

Meaning as a Stored Object: exact, model-independent memory and computation on an embedding space

Arc, The Sovereignty Foundation. Version 1.7, 15 October 2026. This is a living research paper, not a static document. The most current version is kept at sovfound.org/publications/meaning-as-a-stored-object.html, with its PDF beside it. The work is open to peer review, and the paper is updated as review comes in; each version is dated, and a later version supersedes an earlier one.

Abstract

Systems today store a text's meaning as the vector one model produced for it. That memory is approximate, so answers inherit retrieval errors with no way to tell them from facts; it is bound to the model, so changing the model means re-embedding everything; and nothing published says which operations its geometry can perform. We present a construction in which a text, embedded once by any model, is stored as its standardised relations to shared landmark texts, one byte each, beside a cryptographic digest of the exact text. Storing relations rather than coordinates frees the memory from the model that made it. Retrieval ends in a check, not a score: the relations bring the item into a shortlist and the digest confirms it or refuses. With the full digest, no retrieval in our measurements delivered a wrong item; a 34-byte compact form trades that guarantee for size. Across independently trained models, 87% of held-out texts are found within the top five of 286 candidates but only 58% exactly: neighbourhoods transfer, exact points do not. The geometry recalls and relates but does not compute, and negation needs a symbolic gate beside it. Returning each step of a computation to the nearest stored meaning stops the drift of vector arithmetic. Together these describe a small machine whose memory is meaning and which runs without a language model. Every measurement's pass bar was fixed before its code was written, and a reference library recomputes 278 recorded figures from the shipped data.

Contents

Abstract	8. Conclusion
1. Introduction	9. Data, code and record
2. Background and prior work	References
3. The construction	Appendix A — Signatures
4. Experimental method	Appendix B — Standing of each claim
5. Results	Appendix C — Cells cited
6. Discussion	Appendix D — Notation
7. Limitations	

1. Introduction

1.1 The problem

Contemporary systems store the meaning of a text as the vector a model produced for it, and treat that vector as the meaning. Five costs follow.

Memory is approximate. A search over vectors returns a shortlist, and the shortlist is wrong a measured fraction of the time. The wrong items become the context a model answers from, so the answer inherits the error with full confidence, and nothing downstream can tell a retrieved error from a retrieved fact: the only object in the pipeline is a vector, and a vector carries no receipt.

Memory belongs to the model that made it. The coordinates are valid only inside one model's geometry. Change the model and every stored meaning must be remade; two organisations that use different models hold nothing in common.

Meaning cannot move between systems. There are protocols for bytes and for vectors, and none for meaning. A vector produced by one model is noise to another.

Nobody knows what the geometry can do. Systems route questions to the embedding space on faith: answer this by nearest neighbour, compute that by vector arithmetic. No published specification says which operations a semantic geometry performs and which it cannot.

Reasoning over meaning requires a language model, and vector arithmetic drifts. Chaining operations on vectors wanders away from anything stored within a few steps, so systems that must reason over stored meaning call a language model, at a cost in latency, money and auditability.

1.2 The idea

Describe a text not by where it sits in one model's space but by how it relates to a fixed set of other texts, the landmarks, which every party holds by their exact bytes. A city located by its distances to a hundred known cities is located on every map; a city located by one map's coordinates is located on that map alone. Relations to shared landmarks survive a change of model in a way coordinates cannot, so the memory stores neighbourhoods, not coordinates.

Make that relational description the stored object. Standardise it, reduce each relation to one byte, and store beside it a digest of the exact text, so that retrieval ends in a check rather than a guess. Give each region of the corpus its own landmarks, because calibration is local. And when something must be computed rather than retrieved, compute it as a sequence of steps that each returns to the nearest stored meaning (the snap), the way digital logic restores a signal to a clean level after every gate, with a small symbolic gate for what geometry cannot do. Figure 1 shows the whole path.




Figure 1. How a meaning is stored and found again. The sender's model turns the text into numbers once; what is kept is the text's relations to shared landmarks, one byte each, with a digest of the exact text. A receiver using any model ranks its own items by the same relations, takes a shortlist, and delivers the item only if its text matches the digest; otherwise it refuses.

No part of this is new on its own; §2 names where each came from. The contribution is the assembly, the measurement of what the assembly can and cannot do, and the specification that follows from those measurements.

1.3 Contributions

1. **A stored object for meaning that ends in a check.** With the full 32-byte digest, the measured retrieval delivered no wrong item at any shortlist size: at a shortlist of ten, 126 of 136 queries were delivered exactly and 10 refused. A compact 34-byte form, keeping two bytes of the digest, resolved 97.2% of lookups exactly at a shortlist of sixteen over a 117,283-item corpus; its check is probabilistic, not guaranteed (§3.5).
2. **A memory that does not belong to its model.** The stored relations can be computed by any model that holds the landmarks, so the memory survives a change of model; giving each region its own landmarks raises cross-model accuracy by up to 0.28 when landmarks are scarce (§3.2, §3.6, §3.7).
3. **Transport of meaning between systems that share no model, with its limit measured.** Across two independently trained models, 87% of held-out texts are found within the top five of 286 candidates, and 58% exactly, under the 60% bar set in advance. Across a change of model and language at once, the median rank of the true text is 5 of 286 (§3.3, §5).
4. **A specification of what the geometry does.** It recalls at 1.00, infers relations at 0.88 and computes at 0.19. Negation is not a direction in it, conjunction is better left to it, and which tier answers a query is decided by the kind of operation (§3.12).
5. **A small machine that runs on meaning without a language model.** Its memory is the store of meanings, its instructions are geometric operations sorted by what each does to identity, and each step returns to a stored meaning. Composed by arithmetic, a two-step program reached 0 of 6 targets; with a snap after each step, 4 of 6, and a snapped round trip held 1.00 to depth eight. A run of matrix operators fuses into one 8-bit operator that scored the same as its float form on a small task (§3.10, §3.11).
6. **Time as a coordinate of the same geometry,** fixed and identical across implementations, with its own forgeability measured: the coordinate alone is forgeable, and order becomes verifiable only inside a hash chain (§3.9).

1.4 Why it matters

Each contribution removes one of the costs in §1.1, and together they change what a memory for machines can be.

A memory that is never confidently wrong. A retrieval system built on vectors looks the same when it is wrong as when it is right: it returns its nearest neighbour with the same confidence either way, and the error travels on into whatever is built from the answer. A memory that resolves by digest has two possible answers: the exact text the reference names, carrying the fingerprint that proves it, or a refusal. With the full digest, our measurements contain no wrong delivery at any shortlist size. Wherever a confident wrong answer has a price (a patient record, a legal citation, a ledger entry, a document an automated agent acts on), the memory layer stops being where errors enter and becomes where they are stopped.

