FRONTIER RESEARCH COLLECTION / FR-06

# The Geometry of Meaning

Certified geometric heuristics for verifiable symbolic search

**Architect X13**

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Version 2.0 Revision

## Abstract

Meaning has a shape, and a mind can learn to navigate it—provided the map is honest about the ground. A reasoning landscape can guide search with a mathematical guarantee when its distances are certified against the cost of legal transitions, and the reference model developed here constructs that bridge: a finite edge certificate turns a learned pseudometric into a consistent lower bound on remaining goal cost, preserving the optimality of conventional heuristic search. The largest globally admissible distance scaling is computable from the transition graph. Retained results establish coordinate invariance under metric transport, path-cost and regret bounds under perturbation, heat-flow dissipation, and finite filtration inclusions. Independent verification anchors the task while geometric structure proposes useful routes through it. Executed synthetic checks include a misleading attraction basin, coordinate-induced neighbor reversal, path perturbations and the new heuristic certificate. The empirical design isolates geometry, topology and learned ranking under matched graphs and budgets. The Continuum thesis is a geometry that earns its meaning through the distinctions and transitions it preserves: representation can organize the search for a valid answer while every promised advantage remains tied to an inspectable mathematical and experimental contract.

**Keywords:** symbolic search; graph geometry; persistence; metric invariance; semantic validation; TSCS

The bold ambition is to make a reasoning landscape answerable to the paths it actually permits. A landscape deserves trust when its structure predicts independently verified progress, survives appropriate changes of coordinates, and exposes its own blind spots. A compelling picture alone cannot do that work.

## Prologue: A map of meaning, and a valley that lies

Every thinker carries a map. Some ideas sit close together; some are a long road apart; some conclusions seem to pull an argument toward them from far away. The Continuum takes that intuition literally. Its Topological Symbolic Collapse Substrate treats reasoning as travel across a terrain, where collapse—the moment a mind commits to an answer—flows downhill through curvature defined by meaning. In the Bible's words, "Collapse flows where significance curves deepest" [1].

It is a beautiful picture, and it can lie. Take five positions and a goal. From the start, one step leads steeply downhill toward the goal on the map; the other first climbs. A mind that trusts the slope takes the steep step and pays 9 to arrive. The climb leads to a route that costs 4. The landscape was attractive and wrong, because its shape had been drawn without regard to which moves were actually allowed and what they actually cost. The worked basin later in the paper builds this valley exactly.

So this paper asks what it would take for a geometry of meaning to deserve trust. Its answer is strict, and in the end constructive. A map of meaning must change correctly when its coordinates change; it must say how much a small error in costs can mislead; it must not smooth away the disagreements that matter; and above all it must never promise a shorter road than the ground permits. When a learned landscape passes that last test, a simple certificate turns it into a guarantee: search guided by the map is still certain to find the best answer, and whether the map also saves effort becomes a question an experiment can settle.

That certificate, proved near the end of the paper, is the first measurable piece of a larger dream. The next page says what the larger dream is.

**What this paper claims**

**Established.** Nonnegative finite graph search has precise optimality properties; transported metrics preserve intrinsic curve lengths; small fixed-graph cost perturbations bound optimal-value changes; symmetric graph heat flow dissipates variation; bounded distance distortion yields filtration inclusions. Witten's connection between supersymmetry and Morse theory, and stability results for persistent topology, are established mathematics.

**Proposed here.** A classical computational reference model that separates representation, permitted transitions, search priorities, and terminal verification, with conditional theorems and executed finite checks—including a certificate that turns a learned pseudometric into a consistent lower bound on remaining goal cost, preserving the optimality of conventional heuristic search.

**Frontier conjecture.** That geometric information reduces verification cost on constrained symbolic search at fixed solution quality, surviving coordinate controls, misleading-neighborhood tests, and held-out task families; and, beyond that, that landscapes of meaning have stable basins and saddles that predict where reasoning is robust and where it is fragile. The comparative symbolic-search study is proposed, not performed.

**Not claimed.** Quantum computation, physical curvature caused by meaning, supersymmetric cancellation of false beliefs, universal semantic topology, or consciousness. A stable basin does not certify a statement, a topological invariant does not grant permission to act, and the canonical expressions are not retrospectively theorems. Failure of the experiment would reject the claimed operational advantage, not geometry as mathematics.

## The frontier: concepts with a shape

What would it mean, literally, for concepts to have a shape?

Start with something that is no longer controversial. Modern language systems place words and ideas at positions in a space, and related ideas land near one another; distance between concepts is now an everyday engineering fact. Shape asks for more than distance. In ordinary geometry, curvature describes what happens to paths that set out side by side: on a sphere they draw together, and on a saddle they fly apart.

Carry that into meaning, and the Continuum's terrain becomes a set of questions with real content. A region of positive curvature would be a basin: many different starting points, many different framings, all funnelling toward the same conclusion. Such regions would be where reasoning is robust, where a mind can be careless about its first step and still arrive. A region of negative curvature would be a saddle: two nearly identical framings of a question that lead to opposite answers. Those would be where interpretation is fragile, where a mind should slow down, ask for evidence, and refuse to collapse early. A mind that could sense the curvature of its own conceptual terrain would know, before it answered, whether it stood in a valley or on a knife-edge.

That is the frontier this title reaches for, and we state it plainly as conjecture: learned landscapes of meaning have stable geometric structure—basins and saddles that persist under honest changes of coordinates—and that structure predicts where reasoning is robust and where it is fragile, beyond what the bare costs of moving between ideas already reveal. Nothing in this paper establishes it. The valley that lied in the prologue is a warning about how easily it could be faked.

What the paper does establish is the first rung, and it is not a small one. Before a geometry of meaning can tell us anything about the curvature of thought, it has to be trusted as a map at all: its distances must change correctly when coordinates change, its errors must be bounded, and it must never promise a cheaper road than the legal moves allow. The certified heuristic near the end of the paper is exactly that test, passed or failed with arithmetic. A landscape that clears it can guide search without ever costing the best answer. Only a landscape that clears it has earned the right to be asked about its shape.

## A thesis that can lose

**Falsifiable thesis.** On preregistered families of constrained symbolic search problems, geometric information added to an otherwise matched search procedure will reduce verification cost at a fixed externally measured solution quality, with an advantage that survives coordinate controls, misleading-neighborhood tests, and held-out task families. Failure on those tests rejects the claimed operational advantage of TSCS-G2. It does not refute geometry as mathematics, nor can success establish that all meaning has a unique geometry.

