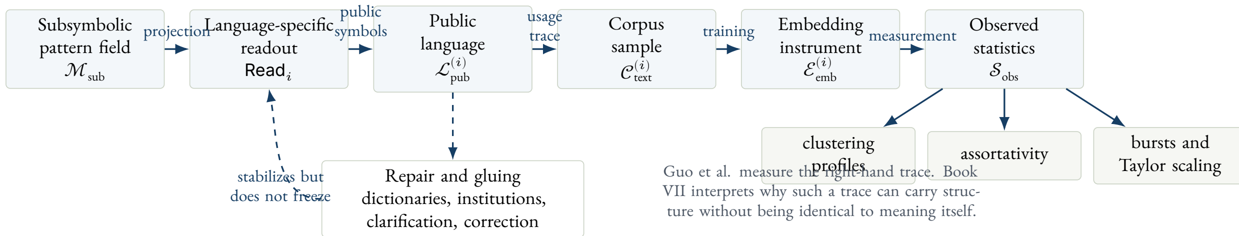

Figure 1. Embedding space as operational shadow with public-language repair. The empirical object in Guo et al. is downstream of language-specific readout, public-language stabilization, corpus formation, and embedding instrumentation. Book VII enters on the left-hand interpretive side; the anchor paper measures the right-hand statistical trace.

would require additional argument.

Second, the fact that embedding regularities are only shadows does not make them trivial. A shadow can still be diagnostic. If no structured relation existed between subsymbolic pattern, language readout, public use, and learned representation, then the appearance of stable cross-method and cross-language regularities would be surprising. In Book VII terms, the success of distributional methods is itself evidence that meaning has a subsymbolic and distributionally recoverable component [2].

There is also a negative version of the same point. If the observed regularities evaporate when the corpus changes, when contextual embeddings are used, when languages with different morphology are inspected, or when polysemy is separated into sense-level trajectories, then the operational shadow is weaker than the present note hopes. That would not refute Book VII. It would refute this paper's usefulness as a strong anchor for Book VII. This is the right scale of vulnerability for a Research Note.

Proposition 3.1 (Operational-shadow bridge). *For this Research Note, an embedding regularity may be used as evidence for structured semantic/readout organization only under the following translation:*

$$\textit{embedding statistic} \Rightarrow \textit{public readout trace},$$
$$\textit{public readout trace} \not\Rightarrow \textit{direct identity with meaning}.$$

The first arrow is the empirical bridge supplied by Guo et al. The blocked second arrow is the overclaim this note refuses.

Scope. Proposition 3.1 is **[Metaphorical]** as a bridge proposition. It states how the present note reads the external paper; it is not a theorem of Book VII.

4. BOOK VII BRIDGE: READOUT FAITHFULNESS, NAMING, TRANSLATION, DRIFT

With the measurement discipline fixed, the Book VII bridge can be stated cleanly. The anchor paper matters to Category τ because its four regularities are regularities of the public readout trace. Book VII gives a vocabulary for why such a trace should have structure: language is a readout functor; readouts can be more or less faithful; subsymbolic patterns precede labels; naming makes patterns portable; translation preserves kernel content while losing carrier-local accretions; public language repairs and stabilizes shared readouts; and meaning drift is the perturbation of a living readout map.

The key enrichment is that the bridge is no longer merely "embeddings shadow meaning." The sharper statement is:

$$\mathcal{M}_{\text{sub}} \xrightarrow{\mathsf{Read}_i} \mathcal{L}^{(i)}_{\text{pub}} \longrightarrow \mathcal{C}^{(i)}_{\text{text}} \longrightarrow \mathcal{E}^{(i)}_{\text{emb}} \longrightarrow \mathcal{S}_{\text{obs}}.$$

Guo et al. measure the right-hand side of this chain. Book VII explains why the left-hand side can leave structured shadows there. The bridge becomes scientifically useful only if every claim remembers its position on the chain.

4.1 Layer discipline: six objects, six claim types

The layer chain is not decorative notation. It is the claim grammar of the note. Each layer licenses a different kind of sentence, and many tempting overclaims arise by moving a sentence to the wrong layer. The subsymbolic pattern field licenses claims about prelinguistic or prelabelled structure, but Guo et al. do not measure that field. The readout functor licenses claims about projection and possible loss, but it is not reconstructed directly from the embedding. Public language licenses claims about shared symbols, repair, drift, and institutional stabilization. The corpus licenses claims about recorded usage under sampling bias. The embedding instrument licenses claims about learned geometry. The observed statistics license claims about frequency, clustering, activity, and scaling.