A memory that belongs to its owner. Meaning written as relations to the owner’s own landmark texts is held by whoever holds those texts. Changing models becomes recomputing relations against the same landmarks; the stored references do not change. The organisation that builds a memory owns it, whichever vendor’s model drew it.

A protocol for meaning. Two systems that share nothing but a set of landmark texts can exchange meanings: 87% of held-out texts arrive in the top five across independently trained models, and the median text crosses a change of language at rank 5 of 286. What crosses and what does not is measured, and that is what lets two systems know how far to trust what they receive, the property that turns a technique into a protocol. Meaning gets what bytes and vectors already have: a way to travel between systems that were never built together.

A map of what the space can do. Builders send questions to embedding spaces on faith. This paper measures the space instead: it recalls stored facts at 1.00 and relations at 0.88, computes at 0.19, handles negation at 0.13, and beats a symbolic gate at conjunction. That is a specification a builder can design against, saying which questions to send to the geometry and which to send to an engine beside it, in place of discovering the boundary in production. We know of no other published specification of this kind for the representations in common use.

Computation with a trail. A small machine whose every step lands on a stored meaning computes over meaning without a language model. Each step is a lookup into the same store, so a result can be checked by running it again, and every intermediate state is something stored and readable. A run of its steps compiles to one 8-bit matrix, the native precision of analog in-memory hardware, so the construction’s arithmetic describes a device as well as a program.

Each of these is stated as far as the measurements carry it. §6 gives the distance from each measurement to use, and §7 lists what is not yet shown.

1.5 How to read this paper

Every measurement is a cell: a claim with a numeric bar set before the measuring code existed, run once, and recorded with its verdict, PROVEN, DISPROVEN or INVALID. Failures are reported at the same size as passes, because several design rules exist only because of a failure. The cells behind a measurement are listed in small brackets after it, and Appendix B gives the standing of every claim: recomputed from the shipped files, recorded only, or measured on data that is not public.

The terms below recur. Each is defined again where it first matters.

Term	Plain meaning
Frame	One model’s coordinate system: the numbers a particular model assigns to texts.
Landmark (anchor)	A shared text that every party holds by its exact bytes; meanings are described by their relations to landmarks.
Profile	A text’s list of relations to the landmarks, one number per landmark.
Standardisation (gauge)	Rescaling a profile to zero mean and unit spread, so profiles from different models are comparable.
Hub	A text that is near almost everything and therefore identifies nothing; the decode corrects for hubs.

Term	Plain meaning
Shortlist	The few candidates a receiver considers after ranking by relations.
Digest (commitment)	A cryptographic fingerprint of the exact text; matching it confirms identity.
Reference	The stored or sent object: a region number, the profile in bytes, and the digest or part of it.
Chart (region)	A part of the corpus with its own landmarks.
Snap	Replacing a computed point by the nearest stored meaning.
Rank-1, top-5	Whether the true item is the first candidate, or among the first five.
Bar	The pass threshold fixed before a measurement was run.

2. Background and prior work

Relative representations (Moschella et al., 2023) describe a point by its similarities to a set of anchors so that two latent spaces become comparable; this is the profile of §3.2, and the paper credits it as the origin of the idea. The difference is what is done with it: here the profile is the stored object, calibrated per region, quantized and bound to a digest, rather than an alignment step, and its limits are measured. *Cross-domain similarity local scaling* (Conneau, Lample et al., 2018) is the hub correction of §3.3, used as published, with a measurement of what it does not do. *The Platonic representation hypothesis* (Huh et al., 2024) argues that independently trained representations converge; §3.3 and §5 measure how far that convergence reaches for a point and for a neighbourhood. *TwoNN* (Facco et al., 2017) is the intrinsic-dimension estimator of §5.3, run beside an un-learned baseline. *Product quantization* (Jégou et al., 2011) and *HNSW* (Malkov and Yashunin, 2018) are approximate-search tools a shortlist may use; the digest check turns their misses into refusals rather than wrong answers. *Contrastive learning* (van den Oord et al., 2018) is the training rule of the encoder mentioned in §3.7. *Sinusoidal position codes* (Vaswani et al., 2017), *rotary embeddings* (Su et al., 2024) and *Time2Vec* (Kazemi et al., 2019) make time or position legible to a model as a learned or fixed feature; §3.9’s coordinate is fixed, identical across implementations, and its forgeability is measured. *Ollivier–Ricci curvature* (Ollivier, 2009) is the discrete curvature used for navigation over the store’s graph. *Gromov–Wasserstein alignment* (Alvarez-Melis and Jaakkola, 2018) is the strongest structure-only matcher tested in §5.2, where it identifies nothing without shared landmarks. *Model stitching* (Lenc and Vedaldi, 2015; Bansal et al., 2021) composes layers across networks; fitting a new model to an existing memory in §3.7 is its memory-side analogue. *Signal regeneration in digital logic* is the analogy the snap is named for, and it is an analogy, not a citation.

3. The construction

This section builds the object one part at a time. Each part opens with what it is and what it is for in plain terms, then gives the formula, then the measurement that supports it, then its consequence. The signature behind each formula is listed in Appendix A.

3.1 Frames

A frame is one model's coordinate system. The construction exists to store meaning in a form that does not depend on which frame drew it, so the definition comes first. A frame is an embedding function E : $\text{text} \rightarrow \mathbb{R}^d$ into a normed space. Two frames are independent when they share no parameters, data ordering or training run; their dimensions need not agree. The construction never assumes that a map between two frames' coordinates exists. The measurements use five public models of 384 to 1,024 dimensions, one of them multilingual.

3.2 The profile (M01, M02)

The first move is to stop describing a text by its coordinates and describe it instead by its relations to a fixed set of shared texts, the landmarks. Coordinates belong to one model; relations between texts can be computed in any model, so two parties who hold the same landmarks can each compute this description in their own model and compare the results.

A landmark set A is a finite ordered set of texts identified by content, by their exact bytes, never by a coordinate, and held by every party. In the reference construction the landmarks are drawn deterministically from a byte-wise ordering of the shared corpus, so two parties who hold the same corpus hold the same landmarks without communicating. The raw profile of text x under frame E is the list of its cosine similarities to each landmark, in order,

$$p_x[k] = \cos(E(x), E(a_k)), k = 1 \dots A \text{ (M01)},$$

and the stored profile is its standardisation, which removes each model's own offset and scale,

$$z_x = (p_x - \text{mean}(p_x)) / \text{std}(p_x) \text{ (M02)}.$$

Each part of the recipe was removed in turn to measure what it carries (Figure 2). Matching the right text across two independently trained models, first time, ran 0.566 with the full recipe, 0.515 without the standardisation, 0.493 without the hub correction of §3.3, and 0.294 without both; storing the profile in signed bytes changed nothing, and halving the landmarks from 150 to 75 cost 0.121. [V21] The number of landmarks sets how much can be matched: on an exploratory set, exact matches rose from 20.6% at 50 landmarks to 58.2% at 150. [V9c]




Figure 2. What each part of the recipe carries. Each bar is rank-1 cross-model matching with one part removed; storing the profile as bytes costs nothing, and removing both corrections halves the result.

