

AEGIS INSIGHT

See the Shape of Knowledge

A Technical Vision for Builders and Architects

Cedrus Strategic LLC
January 2026

What If Your AI Could See the Shape of Knowledge?

Not just retrieve documents. Not just return the top-K chunks with the highest cosine similarity. What if your AI system could understand that a topic contains three distinct narrative clusters, each with different suppression signatures and coordination patterns, and that certain claims bridge these clusters in ways that suggest synthesis opportunities?

This is what Aegis Insight provides: epistemic topology as a queryable service.

Aegis Insight is open-source infrastructure that transforms document corpora into multi-dimensional knowledge graphs, then exposes that structure through APIs designed for AI system integration. Rather than treating knowledge as a flat collection of text chunks, Aegis Insight extracts the relational architecture of claims, entities, sources, and narratives—then makes that architecture accessible to the systems you build.

The Core Capability: Perspective Clustering

The most powerful endpoint in the system is `get_perspectives`. This is not a search function. It is a topology function. You provide a topic, and Aegis Insight returns the structure of perspectives that exist around that topic—each cluster with its own detection scores, representative claims, and position label.

What the Response Actually Contains

Consider a query about AI development risk—a domain where serious researchers hold genuinely different positions. The system returns a response structure that reveals the epistemic landscape:

```
{
  "topic": "AI existential risk",
  "clusters": [
    {
      "cluster_id": "cl_001",
      "label": "acceleration_optimist",
      "size": 412,
      "position": "dominant",
      "coordination_score": 0.45,
      "suppression_score": 0.08,
      "representative_claims": [
        {
          "text": "AI development should proceed rapidly; benefits outweigh speculative risks",
          "type": "PRIMARY",
          "source": "Andreessen, 2023"
        }
      ]
    },
    {
      "cluster_id": "cl_002",
      "label": "alignment_concerned",
      "size": 287,
      "position": "institutional",
      "coordination_score": 0.15,
      "suppression_score": 0.85,
      "representative_claims": [
        {
          "text": "AI development should be carefully monitored; risks outweigh benefits",
          "type": "SECONDARY",
          "source": "Bostrom et al., 2014"
        }
      ]
    }
  ]
}
```

```

  "coordination_score": 0.52,
  "suppression_score": 0.15,
  "representative_claims": [
    {
      "text": "Alignment problem poses existential risk requiring coordinated
slowdown",
      "type": "PRIMARY",
      "source": "Future of Life Institute"
    }
  ]
},
{
  "cluster_id": "cl_003",
  "label": "capability_skeptic",
  "size": 89,
  "position": "academic",
  "coordination_score": 0.12,
  "suppression_score": 0.35,
  "representative_claims": [
    {
      "text": "Current architectures cannot achieve general intelligence; risk
discourse is premature",
      "type": "PRIMARY",
      "source": "Marcus et al."
    }
  ]
},
"bridging_claims": [...],
"topology_summary": {
  "cluster_isolation": 0.67,
  "cross_citation_rate": 0.12
}
}
  
```

This response tells you something no traditional RAG system can: the shape of the discourse. You have three distinct clusters—optimists pushing acceleration, institutionally-backed alignment researchers urging caution, and academic skeptics questioning the entire premise. Each cluster has different coordination and suppression signatures. The moderate isolation score and 12% cross-citation rate indicate these communities do engage with each other, unlike more polarized domains.

This is actionable intelligence for AI systems. Your RAG pipeline can now decide how to present information with awareness of the epistemic context. Your synthesis can acknowledge multiple perspectives with appropriate uncertainty. Your retrieval can surface claims from underrepresented clusters when appropriate.

Bridging Claims: Where Perspectives Touch

The bridging_claims field contains something unique: claims that cite or connect multiple clusters. These are points of potential dialogue, synthesis opportunities, or simply claims that don't fit neatly into polarized positions.