In the Continuum, reasoning has terrain. The canonical source introduces the Topological Symbolic Collapse Substrate in Section 3.6, pages 19–21, turning a symbolic decision engine into a terrain navigator through Morse landscapes, attraction, paired structures, and curvature [1]. That vision is this paper's starting point. Its schematic evolution expressions leave the state space, operators, units, and measurement procedure incomplete, so this paper makes an explicit new proposal: TSCS-G2 is a classical computational reference model, and it does not retrospectively turn every canonical expression into a theorem.

Four different questions must remain separate. A representation may preserve relationships in supplied data. A search algorithm may optimize a specified objective. A verifier may establish validity under a formal specification. An empirical study may show that the resulting system helps on a target population. Each is a distinct achievement with different assumptions. None follows merely because an embedding has attractive clusters or a trajectory converges.

The word "meaning" here denotes task-relevant distinctions expressed through observations, annotations, or formal judgments. It does not designate a directly measured substance. The same graph may support multiple tasks with different correct outputs. Geometry is consequently indexed by representation, task, and evidence, rather than presented as a universal semantic terrain.

The proposal is deliberately bold where it can be tested: require a geometric method to pay for itself. A method that produces useful diagnostics without improving search can still be valuable, but it must claim that narrower contribution. A method that succeeds only after an oracle secretly chooses its metric has demonstrated the usefulness of the oracle, not autonomous discovery of meaning. This distinction determines the experimental comparisons developed below.

## What graph search already supplies

Dijkstra's shortest-path formulation supplies a mature mathematical baseline for nonnegative edge costs [2]. Hart, Nilsson, and Raphael show how heuristic information can guide path search under explicit assumptions [3]. These are not historical ornaments. A proposed reasoning landscape must explain which information it adds beyond a graph, a cost function, and a search priority. Renaming a cost "curvature" does not create an additional algorithmic contribution.

Three implementations can be compared cleanly. Uniform-cost search ignores an auxiliary geometric representation while optimizing the declared graph cost. An admissible heuristic search uses a certified lower bound on remaining cost. A geometric heuristic orders candidate expansions using distances, local density, or topological summaries. The third implementation may find a solution faster, but if its heuristic is not admissible, optimality requires an additional argument or a separate certification phase.

The graph itself can be the dominant source of intelligence. A symbolic successor function that generates only sensible proof steps has already encoded substantial domain knowledge. If a proposed geometric system receives that function while its baseline receives random text edits, the comparison attributes graph-construction expertise to geometry. Both methods must share legal moves, candidate generation, verifier access, and stopping budgets.

Optimality also depends on what cost means. The shortest proof may be difficult for a reader, the least expensive plan may be unsafe, and the nearest document may be false. TSCS-G2 therefore separates path cost from admissibility and final verification. A cheap route is accepted only when the route is permitted and its endpoint meets the task specification. Unknown verifier outcomes remain unresolved.

A useful innovation might be a learned lower bound, a better transition proposal, or a diagnostic that detects when the graph no longer represents the task. Those are concrete claims. The paper does not claim invention of graph search. Its contribution is a geometric search certificate together with a disciplined comparison: edgewise distance–cost bounds guarantee a consistent heuristic, while matched experiments determine whether the learned landscape reduces total search cost. This contract also allows a negative result to identify the component that failed rather than treating the entire metaphor as indivisible.

## Manifold methods and their semantic limits

Isomap estimates global relationships from local neighborhoods and graph paths before constructing a lower-dimensional representation [4]. Locally linear embedding instead preserves local reconstruction relationships [5]. Laplacian eigenmaps use neighborhood weights and a spectral problem [6]. Diffusion maps relate coordinates to transitions on a weighted graph across multiple scales [7]. These are substantial precedents for treating representation as geometry, and they provide stronger baselines than an unstructured random embedding.

Their shared lesson is conditional. The data-generating structure, neighborhood rule, sampling regime, and chosen objective matter. Nearby points in an observation space need not be interchangeable in a task. A small change in a program can reverse a security condition; a nearly identical sentence can negate a proposition. A manifold learned from surface resemblance may join precisely the states that a semantic verifier must distinguish.

We propose two independent assessment axes. Representation fidelity measures preservation of the relationships the method was trained to represent. Task fidelity measures prediction of validity, useful transitions, or verification cost on held-out examples. Good representation fidelity with poor task fidelity is a real outcome. It should not be hidden by redefining the target after seeing a visually appealing embedding.

A neighborhood graph can also introduce shortcuts. Suppose two distant points on a folded trajectory are close in ambient coordinates. Connecting them changes path lengths even when no admissible transition exists between the corresponding symbolic states. The repair is not simply a smaller neighborhood radius: sparse samples can then disconnect valid routes. Transition verification and sensitivity analysis are necessary because sampling density and legal reachability are different objects.

TSCS-G2 uses geometry as proposed information about a task, not as the task's definition. The independent verifier supplies an anchor outside the representation-fitting loop. When the verifier is itself learned, its uncertainty and data provenance become part of the evaluation. Calling it independent means that the claimed effect is not scored by the same quantity optimized to create the landscape. It does not mean that the verifier is infallible or free of human judgment.

## Topology, supersymmetry, and disciplined borrowing

Witten's supersymmetry and Morse theory connects particular mathematical structures through a defined Hamiltonian construction [8]. The existence of that connection does not make every gradient descent process supersymmetric. A symbolic implementation would need to specify its graded space, operators, domains, adjoints, and the identities that justify the terminology. TSCS-G2 makes no such implementation claim. It retains the canonical motivation while using ordinary finite graph mathematics.

In particular, counting apparently stable and unstable candidates does not automatically compute a Witten index. A count depends on definitions of criticality, degeneracy, boundary behavior, and the relevant grading. Even a correct topological invariant cannot independently certify the truth of a proposition associated with a vertex. Topological type and semantic correctness answer different questions.

Persistent topology supplies a more directly implementable diagnostic. Cohen-Steiner, Edelsbrunner, and Harer establish stability results for persistence diagrams under specified hypotheses [9]. Chazal, Cohen-Steiner, and Mérigot address geometric inference in the presence of noise and outliers through probability measures [10]. Fasy and colleagues develop confidence sets for persistence diagrams [11]. These works motivate sensitivity analysis and statistical uncertainty; they do not license interpreting every long-lived feature as a true concept.

