



TOPOLIFT AI

PRECISION INTELLIGENCE. BUSINESS READY.



The Trust Layer for Enterprise AI

How Topological Ontologies contain confabulation at the source — and make sovereign, decision-grade AI a reality.

33.86×

context compression

~97%

fewer tokens, same or better evidence

01 · THE PROBLEM

Confabulation Is a Context Defect, Not a Model Defect

Enterprises cannot deploy LLMs into consequential decisions because the models confabulate — they generate confident, fluent answers unmoored from the organization's actual data. In a consumer chat, that is an annoyance. In underwriting, clinical planning, or an agentic workflow, it is a liability that compounds, because every downstream step inherits the fabrication as fact.

The clinical term is precise: a confabulating patient is not lying and is not malfunctioning — they are constructing a coherent narrative in the absence of the memory that should have constrained it. This is exactly what an LLM does when the specific structure of an organization's reality is absent from its context, or present only as flat, disconnected text. It fills the gap with the statistically typical answer.

The missing link in today's tools

The industry has treated confabulation as a model defect and thrown fine-tuning, retrieval, and bigger models at it. But retrieval-augmented generation treats context as **volume** when the problem is **structure** — more tokens of unstructured evidence widen the space of plausible narratives rather than narrowing it. Semantic ontology platforms constrain what the AI may *look at*; they do not bound its inferential moves. Guardrail tools inspect outputs *after* the model reasons. Nothing on the market constrains the reasoning itself.

“A model cannot be talked out of confabulating. It must be given a context whose structure makes confabulation the harder path.”

That is the layer TopoLift built: **structural context constraint** — the context handed to the model mathematically bounds the set of reasoning trajectories it can take, and every output is attributable to the elements of context that licensed it.

02 · THE SOLUTION

Topological Ontologies: Structure Discovered, Not Declared

A Topological Ontology is a knowledge structure computed directly from an organization's data using combinatorial topology, graph theory, and network analysis. Rather than asking experts to declare the schema of the business, TopoLift measures it. The ontology is the shape of the data itself.

From records to hypergraphs

The construction treats heterogeneous enterprise data — tables, clinical notes, process rules, decisions, outcomes — as points and relations in a high-dimensional space, and yields a **hypergraph**: a structure in which a single edge may bind any number of variables. Real enterprise phenomena are rarely pairwise. A property's fire-severity risk is not driven by smoking, or absent detectors, or slow response — it is driven by the **joint occurrence of all three**. In a topological ontology, that three-way interaction is a single hyperedge: a first-class object, discovered because the mathematics detected the joint effect, not because an underwriter thought to encode it.

Reasoning atoms: the compiled unit of trust

The discovered structure is compiled into **reasoning atoms** — governed units of evidence carrying five things together, inseparably:

Signal	the predictive evidence itself
Topology	position, weights, and hyperedge memberships within the network
Provenance	the exact source records that produced it
Uncertainty	calibrated confidence in the signal
Governance	who may use it, under what policy

This co-compilation matters. In retrieval architectures, relationships, lineage, and confidence are bolted on after text is fetched — and each seam is a place where fabrication leaks in. In a topological ontology there are no seams: an atom that reaches the model's context arrives already structured, already sourced, already weighted. And because atoms are dense, structured representations rather than prose, they are radically compact: **33.86× less context — roughly 97% fewer tokens** — for the same or better task evidence. Nothing semantically load-bearing is discarded; only the linguistic packaging the model never needed to re-derive.

03 · TOPOLOGICAL ONTOLOGIES IN DEPTH

How Topology Contains Confabulation

The ontology is not a schema and not a graph database — it is a global, quantitative characterization of the organization's data: persistent features, influence pathways, communities of interacting factors, and the centrality of each element within the whole. Where a semantic ontology records that things are related, a topological ontology records how much, in what combinations, and with what reach. And it is a living structure: the topology re-computes as the underlying data changes, so the ontology never drifts from the business it measures. Structural context constraint then operates through four reinforcing effects.

A bounded reasoning manifold

When the context consists of reasoning atoms, the model is no longer navigating the open sea of everything fluent English permits. The atoms and their hyperedges define a manifold of admissible inferences: factor A may connect to outcome C only through pathways that exist in the topology, with strengths the topology specifies. Confabulation would require the model to assert a connection *against* an explicit, weighted structure sitting in its context — a far higher bar than inventing a connection into a vacuum. The empty space that confabulation fills has been pre-filled with measured reality.

Atomic traceability

Every output can be traced to the specific atoms — and through their provenance, the specific source records — that licensed it. This is stronger than citation: a citation says a document exists that is consistent with the claim; atomic traceability says the claim is a *function of identified evidence*, and exposes the function. Verification becomes mechanical: reviewers do not ask whether an answer sounds supported, they inspect whether the supporting atoms entail it. A step with no atom behind it is visible as exactly that.

