Semantic Discovery for Patent Landscape Analysis

by Nick Clark | Published March 27, 2026 | PDF

Patent landscape analysis determines the intellectual property terrain surrounding a technology area. Current approaches combine keyword search with classification-code filtering, which confines discovery to the vocabulary and taxonomy of existing patent classification systems. Semantic discovery provides governed traversal through patent corpora that follows technical concepts across classification boundaries, enabling landscape mapping that discovers relevant prior art in unexpected classifications and jurisdictions.

The classification boundary problem in patent search

Patent classification systems organize inventions by technical domain. A search within CPC class G06F returns computing patents. A search within H04L returns networking patents. But inventions increasingly span classification boundaries. A patent describing a machine learning approach to network

optimization lives at the intersection of G06F and H04L, and may be classified in either, both, or neither in ways that depend on the examiner's judgment.

Prior art that is semantically relevant but classified in an unexpected domain is systematically missed by classification-bounded search. A mechanical engineering patent describing a structural optimization algorithm may be relevant prior art for a software patent describing an analogous computational optimization, but no classification-based search will surface the connection.

Freedom-to-operate analysis requires comprehensive landscape coverage. Missing relevant prior art in an unexpected classification does not reduce the patent risk. It increases it, because the risk exists but is undetected.

Why keyword expansion does not solve cross-domain discovery

Patent search tools expand keywords through synonym lists and semantic similarity. This broadens the vocabulary but does not cross domain boundaries. Expanding the keyword "neural network" to include "artificial neural network" and "deep learning" stays within the machine learning vocabulary. It does not discover that a specific signal processing technique described in telecommunications patents uses a mathematically equivalent approach under a completely different name.

The discovery challenge is not vocabulary breadth within a domain. It is conceptual connection across domains. This requires traversal through semantic relationships between ideas, not expansion of keyword lists.

How semantic discovery addresses patent landscape analysis

Semantic discovery treats the landscape analysis as a persistent discovery object that carries the technical concept under investigation, the accumulated prior art findings, and the evolving understanding of the technology space. Traversal proceeds through semantic neighborhoods of technical concepts rather than through classification hierarchies or keyword lists.

A traversal that begins with a specific machine learning technique can follow semantic connections into signal processing, control theory, or optimization mathematics when the conceptual relationships warrant it. The discovery object's persistent state ensures that cross-domain traversal remains governed by the original technical question rather than drifting into irrelevance.

Trust-scoped resolution enables differentiated treatment of granted patents, published applications, provisional filings, and non-patent literature. The landscape analysis can weight granted patents more heavily than applications, or scope traversal to specific jurisdictions, depending on the analytical purpose. A freedom-to-operate analysis scopes to granted patents in target jurisdictions. A technology scouting analysis may scope broadly across all document types.

The discovery object's traversal lineage provides the documentation that patent landscape reports require. Every identified reference can be traced back through the traversal path that surfaced it, and the semantic connection to the technology under investigation is documented in the lineage.

What implementation looks like

An IP department deploying semantic discovery provides analysts with persistent landscape objects that accumulate findings across research sessions. A landscape analysis that spans weeks maintains its state, with each session extending the traversal from the accumulated knowledge rather than re-searching established terrain.

For patent prosecution, semantic discovery enables prior art search that crosses classification boundaries, reducing the risk of overlooked references during prosecution and strengthening the resulting patent by addressing a broader range of prior art during drafting.

For technology licensing and acquisition, semantic discovery provides landscape coverage that identifies licensing opportunities in adjacent technology domains, surfacing patents whose technical concepts overlap with the target technology even when the classification and vocabulary suggest unrelated fields.

Semantic Discovery All 21 steps →

Search, inference, and execution as one governed step.

Primary Technical Disclosure
○ Governed Semantic Discovery: Search, Inference, and Execution Through Adaptive Traversal
Secondary Technical
○ The Adaptive Index as Unified Search-Inference-Execution Substrate○ Three-in-One Traversal: Search, Inference, and Execution in a Single Step○ The Discovery Object: A Traversal-Native Semantic Agent○ Post-PageRank Semantic Ranking: Relevance Through Governed Traversal○ Persistent Semantic State: Eliminating Prompt Reconstruction○ Traversal Lineage as Index Evolution Signal○ Anchor Semantic Neighborhood Publication○ Inference-Time Execution Control as Traversal Primitive○ Anchor Self-Organization Under Entropy and Load Pressure○ Alias Resolution as Navigational Traversal○ Three Discovery Operating Modes: Human Search, Agent Reasoning, Answer Synthesis○ Model-Agnostic Semantic Discovery○ Affect-Modulated Discovery Traversal○ Confidence-Gated Discovery Traversal○ Integrity-Tracked Traversal Drift Detection○ Biological Identity-Scoped Access During Discovery○ Rights-Grade Anchor Governance for Content Discovery○ Forecasting-Shaped Discovery Traversal○ Capability-Constrained Anchor Accessibility○ Collaborative Multi-Object Discovery Traversal
Applications (General)
○ Enterprise Knowledge Management Through Governed Traversal○ AI-Native Search That Replaces PageRank With Contextual Relevance○ Semantic Discovery for Scientific Research○ Semantic Discovery for Legal Case Research● Semantic Discovery for Patent Landscape Analysis○ Semantic Discovery for Medical Literature Search○ Semantic Discovery for Competitive Intelligence○ Semantic Discovery for Regulatory Compliance Search
Applications (Specific)
○ Google Search Retrieves Results, Not Understanding○ Perplexity Answers Questions Without Discovery State○ Elasticsearch Indexes Documents, Not Discovery○ Algolia Optimizes Relevance Without Discovery State○ Pinecone Finds Vectors, Not Understanding○ Weaviate Stores Semantics Without Discovery Governance○ You.com Answers Questions but Does Not Govern Discovery○ Brave Search Built an Independent Index Without Governed