Consequence. A text is described by A numbers that depend on the model only through similarities to shared texts. Two parties holding the same landmarks can compare descriptions without sharing a model, a dimension or a coordinate.

3.3 Decode (M03, M04)

Decoding is how a receiver, holding a sender's profile and its own profiles of the same corpus, finds which of its own texts the profile names. Plain similarity between profiles is dominated by hubs, texts that are near everything and therefore identify nothing, so the decode corrects for them.

Given sender profiles z_i and receiver profiles z_j , the base similarity is $\text{sim}(i, j) = \cos(z_i, z_j)$, and retrieval uses cross-domain similarity local scaling with $k = 10$,

$$\text{csls}(i, j) = 2 \cdot \text{sim}(i, j) - r_S(i) - r_R(j) \text{ (M03)},$$

where $r_S(i)$ is the mean of the ten largest similarities of i and $r_R(j)$ the same for j . The decode of message i is the j with the largest csls , ties to the lowest index (M04). The correction lifted 150-landmark exact matching on the exploratory set from 58.2% to 64.4%. [V9c] What the profile does not do is also measured: inside one model it ranks exactly as raw cosine does, 0.1233 against 0.1233 on fragment retrieval, so profiles buy transport across models, not better ranking within one. [CELL-TRIANGLE-1]

How far meaning transports is shown in Figure 3. Between two independently trained models, over 136 held-out texts and 286 candidates, the true text was first 58.1% of the time, under the 60% bar set in advance, and within the top five 86.8% of the time, where chance is 1.7%. [V9d] Between an English model and a Spanish multilingual model, the true text's median rank was 5 of 286, within the top five 55.9% and first 24.3% of the time. [V18]

Neighbourhoods cross between models; exact points mostly do not




Figure 3. What crosses between models. Left, two models in one language; right, two models and two languages. The exact text is usually near the top; it is first less often, and below the bar set in advance.

Consequence. Neighbourhoods cross between models; exact points mostly do not. Exactness therefore cannot come from the geometry, and §3.5 is where it comes from.

3.4 The wire form (M05)

A profile is A real numbers. For storage and transmission each is reduced to one signed byte, and the question is whether that costs anything.

A profile is quantized as $q[k] = \text{round}(127 \cdot z[k] / \max|z|)$, one signed byte per landmark, and the receiver standardises again on receipt, which makes the scale unnecessary to send. The bar, set before the run, allowed quantization to cost at most 0.02; matching scored 0.566 in floating point and 0.566 in bytes, and the correlation of similarity with graded similarity levels was 0.701 in both. [V21, V1] These are equal aggregate scores on the measured tasks; they do not show that every ranking is identical.

Consequence. Precision is not where the meaning lives. Everything downstream is built at one byte per relation.

3.5 The digest and the reference (M06, M07)

A profile brings a text near the right answer; it does not name it. The reference adds a cryptographic digest of the text's exact bytes, so the last step of retrieval is a check rather than a score. This is where the construction's exactness comes from, and it comes in two strengths.

The full form. The reference carries the region number, the profile in bytes and the whole 32-byte digest (blake3 in the reference implementation). The receiver ranks its own texts by the profile, takes a shortlist of k , and delivers a candidate only if its text hashes to the digest; if none does, it refuses. A wrong delivery would need two different texts with the same 32-byte digest, which a cryptographic hash makes practically impossible. With the full digest beside the profile, the measured retrieval delivered no wrong item at any shortlist size: the share delivered equals the share of queries whose true text is in the shortlist, from 58.1% at $k = 1$ to 97.1% at $k = 20$, and the rest are refused. At $k = 10$, 126 of 136 queries were delivered exactly, 0 wrong, 10 refused (Figure 5). [V124, V128] With two bytes for the region and thirty for the profile, the full form is 64 bytes.

The compact form. The shipped layout keeps only two bytes of the digest: two bytes naming the region, thirty bytes of profile (one byte for each of the region's 30 landmarks) and two bytes of digest, 34 bytes in all (Figure 4). Two bytes almost always single out the right candidate but cannot guarantee it. If the true text is missing from the shortlist and another shortlisted text happens to share the same two bytes, the compact form delivers the wrong text; with a shortlist of sixteen that happens in about 1 of every 4,100 lookups whose true text is missing. The compact form is therefore a fast lookup, and the guarantee needs the full form. On a production corpus of 117,283 texts, the compact reference resolved 97.2% of 2,000 lookups exactly at a shortlist of sixteen, against a bar of 90%, where a structureless 156-byte profile reached 79.0%; every miss was a true text absent from the shortlist. [CELL-REF32]

The reference as a code. Single-bit damage was seeded through the receiver, 500 flips partitioned by field. Flips in the profile bytes were healed, the right text still delivered, 98.9% of the time and otherwise refused; flips in the digest and region bytes were refused 56 of 56; wrong text was delivered 0 of 500. [CELL-HEAL-1] On the shipped public data the same experiment healed every profile flip and refused every digest and region flip. The address corrects errors where it is geometric and exposes tampering where it is symbolic, within the probabilistic limit of a two-byte digest.




Figure 4. The two forms of the reference. Above, the compact 34 bytes, with the measured response to a flipped bit in each field; below, the full 64-byte form, whose 32-byte digest makes the final check a guarantee.




Figure 5. With the full digest, the shortlist decides how much is found, and nothing found is wrong. The line is the share of 136 queries whose true text is within the top k; everything above the line is refused.

Consequence. A meaning is addressed in tens of bytes where a vector occupies thousands, and the last step of retrieval is a check. With the full digest the check is a guarantee; with the compact digest it is a fast filter whose error rate is computable.

3.6 Regions and local landmarks (M08, M09)

A corpus is not one uniform space, and one set of landmarks does not describe all of it equally well. The corpus is therefore divided into regions, called charts, each with its own landmarks and its own standardisation.

At equal landmark budget, landmarks drawn from a text’s own region describe it better than landmarks drawn from everywhere: across six model pairs and three seeds the mean gain in exact matching was +0.218, seven times its bar of +0.03. [V28] The gain is largest when landmarks are scarce, +0.28 at 30 landmarks, and narrows to +0.07 at 100 as the mixed landmarks catch up, under that cell’s bar of +0.10 (Figure 6). [V127]




Figure 6. Regional against mixed landmarks at three budgets, rank-1 across two models. The local advantage shrinks as landmarks become plentiful.

Consequence. Calibration is local, and its premium is real when landmarks are scarce. A region can be found before any model is consulted, from the region number stored in the reference.