This discipline is what lets the note be ambitious without becoming loose. When we say that semantic space has a shape, we mean that the measured public trace has shape. When we say that the result is relevant to readout geometry, we mean that such a trace is where a structured readout should become visible after projection through public language and corpus sampling. When we say that the result is relevant to Book VII faithfulness, we mean that the trace may preserve distinctions well enough to be a probe. We do not mean that every embedding neighborhood is a kernel neighborhood, or that every statistically active region is a natural kind.

The same discipline also protects the external paper. Guo et al. should not be asked to carry metaphysical weight that its method was not designed to carry. Its achievement is empirical: it constructs measurement surfaces on which vocabulary placement, neighborhood structure, word birth, and fluctuation scaling can be studied. The tau-facing work begins after that surface is in hand. It asks which parts of the measured shape are stable under instrument changes, which parts depend on corpus and language carrier, which parts are affected by public repair, and which parts can be reproduced by explicit shared-code dynamics.

This is also why failure modes matter. If a future embedding family erases one regularity but preserves another, the result should not be flattened into success or failure. It should be assigned to the correct layer. A failure under corpus change may be a sampling result. A failure under contextual embeddings may be a polysemy result. A failure under non-English word-origin dating may be a carrier or historical-record result. A failure of a tau-facing model to reproduce Taylor-law scaling may be a mechanism result. The layer chain turns disappointment into information.

4.2 Language as readout functor

Book VII's readout functor formalism treats language as a projection from kernel or pattern structure into communicable expression [2]. In its simplest form:

$$R\colon \mathcal{K}_\tau \longrightarrow \mathcal{E}\mathrm{xpr}.$$

Different languages, notations, genres, and cultural practices are different readout functors. No single readout is privileged; each projection preserves some structure and loses some structure. Faithfulness is the quality criterion: a readout is more faithful when it distinguishes distinctions that the source layer distinguishes.

This matters for Guo et al. because embedding regularities are not discovered in the private heads of speakers. They are discovered in the public expression layer. Frequency assortativity says that some regions of public expression become frequency-rich. Clustering velocity says that public expression has multi-scale neighborhoods. Word-birth dynamics say that some regions of public expression become active sites of naming. Taylor's law says that those activities fluctuate with measurable scaling.

The readout-functor reading is therefore:

$$R_i(\text{patterns}) \rightsquigarrow R_i\text{-readouts},$$
$$R_i\text{-readouts} \rightsquigarrow \text{vocabulary regions}.$$

The arrow is not equality. It is projection. That is why the embedding is a shadow and also why the shadow is meaningful.

4.3 Readout faithfulness as empirical discipline

The central Book VII enrichment is faithfulness. In Book VII, readout faithfulness is not a slogan about eloquence, precision, or the cultural status of a language. It is the local criterion that a readout should preserve distinctions that matter at the source layer: if the source distinguishes two patterns, a faithful readout should not collapse them without leaving any recoverable trace [2]. Conversely, a readout can be unfaithful in a particular region when different source distinctions are forced through one public label, when polysemy overloads a single word, or when a community lacks a stable vocabulary for a difference it nevertheless experiences.

Guo et al. do not test this theorem directly. They cannot, because their instrument lives at the right-hand side of the measurement chain:

$$\mathcal{C}^{(i)}_{\text{text}} \longrightarrow \mathcal{E}^{(i)}_{\text{emb}} \longrightarrow \mathcal{S}_{\text{obs}}.$$

What they can test is whether the public trace behaves as if it preserves structured distinctions. Frequency assortativity suggests that usage is not evenly smeared over the trace. Clustering velocity profiles suggest that neighborhoods remain visible across scales. Word-birth persistence suggests that new public names arise near regions already under naming pressure. Taylor-law scaling suggests that the fluctuation of word creation is not merely white noise. These are not proofs of readout faithfulness, but they are operational probes of it.

This point is easy to overstate, so the note keeps the direction one-way. A stable embedding regularity may indicate that the public readout preserves some structure. It does not identify the preserved structure with kernel identity, and it certainly does not rank languages by metaphysical adequacy. Book VII explicitly rejects that kind of ranking: different readouts can be faithful in different regions and lossy in different ways. A mathematical notation can preserve formal distinction while losing affective nuance. A poem can preserve affective and imaginal distinction while refusing formal resolution. A technical vocabulary can stabilize a region for experts while excluding ordinary speakers. Faithfulness is local, task-bound, and auditable only relative to a stated source distinction.

Proposition 4.1 (Operational faithfulness probe). *For the purposes of this Research Note, an embedding-space regularity counts as an operational probe of readout faithfulness only when it satisfies three conditions:*

1. *it is computed from a specified corpus and embedding instrument;*
2. *it survives at least one reasonable change of measurement choice;*
3. *it is interpreted as a public trace, not as kernel identity.*

Under those conditions, the regularity may support the claim that a public-language readout preserves structured distinctions in the observed region.