Traversal○ Kagi Charges for Better Results, Not Governed Discovery○ Metaphor Systems Predicts Links but Does Not Govern Traversal○ Glean Indexes Enterprise Knowledge Without Governing Its Discovery○ Coveo Personalizes Retrieval, Not Discovery Governance
Semantic Discovery overview →
AQ
deterministic
autonomy

Legal

Last updated: 2026-03-03

-
- Inventive Steps

- Licensing
- Patents
- Articles
- Legal
- Opportunities
- Sitemap

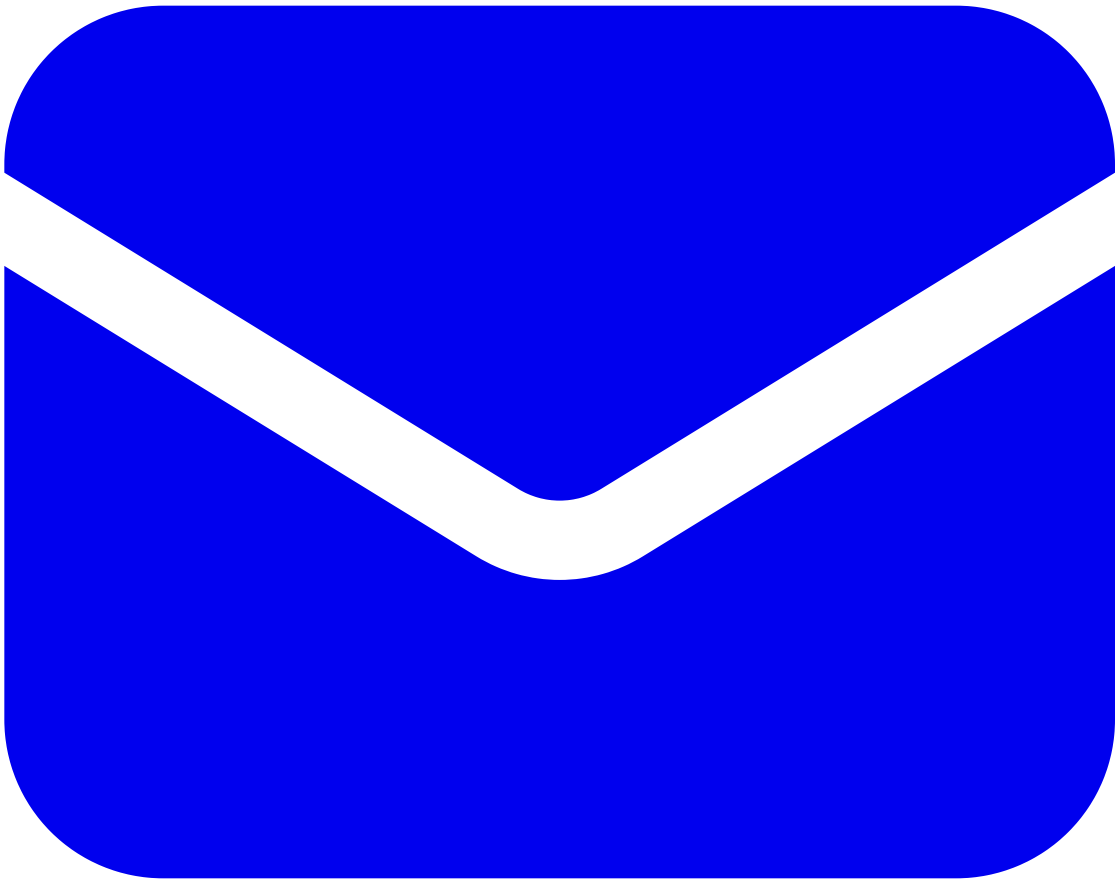

-
- nick@qu3ry.net
- 72 28 14 36 01

Invented by Nick Clark | Founding Investors: Devin Wilkie