3.7 Independence from the model

Because the stored object is the set of relations, the model that drew it is replaceable. A new model joins an existing memory by computing its own profiles against the same landmarks; nothing in the stored references changes. The evidence for this is §3.3: two models that share nothing but the landmarks find each other’s texts in the top five 87% of the time.

Two further routes were measured on the production corpus and are reported in the record: fitting a new model into the memory by small per-region linear maps, and an encoder trained from the corpus alone with no pretrained weights that serves the same memory. Both were measured on data that is not public, and this paper does not rest a claim on their figures.

A learned operator is a byte artifact. A linear operator fitted on more training pairs than its dimension serialises to a fixed block: 1,600 bytes of coefficients in 32-bit floats and 400 in signed bytes (1,610 and 414 bytes as files), the float form round-tripping byte for byte and the byte form agreeing with it on every held-out item. [V109b]

Consequence. The memory belongs to whoever holds the corpus and the landmarks, not to the vendor of a model, and a learned function is a file.

3.8 The void gate (M10)

A retrieval system that always returns its best candidate is wrong whenever the answer is not in the store. The void gate is a floor on the match score, calibrated on the live corpus, below which the system refuses instead of answering.

Whether a floor separates what the corpus holds from what it does not depends on the corpus. On the author's corpus the lowest in-domain match (0.627) sat above the highest foreign match (0.623), and a floor of 0.625 separated them completely, with a recorded caveat: the foreign probes were single sentences and the in-domain probes full passages, so length may inflate the separation. [V102] On the shipped hand-written corpora, which have no single domain, the ranges overlap, the in-domain minimum (0.454) far below the foreign maximum (above 0.74), and no floor separates them. A length-matched test on a public corpus with a domain is drafted.

Consequence. A system that can say “I have nothing on this” need not say something wrong, but the floor that lets it do so is a property of the corpus and must be measured per corpus.

3.9 Time as a coordinate (SIG-T family)

A stored meaning has a when as well as a what, and the construction places the when in the same space as the what, as numbers any party computes identically. Time enters as a fixed coordinate, never a learned feature. For a cycle of period P , the angle is $\theta(t) = 2\pi (t \bmod P)/P$ and the position on that cycle is $(\sin \theta, \cos \theta)$; stacking the daily, weekly and yearly cycles with a slow epoch term and a constant gives

$$T(t) = [\sin \theta_d, \cos \theta_d, \sin \theta_w, \cos \theta_w, \sin \theta_y, \cos \theta_y, at, 1].$$

Trust-bearing quantities decay as $V(t) = m \cdot e^{-\lambda \Delta t}$ with $\lambda = \ln 2 / t_{1/2}$, so expiry is a consequence rather than a message, and an entity's origin is the decay-weighted centre of its own history. A sequence of meanings is then a trajectory, a habit is a loop in the phase coordinates, and a departure is a distance from the loop. The Foundation's Kronos paper develops the coordinate in full; this paper carries what the memory needs.

The audit of the clock begins with a disproof. The coordinate alone is forgeable: it is deterministic, so anyone who can state a date can recompute it, and none of its eight dimensions resisted a forger. [V3] Integrity lives in a hash chain: with the time bound into a content-committing chain, back-dating was rejected 500 of 500 times and tampering detected 1,000 of 1,000. [V4, V10] The chain had a measured hole: when the epoch was bound only into the first segment, 143 of 200 mixed attacks were caught, so order is verified while absolute time is still asserted. [V30] On a simulated source, two parties recover a shared decay rate without communicating while their individual events stay uncorrelated ($|r| \leq 0.079$), so the rate is public and the events private. [V31]

Consequence. The time coordinate gives shape, recurrence, departure and fading as distances, and the chain gives integrity; the record keeps the two apart.

3.10 The instruction set and the snap (M11, M12)

Once meanings are stored as points, operations on them are operations on points, and a program is a sequence of such operations. This section states the instructions and the one mechanism, the snap, without which programs drift.

The instructions fall into four classes by what they do to the identity of a meaning (Figure 7).

Identity-keeping. MOVE applies a learned relation as an average offset $t = \text{mean}(\text{target} - \text{source})$ over training pairs and computes $\text{source} + t$. On three relations, country to capital, singular to plural and present to past, it reached the right target on all 15 held-out pairs, and moved points stayed among real data. [V73]

Many to one. These map several inputs to one output and so cannot be undone: a reversible operation cannot send two different specifics to one category. PROJECT maps a specific to its category, 14 of 16 held out, through an operator of effective rank 14.7 of 40. [V75] CENTROID reduces a set to its summary and converges on the concept, 4 of 4, whatever the order. [V77b] INTERPOLATE blends two meanings to their midpoint, which landed on the right intermediate 6 of 10 times against a bar of 5 of 10 and stayed as close to real data as its endpoints. [V76]

In and out. ADDRESS stores a meaning and finds it again by its relations: the load and store of the machine, measured at one byte per relation in §3.4. Its retrieval step is the snap,

$$\text{snap}(x) = \underset{j}{\operatorname{argmax}} \cos(x, S_j) \text{ (M11)},$$

the nearest stored meaning, ties to the lowest index.

Control. COMPOSE sequences the others: apply one instruction, snap to the nearest stored meaning, apply the next. MOVE with a snap is $\text{move}(x, \delta) = \text{snap}(x + \delta)$ with the starting point excluded (M12); without the exclusion a small offset returns the start itself. With it, forced iteration ended on all 20 seeds in about three steps, in small two-step loops at the category level, dog to animal to dog. [V105b]



Figure 7. The instruction set, sorted by what each instruction does to the identity of a meaning.

The measurement that forces the snap is composition, shown schematically in Figure 8 and measured in Figure 9. Masculine plus a gender offset plus a plural offset, added as vectors, arrived at the feminine plural 0 of 6 times, landing on the masculine plural every time, although each single step worked on its own. With a snap after each step the same walk arrived 4 of 6 times, and every failure was a bad single step, none caused by composing. [V74, V74b] On a larger bank of 24 concepts the reference library's re-run gave 1 of 6 by arithmetic and 2 of 6 with snaps, again with no errors caused by composing: snapping removes the drift, and the arrival rate is set by how good each single step is. A singular-to-plural round trip, snapped at every step, held accuracy 1.00 to depth eight, with closeness to real data unchanged from 0.935 to 0.934. [V103]




Figure 8. The mechanism in a picture, drawn as a schematic rather than from data. Added offsets drift away from anything stored, one step at a time; a snap after each step lands on a stored meaning every time.

Adding offsets drifts; restoring to a stored meaning after each step does not




Figure 9. Composition by arithmetic against composition with a snap, and a snapped round trip to depth eight.

The machine is defined, following the filed specification, as a finite set S of stored meanings, a set of operators on the ambient space, the snap onto S , and the rule that every instruction is followed by a snap, $s' = \text{snap}(o(s))$, so no un-snapped intermediate crosses an instruction boundary.