Scope. Proposition 4.1 is **[Metaphorical]**. It translates Book VII faithfulness into an empirical discipline for this note; it is not itself a theorem about natural language.

4.4 Subsymbolic label-independence and naming

Book VII's subsymbolic layer, $\mathcal{M}_{\text{sub}}$, is the layer of patterns beneath symbolic naming. A pattern may be present before a word exists for it. Language begins when such patterns receive names sufficiently stable and portable to circulate across agents. This is the conceptual bridge from Guo et al.'s dynamic word-birth result to Book VII.

In that reading, a neologism is not merely a new string. It is a public sign that a community has found a distinction worth carrying. Sometimes the distinction is a new technology; sometimes a social identity; sometimes a scientific classification; sometimes a recombination of older regions. The external paper does not classify each case, but it shows that new words appear in active semantic regions rather than uniformly. That is what what one would expect if naming follows social salience and local readout pressure.

The word "pressure" is important. It prevents a naive one-to-one picture in which every new concept immediately receives one word. In living language, many patterns never stabilize as public names; some receive several names; some receive names only inside subcultures or technical communities; and some names survive after the original pattern has shifted. The empirical birth of a word therefore marks a threshold event in the shared readout system, not the metaphysical beginning of a meaning.

There is an important asymmetry here. Guo et al. study the birth of words, not the birth of meanings. Book VII says meanings can precede words; the paper shows where new words enter the public embedding trace. The bridge is therefore not "new embedding point equals new meaning." It is: new word birth is one measurable public signal that a region of the readout system is active.

This asymmetry is what label-independence buys us. If the subsymbolic layer were nothing but a catalogue of already named concepts, then the birth of new words would be mysterious: there would be no pressure from unnamed or undernamed distinctions. Book VII instead allows the prior existence of patterns that are carried in perception, practice, affect, skill, social tension, or technical work before they become public vocabulary. A new word then functions as a handle introduced into a living public code. It gives a community a way to point, combine, teach, contest, and remember.

The anchor paper's active-region result makes this concrete without overclaiming. New words do not arise uniformly; they appear near regions where recent word birth has already occurred. In Book VII language, this is what one expects when naming is a response to local readout pressure. A scientific field under rapid development, a social movement that needs new self-descriptions, a technology that creates new practices, or a political conflict that polarizes old terms can all create regions where new handles are useful. Guo et al. do not classify the causes of each region, and this note does not pretend that they do. The bridge is structural: public word birth is not random sprinkling over a dictionary; it is clustered entry into an already organized readout trace.

4.5 Language as E2 self-enrichment

Book VII's stronger claim is that language is not merely a layer of labels placed on top of thought. Language is the E_2 self-enrichment of the subsymbolic layer: a system by which patterns become nameable, composable, transportable, revisable, and publicly stored [2]. This is why the embedding-space result matters. A vocabulary is not just a bag of names. It is a record of which distinctions have become portable enough to enter public circulation.

The E_2 reading changes the role of the four anchor regularities. Frequency assortativity becomes a sign that some public names are deeply embedded in repeated communication. Clustering velocity becomes a sign that composable neighborhoods exist at multiple scales. Word-birth persistence becomes a sign that self-enrichment is locally active: once a region begins to acquire new handles, it tends to keep generating them for a time. Taylor-law scaling becomes a sign that this enrichment has fluctuation structure rather than merely a smooth growth curve.

None of this says that Guo et al. have measured E_2 enrichment itself. They have measured downstream traces of vocabulary growth in embedding space. The value for Book VII is that those traces are where the self-enrichment picture should become visible if it has empirical grip. A future tau-facing model should make the enrichment operations explicit: naming, composition, reuse, correction, institutional uptake, translation, and decay. Only then can it be compared honestly with the anchor observables.

4.6 Translation and universality without flattening

The cross-language dimension of the paper is important but limited. The first two regularities are verified across 21 languages beyond English, which suggests that frequency assortativity and clustering profiles are not merely artifacts of English. Book VII gives a disciplined interpretation: cross-language recurrence can indicate shared structural constraints beneath different carriers.

It must not be read as a universal-carrier claim. Translation in Book VII preserves kernel content while losing or transforming carrier-local features: connotation, phonetic rhythm, morphology, grammatical gender, and cultural embedding [2]. Therefore the correct statement is not that all languages share one semantic space. The correct statement is that different languages may project comparable structural neighborhoods while retaining local carrier differences.

This distinction is especially important for languages with different morphological systems, writing systems, and cultural histories. A cross-language embedding regularity is valuable precisely because it appears despite carrier diversity. It is not permission to erase that diversity.