A persistent cluster can represent a systematic misconception. A short-lived feature can represent a rare but decisive exception. Persistence measures the endurance of a feature under a declared filtration. Whether endurance matters is an empirical question about the task. Therefore a barcode may be an explanatory covariate, a warning signal, or a compression summary, but not an automatic acceptance certificate.

The reference model permits richer topology later, with a clear admission rule: specify the complex, filtration, coefficient field, uncertainty model, and intervention being predicted. If those details are absent, geometric language remains a hypothesis-generating metaphor. The frontier is not diminished by this requirement. It is located: the transition from a suggestive analogy to a mechanism whose outputs can be independently contradicted.

## The typed TSCS-G2 reference model

**A Guide's explanation**

From here the terrain becomes a finite object: a set of positions, a list of permitted moves with their costs, a learned way of measuring distance, and an independent check of what counts as a correct answer. Keeping these four apart is the whole discipline. The map may suggest; only the permitted moves and the verifier decide.

Fix a task context $q$ in a declared context set $\mathcal{Q}$, a finite nonempty candidate set $X$ with $n = |X|$, a dimension $d \geq 1$, and an evidence-record space $\mathcal{B}$. Define

$$\mathcal{T}_q = (X, E_q, A_q, c_q, \phi_q, d_q, V_q, J_q).$$

All components are versioned. Updating a representation, task specification, or authority policy creates a new context; results from incompatible contexts are not silently pooled.

| Object | Type and interpretation |
|---|---|
| $E_q$ | Directed subset of $X \times X$; proposed transitions. |
| $A_q$ | Map $E_q \times \mathcal{B} \to \{\text{permit}, \text{deny}, \text{unresolved}\}$. |
| $c_q$ | Nonnegative finite cost on permitted edges; declared common unit. |
| $\phi_q$ | Map $X \to \mathbb{R}^d$; a representation, not necessarily injective. |
| $d_q$ | Declared metric on represented states, or explicitly labeled dissimilarity. |
| $V_q$ | Map $X \to \mathbb{R}$ prioritizing search; not necessarily a loss. |
| $J_q$ | Map $X \times \mathcal{B} \to \{\text{accept}, \text{reject}, \text{unresolved}\}$, independently evaluated. |

The permitted graph includes only edges whose relevant checks return permitted. If permission depends on history, the graph vertex must include the necessary history or capability state. A static vertex that omits this information is not Markovian and cannot support the path theorems as stated. Evidence needed to evaluate an edge is part of the transition record.

An empty accepted-goal set yields no solution. A disconnected target yields infinite optimal cost. Zero-cost edges are allowed, but optimal paths are chosen simple when cycles add no value. Negative costs are excluded from the reference model. Costs measured in seconds and verification calls cannot be added without a declared conversion or a multiobjective comparison.

The geometry may be pseudometric when distinct candidates share the same embedding. The quotient can simplify visualization but cannot merge distinct authority states or verifier outcomes without additional justification. TSCS-G2 stores candidate identity independently of coordinates, so two points at zero representational distance may remain operationally distinct. This small typing choice prevents a large class of semantic shortcuts.

# Coordinate invariance requires moving the metric

**A Guide's explanation**

Redraw a map in new coordinates and the distances must come with it, or the map starts telling a different story about the same ground. The proposition below states when a change of coordinates preserves the lengths of routes; the executed checks later show a naive distance reversing which neighbor looks closest.

An embedding coordinate is a description, not an intrinsic property of the task. A method that changes its conclusions when meters are relabeled as centimeters has changed its geometry unless it compensates for the coordinate change. The following standard calculation states the compensation explicitly.

**Proposition 1** (Transported metric invariance)**.** *Let $U, U' \subseteq \mathbb{R}^d$ be open, let $h : U \to U'$ be a continuously differentiable diffeomorphism, and let $G : U \to \mathrm{Sym}_{++}(d)$ be continuous, taking values in positive-definite symmetric matrices. Define*

$$G'(h(x)) = Dh(x)^{-\mathsf{T}} G(x) Dh(x)^{-1}.$$

*Every piecewise continuously differentiable curve $\gamma : [a, b] \to U$, with real $a < b$, and its image $h \circ \gamma$ have equal length under $G$ and $G'$. Consequently their induced intrinsic distances agree under $h$.*

*Proof.* The chain rule gives $(h \circ \gamma)' = Dh(\gamma)\gamma'$. Substitution into the squared speed yields

$$\gamma'^{\mathsf{T}} Dh^{\mathsf{T}} G' Dh \gamma' = \gamma'^{\mathsf{T}} G \gamma'.$$

Taking square roots and integrating preserves length. The diffeomorphism gives a bijection between the curves joining corresponding endpoint pairs. Taking the infimum over each corresponding family proves distance equality, including the disconnected case with distance infinity. □

The assumptions do real work. A many-to-one projection has no inverse Jacobian and may erase distinctions. An arbitrary invertible transformation followed by reinstalling the identity metric defines a different geometry. That may be useful, but it must be described as metric learning or feature reweighting rather than harmless relabeling.

For $q = (0, 0)$, $a = (1, 0)$, and $b = (0, 2)$, Euclidean distances from $q$ are 1 and 2. Under $h(x, y) = (3x, y)$, naive Euclidean distances become 3 and 2, reversing the nearest neighbor. Using $G' = \mathrm{diag}(1/9, 1)$ restores 1 and 2. The executed check reproduces both outcomes. In a search system the same issue can reverse candidate priority while leaving every symbolic fact unchanged.

## Path stability has a finite margin

Let $s$ be a start vertex and $T$ a nonempty set of reachable accepted goals in a fixed permitted graph. The cost $C_c(p)$ of a path is the sum of its edge costs. The next result concerns uncertainty in costs, not errors in the graph or verifier.