Weak-chain detection in multi-step workflows

Agentic systems magnify confabulation because errors propagate: a fabricated premise in step two becomes trusted input by step seven. Each link in a reasoning chain has a measurable structural footing — the weight and confidence of the atoms it rests on. Chains whose footing falls below threshold are flagged, re-grounded, or routed to a human before their conclusions enter downstream steps. The same instrumentation detects a subtler threat: injected instructions and anomalous premises arriving through tool outputs stand out as structural discontinuities rather than blending into a stream of undifferentiated text.

Constraint as efficiency, efficiency as sovereignty


Because the constraining structure is compact, smaller and local models perform tasks that previously demanded frontier-scale context windows. A 33.86× reduction in required context is simultaneously a cost reduction, a latency reduction, and a sovereignty enabler: sensitive data need never leave local infrastructure to obtain frontier-grade reliability, because reliability is supplied by the structure, not by the model's size.

Constraint, correctly engineered, is not a tax on capability. It is where the capability comes from.

04 · THE ARCHITECTURE

The Context Constrained Transformer

Everyone else uses the LLM as the reasoning engine and hopes. We use it as a constrained component inside a topologically governed pipeline — where the structure, not the model, guarantees reliability.



Model-agnostic by design

The interface between our layer and any model is a formalized **atomic context contract**: exactly what structured context goes in, exactly what traceable output must come out. Swapping model families — closed frontier APIs, open self-hosted models, or a customer's own proprietary models — is a configuration change, not a re-architecture. For sovereignty-constrained buyers, the entire stack runs on-premise: because reliability is supplied by the structure rather than by model size, compressed reasoning atoms let small, local models deliver frontier-grade results. **Sovereign AI stops being a compromise.**

05 · THE EVIDENCE

Benchmark Results

The argument is architectural; the evidence is empirical. In two independent due-diligence benchmarks, TopoLift-augmented agents were evaluated against a frontier baseline (GPT-5.1) on domain reasoning tasks with expert-designed rubrics.

33.86×

context compression,
same or better task evidence

~97%

fewer tokens than
raw-text retrieval

Insurance underwriting

On a property underwriting scenario requiring risk-driver identification, synergy mapping, and actionable conditions, the TopoLift-augmented agent outperformed the baseline by a wide margin — with near-full marks in structural reasoning where the baseline scored near zero. The decisive gap was combinatorial: the topological agent surfaced multi-factor risk synergies (freeze–thaw climate × older roof; smoking × absent detectors × slow fire response) that the baseline, reading the same facts, never connected.

Clinical rehabilitation planning

On an identical stroke-patient case, the TopoLift-augmented agent achieved a perfect plan-quality score and a clear overall margin over a baseline that was already clinically competent. Expert reviewers judged the plan indistinguishable from expert therapist reasoning — because it reasoned across domains, linking cardiometabolic comorbidities to therapy tolerance and cognitive status to home-environment safety.

06 · TWO ONTOLOGICAL PHILOSOPHIES

Declared vs. Discovered

Palantir deserves genuine recognition: its Ontology proved that grounding AI in an explicit model of the enterprise pays enormous dividends, and TopoLift shares that conviction. But the prevailing architecture has a ceiling built into how the ontology comes to exist. The semantic approach asks experts to declare the structure of the business; the topological approach measures it from the data.

Dimension	Semantic Ontology (Foundry / AIP)	Topological Ontology (TopoLift)
Construction	Human-defined: teams map datasets to object types and links by hand.	Machine-discovered: structure computed directly from the data, then governed.
Relationships	Pairwise, typed links (A relates to B).	Hyperedges: n-ary interactions capturing joint, combinatorial effects.
Weighting	Largely categorical; influence strength is not intrinsic.	Every edge carries measured influence — centrality, persistence, propagation.
Context for LLMs	Objects and embeddings retrieved as text; retrieval and reasoning are separate.	Compiled reasoning atoms constrain generation directly — 33.86× compression.
Confabulation	Mitigated after generation via governance and citation prompting.	Contained at the source: outputs must trace to atoms; weak chains flagged before they propagate.
Footprint	Full-platform commitment; frontier-scale context volumes.	A lightweight layer above existing stacks; sovereign models become viable.
Time to value	Multi-quarter ontology-building engagements.	Ontology discovered in the first pass; curation refines rather than creates.

These paradigms are not mutually exclusive: TopoLift is a **layer, not a platform**. A topological ontology can be discovered over data already living in a semantic ontology — inheriting its governance while contributing the quantitative, higher-order structure the semantic layer cannot express. Organizations invested in Foundry lose nothing by adding measured topology above it. Their success is our distribution.

07 · THE MOAT

Not a Wrapper — and Why Better Models Make Us Stronger

A wrapper passes prompts to a model. We built the layer the model reasons inside — on mathematics, not prompt engineering. And if the base models improve tomorrow, our customers' answers get better while our layer remains the only thing standing between that new power and a real decision.