The correct analogy is not a single coordinate chart imposed on all languages. It is a family of charts whose transition behavior is partly structured and partly lossy. Book VII's translation picture allows this middle position. If two languages show comparable frequency assortativity or clustering velocity, that may indicate common constraints on communicable semantic organization. But the carrier-level realization can remain deeply different: compounds, inflectional systems, scripts, idioms, register boundaries, and culturally saturated terms can all reshape the public readout without destroying the possibility of structural comparison.

This is also where Book VII's translation-as-functor picture is helpful. Translation is not a lookup table between word tokens. It is a structured transport from one public readout into another, mediated by an intended kernel content and limited by carrier-local losses [2]. Some distinctions transport well. Others leak rhythm, register, moral valence, historical allusion, or technical convention. A cross-language embedding regularity therefore has a precise status in this note: it suggests that the public traces of different languages may be constrained by comparable source-side or communication-side structure, while the actual carrier maps remain lossy.

A future empirical refinement should preserve that middle position. It should not ask only whether a headline regularity exists in each language. It should ask how the regularity changes with morphology, compounding, script, genre, and translation history. If clustering profiles are stable across such changes, the readout-geometry bridge becomes stronger. If the profiles shift systematically, that is not a failure; it is evidence about where carrier-local structure enters the public trace. The right question is not whether all languages reduce to one semantic space. It is how families of public readouts preserve, distort, and repair comparable distinctions.

4.7 Drift as living readout

Book VII formalizes meaning drift as perturbation of a time-dependent readout functor:

$$\Phi_t : \mathcal{K}_\tau \longrightarrow \mathcal{L}_t.$$

Natural languages are living maps. The kernel or source structure is not itself drifting in the same way that public usage drifts; rather, the community-maintained readout changes over time. Words narrow, broaden, shift, improve, degrade, and acquire new connotations.

Guo et al. do not directly measure this full drift process. Their dynamic object is the birth of new words, not contextual movement of existing word senses. This limitation is productive. It tells us where the next research surface lies. Static embeddings plus dated vocabulary birth give one side of the living-map picture: where new public representatives appear. Contextual and diachronic embeddings would be needed to inspect how existing readout maps move after birth.

This also clarifies why polysemy is not an incidental technical problem. In Book VII terms, a polysemous word is a shared label through which several pattern regions are read out, often with context deciding which path is active. A static embedding collapses those paths into one coordinate. That collapse can still be useful for large-scale statistics, but it is the wrong instrument for fine-grained drift. A successor study could use the same anchor paper as a baseline and then ask how contextual sense trajectories refine, split, or contradict the static semantic-space picture.

The bridge from Guo et al. to drift therefore has two faces. Word birth is a discrete public event: a new label becomes visible in the corpus. Semantic drift is a continuous or punctuated movement of an existing label's readout path. The two processes interact. A region whose existing labels no longer carry emerging distinctions may generate neologisms. A new word may later drift away from its original technical, social, or ironic use. Conversely, repair mechanisms may pull a drifting word back toward a stabilized public definition. The anchor paper directly studies the first face and only indirectly points toward the second.

This distinction gives the future program a sharper empirical question: when new words enter an active region, do neighboring old words move? If contextual or diachronic embeddings show that existing words split into new sense trajectories around the same periods in which neologisms appear, then the Guo-style word-birth surface and the Book VII drift surface begin to meet. If they do not, then word birth may be a more independent cultural innovation process than the present bridge expects.

4.8 Public language, repair, and gluing

Living readouts do not drift freely. Public language has repair infrastructure: dictionaries, schooling, legal definitions, editorial standards, scientific nomenclature, translation practices, usage guides, conversation, correction, and institutional memory. Book VII treats these as gluing-condition enforcement for language: local uses must cohere enough to remain shareable across contexts [2]. Without such repair, language fragments into private or local codes. With too much repair, it freezes and fails to absorb new distinctions.

This repair layer is not directly measured by Guo et al., but it sits between public language and corpus formation in Figure 1. It helps explain why frequency-rich regions can be stable. A word becomes frequent not only because it is used often but because communities agree well enough on how it may be reused. Technical vocabularies, legal terms, and standardized educational language are extreme cases: they spend social energy to keep readouts fixed. Open cultural language is looser; it allows more drift, slang, metaphor, irony, and contestation.

The research opportunity is to treat repair as a measurable gradient. A future sprint could compare ordinary corpora, technical corpora, legal corpora, and social-media corpora under the same embedding and clustering procedures. The prediction is not a theorem, but a useful test target: highly repaired domains should show sharper stabilization of terms, lower uncontrolled drift, and possibly different word-birth statistics than open cultural domains. If the same Guo-style regularities persist across the gradient, the readout-geometry bridge strengthens. If they fragment by domain, repair becomes a major missing variable.