Consider a claim like: "Targeted regulation of frontier models may address near-term risks without impeding beneficial AI development." This claim bridges the gap between accelerationist

and alignment-concerned clusters. It acknowledges both the value of continued development and the legitimacy of safety concerns. The system identifies these automatically by analyzing citation patterns and semantic positioning.

For builders, bridging claims offer a starting point for nuanced responses. Rather than presenting a false binary, your system can surface the claims that acknowledge complexity. This is particularly valuable in contested domains where users benefit from seeing that reasonable middle positions exist.

The Seven-Dimensional Extraction Pipeline

Perspective clustering depends on rich metadata. Aegis Insight achieves this through a seven-dimensional extraction pipeline, where each document chunk passes through specialized extractors:

1. **Entity Extraction:** People, organizations, places, concepts, events—with alias resolution to connect variant names to canonical entities.
2. **Claim Extraction:** Factual assertions with confidence scoring and claim type classification (PRIMARY, SECONDARY, CONTEXTUAL, META).
3. **Temporal Extraction:** Dates, periods, sequences, relative time references—enabling chronological analysis.
4. **Geographic Extraction:** Locations, regions, spatial relationships—enabling geographic filtering.
5. **Citation Extraction:** Source attributions, references, authority claims—building the citation network.
6. **Emotion Extraction:** Sentiment, rhetorical tone, emotional markers—using behavioral analysis frameworks.
7. **Authority Analysis:** Domain expertise indicators, credential markers, institutional affiliation signals.

Each dimension adds analytical capability. Temporal data enables detection of coordinated publication bursts. Geographic data reveals regional clustering. Emotion data identifies rhetorical patterns. Authority data enables credential-weighted analysis.

The result is a knowledge graph where claims carry rich, queryable metadata across multiple axis—not just embeddings for similarity search, but structured data for analytical reasoning.

The Detection Stack: Goldfinger Scoring

Aegis Insight implements three pattern detection algorithms that operate on the knowledge graph:

Suppression Detection

Identifies claims or sources that show characteristics of epistemic marginalization: high-credential sources with low visibility, primary research ignored by meta-claims, citation network isolation despite topical relevance.

Coordination Detection

Identifies claims that show characteristics of synchronized messaging: temporal clustering of publication, linguistic similarity across sources, source network density patterns suggesting organized origin.

Anomaly Detection

Identifies structural patterns that deviate from baseline expectations: unusual claim type distributions, authority-visibility inversions, citation patterns that suggest non-organic dynamics.

Each detector uses "Goldfinger scoring"—a scoring method where multiple independent indicators compound the final score. The principle: coincident patterns suggest non-random causes. A single indicator might be noise; five coincident indicators warrant attention.

Detection scores are calibrated using configurable profiles. Different domains require different baselines. Academic literature has different citation norms than news media. The calibration system allows tuning for domain-appropriate sensitivity.

Model Context Protocol (MCP) Integration

Aegis Insight exposes its capabilities through a Model Context Protocol server, enabling direct integration with AI systems that support MCP. This means your Claude instance, your custom agent, or your LangChain pipeline can query Aegis Insight mid-conversation.

Consider a research assistant that's asked about a contested topic. Before generating a response, the assistant queries Aegis Insight:

```
# Mid-conversation MCP query
perspectives = await aegis.get_perspectives(
    topic="AI_deployment_risk",
    max_clusters=3,
    claims_per_cluster=5
)

if perspectives["topology_summary"]["cluster_isolation"] > 0.7:
    # High isolation = highly contested, present multiple views
    response = generate_multi_perspective_response(perspectives)
```

```

else:
    # Low isolation = convergent views, standard response
    response = generate_standard_response()

```

The AI system now has epistemic awareness. It knows this topic has distinct clusters. It can present multiple perspectives with appropriate hedging. It can surface bridging claims that acknowledge complexity. This is qualitatively different from RAG systems that simply return the top-K chunks.