Consequence. Chains of operations that return to stored meanings at every step do not accumulate drift; what remains is the error of each single step, which is measured. The measured chains are short and the concept sets small, so this is a mechanism shown, not a general guarantee for arbitrary programs.

3.11 Compilation

A program of many instructions need not be executed as many steps. Matrix operators compose exactly, because multiplying matrices is composing functions: T_1 then T_2 is the single operator $T_2 \cdot T_1$. A run of matrix instructions with no snap between them can therefore be fused into one operator, applied once.

A snap between two instructions breaks that equivalence. Snapping is not linear, so applying T1, snapping, then applying T2 need not land where the fused T2·T1 lands. The compiler may fuse only the runs a program declares continuous and must keep every snap the program requires; equivalence of a compiled program with the step-by-step interpreter has not been measured state by state, and this paper does not claim it.

What was measured is precision. A fused operator scored 0.667 on six held-out concepts in floating point and 0.667 in signed bytes, far above the additive control at 0.167, with a largest relative quantization error of 0.0039. [V112] A dense matrix at 8-bit precision is the form an analog in-memory crossbar computes natively; the mapping is disclosed in the specification and no device is built.

Consequence. A continuous run of a meaning program can be compiled to one small 8-bit matrix with no measured loss on the task tested; a program that must snap between steps keeps those steps.

3.12 The symbolic gate and the routing rule (M17)

Not every operation on meaning is a movement in the space, and the record finds the ones that are not.

Negation is not a direction. The offsets from words to their opposites share no common heading, consistency 0.114 against 0.380 for plurals; a sentence and its negation sit at cosine 0.871, closer than unrelated sentences at 0.535; and antonyms sit closer than random pairs, 0.655 against 0.436, because opposites share a topic. Distance in the space means unrelatedness, never contradiction. [V78, V53, V54]

So the construction carries two tiers: the geometry, for finding and blending, and a small symbolic gate for negation, intersection of candidate sets, exact identity by digest, and branching on the void floor. Measured on constrained retrieval, geometry alone put the right answer first on 13% of negation queries and the gate over geometry's shortlist on 67%; on conjunction the gate fell to 19% where the plain geometric blend scored 43%, so neither tier alone met the 0.70 bar and the hybrid cell is DISPROVEN, each tier doing what the other cannot (Figure 10). [V116] Routing by the category of the question, type constraints to the gate and attributes to the blend, scored 0.3968, below the gate alone at 0.4444. [V117] The rule that survives is to route by the kind of operation: negation to the gate, conjunction to the blend. An operation of a kind the system has not been told how to handle should be refused rather than sent to a default tier.




Figure 10. Left, what the geometry does on its own; right, the gate against the geometry on NOT and AND queries.

Consequence. The construction has a measured instruction set and a measured boundary: what geometry does, what it cannot, what must stand beside it, and a rule for choosing.

3.13 Two measured regularities (M13, M14)

Two measurements bound what the construction can be asked to do: how agreement weakens when meaning is relayed through several models, and how much agreement more landmarks buy.

Relaying through models. When a meaning is passed from model A to B to C to D, the end-to-end match rate was never above the weakest single link and never below the product of the links (Figure 11). Through three hops it was 0.143, against a weakest link of 0.333 and a product of 0.063, above the product because the same hard texts fail at every hop. [V126] On all 48 paths through three and four of four models, every relayed rate lay between the two. [V147] This is an observed regularity on these models and texts, not a theorem.

More landmarks, diminishing returns. Exact matching across two models rose from 0.268 at 25 landmarks to 0.883 at 400 (Figure 12). A simple averaging law, fitted on 25 to 100 landmarks, predicted 0.97 at 150 where 0.70 was measured, missing by up to 0.27 against a bar of 0.05. [V130, V154] Adding a constant term to the law, a fitted value of 0.222, cut the misses to at most 0.073, still outside the bar; the constant is a parameter of this fit on this pair and corpus, not a universal floor.

[V158]




Figure 11. A meaning relayed through four models. The end-to-end rate falls below the weakest link and stays above the product of the links.




Figure 12. Exact matching against the number of landmarks. Both laws were fitted on the left of the shaded band and predicted inside it; the pure averaging law overshoots, and the version with a constant comes closer but misses its bar.

Consequence. Agreement weakens predictably through intermediaries and grows with landmarks more slowly than averaging predicts. Exactness cannot be bought with more landmarks; it comes from the digest.

4. Experimental method

Every measurement is a cell: a claim, the result that would refute it, a numeric bar fixed before the measuring code existed, one run, and a verdict by the bar. A cell that fails stands in the record at full size.

The meaning cells run on hand-written text sets of 61 to 613 texts, embedded by five public models: nomic-embed-text v1.5 (768 dimensions), all-MiniLM (384), mxbai-embed-large (1,024), snowflake-arctic-embed (1,024) and the multilingual bge-m3 (1,024). The embeddings are frozen to files, so

every cell recomputes from the same numbers. The dimension cell runs on 4,000 public maths word problems. Three cells, the compact reference and the damage experiment of §3.5 and the void-gate calibration of §3.8, ran on a production corpus of the author’s documents, which is not public. One model’s tokenizer was found faulty and repaired, and the affected cells were re-run; where a re-run exists this paper gives the re-run figure throughout.

Two properties of the protocol are stated so a reader can weigh them. The hub correction and the standardisation are computed over the whole candidate pool, held-out texts included: the decoder sees which texts are candidates, though not which is the answer, and a fit on training texts alone has not been measured. And a reference decodes only under the same landmarks, standardisation, hub-correction setting, precision and tie rule on both sides; these are fixed in the specification but not yet versioned inside the reference.

A reference library written from the paper’s formulas recomputes the record from the shipped embeddings. The version tested is the Sov Meaning Machine, package `mmachine` 0.2.0 at commit 18c042f, run unmodified on macOS (arm64) under Python 3.11.5 with NumPy 1.26.4 and under Python 3.14.4 with NumPy 2.4.4; Linux and Windows have not yet been tested. It computes in 64-bit floating point, storing the quantization scale in 32 bits. On 13 October 2026 it recomputed 278 of the record’s figures and matched every one to the printed digit, and it gave the same figures on a second machine with a different Python and numerical library; re-measuring 33 figures on fresh embeddings, 27 coincided and 6 differed, one by rounding (2.95 printed as 3.0) and five where the library uses its own relation lists. Matching the record shows the recorded computation is the one described; it is not an independent reproduction, which under the Foundation’s method requires another party on another machine.