4.9 Distributional structure as evidence

Book VII's discussion of distributional semantics and large language models is useful here only if kept in its disciplined form. The claim is not that a statistical model understands meaning. The claim is that the success of distributional methods reveals that some component of linguistic meaning is encoded in public statistical structure [2]. Guo et al.'s paper is a cleaner anchor for this claim than a debate about artificial minds, because it uses embeddings as instruments to study vocabulary evolution rather than as agents.

The conclusion is modest: distributional structure is evidence, not identity. It is evidence because words with related roles in the public readout system leave related traces in corpora. It is not identity because the corpus does not contain the whole pattern field, the embedding does not contain the corpus, and the coordinate system does not contain the act of meaning.

This is why Guo et al. are unusually well positioned as an anchor for our sixth Research Note. The paper does not require us to enter the question of machine consciousness or artificial understanding. It asks a cleaner scientific question: if one treats embeddings as spaces, do vocabularies occupy and grow inside those spaces in regular ways? That question can be answered with statistics, compared across instruments, and converted into future tests. It is the level at which Book VII's distributional claim can meet empirical work without becoming inflated.

Book VII's proposition about distributional and LLM evidence should be read in the same modest key. Distributional success is evidence that public text contains recoverable traces of subsymbolic organization; it is not evidence that the model has the whole subsymbolic layer, and it is not evidence that the model understands as a human agent understands [2]. This Research Note inherits the first part and refuses the second. Embedding spaces can reveal structure because language use is not arbitrary. They remain instruments because they are trained on traces rather than living the full readout relation.

The polysemy limitation is the clearest warning. Static embeddings collapse many contextual paths into one coordinate. That collapse may be acceptable for large-scale spatial statistics; it is not acceptable for a fine theory of meaning repair, connotation, metaphor, or drift. A high-value successor study would compare static embeddings with contextual embeddings on the Guo observables. If both instruments see the same large-scale regularities, the operational-shadow bridge gains robustness. If contextual methods reveal that static neighborhoods average incompatible senses, then the bridge must be narrowed to vocabulary-level public traces rather than semantic trajectories.

5. BOOK VI BRIDGE: SHARED CODES AND CULTURAL EVOLUTION

Book VII explains why words are readouts of structured patterns. Book VI adds the population-level mechanism: language is a shared code. In Book VI, a language is a triple of grammar, encoding, and decoding maps maintained across a population of agents [1]. Communication succeeds when agents encode and decode sufficiently aligned content through a shared signal channel.

This is the correct home for the cultural dimension of Guo et al. A vocabulary is not a private list. It is a public code under social maintenance. Words are introduced, reused, forgotten, standardized, repurposed, translated, and contested. A new word is therefore an innovation in a shared code. It is not merely an individual's invention; it becomes linguistically real when the community can receive, decode, and reuse it.

This makes the embedding-space result sociological as well as semantic. The new word is a point in a learned vector space only after it has survived enough social circulation to enter the corpus. Before that, it may be an individual coinage, a joke, a local technical term, or a failed proposal. Book VI's shared-code picture explains why this threshold matters: code innovations become durable only when enough agents can coordinate around them. The empirical birth of a word is therefore already filtered by collective uptake.

Book VI describes cultural evolution as PPAS on shared codes. The components line up naturally:

- the shared code is the evolving cultural genome;
- innovation introduces variants such as new words, new techniques, and new distinctions;
- recombination occurs through cultural exchange, translation, migration, and media circulation;

External observable	Book VII/VI concept	Allowed manuscript claim	Forbidden overclaim	Future empirical test
Embedding spaces show non-random structure.	VII.T10, VII.D52	The public readout trace can act as an operational faithfulness probe.	Embedding coordinates identify kernel positions or meaning itself.	Compare observables across corpora, embeddings, and partitions.
Frequency-rich words cluster together.	VII.D52; public language	Readout density is non-uniform and partly stabilized by use.	Frequency equals semantic centrality or metaphysical priority.	Control for part of speech, corpus domain, and register.
Clustering profiles are hierarchical.	VII.T10; presheaf shadow	Public vocabulary neighborhoods have multi-scale operational structure.	The empirical hierarchy is the Book VII presheaf itself.	Test robustness under clustering method and embedding family.
New words arise in active regions.	VII.T20, VI.T39	Word birth is a public signal of naming pressure in a shared code.	The paper measures the birth of meanings themselves.	Repeat with non-English dated vocabularies.
Cross-language spatial regularities recur.	VII.R25; no privileged readout	Different carriers may preserve comparable public-trace constraints.	All languages share one universal semantic carrier.	Compare morphology, script, genre, and translation history.
Static embeddings collapse polysemy.	VII.R26, VII.P14	One vector can hide several readout paths while preserving coarse traces.	Static embeddings fully measure semantic drift.	Compare contextual, diachronic, and sense-aware embeddings.
Taylor-law scaling appears in word creation.	VII.R26; Book VI PPAS	Fluctuation scaling is a calibration target for future dynamics.	Tau already predicts the fitted exponents.	Ask whether tau-facing models reproduce scaling without manual tuning.
DPP broadly matches the tested package.	Book VI cultural PPAS	DPP is a benchmark model family for comparison.	DPP is PPAS or a full causal theory.	Build and compare an explicit tau-facing shared-code model.