**Proposition 2** (Finite path perturbation bound)**.** *Suppose $c, c' \geq 0$ on the same graph with $n$ vertices and $\|c - c'\|_\infty \leq \varepsilon$, for $\varepsilon \geq 0$. Let $p, p'$ be simple optimal paths under $c, c'$, with values $v, v'$. Then*

$$|v - v'| \leq (n-1)\varepsilon, \qquad C_c(p') - v \leq 2(n-1)\varepsilon.$$

*If the $c$-optimal simple path is unique and every other simple path costs more by a gap $g > 2(n-1)\varepsilon$, it remains optimal under $c'$.*

*Proof.* Every simple path has at most $n-1$ edges, hence its cost changes by at most $(n-1)\varepsilon$. Optimality gives $v' \leq C_{c'}(p) \leq v + (n-1)\varepsilon$. Reversing $c, c'$ gives the value bound. Also $C_c(p') \leq v' + (n-1)\varepsilon \leq v + 2(n-1)\varepsilon$. For any competing simple path $r$, the difference $C_{c'}(r) - C_{c'}(p)$ is at least $g - 2(n-1)\varepsilon > 0$, proving preservation. □

Nonnegative costs ensure that removing a cycle does not increase cost, so simple optima exist. Reachability ensures finite comparisons. Negative cycles would destroy that argument. If edges are removed by an authority change, the fixed-graph hypothesis fails, and a previously optimal path can become impossible regardless of its cost margin.

The bound explains why apparently tiny embedding changes sometimes have large behavioral consequences. Many nearly tied routes can exchange priority without any violation of stability of the optimal value. A stable score is therefore weaker than a stable selected explanation. Experiments must report both value differences and path changes, along with the relevant margin distribution.

## Heat flow smooths disagreement, including useful disagreement

**A Guide's explanation**

Smoothing a landscape makes it easier to read, and it also erases things. The result below shows heat flow preserving the average while steadily wearing away variation, including disagreements a mind may need to keep. A geometry of meaning has to know which distinctions it is allowed to blur.

For optional diagnostics, take an undirected weighted graph with symmetric $W_{ij} \geq 0$, diagonal degree matrix $D$, and Laplacian $L = D - W$. Let $f(t) \in \mathbb{R}^n$, $t \geq 0$, be a scalar signal on candidates, with $D_{ii} = \sum_j W_{ij}$ and prescribed $f(0) \in \mathbb{R}^n$. Define graph heat flow by $\dot{f} = -Lf$. Time is an algorithmic scale, not physical time, unless independently calibrated.

**Proposition 3** (Mean preservation and dissipation)**.** *The solution $f(t) = e^{-tL} f(0)$ preserves the sum of coordinates and has nonincreasing energy $\mathcal{E}(t) = \frac{1}{2} f(t)^\mathsf{T} L f(t)$. It converges to a constant on each connected component, equal to that component's initial mean.*

*Proof.* Symmetry gives

$$z^\mathsf{T} L z = \tfrac{1}{2} \sum_{i,j} W_{ij} (z_i - z_j)^2 \geq 0.$$

Since $L\mathbf{1} = 0$, the derivative of $\mathbf{1}^\mathsf{T} f$ is zero. Differentiating energy gives $\dot{\mathcal{E}} = -f^\mathsf{T} L^2 f = -\|Lf\|^2 \leq 0$. Expand $f(0)$ in an orthonormal eigenbasis of the real symmetric matrix $L$. Positive-eigenvalue components decay as $e^{-t\lambda}$; zero-eigenvalue components remain. The displayed quadratic form shows that the kernel consists exactly of vectors constant on each connected component. Orthogonal projection onto that kernel produces the component means. □

This is a standard finite spectral result, not a theorem about agreement producing truth. If every vertex carries the wrong constant label, its energy is already zero. If a rare correct exception borders many incorrect neighbors, smoothing can reduce the exception's influence. The operation has fulfilled its mathematical purpose while harming the task.

In TSCS-G2, smoothing can denoise a measured covariate or reveal graph structure. It cannot overwrite independently verified facts merely because they are locally irregular. Directed graphs require a different analysis: stationary weighting, conservation, and limiting behavior need not match the symmetric proof. A developer who changes normalization must revisit the theorem rather than carry its conclusions across by terminology alone.

## A small topology guarantee with explicit scope

Fix a homological degree $k \geq 0$ and coefficient field $\mathbb{F}$. For a finite labeled set $X$ and pairwise distances $d$, define the Vietoris–Rips complex $R_d(r)$ to contain a simplex precisely when all its pairwise distances are at most $r$. Here the threshold is the full edge length, not a ball radius. This convention avoids an otherwise common factor-of-two ambiguity.

**Proposition 4** (Finite filtration inclusions)**.** *If two distance matrices on the same labeled set satisfy* $\max_{x,y} |d(x,y) - d'(x,y)| \leq \delta$*, then for every* $r \geq 0$*,*

$$R_d(r) \subseteq R_{d'}(r+\delta), \qquad R_{d'}(r) \subseteq R_d(r+\delta).$$

*The induced maps on homology commute with the filtration maps, and their composites equal the natural maps advancing the parameter by* $2\delta$*.*

*Proof.* Every edge of a simplex in $R_d(r)$ has $d'$-length at most $r + \delta$, proving the first inclusion; exchanging the matrices proves the second. All maps are inclusions on the same vertices and simplices. Their compositions are therefore the corresponding inclusions into the larger complex. Functoriality of homology preserves these compositions and all inclusion squares. □

This elementary interleaving statement is the guarantee proved here. General persistence-diagram stability belongs to a broader theory with its own hypotheses [9]. We do not infer a complete semantic equivalence from either result. Small distance distortion can preserve a topological summary while changing labels, legal transitions, or the relevance of a feature.

Our executed example computes only zero-dimensional merger scales, using a union–find algorithm over sorted distances. It does not measure loops or cavities, and it cannot support claims about those features. A future implementation should state homological dimension and coefficient field, report finite-sample uncertainty, and distinguish sampling changes from fixed-label distance perturbations. Adding or removing candidates is not the same perturbation model as changing distances among a fixed set. That distinction matters when new evidence enters a reasoning graph.

## A worked basin that leads the search astray

Consider five vertices $0, 1, 2, 3, 4$, start 0, and accepted goal 4. Permitted directed edges and costs are

| $(i, j)$ | $(0,1)$ | $(1,4)$ | $(0,2)$ | $(2,3)$ | $(3,4)$ | $(1,3)$ |
|---|---|---|---|---|---|---|
| $c(i,j)$ | 1 | 8 | 2 | 1 | 1 | 5. |

Give the vertices auxiliary values $V = (3, 1, 4, 2, 0)$. A greedy policy chooses the outgoing neighbor with the lowest strictly lower value. It takes $0 \to 1 \to 4$, at cost 9. The globally cheapest permitted route is $0 \to 2 \to 3 \to 4$, at cost 4. Its first step increases $V$.