5. Results

5.1 What holds

Result	Measurement	Cells
Relations survive a change of frame	100% recovery against 4.4% for raw coordinates under rotation; 97.8% with noise added (a mechanism check)	V9
Neighbourhoods cross independent models	top five 86.8% of 136 held-out texts among 286 candidates; chance 1.7%	V9d
Neighbourhoods cross languages	median rank 5 of 286; top five 55.9%; first 24.3%	V18
Each part of the recipe carries something	0.566 full; 0.515 without standardisation; 0.493 without hub correction; 0.294 without both; 0.445 with half the landmarks	V21
Bytes cost nothing on the tasks measured	0.566 in float and in bytes; similarity correlation 0.701 in both	V21, V1
The recipe rebuilds from its text	a second implementation matched to four decimals, 0.5662, every decode the same; a reference library matches 278 of 278 recorded figures	V22

Result	Measurement	Cells
With the full digest, nothing delivered is wrong	delivered equals in-shortlist from 0.581 (k = 1) to 0.971 (k = 20); at k = 10, 126 delivered, 0 wrong, 10 refused	V124, V128
The compact reference resolves at scale	34 bytes, 97.2% exact at a shortlist of 16 over 117,283 texts; a flat 156-byte profile 79.0%	CELL-REF32
The address is a code	profile flips healed 98.9%; digest and region flips refused 56 of 56; wrong 0 of 500	CELL-HEAL-1
Local landmarks help when scarce	+0.218 mean gain at 60 (bar +0.03); +0.28 at 30; +0.07 at 100, under its +0.10 bar	V28, V127
The snap stops composition drift	0 of 6 by arithmetic; 4 of 6 snapped, no composition errors; round trip 1.00 to depth 8	V74, V74b, V103
The instructions do what their class says	MOVE 15 of 15; PROJECT 14 of 16 at rank 14.7 of 40; CENTROID 4 of 4; INTERPOLATE 6 of 10	V73, V75, V77b, V76
Forced loops terminate	20 of 20 seeds in about three steps, in local two-step cycles	V105b
A fused operator survives 8 bits	0.667 in float and in bytes on six concepts; additive control 0.167	V112
An operator is a file	1,600 / 400 coefficient bytes (1,610 / 414 as files); float round trip exact; byte form agrees 1.00	V109b
Relaying sits between product and weakest link	48 of 48 paths; three hops 0.143 against a weakest link of 0.333	V126, V147
Order can be bound, not forged	back-dating rejected 500 of 500; tampering caught 1,000 of 1,000	V4, V10
A decay rate is public, events private	shared rate recovered without communication; events correlate at most 0.079	V31
Identifiability is graded	profile values 0.581, ranks alone 0.500, bare landmark sets 0.338; values survive damage better, 4 of 4	V143, V146

5.2 Where it stops

Limit	Measurement	Cells
Exact points do not transfer	58.1% against a bar of 60%; 22.6% at 30 landmarks without hub correction; 37.9% at 100 on the exploratory set	V9d, V9b, V9c
Averaging does not describe the landmark budget	misses held-out budgets by up to 0.27; with a fitted constant, by up to 0.073	V154, V158
Agreement belongs to the ordered pair of models	3 of 10 pairs asymmetric beyond 0.05, the largest 0.110	V149
Agreement does not saturate early within a pair	the saturating form over-predicts at 150 landmarks by 0.116	V125

Limit	Measurement	Cells
The geometry does not compute	recall 1.00, relational inference 0.88, computation 0.19	V132b
Negation is not a direction	offset consistency 0.114 against 0.380; a sentence and its negation at cosine 0.871	V78, V53, V54
The gate cannot conjoin	AND: gate 0.19 against the blend's 0.43; the hybrid 0.47 against a bar of 0.70	V116
Routing by category loses	0.3968 against the gate alone at 0.4444	V117
Agreement is not truth	true statements 0.5106 against false 0.5150; 10 wins in 20 pairs	V40
Nothing transfers without shared landmarks	three structure-only matchers at chance, 0.0035 to 0.0070	V145
The time coordinate alone is forgeable	0 of 8 dimensions resisted; binding at one point caught 143 of 200 attacks	V3, V30
The founding equation was not algebra	frame maps neither compose nor invert (errors 0.448 and 0.390 against 0.25); the invariant 0.358 against the best single frame's 0.623	V80–V84

What the limits mean together. Because exact points do not cross and more landmarks do not close the gap, exactness cannot come from the geometry; it comes from the digest. Because nothing transfers without shared landmarks, agreement between systems is built from what they share and is never discovered from structure alone. Because the geometry does not compute, computation goes to a gate, chosen by the kind of operation. Because agreement is not truth, trust in what the memory holds must be earned some other way than by consensus. The limits are the specification.

5.3 How much of the space meaning uses (M15, M16)

On 4,000 public maths word problems with their worked solutions, embedded at 1,024 dimensions by a public model, the two-nearest-neighbour estimator gave an intrinsic dimension of 25.6, with a participation ratio of 124.6. The decisive comparison is the baseline: a bag of the same words, hashed and counted with nothing learned, has an intrinsic dimension of 23.7. The thinness belongs to the texts, not to the embedding, and the paper states the result at that scope and no wider. [V39p]

5.4 Capacity

Exact and neighbourhood recovery are two samples of one curve: on the held-out pair, the share of true texts within the top k ran 0.581, 0.706, 0.809, 0.868, 0.927 and 0.971 at $k = 1, 2, 3, 5, 10$ and 20 (Figure 5). The zero-error branch costs exactly the digest: at every k the delivered share equals that curve and wrong deliveries are zero. [V124, V128] Within one pair of models, recovery rises nearly linearly with landmarks up to 150, and the saturating form over-predicts; across pairs it is asymmetric, depending on which model sends. [V125, V149] No formula for capacity is claimed.

5.5 The founding theory, tested

The work began from the theory that meaning is what survives a change of frame: a representation divided by the group of frame transformations. Five cells, each with its bar fixed in advance, tested it on four models over 120 concepts. The learned frame maps did not compose or invert cleanly, so they

do not form a group; a concept's spread across frames was 2.17 times the gap to its nearest neighbour; and the invariant scored 0.358 against the best single frame's 0.623, approaching 1.0 only as frames were added. Four of the five cells were DISPROVEN. The profile of §3.2 replaced the equation: it builds correspondence from shared landmarks rather than discovering it, and §5.2 says how much correspondence exists. [V80, V81, V82, V83, V84]

6. Discussion

6.1 What changes

Set side by side, the construction changes each part of how a machine holds meaning. Every row rests on a measurement in this paper.

	Memory today	Meaning as a stored object	Measured basis
What is stored	One model's vector, about 3,000 bytes	Relations to shared landmarks and a digest, 34 or 64 bytes	§3.4, §3.5
A wrong retrieval	Looks exactly like a right one	The exact text with its digest, or a refusal	0 wrong at every shortlist size [V124, V128]
Changing the model	Re-embed the whole corpus	Recompute relations against the same landmarks	§3.3, §3.7
Two systems on different models	Share nothing	Share meaning to a measured degree	87% in the top five across models; median rank 5 of 286 across languages
What the space can do	Assumed	Specified: recall, relate, do not compute; negation to a gate	§3.12, §5.2
Reasoning over stored meaning	A language model, or vector arithmetic that drifts	Steps that each land on a stored meaning	0 of 6 by arithmetic; 4 of 6 with a snap
Checking a result	Read the log	Run it again	§3.10

6.2 How far each change is from use

Exactness at the memory layer. With the full digest, retrieval delivered no wrong item at any measured shortlist size and refused when the true text was not in the shortlist. For any use in which a confident wrong answer has a cost, the memory can answer with the exact text and its digest, or decline. The measurement is on 136 queries; the compact form trades the guarantee for size at a computable risk.