Table 2. Book VII enrichment matrix. Each row states how an external observable may enter the manuscript, what it must not be made to claim, and which future empirical test would sharpen the bridge.

- selection is social feedback: useful, resonant, or institutionally supported forms spread;
- drift is stochastic loss or local fixation of variants.

Guo et al.'s active semantic regions are empirical shadows of this process. When a region is culturally active, new words tend to appear near recent additions. When a region cools, creation slows while existing words may keep accumulating frequency. The paper's conclusion even sketches a possible cycle: little creation and little use; high creation and limited use; later use following earlier creation; and gradual cooling. That is the kind of phase structure a future shared-code model should be able to address.

Those phases also give the note a public-facing story without weakening the technical claim. A culture does not merely add words one by one. It opens zones of attention. New practices, devices, conflicts, disciplines, and institutions create regions where distinctions proliferate. Once the region has stabilized, usage may keep growing even while naming slows. This is why the paper's distinction between creation and frequency is so valuable: it lets us separate the birth of new readouts from the later consolidation of old ones.

DPP enters here as a useful foil. It generates embedding-like point sets by placing new entities with local spatial context, pushing them away from dense areas and seeding new semantic regions. Cultural PPAS is not DPP. But the comparison is productive because both focus on the same qualitative package: cumulative advantage, active regions, bursts, lulls, and heterogeneous growth. A tau-facing model should not merely redescribe those features. It should eventually make them testable.

The minimal tau-facing version would have to represent at least four ingredients. It would need a space of candidate distinctions or patterns. It would need agents who encode and decode through a shared but imperfect code. It would need local pressure for new names when existing readouts fail to carry emerging distinctions. And it would need social feedback so that some innovations spread while others disappear. Only after those ingredients are specified should the model be compared to DPP on the same observables.

Remark 5.1 (DPP and cultural PPAS). DPP is an external generative model in the anchor paper. Cultural PPAS is the Panta Rhei framework's internal model family for shared-code evolution. This note compares them only at the level of future research targets: which observable semantic-space statistics should a tau-facing shared-code model reproduce?

6. FROM ANCHOR PAPER TO MINI RESEARCH PROGRAM

The preceding sections establish the bridge. This section states the work it opens. A useful anchor paper should not merely decorate the framework; it should create obligations. Guo et al. create four such obligations for the Panta Rhei language program: an observable ledger, a readout-faithfulness ledger, a model-comparison ledger, and an editorial ledger.

Observable ledger. The first obligation is to preserve the empirical observables in a form that future work can actually test. For this note, the observables are not "meaning" or "semantic space" in the abstract. They are concrete: frequency-neighbor correlations, clustering-velocity profiles, counts of new words in fixed regions and time windows, temporal persistence of active regions, and Taylor-law mean-variance scaling. Each observable lives at a specific layer of the measurement chain. Frequencies and word births belong to public corpus and lexicon data. Hypercone counts belong to the normalized embedding instrument. Scaling exponents belong to statistics computed after region and time-window choices have already been made.

That layered view is crucial. It prevents the framework from treating a statistic as if it were a direct metaphysical signal. It also makes the future audit cleaner. If a result changes, we can ask where it changed: corpus choice, embedding family, dimensional normalization, partition, clustering rule, dated-vocabulary source, or statistical estimator. A tau-facing language model should eventually be evaluated against this ledger, not against a loose memory of the paper's headline claims.

Readout-faithfulness ledger. The enrichment adds a second obligation: record which public-trace regularities are being used as probes of readout faithfulness and which are not. A faithfulness claim always has the form "this public trace appears to preserve this kind of distinction under these measurement choices." It should never have the form "this embedding coordinate is the meaning" or "this language is more metaphysically faithful than that language."

The ledger should therefore track three classes of evidence. First, there are robustness checks: does the regularity survive changes in embedding family, corpus, partition, and clustering procedure? Second, there are carrier checks: does it survive morphology, script, genre, and translation history? Third, there are repair checks: does it behave differently in domains where public language is heavily stabilized, such as legal, scientific, or technical corpora? A future model that ignores all three classes may still be interesting, but it will not have earned the Book VII faithfulness bridge.

Model-comparison ledger. The third obligation is to keep DPP as a real comparator. A weak bridge would praise Guo et al., borrow their language, and then retreat into framework terms. A stronger bridge asks whether the Panta Rhei model family can meet the same observable bar. That means preserving DPP as an external baseline even where its interpretation differs from ours.