What You Can Build With This

The infrastructure Aegis Insight provides enables several categories of applications:

Epistemic-Aware RAG Systems

Replace naive chunk retrieval with topology-aware retrieval. Weight sources by epistemic context. Surface underrepresented perspectives when suppression scores are high. Flag coordinated messaging for user awareness. Your RAG system becomes a tool for understanding, not just regurgitation.

Research Synthesis Tools

Build tools that help researchers navigate contested domains. Ingest a corpus of papers on a controversial topic. Show researchers where the cluster boundaries are. Identify which claims bridge perspectives. Highlight which sources are marginalized despite high credentials. Enable synthesis across epistemically fragmented research.

Narrative Intelligence Platforms

Monitor how narratives evolve. Track coordination patterns over time. Identify when new claims enter the discourse and which clusters adopt them. Build dashboards that show narrative topology changing in real-time. This is infrastructure for understanding information environments.

Multi-Perspective Chatbots

Deploy AI assistants that don't just answer questions, but acknowledge epistemic complexity. When asked about contested topics, the assistant can say: "There are three main perspectives on this topic" and provide representative claims from each. This is more honest and more useful than systems that present a single synthesized answer.

Domain-Specific Knowledge Bases

Build knowledge graphs for specialized domains—declassified documents, academic subfields, organizational archives. The extraction pipeline handles the heavy lifting. You get structured, queryable knowledge with pattern detection built in.

A Synthesis Workflow: Reconstructing Scattered Knowledge

Consider a concrete example. You're researching unconventional propulsion physics—a domain where relevant research is scattered across declassified government documents, patent filings, academic papers from multiple decades, and technical reports from various organizations.

Traditional RAG would struggle here. The relevant documents don't share vocabulary. Citation networks are sparse because institutional barriers prevent cross-citation. Embeddings might find surface-level similarity but miss deep conceptual connections.

With Aegis Insight, you ingest the scattered data set. The extraction pipeline identifies claims about experimental parameters, theoretical frameworks, and reported results. The graph builder creates connections based on entity overlap, citation patterns, and semantic similarity.

Now you query: "Where do the experimental parameters converge across sources?" The system can identify claims from different sources that reference similar technical specifications, even if they use different terminology. The bridging claims surface connections that would require deep domain expertise to identify manually.

This is not search. This is epistemic reconstruction—reassembling fragmented knowledge into queryable structure.

Technical Requirements and Architecture

Aegis Insight is designed to run on consumer hardware. The complete system operates effectively on a machine with a modern GPU (RTX 3060+ or better), 32GB RAM, and reasonable SSD storage.

Component Stack

Component	Purpose
Neo4j 5.x	Knowledge graph storage and traversal
PostgreSQL + pgvector	Vector embeddings and similarity search
Ollama	Local LLM inference for extraction and analysis
FastAPI	REST API and MCP server
D3.js	Interactive graph visualization

Deployment

The entire system deploys via Docker Compose. A single command brings up all components, pre-configured and ready for use:

```
git clone https://github.com/Eleutherios-project/aegis-insight.git
cd aegis-insight
docker-compose up -d
```

No cloud dependencies. No API keys required for core functionality. Complete data sovereignty—your documents never leave your infrastructure.

Getting Started

For builders ready to experiment:

8. **Clone and deploy:** Use Docker Compose to bring up the complete stack.
9. **Ingest a data set:** Use the web interface or CLI wizard (parallelism for multi-GPU) to import your documents.
10. **Explore the graph:** Use the visualization interface to understand your knowledge structure.
11. **Query the APIs:** Use the REST endpoints or MCP server to integrate with your applications.
12. **Build on top:** The infrastructure is open source. Extend, customize, integrate.

Full documentation, API references, and example integrations are available in the repository.

Build On It

Aegis Insight exists because the tools for understanding knowledge structure should be accessible to everyone building AI systems. The infrastructure is ready. The APIs are documented. The source is open.

What will you build when your AI can see *the shape of knowledge*?

— — —

Aegis Insight — Open Source Epistemic Infrastructure
github.com/Eleutherios-project/aegis-insight

Cedrus Strategic LLC • cedrusstrategic.com