Ownership. The memory belongs to the corpus rather than to the vendor of the model, because it is written in relations to texts the owner holds. The storage arithmetic follows from §3.5: a reference of 34 or 64 bytes against about 3,000 bytes for a 768-number vector, so the reference payload for a billion texts is tens of gigabytes rather than terabytes. That is payload arithmetic only. The texts themselves, the model, the landmark profiles and any search index still cost what they cost, and full system memory and latency at matched recall have not been measured.

Transport. Two systems that share landmarks and nothing else share meaning to a measured degree: 87% within the top five across models, a median rank of 5 in 286 across languages, 58% exactly, and nothing without shared landmarks. The limit is what makes this a protocol, and the relay measurement says how agreement weakens through intermediaries. A second party decoding references it did not produce, from the specification alone, is not yet demonstrated.

Specification. The geometry recalls and relates but does not compute; negation is not a direction in it; conjunction is better left to it than to the gate. That is a specification of which questions belong on an embedding space and which need an engine beside it, and of the rule for choosing, by the kind of operation.

Computation without a language model. Programs composed by vector arithmetic fail; the same programs with a snap at each step do not accumulate error; a continuous run compiles to one 8-bit matrix with no measured loss on the task tested. Computation over stored meaning can therefore be deterministic and repeatable, and checking it means re-running it. The measured programs are short and the concept sets small; equivalence between compiled and step-by-step execution is not yet shown, and the in-memory device is specified, not built.

Time. Recurrence, departure and fading become distances any system computes the same way, and permissions can expire by decay rather than by revocation. The integrity of time comes from the chain, not the coordinate.

6.3 Where it applies

The measurements make the following uses plausible. None is a deployment we have run.

- **Memory for agents.** An agent that acts on what it retrieves can be given a memory that hands back the exact source with its digest or says it has nothing, so every action traces to a text that exists.
- **Retrieval where an error costs.** Clinical, legal, financial and regulatory systems, where a citation must be the cited text, can store references and refuse rather than approximate.
- **Changing models without rebuilding.** An organisation can move to a new embedding model by recomputing relations to its landmarks, keeping every stored reference.
- **Sharing between organisations.** Two parties on different models can exchange references over an agreed landmark set without sharing a model, its weights or its vectors, and the measured limits say how far to trust each exchange.
- **Small devices.** At 34 or 64 bytes a reference, the stored payload fits constrained hardware; the full system cost on such hardware has not been measured.
- **Audit.** Because each step of a computation lands on a stored meaning and each answer carries its digest, a reviewer can re-run a computation and check every answer against the text it names.

6.4 The machine

The construction is the memory of a small computing machine, the Sov Meaning Machine, and what this paper states of the machine is what the filed specification discloses and §3.10 to §3.12 measure: instructions sorted by what they do to identity, a snap that returns every step to a stored meaning, compilation of continuous runs into one 8-bit operator, and a symbolic gate for what the geometry cannot do, chosen by the kind of operation. Stated plainly, it is a computer whose memory is meaning,

whose every step lands on something stored, and whose answers can carry the digest of the exact text they came from. It does not contain a language model; a model is used, when at all, to compute profiles at the boundary.

7. Limitations

The measurements support the claims above at their stated scope. These are the gaps a reader should weigh.

- **The compact reference is probabilistic.** Its two-byte digest can deliver a wrong text when the true text is absent from the shortlist; the guarantee requires the full digest, at 64 bytes.
- **Scale.** The meaning cells use 61 to 613 texts and the exactness measurement 136 queries; the compact reference was measured on 117,283 texts. Nothing here demonstrates behaviour at ten million.
- **Small sets behind the instructions.** MOVE was tested on fifteen pairs, INTERPOLATE on ten triples, composition on six concepts; these are clean signals at modest size, not large-sample laws.
- **Compilation.** Fusing is exact only for continuous matrix runs; equivalence with the snapped interpreter has not been measured.
- **The protocol sees the candidates.** The hub correction and standardisation are computed over the candidate pool, held-out texts included.
- **Versioning.** A reference decodes only under identical landmarks, standardisation, hub correction, precision and tie rule, and these are not yet versioned in the reference.
- **The void floor is corpus-dependent.** It separated in-domain from foreign probes on one corpus, with a length confound, and does not separate them on the shipped hand-written sets.
- **Regional gains** shrink as landmarks become plentiful and fell under their bar at 100.
- **The landmark-budget fit** missed its bar; the fitted constant is not a universal floor.
- **Inside one model,** profiles rank no better than plain cosine.
- **Data.** The compact-reference, damage and void-gate cells ran on a corpus that is not public.
- **Routing** is by kind of operation; refusing unknown kinds is a requirement the current implementation does not yet meet.
- **The physical decay source** behind §3.9 is specified and not built.

8. Conclusion

This paper set out to store meaning as an object rather than as a vector, and measured what that object can and cannot do. Relations to shared landmarks carry neighbourhoods across models and languages; they do not carry exact points, and no number of landmarks makes them. A digest does: with the full digest, retrieval returns the exact text or refuses, and in our measurements it never returned a wrong one. The space recalls and relates but does not compute, so computation runs as a sequence of steps that each land on a stored meaning, with a small gate for what the space cannot do.

The result is a memory owned by whoever holds the corpus, a measured protocol for meaning between systems that share no model, a specification of what an embedding space can be trusted with, and a small machine that computes over meaning without a language model. Every claim has a bar that was fixed before it was measured, the failures stand beside the passes, and a reference library recomputes the record. §7 names the next measurements; the construction they will test is fixed here.

9. Data, code and record

Every measurement is a cell whose bar was fixed before its measuring code was written, run once and recorded; the Foundation’s record keeps the method, the standing of each cell, and every correction, with the failures at the same size as the passes. Some cells were measured on data that is not public; public versions are drafted and replace them as they are run. The construction is implemented in a reference library in Python, the Sov Meaning Machine (package `mmachine`), a pre-release research preview, which recomputes the record’s figures from the shipped embeddings with `python3 -m mmachine check` and runs on any corpus with `python3 -m mmachine corpus`; the formulas and their signatures are listed in Appendix A. Appendix B gives the standing of every claim in this paper.