The comparison should be staged. First, reproduce DPP's qualitative success on the anchor observables. Second, define a minimal shared-code model with agents, readout pressure, innovation, social uptake, and drift. Third, compare both models on identical measurement surfaces. Fourth, inspect the failure modes. Does the shared-code model fail to create persistent active regions? Does it create bursts but not Taylor-law scaling? Does it reproduce frequency assortativity only by baking frequency preference into the generative rule? Those failures would be research results, not embarrassments. They would tell us which parts of Book VI and Book VII still need formal mechanism.

Editorial ledger. The fourth obligation is editorial. This Research Note should make the public claim simple without making it simplistic. The public sentence is: "semantic space has a shape." The audit sentence is longer: learned embedding spaces, constructed from public language traces, show structured regularities in vocabulary placement and word birth; those regularities can be read as operational shadows of Book VII readout geometry and Book VI shared-code dynamics, provided that embeddings are not identified with meaning and the dynamic results remain bounded by their English historical data.

That two-sentence pairing should guide any later publication version. The first sentence gives the note its reach. The second sentence gives it its honesty. If the draft loses the first sentence, it becomes a technical memo with little public force. If it loses the second, it becomes an overclaim. The final publication version should keep both together.

What would count as progress. Progress after this draft has three forms. A prose improvement would sharpen the Book VII bridge around readout functors, translation, repair, faithfulness, and drift. A computational improvement would reproduce the anchor statistics from the archived data and test instrument dependence. A modelling improvement would build the first explicit shared-code simulation and compare it against DPP. The sixth Research Note becomes important if it can hold these three forms together: faithful external science, disciplined tau interpretation, and future tests that can genuinely fail.

7. FUTURE TEST SURFACES

The strongest ending for this note is not triumph but testability. If Guo et al. are a serious anchor, then they should change what future Panta Rhei language work must be able to inspect. The following surfaces are the first ones.

Reproduction of headline regularities. The paper archives data and code in Dryad and supplementary material in Figshare [3, 5]. A future computational sprint should reproduce the headline statistics: frequency assortativity, clustering profiles, active-region co-movement, and Taylor-law slopes. The goal is not to challenge the paper by default. It is to make the research note's anchor auditable inside our own workflow.

The reproducibility target should be concrete. The first pass should rebuild the cleaned vocabulary spaces, reproduce the unit-hyperball normalization, recompute the hypercone-region counts, and regenerate the headline plots or statistics used in the paper. The second pass should vary the instrument: alternative embeddings, alternative corpora, and alternative partitions. The third pass should record which conclusions survive those changes. That is the point at which this note can mature from a conceptual bridge into a computational research program.

Static versus contextual embeddings. The paper explicitly uses static embeddings. A natural follow-up is to ask which regularities survive when contextual embeddings or diachronic embeddings are used. If contextual methods sharpen the same structure, that would strengthen the operational-shadow bridge. If they dissolve it, then the bridge may be more instrument-dependent than the present note assumes.

The deeper Book VII version is a contextual polysemy benchmark. Choose regions where static embeddings suggest high activity and then ask whether contextual embeddings reveal one coherent sense trajectory, several overlapping trajectories, or an unstable mixture. The first outcome would support the coarse public-trace interpretation. The second would refine it: the static point is not wrong, but it hides multiple readout paths. The third would weaken the anchor by showing that the observed neighborhood is mostly an artifact of sense collapse.

Readout-faithfulness robustness. A future sprint should explicitly turn Proposition 4.1 into a benchmark. For each regularity, record the corpus, embedding family, normalization, partition, and statistic. Then vary one choice at a time. The question is not whether every numerical value remains fixed. The question is whether the qualitative distinction being read out remains visible. If frequency assortativity appears only in one embedding family, it is a weak faithfulness probe. If clustering profiles survive several instruments, they become stronger evidence that the public trace preserves multi-scale semantic organization.

Non-English dynamics. The first two regularities have cross-language support, but word-birth dynamics are English-only in the paper's historical timing analysis. A high-value future test would use dated vocabularies in other languages to ask whether active-region persistence and Taylor-law behavior generalize. Failure to generalize would not refute Book VII, but it would weaken the claim that Guo et al.'s dynamic result is a broad language-evolution anchor.

Public-language repair gradient. The repair/gluing bridge suggests another test: compare domains with different stabilization pressure. Legal and technical corpora should contain terms whose public readout is actively repaired by definitions, standards, and institutions. Open cultural corpora should contain looser, more metaphoric, and more contested usage. Running Guo-style spatial statistics across such a gradient would ask whether repair produces sharper clusters, slower drift, or different word-birth patterns. This would move the note from a general statement about public language to an empirical question about how public language keeps itself coherent.