What we built

- **Topological ontology construction** — proprietary algorithms that discover multi-way hypergraph structure in enterprise data, versus the declared, pairwise ontologies of incumbent platforms. Applied topology, not embeddings-plus-a-graph-database.
- **The reasoning atom** — our core data structure. Atomic traceability and gated outputs are only possible because signal, topology, provenance, uncertainty, and governance are compiled together. You cannot retrofit this onto RAG.
- **Weak-chain detection** — reasoning paths are scored against the hypergraph's structure; chains crossing thin regions are flagged, re-grounded, or routed to a human before their conclusions enter downstream steps. The same instrumentation surfaces injected instructions and anomalous premises as structural discontinuities.
- **The atomic context contract** — the formalized interface that makes the system model-agnostic.

Three layers of defense

- **The math is the moat.** Topological and hypergraph methods at this depth are a rare specialization — years of research, not a weekend of tooling. Model releases don't touch this layer; we are the first beneficiary of every model improvement.
- **A compounding data asset.** Every deployment builds a customer-specific ontology that deepens with use — each gated decision, traced answer, and flagged weak chain enriches the topology. Copying our code copies none of it.
- **Compliance lock-in.** Once TopoLift is the system of record for *why* decisions were made, we are audit infrastructure, not a tool. Nobody rips out their audit trail.

The better the models get, the more the trust gap — not the capability gap — becomes the binding constraint on enterprise AI. **That gap is our business.**

08 · MARKET & BUSINESS MODEL

Who We Serve, and How We Capture Value

Our ideal customer is a mid-to-large enterprise making high-stakes decisions in complex operational environments — where a confidently wrong AI answer has real financial or safety consequences. We sell to the executive accountable for the decision, not the data science team.

- **Existing Foundry customers** — they've invested in semantic ontology and hit its ceiling. TopoLift sits above their stack: an expansion, not a rip-and-replace. Our fastest sales motion.
- **Sovereignty and data-residency buyers** — government-adjacent, regulated finance, healthcare. They cannot send data to closed APIs and need every AI-assisted decision auditable. Our premium segment.
- **Mid-size operators in project-driven industries** — firms running large capital projects who face the same go/no-go risk but can never afford a platform-scale engagement. Gated scenario runs priced against project value make us accessible where incumbents never will be.

Pricing: outcomes and the ontology — never seats or tokens

Platform license (annual)	Priced on the Topological Atomic Ontology — the customer's living structural asset. The recurring core.
Outcome-based consumption	Gated scenario runs priced per gate, anchored to the value of the decision at stake — the cost of being wrong, not the cost of compute.
Pilot-to-expand	Paid pilots prove value on one decision domain, then expand across the ontology.
Sovereignty premium	On-premise, model-agnostic deployments carry a premium — these buyers have the fewest alternatives and the highest stakes.

We deliberately decoupled pricing from inference. Per-token pricing commoditizes as models get cheaper; our billable assets — the ontology and the gated decision — appreciate with use.

09 · RISK PHILOSOPHY

Risk Is the Product — and Risk Is a Dial

AI risk isn't a compliance checkbox we manage around the product. It is the product.

- **Confabulation** — architected against, not disclaimed. Structural context constraint, weak-chain detection, and gated outputs catch confident-but-wrong answers before they reach a decision-maker.
- **Model unpredictability** — the model is never the guarantor; the topology is. A model behaving unexpectedly is contained by structure, flagged by weak-chain scoring, and blocked at the gate.
- **Data privacy and sovereignty** — customer data never has to leave customer infrastructure. Compressed reasoning enables fully on-premise deployment. Privacy isn't a policy; it's a deployment architecture.
- **Risk is a dial, not a switch** — customers control how tightly the topology constrains the model, from fully gated decisions to exploratory reasoning, and the setting is always visible. Assuming more creative freedom is an informed, auditable choice — never a hidden default. The constraint is adjustable; the audit trail is not. Atomic traceability and weak-chain visibility never turn off.

From answers to outcomes

A chatbot answers questions. TopoLift delivers **gated decisions**: data assembled, structure verified, reasoning checked for weak chains, answer traced to source, decision gated for sign-off. The customer doesn't buy AI assistance — they buy a resolved decision with an audit trail. And as agentic systems assume multi-step operational work, the CCT architecture is what makes them safe to ship: agents fail in production because errors compound silently across steps; reasoning atoms and weak-chain detection constrain every step, with each hop grounded in verified structure and traceable end to end.

“Chatbots automate the conversation. Agents automate the task. TopoLift automates the trusted outcome — the only kind an enterprise will actually deploy.”



Most AI generates answers. TopoLift builds understanding.

Understanding is the only reliable constraint on a machine
that would otherwise say anything.

TOPOLIFT AI

Dr. Joseph Geraci · Founder & Executive Chairman
Joseph.Geraci@topolift.ai

Benchmark figures reflect internal due-diligence evaluations.

Palantir, Foundry, and AIP are trademarks of Palantir Technologies Inc. This document is not an offer of securities.