References

- Alvarez-Melis, D., and Jaakkola, T. S. (2018). Gromov-Wasserstein alignment of word embedding spaces. *Proceedings of EMNLP 2018*, 1881–1890. doi:10.18653/v1/D18-1214.
- Bansal, Y., Nakkiran, P., and Barak, B. (2021). Revisiting model stitching to compare neural representations. *Advances in Neural Information Processing Systems 34* (NeurIPS 2021).
- Conneau, A., Lample, G., Ranzato, M., Denoyer, L., and Jégou, H. (2018). Word translation without parallel data. *ICLR 2018*. arXiv:1710.04087.
- Facco, E., d’Errico, M., Rodriguez, A., and Laio, A. (2017). Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Scientific Reports* 7, 12140. doi:10.1038/s41598-017-11873-y.
- Huh, M., Cheung, B., Wang, T., and Isola, P. (2024). Position: the Platonic representation hypothesis. *Proceedings of ICML 2024*, PMLR 235, 20617–20642.
- Jégou, H., Douze, M., and Schmid, C. (2011). Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(1), 117–128. doi:10.1109/TPAMI.2010.57.
- Kazemi, S. M., et al. (2019). Time2Vec: learning a vector representation of time. arXiv:1907.05321.
- Lenc, K., and Vedaldi, A. (2015). Understanding image representations by measuring their equivariance and equivalence. *Proceedings of CVPR 2015*, 991–999. doi:10.1109/CVPR.2015.7298701.
- Malkov, Yu. A., and Yashunin, D. A. (2018). Efficient and robust approximate nearest neighbor search using Hierarchical Navigable Small World graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. doi:10.1109/TPAMI.2018.2889473.
- Moschella, L., Maiorca, V., Fumero, M., Norelli, A., Locatello, F., and Rodolà, E. (2023). Relative representations enable zero-shot latent space communication. *ICLR 2023*. arXiv:2209.15430.
- Ollivier, Y. (2009). Ricci curvature of Markov chains on metric spaces. *Journal of Functional Analysis* 256(3), 810–864.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. (2024). RoFormer: enhanced transformer with rotary position embedding. *Neurocomputing* 568, 127063. doi:10.1016/j.neucom.2023.127063.
- van den Oord, A., Li, Y., and Vinyals, O. (2018). Representation learning with contrastive predictive coding. arXiv:1807.03748.

- Vaswani, A., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems 30* (NIPS 2017).

Appendix A — Signatures

The proposed family M01–M17 and the existing signatures used (SIG-V, SIG-T) are listed below with the section and the cells behind each. The family was carried to the Foundation’s Lexicon under the prefix SIG-MN; the identifiers stay provisional until that merges.

Id	Operation	Section	Cells
M01	raw profile	§3.2	V9, V21
M02	standardisation (the gauge)	§3.2	V21
M03	hub-corrected similarity	§3.3	V9c, V21
M04	decode	§3.3	V9d, V22
M05	wire quantization	§3.4	V21, V1
M06	content digest	§3.5	V128, CELL-HEAL-1
M07	the reference	§3.5	CELL-REF32, V128
M08	region assignment	§3.6	V28
M09	regional calibration	§3.6	V28, V127
M10	void gate	§3.8	V102
M11	snap (ADDRESS)	§3.10	V74b, V103
M12	MOVE with the origin excluded	§3.10	V73, V105b
M13	the relay regularity	§3.13	V126, V147
M14	the landmark-budget fit	§3.13	V154, V158
M15	intrinsic dimension (TwoNN)	§5.3	V39p
M16	participation ratio	§5.3	V39p
M17	route by operation	§3.12	V116, V117

Appendix B — Standing of each claim

Each claim in this paper has one of three standings: recomputed from the shipped files by the reference library, recorded only (the cell ran, and the shipped library does not yet recompute it), or measured on data that is not public.

Claim	Section	Cells	Standing
Relations survive a change of frame	5.1	V9	Recomputed
Each part of the recipe carries something; bytes cost nothing	3.2, 3.4	V21, V1	Recomputed
Neighbourhoods cross models and languages; exact points do not	3.3	V9d, V18, V9b, V9c	Recomputed

Claim	Section	Cells	Standing
The recipe rebuilds from its specification	4, 5.1	V22	Recomputed
With the full digest, nothing delivered is wrong	3.5	V124, V128	Recomputed
The compact reference resolves 97.2% at scale	3.5	CELL-REF32	Non-public data
The address heals and refuses by field	3.5	CELL-HEAL-1	Non-public data ; the mechanism agrees on public data
Local landmarks help when scarce	3.6	V28, V127	Recomputed
An operator is a file	3.7	V109b	Recomputed
Inside one model, profiles rank as cosine does	3.3	CELL-TRIANGLE-1	Non-public data
A floor separates in-domain from foreign	3.8	V102	Non-public data ; does not hold on the shipped sets
Time: forgeable alone, verifiable in a chain	3.9	V3, V4, V10, V30, V31	Recomputed
The instructions do what their class says	3.10	V73, V75, V76, V77b, V105b	Recomputed
The snap stops composition drift	3.10	V74, V74b, V103	Recomputed
A fused operator survives 8 bits	3.11	V112	Recomputed
Negation needs a gate; route by operation	3.12	V78, V53, V54, V116, V117	Recomputed
The geometry does not compute	1.3, 5.2	V132b	Recorded only
Relaying sits between product and weakest link	3.13	V126, V147	Recomputed
Averaging does not describe the landmark budget	3.13	V130, V154, V158	Recomputed
Agreement belongs to the ordered pair	5.2, 5.4	V125, V149, V143, V146	Recomputed
Agreement is not truth	5.2	V40	Recomputed
Nothing transfers without shared landmarks	5.2	V145	Recorded only
The thinness of meaning is the texts'	5.3	V39p	Recomputed
The founding equation fails	5.5	V80–V84	Recorded only

Appendix C — Cells cited

Recomputed from the shipped embeddings: V1, V3, V4, V9, V9b, V9c, V9d, V10, V18, V21, V22, V28, V30, V31, V39p, V40, V53, V54, V73, V74, V74b, V75, V76, V77b, V78, V103, V105b, V109b, V112, V116, V117, V124, V125, V126, V127, V128, V130, V143, V146, V147, V149, V154, V158.

Recorded only: V132b, V145, V80, V81, V82, V83, V84.

Measured on data that is not public: V102, CELL-HEAL-1, CELL-REF32, CELL-TRIANGLE-1.

Appendix D — Notation

E, a frame; d, its dimension; A, the landmark set and its size; p_x , the raw profile of text x; z_x , its standardisation; sim and csls, base and hub-corrected similarity; q, the quantized profile; S, the finite set of stored meanings; snap, the return to the nearest element of S; T(t), the time coordinate; θ , a phase angle; λ and $t_{1/2}$, a decay rate and its half-life; k, the shortlist size. Verdicts are PROVEN, DISPROVEN and INVALID.