

Latent-Sense Technologies

With you every step of the way.

rxMaps:
Upstream Semantic Infrastructure for
AI Reasoning Systems

📍 112 -970 Burrard Street,
Unit 1330, Vancouver,
BC V6Z 2R4
Canada.

📍 Suite 111,
1201 North Market Street,
Wilmington, DE 19801,
United States.

🌐 www.latentsense.ai

✉ inquiries@latentsense.ai

rxMaps: Semantic Infrastructure for AI Reasoning



From Data Storage to Semantic Infrastructure

Modern cloud and data providers excel at storage, retrieval, and compute, but none provide the semantic normalization, meaning-preserving structure, or constraint-aware reasoning required for advanced AI systems. As a result, today's AI operates over flat text chunks and probabilistic embeddings, not true semantic understanding.

A semantic infrastructure is the foundation that prepares, structures, and connects meaning across all enterprise data; enabling AI systems to reason with consistency, logic, and traceable understanding.

rxMaps semantic infrastructure provide:

- a normalized semantic space,
- with explicit relationships,
- constraint-aware knowledge, and
- provenance-preserved structure.

rxMaps: Upstream Semantic Infrastructure

Beyond standard storage and compute, rxMaps provides the essential semantic normalization layer for advanced AI. Unlike probabilistic RAG or vector search, rxMaps operates upstream to transform raw data into a deterministic, long-term semantic substrate. It establishes the cross-document framework and auditable logic required for true, defensible Cognitive AI.

Cloud & Data Service Providers



Raw Data Storage



Pre-Storage Semantic Structure




No Semantic Space Constraints



Semantic Relationships Algebra



Isolated Documents



Cross-Document Binding



No Drift Detection



Drift & Conflict Detection



Session-Based Retrieval



Persistent Semantic Memory



Probabalistic Search



Deterministic & Auditable Graph

With rxMaps

Raw Data

Raw corpora, Documents, Unstructured and semi-structured data sources

Resides in S3 | ADLS | GCS | Data Lakes| enterprise stores | source systems

LST Reasoning Infrastructure - rxMaps

Reasoning Fabric Formation

rxMaps as the semantic fabric substrate enables:

- semantic fabric assembly and mapping
- semantic persistence
- relationship integrity preservation
- reasoning-ready structure for downstream model use

Reasoning Representation Engineering

rxMaps as enterprise structured semantic memory provides:

- semantic operations for querying, exploring, extracting
- semantic structure and relationships storage
- semantic structure artefacts
- knowledge artefacts

Data Engineering

The data engineering layer primarily prepares, transforms, and organizes data so that downstream systems (ML models, analytics dashboards, and AI agents) can operate on it.

This is where enterprises transform raw organizational data into structured assets that can be used for analytics, machine learning, and AI applications.

rxMaps become engineered semantic artefacts residing in GraphDBs, RAGs and Datalake artifacts with:

- Context gain: document meaning is retained during Extract-Transform-Load
- Unified Knowledge: Data is unified across datasets through semantic structure
- Causal understanding: Systems detect relationships and semantic structures
- Reasoning layer: Infers implications across records.

Integrates with Databricks | Snowflake | Palantir | Dataiku | Alteryx | DataRobot

Model Infrastructure (Pre-Training)

Pre-training, representation learning, model optimization, training orchestration etc.

rxMaps enable:

- training-data semantic preparation
- tripletization with intact latent semantics
- reasoning-ready representation before training

Integrates with OpenAI | Anthropic | Meta | Google DeepMind | Cohere | Hugging Face

Model Infrastructure (Inference Layer)

Runtime Inference

Runtime inference infrastructure serving:

- context execution,
- memory access,
- tool use,
- generation-time orchestration

Reasoning Harness

rxMaps and ReX function as a reasoning harness for post-training / inference-time cognition and control, through:

- Evidence-first decision validation
- Agent reasoning verification
- Semantic structuring for reasoning traces
- Traceability across prompts, evidence, memory and decisions

Integrates with OpenAI APIs | Anthropic APIs | Fireworks AI | vLLM-based serving stacks

Application Layer

Enterprise AI products, copilots, AI SaaS, agentic applications

Microsoft Copilot | Salesforce Einstein | ServiceNow AI | Sierra.