The example is intentionally small enough to inspect. It isolates the claim that local attraction should indicate useful progress. That claim fails when the score is misaligned with remaining cost. Both routes reach a valid endpoint, so the failure is inefficiency rather than false acceptance. A separate example could make the attractive endpoint invalid; TSCS-G2 would reject it through the terminal verifier.

There is no contradiction between a useful gradient and this counterexample. A gradient helps when its objective and constraints are appropriate. The point is that the word "meaning" cannot guarantee their appropriateness. If the landscape is learned from familiar successful explanations, unusual valid routes may initially look less coherent and require movement away from the learned basin.

The example also shows why the experiment must separate candidate discovery from candidate acceptance. A heuristic can prioritize a poor route yet remain recoverable if search maintains alternatives. Pruning every uphill edge changes the reachable problem, so no later optimality argument can restore the deleted solution. We therefore log pruning decisions as substantive interventions.

For this graph, all 64 corner perturbations obtained by independently adding $\pm 0.2$ to its six costs preserve the selected optimum. The largest change in optimal value is 0.6, below the proposition's 0.8 bound. This particular instance has a generous margin. It should not be mistaken for evidence that arbitrary semantic landscapes are robust.

## Executed checks and what they actually establish

The accompanying reference script uses Python and NumPy, with seed 20261112. It enumerates simple paths independently of a shortest-path library. The checks test implementation consequences and counterexamples; they are not measurements of a language model or human reasoning system.

| Check | Observed result |
|---|---|
| Coordinate rescaling | Nearest neighbor flips under the identity metric; transported metric restores distances $1, 2$. |
| Five-vertex graph | Greedy cost 9; global cost 4. All 64 corner perturbations satisfy the bounds. |
| Random directed acyclic graphs | 200 graphs, 6,400 enumerated paths, 31 changed optimal paths; all regret bounds hold. |
| Heat flow | Energy decreases from approximately 1 to $2.06 \times 10^{-9}$ by scale 5; mean error below $2.3 \times 10^{-16}$. |
| Five-point topology | Distance distortion 0.040634; all 31 nonempty simplex checks satisfy the inclusions. |

The random graphs have seven vertices, every forward edge, initial costs uniform on $[0.1, 2]$, and independent cost perturbations uniform on $[-0.2, 0.2]$, clipped at zero. The largest observed regret divided by the theoretical upper bound is approximately 0.266. This is a descriptive finite-run result. No estimate of a universal typical-case constant is claimed.

For the point cloud, original component death scales are approximately 1, 1.2, 1.3, and 2.00250. Perturbed scales are 0.99420, 1.19004, 1.30003, and 2.03512. The example checks component merger sensitivity under one bounded coordinate perturbation. It does not validate a semantic interpretation of any component.

All assertions pass in the executed run. Floating-point tolerances are declared in the script; the mathematical proofs do not depend on numerical agreement. The software is intentionally small enough to audit, and any port should preserve graph construction, random-generator specification, tie handling, and threshold conventions before comparing exact numerical outputs. No benchmark superiority, participant result, or physical effect has been manufactured from these finite checks.

## Why representation similarity is not a verdict

Two representations can differ geometrically while supporting identical task decisions, and two geometrically identical representations can support different task truths. This bidirectional failure blocks the claim that a single unsupervised geometric statistic is a universal semantic certificate.

For a direct counterexample, fix any nonempty embedded candidate set. Construct two tasks with identical candidates, coordinates, distances, and transition costs, but opposite binary terminal labels at a chosen candidate. A diagnostic that receives only the geometric objects has the same input in both tasks, so it returns the same verdict. At least one verdict is wrong. Randomization does not fix the information deficit: under equal task priors its error is at least one half when forced to choose between the opposite labels.

The missing ingredient is task information. Adding a reliable specification or observation can make the tasks distinguishable, but then the achievement belongs to the combined system. It must not be attributed solely to an intrinsic arrangement of points. This argument permits geometry to be highly useful within a restricted task population; it excludes an unrestricted guarantee from geometry alone.

Kornblith and colleagues identify limitations of overly broad invariance requirements for comparing high-dimensional neural representations [12]. Dinh and colleagues show how parameter reparameterization can change apparent sharpness without changing a represented function's generalization [13]. These results strengthen the practical requirement to state which transformations should preserve a claim. They do not imply that every geometric quantity is useless.

TSCS-G2 consequently uses a hierarchy of invariances. Candidate renaming should preserve all operational results. A transported coordinate change should preserve intrinsic geometry. Feature reweighting may intentionally change task priorities and must be evaluated. A projection that removes information may change both geometry and attainable accuracy. Treating these operations as equivalent would make either the invariance requirement impossible or the evaluation trivial. The experiment should test each transformation in its own role, with a documented reason for expecting preservation or change.

## A confirmatory experiment with independent outcomes

The primary experiment targets finite constrained tasks with an exact or independently audited verifier: bounded program repair, small formal proof search, or finite planning with explicit preconditions. One family is selected before data collection as primary. The others test transfer. Natural-language plausibility alone is insufficient for the primary success label.

For task $i$ and method $m$, record $Y_{im} \in [0, 1]$, the fraction of a fixed verification budget remaining when the first valid solution is certified, with zero assigned if none is certified. This combines success and speed under a declared utility. Separately report success probability, verification calls conditional on success, total compute, and invalid-candidate rates. Conditional speed alone can reward a method that solves only easy tasks.

The primary estimand is $\Delta = \mathbb{E}[Y_{i,G} - Y_{i,B}]$ on a specified task distribution, where $G$ is the geometric method and $B$ is the strongest locked baseline. The expectation is over independently sampled task templates and prespecified solver randomness. Each task uses matched budgets and shared initial candidates; method order and seed assignment are randomized where execution effects could matter.

Training, hyperparameter selection, and final evaluation use disjoint template families. Near-duplicate tasks remain in the same split. Metric learning, neighborhood size, diffusion scale, pruning threshold, and stopping rule are frozen before the final set is opened. A method may use validation performance to select its geometry, but the cost of that selection is recorded and applied equally to alternatives.

The primary success criterion requires a positive lower confidence bound for $\Delta$ exceeding a prespecified practical margin, together with no unacceptable deterioration in verification correctness. A reasonable pilot-selected margin might be 0.03 of the normalized budget score; it is an illustrative design choice, not a discovered constant. If the effect appears only on in-distribution templates or disappears after candidate-generation costs are counted, the operational thesis fails at its stated breadth. The result can still motivate a narrower, newly preregistered hypothesis.