Tau-facing shared-code models. The central modelling task is to construct a tau-facing shared-code model that can be compared against DPP. Such a model should specify how naming pressure, carrier-local constraints, social feedback, and cultural recombination produce public vocabulary change. It should then be tested against the same statistics Guo et al. use: frequency assortativity, clustering-velocity profiles, activity persistence, and Taylor-law scaling.

The first benchmark should not be ambitious. It should begin with a toy shared-code process whose outputs are explicit word-birth events in a structured space. If the toy model cannot produce non-uniform active regions or any mean-variance scaling, that failure is informative. If it can, the next step is to ask whether the mechanism is genuinely tau-facing or merely a disguised version of DPP. The model should earn its framework connection by exposing readout, code sharing, innovation, selection, and drift as separate inspectable components.

Falsification criteria for this anchor. The anchor would weaken in at least four ways. First, the spatial regularities could fail under reasonable embedding alternatives. Second, the cross-language findings could collapse once morphology, corpus size, or register are controlled more aggressively. Third, English word-origin timing could turn out to encode dictionary or etymological artifacts rather than semantic-region activity. Fourth, a model with no semantic or social structure could reproduce the whole package as a statistical artifact. Those are not threats to avoid. They are the tests that make the anchor useful.

Possible weakening results. The readout-functor interpretation would be weakened if the observed regularities disappear under reasonable alternative corpora, if they depend almost entirely on embedding method artifacts, if language families with different morphology systematically fail to show comparable spatial structure, or if contextual embeddings reveal that static clusters mainly average unrelated senses. These would not disprove Category τ, but they would reduce the evidential value of this anchor paper.

8. HONEST NON-CLAIMS

This Research Note is intentionally narrow. It makes no claim that the Proceedings B paper proves Category τ. It makes no claim that word embeddings are meaning. It makes no claim that all languages collapse into one universal semantic carrier. It makes no claim that Book VII predicts the Taylor-law exponents or that Book VI already contains a model equivalent to DPP. It also makes no claim that the anchor paper proves readout faithfulness for natural language as a whole, that public-language repair has already been measured, that LLMs understand, or that languages can be ranked by their metaphysical adequacy.

It also does not claim that the paper fully studies semantic drift. Guo et al. study new-word creation and usage patterns using static embeddings. Book VII adds the broader living-map picture, but the paper itself does not measure the full perturbation of word meanings over time. The correct bridge is therefore from word-birth dynamics to naming pressure and shared-code activity, not from word-birth dynamics to complete semantic drift.

Proposition 8.1 (Semantic-space anchor)**.** *Guo et al. provide a calibrated empirical anchor for Book VII's readout-functor picture in the following sense: their embedding-space regularities are consistent with structured public semantic/readout organization; they give operational probes of readout faithfulness, naming pressure, translation robustness, drift limits, and shared-code dynamics; and they supply quantitative test surfaces for future tau-facing language models. They do not, by themselves, establish the ontology, formal machinery, or predictions of Category τ.*

Scope. Proposition 8.1 is **[Metaphorical]**. It is the interpretive claim of this Research Note, not a formal theorem.

The intended contribution is therefore smaller and stronger than a confirmation claim. The anchor paper shows that public vocabularies leave a structured geometric trace. Book VII explains why such a trace is the kind of place one should look if language is a readout of structured subsymbolic patterns, if translation is lossy but structured, if public language repairs itself, and if drift is the perturbation of a living readout. Book VI explains why the trace should evolve through shared-code dynamics. The next task is to make those bridges computationally auditable.

ACKNOWLEDGEMENTS

We thank the authors of the anchor paper for making the data and supplementary materials publicly available, which makes this kind of calibrated bridge note possible.

REFERENCES

[1] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book VI: Categorical Life*. Panta Rhei Research, 2nd edition, 2026. Life as Self-Decoding Distinctions; ISBN 9798250508919 and 9798250511308; DOI 10.5281/zenodo.19553667.

[2] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book VII: Categorical Metaphysics*. Panta Rhei Research, 2nd edition, 2026. The Final Self-Enrichment; ISBN 9798250509282 and 9798250511414.

[3] Xingzhi Guo, Sergiy Verstyuk, Haochen Chen, Baojian Zhou, and Steven Skiena. Data and code from: Statistical structure and the evolution of languages, 2026.

[4] Xingzhi Guo, Sergiy Verstyuk, Haochen Chen, Baojian Zhou, and Steven Skiena. Statistical structure and the evolution of languages. *Proceedings of the Royal Society B*, 293:20252374, 2026. Received 2025-01-31; accepted 2026-01-26.

[5] Xingzhi Guo, Sergiy Verstyuk, Haochen Chen, Baojian Zhou, and Steven Skiena. Supplementary material from: Statistical structure and the evolution of languages, 2026.