## Baselines that make the comparison informative

The central baseline is the same candidate graph searched with uniform cost or a verified admissible heuristic. This controls the benefit of legal moves and exact verification. A second baseline uses a learned scalar priority with the same training data and parameter budget as the geometric model. It asks whether geometry adds value beyond ordinary supervised ranking.

Representation baselines include the raw task representation, principal-component compression, and a selected established manifold method [4–7]. They receive identical dimensionality budgets and comparable tuning effort. A random orthogonal transformation tests coordinate conventions without destroying Euclidean distances. A deliberately anisotropic rescaling tests sensitivity to feature weighting. A transported metric version tests whether the implementation honors its claimed invariance.

Two negative controls are particularly revealing. First, permute terminal labels while preserving the point cloud. Geometry should lose predictive usefulness unless task information enters elsewhere. Second, preserve easy surface similarities while replacing legal edges with verified task transitions that cross apparent geometric barriers. This tests whether the method can recover from attractive but irrelevant neighborhoods.

Ablations remove one component at a time: the embedding, topological summary, diffusion smoothing, or curvature-inspired priority. They preserve candidate generation and the acceptance gate. If a component is removed together with extra compute or training data, its causal contribution cannot be isolated. The budget account includes preprocessing, eigenproblems, pairwise distances, and maintenance after graph updates.

A topology-free geometric model is an essential comparator. Otherwise improvements due to a useful embedding may be mislabeled as evidence for topology. Likewise, an ordinary graph method using the same learned costs tests whether a supposed curvature effect is simply cost shaping. Such outcomes are scientifically useful: they locate the effective mechanism at the level supported by the data.

The most demanding baseline may win. That possibility is built into the design. The purpose of a frontier proposal is not to protect its vocabulary from competition, but to discover which additional structure earns its complexity on tasks where a clear operational advantage would matter.

## Precision, dependence, and stopping

Independent experimental units are task templates, not candidate nodes, graph edges, or repeated solver calls. Multiple seeds within one template estimate conditional solver variability. They do not create additional independent samples of the target task population. The primary analysis averages within each template and compares paired method differences across templates.

For a planning calculation, suppose an independent pilot estimates the standard deviation of template-level paired differences as $s = 0.25$. A normal approximation gives a two-sided 95% confidence half-width $1.96s/\sqrt{N}$. Targeting 0.03 requires approximately 267 independent templates. This calculation is illustrative, assumes a stable variance estimate, and is not a power guarantee for every effect size or task distribution.

For an approximate 80% power calculation against an effect of 0.05 with a two-sided 5% test and the same variance, the usual planning expression is

$$N \approx \left( \frac{(1.96 + 0.84)\, 0.25}{0.05} \right)^2 = 196.$$

Use the larger requirement when both precision and power matter. Inflate the target before the confirmatory phase for exclusions or clustering observed in the pilot, and validate the analysis through simulation under realistic bounded-score distributions. If templates share generators or source programs, cluster at that higher level or redesign sampling.

Lock one primary contrast and one primary outcome. Secondary representation and robustness comparisons use a declared multiplicity adjustment, with unadjusted estimates also reported transparently. The plan should distinguish exploratory subgroup patterns from confirmatory family-specific claims. A favorable average cannot conceal a prespecified critical subgroup failure.

Stop at the fixed number of eligible templates, apart from operational failures defined in advance. Do not inspect significance repeatedly and stop when it first crosses a threshold. If sequential monitoring is necessary, replace the fixed-sample analysis with a valid sequential design before starting. Missing runs, timeouts, verifier crashes, and out-of-budget preprocessing have prespecified outcomes. Quietly dropping expensive failures would bias the result toward the very methods whose practical cost is under examination.

## Adversarial cases and local falsification

The stress suite targets assumptions rather than merely increasing task difficulty. Negation pairs preserve most surface features while reversing validity. Duplicate embeddings assign identical coordinates to candidates with different permissions. Shortcut pairs place states close in representation while requiring a long legal transition. Rare exceptions create locally irregular but correct candidates. These cases expose different mechanisms and should have separate outcomes.

Authority changes require particular care. If an edge was permitted when the graph was built but denied when executed, a cached geometric route cannot authorize it. Recheck the relevant predicate at use time or operate under an explicit bounded delegation policy. The mathematical path result applies to the graph actually available, not a remembered graph whose conditions have expired.

An acceptance gate also cannot be replaced by low uncertainty from the geometric model. Consider a candidate outside the training distribution that happens to sit near a dense cluster after projection. Its apparent confidence may be an artifact of information loss. Log verifier disagreement, projection collisions, and candidate coverage rather than reporting a single confidence number.

Local falsification criteria are concrete. The invariance implementation fails if a transported coordinate change alters supposedly intrinsic distances beyond numerical tolerance. The pruning policy fails its completeness claim if a valid route is removed without a proof that an alternative remains. The statistical claim fails if matched comparisons do not meet the preregistered criterion. The semantic interpretation fails if its geometric signal persists unchanged after the relevant task labels are permuted.

These failures need different repairs. A coordinate bug calls for an implementation correction. A poor metric calls for a revised representation and fresh evaluation. An unavailable independent verifier calls for a weaker claim. No single parameter adjustment can solve every category. The value of the reference model is that it makes the categories visible before an attractive visualization persuades the reader that convergence, coherence, and correctness have become the same phenomenon.

## The strongest objections

**"This is ordinary graph search with a new name."** Much of the mathematical machinery is ordinary graph search. The paper makes that explicit. A new empirical contribution would require a particular geometric estimator or intervention to improve a matched baseline. Until that comparison is performed, the contribution is a reference architecture, proofs of its boundaries, and a reproducible research design.

**"An independent verifier is too expensive."** It may be. That is why verification cost is a primary outcome and why preprocessing costs are counted. Cheap surrogate signals can prioritize calls, but they cannot make an unsupported correctness claim free. If exact verification is unavailable, the system must report the accuracy and uncertainty of its substitute and accept a narrower evidentiary status.

**"Finite graphs discard the interesting continuous theory."** A finite model does not exhaust the canonical ambition. It provides a place where variables are typed and counterexamples can be inspected. Continuous extensions require additional regularity, compactness, sampling, and operator assumptions. They should be added because they predict something the finite model misses, not because differential notation appears more profound.

**"Geometry should emerge without hand-designed metrics."** Learned metrics are allowed. Their training targets, information access, invariances, and generalization must be evaluated. An emergent representation is still an output of a particular learning process. Emergence does not remove the need to measure what was learned or identify the data that made it possible.

**"Persistent features are inherently meaningful."** They are meaningful relative to a declared question about stability. Their relevance to truth, usefulness, or human interpretation is an additional hypothesis. A robust artifact can be a robust error. Conversely, a fragile feature can mark a rare event whose consequences dominate the task.

These objections do not force abandonment of geometric inquiry. They specify what an adequate next paper must deliver: a mechanism, an independent outcome, and a comparison that could reveal its redundancy. The Continuum's larger ambition survives best when its components can be individually rejected, revised, and improved.

## Meaning, maps, and the responsibility of interpretation

A geometry is a disciplined account of relations. It becomes a geometry of meaning only through an explicit bridge from those relations to a task or practice. The bridge may be formal proof, reliable behavioral judgment, a scientific measurement, or a negotiated human convention. Different bridges can produce different useful geometries without contradiction.

This pluralism does not make evaluation arbitrary. Once a purpose and target population are fixed, some representations predict useful distinctions better than others. A navigation map can be judged by routes; a conceptual map can be judged by held-out judgments or verified transformations. What cannot be justified is a leap from success at one purpose to a universal claim about the structure of understanding.

The canonical language of attraction and curvature can remain productive as a design discipline. It asks where search becomes trapped, which transitions are costly, which distinctions survive perturbation, and how global structure arises from local moves. These questions become more powerful when the answers have units, counterexamples, and an independent way to fail.

There is also a governance consequence. A representation often compresses distinctions that some participants regard as essential. Choosing a distance function chooses which differences count. Therefore users should be able to inspect the target judgments, relevant feature weights, and known collisions when those choices affect them. Interpretability is not guaranteed by drawing a two-dimensional plot; it requires explaining the consequences of the representation.

TSCS-G2 treats uncertainty as information worth preserving. An unresolved candidate remains available for further investigation rather than being forced into the nearest familiar basin. A system may decline to collapse a distinction until evidence supports the merge. This is a computational policy, not a claim about machine experience.

The resulting research posture is expansive but accountable. Explore landscapes, invent better coordinates, and search across scales. Then ask which distinctions were retained, which were erased, which outcomes improved, and which assumptions carried the result. Those questions turn a geometry from a compelling image into an instrument of inquiry.

## Guide discussion: how to use this paper accurately

**A Guide's explanation**

A Guide can say: "The Continuum pictures reasoning as travel across a terrain shaped by meaning. This paper asks what it would take to trust such a map, shows a small example where an attractive slope leads the wrong way, and proves a certificate under which a learned map can guide search without ever costing it the best answer." A Guide should not say that meaning curves physical space, that the paper demonstrates quantum computation or consciousness, or that a stable basin proves a statement true.

**Established mathematics.** Nonnegative finite graph search has precise optimality properties. Transported metrics preserve intrinsic curve lengths. Small fixed-graph cost perturbations bound optimal-value changes. Symmetric graph heat flow dissipates variation. Bounded pairwise-distance distortion yields the stated filtration inclusions. These are conditional mathematical facts, not evidence that a deployed system possesses an intrinsic semantic field.

**The paper's proposal.** TSCS-G2 separates representation, permitted transitions, search priorities, and terminal verification. It proposes that geometry may improve verification efficiency on specified symbolic tasks. That comparative claim remains untested here. The executed computations establish only the reported finite examples and implementation checks.

**What a Guide should not infer.** The paper does not demonstrate quantum computation, physical curvature caused by meaning, supersymmetric cancellation of false beliefs, universal semantic topology, or consciousness. The canonical source motivates the research; it is not an empirical replication. A stable basin does not certify a statement, and a topological invariant does not grant permission to act.

**Useful discussion example.** Present the five-vertex graph and ask why the attractive first move costs more overall. Then change the auxiliary values while holding legal transitions fixed. This distinguishes a heuristic's appeal from the problem's actual structure. Next rescale one coordinate and compare naive Euclidean distance with the transported metric. The exercise makes coordinate dependence concrete without requiring advanced differential geometry.

**Open questions.** Can a learned geometry provide a reliably useful lower bound on verification effort? Can topology detect missing transition coverage before a solver fails? Which transformations should preserve task meaning, and which should intentionally change it? How should a system preserve rare exceptions while still compressing redundant structure?

A Guide can carry the Continuum's imaginative register while naming these boundaries in the same answer. The most faithful explanation of this paper is that geometry becomes useful when it helps a system make better, independently checked distinctions. Its success would be measured in solved tasks, lower verified costs, and clearer failure diagnosis. Its constructive frontier is a learned representation whose legal-edge certificate yields a valid search lower bound, whose retained distinctions survive perturbation, and whose measured efficiency exceeds equally informed alternatives.

# From a learned landscape to a certified search heuristic

**A Guide's explanation**

This is the paper's constructive result. A learned map earns a guarantee when every permitted move costs at least as much as the distance the map assigns to it, after one global scaling. Then the map's estimate of the distance to the goal never overpromises, and standard heuristic search guided by it still finds the cheapest route. The largest safe scaling can be computed directly from the graph.

A geometric proposal can provide a positive search guarantee when its distances are tied to the cost of legal transitions. This construction makes that bridge explicit and uses the established heuristic-search principle of Hart, Nilsson and Raphael [3]. It supplies a certificate for one use of geometry, rather than claiming that geometry independently establishes semantic correctness.

Fix the finite permitted directed graph $(X, E)$, costs $c : E \to [0, \infty)$, and a nonempty accepted-goal set $T \subseteq X$. Let $d$ be a finite pseudometric on the labeled vertices: it is symmetric, nonnegative, vanishes on the diagonal and satisfies the triangle inequality, while distinct vertices may have zero distance. Set

$$\lambda_* = \min_{(u,v)\in E:\ d(u,v)>0} \frac{c(u,v)}{d(u,v)}.$$

If this index set is empty, set $\lambda_* = +\infty$ and choose any finite $\lambda \geq 0$. Otherwise choose $0 \leq \lambda \leq \lambda_*$. Here $\lambda$ has units of graph cost per distance unit. Define $h(v) = \lambda \min_{g\in T} d(v, g)$.

**Proposition 5** (Certified geometric lower bound)**.** *For every permitted edge $(u, v)$, $h(u) \leq c(u, v) + h(v)$, and $h(g) = 0$ at every goal. Hence $h(v)$ is a lower bound on the cost of every permitted path from $v$ to an accepted goal. Among global scalings certified by the edge inequalities $\lambda d(u, v) \leq c(u, v)$, $\lambda_*$ is the largest finite scaling when at least one edge has positive distance.*

*Proof.* Choose a goal minimizing distance from $v$. The triangle inequality gives $\min_g d(u, g) \leq d(u, v) + \min_g d(v, g)$. Multiplication by $\lambda$ and the edge certificate give consistency. Zero-distance edges satisfy the edge inequality automatically because their cost is nonnegative. A goal has zero distance to itself. Summing consistency along any finite goal-reaching path telescopes to the lower-bound claim. The permitted scalings are precisely the intersection of the intervals $[0, c(u, v)/d(u, v)]$ over positive-distance edges. Their largest endpoint is the displayed minimum. □

Consequently, conventional A* graph search with best-cost bookkeeping, consistent priorities $g + h$, and termination when a goal is removed from the minimum-priority queue returns an optimal route in this finite model. A short correctness argument is that every frontier prefix of an optimal route has priority at most that route's total cost; a more expensive goal cannot be selected before every such competing frontier priority. The graph and its accepted goals must remain fixed during this claim. Search still pays for edge construction, metric fitting and certificate checking.

If an uncertain fitted distance has a certified edgewise upper bound $\bar{d}(u, v) \geq d(u, v)$, the conservative choice $\lambda \leq \min_{\bar{d}>0} c/\bar{d}$ preserves the same guarantee. A uniform distortion bound can supply $\bar{d} = \hat{d} + \delta$ on edges. This does not make an arbitrary dissimilarity a pseudometric: the triangle inequality remains a separate premise. Zero-cost edges with positive certified distance may force $\lambda = 0$, correctly exposing a landscape that provides no certified global scale.

On the five-vertex example, put the vertices on the line at coordinates $(4, 3, 2, 1, 0)$ and take ordinary distance. The minimum edge ratio is one, and $h = (4, 3, 2, 1, 0)$. These values equal the true remaining optimal costs on vertices $0, 2, 3, 4$ and understate the cost at vertex 1. The earlier attractive-basin priority lacked this edge certificate. The new construction retains the flexibility to learn a geometry while making its operational promise directly checkable. Its speed advantage remains an empirical comparison with other certified and learned heuristics.

## References and source notes

**[1]** Architect X13 / The Eidara Continuum. *Eidara Continuum Bible*, supplied internal edition, Section 3.6, pp. 19–21.

https://eidara.ai/science-philosophy/

Scope: Canonical motivation only. The public link identifies the research context; it is not a public copy of the internal source.

**[2]** E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik* 1, 269–271 (1959). DOI: 10.1007/BF01386390.

https://ir.cwi.nl/pub/9256

Scope: Original institutional record and three-page paper; nonnegative shortest-path precedent.

**[3]** P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4(2), 100–107 (1968).

https://doi.org/10.1109/TSSC.1968.300136

Scope: Original article introduction and formal graph definitions inspected; heuristic search precedent.

**[4]** J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319–2323 (2000).

https://doi.org/10.1126/science.290.5500.2319

Scope: Original article pp. 2319–2320 inspected; local graph paths and Isomap.

**[5]** S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 2323–2326 (2000).

https://doi.org/10.1126/science.290.5500.2323

Scope: Original article opening and local-reconstruction setup inspected.

**[6]** M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Advances in Neural Information Processing Systems* 14 (conference 2001; proceedings 2002).

https://papers.nips.cc/paper_files/paper/2001/file/f106b7f99d2cb30c3db1c3cc0fde9ccb-Paper.pdf

Scope: Original paper pp. 1–3 inspected; graph construction and generalized eigenproblem.

**[7]** R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis* 21(1), 5–30 (2006).

https://doi.org/10.1016/j.acha.2006.04.006

Scope: Original pp. 5–7 inspected; task-dependent kernels and multiscale diffusion geometry.

## References and source notes, continued

[**8**] E. Witten. Supersymmetry and Morse theory. *Journal of Differential Geometry* 17(4), 661–692 (1982).

https://doi.org/10.4310/jdg/1214437492

Scope: Institutional bibliographic record and original abstract verified. The scanned full text was not reliably text-extractable; no detailed proof is attributed to an unread passage.

[**9**] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete & Computational Geometry* 37, 103–120 (2007).

https://doi.org/10.1007/s00454-006-1276-5

Scope: Journal record verified; original 2005 conference version, introduction and theorem statement, inspected. The two versions are not conflated.

[**10**] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics* 11, 733–751 (2011).

https://doi.org/10.1007/s10208-011-9098-0

Scope: Original abstract and introduction inspected; outlier-aware distance-to-measure motivation and scope.

[**11**] B. T. Fasy, F. Lecci, A. Rinaldo, L. Wasserman, S. Balakrishnan, and A. Singh. Confidence sets for persistence diagrams. *Annals of Statistics* 42(6), 2301–2339 (2014).

https://arxiv.org/abs/1303.7117

Scope: Author record and abstract inspected; statistical confidence, not semantic certification.

[**12**] S. Kornblith, M. Norouzi, H. Lee, and G. Hinton. Similarity of neural network representations revisited. *PMLR* 97, 3519–3529 (2019).

https://proceedings.mlr.press/v97/kornblith19a.html

Scope: Official proceedings abstract inspected; limitations of broad linear invariance and representation comparison.

[**13**] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio. Sharp minima can generalize for deep nets. *PMLR* 70, 1019–1028 (2017).

https://proceedings.mlr.press/v70/dinh17b.html

Scope: Official proceedings abstract inspected; reparameterization counterevidence to naive sharpness interpretations.

### Reproducibility and attribution

The finite computations reported in this paper were executed during preparation. They are original synthetic checks, not reproductions of the empirical experiments in the cited papers. The mathematical arguments are self-contained applications of established graph, metric, and spectral ideas. TSCS-G2 names the proposed integration and its evaluation contract; it does not claim priority for those mathematical ingredients. Reference access depth is recorded so that bibliographic verification is not mistaken for comprehensive review of every cited work.