

Personal Cognitive Asset: How Per-User Lineage Re-Weights the Same Substrate

Anchor-level personalization without platform lock-in. Each user maintains a portable lineage layer whose accumulated traversal history biases anchor neighborhood scoring on their subsequent traversals, modulating preference among admissible candidates without altering admissibility itself.

Mechanism

Two people can present the same credentials, issue the same query, and mean entirely different things. The adaptive index already scopes what each requester is permitted to see: an anchor publishes its semantic neighborhood to an arriving discovery object under that object's policy profile and biological-identity credential, so access is personalized. What that access-control scoping does not capture is meaning. The same admissible neighborhood, ranked the same way for everyone, cannot distinguish the user who always means their own pet when they say a name from the user who means the public figure. The personal lineage layer adds that missing axis. Each user maintains a per-user record of their own accumulated traversals, and that record biases how an anchor ranks the candidate transitions it offers to a discovery object carrying that user's identity, without changing which transitions are admissible.

The layer is a weighting overlay applied at the moment an anchor publishes its neighborhood to a discovery object. The discovery object carries, in its policy reference field, a reference to its originating user's personal lineage layer: a structured subgraph of that user's prior admitted transitions, scored by recency, repetition, and the satisfaction of the traversals they belonged to. At each anchor, the candidate transition set produced by the search step is reweighted against this subgraph. Transitions structurally aligned with neighborhoods the user has recently and repeatedly traversed are elevated; transitions aligned with the user's prior abandonments and backtracks are deprioritized. The underlying anchors, the shared substrate, and every other user's experience are untouched. The substrate stays common; the experience of it becomes personal.

The defining property of the layer is that it is a first-class, user-owned cognitive asset, not an opaque platform profile. It is portable, exportable, and deletable by the user, and its weights decay under a policy-defined half-life so that the layer tracks current preference rather than ossifying around stale history. A user who leaves a platform takes the accumulated value of their traversal history with them; a user who wants to be forgotten deletes it. The personalization is real, and it does not require the user to surrender a behavioral dossier to the host.

The Governance Boundary

The personal lineage layer modulates preference among admissible candidates; it never alters admissibility. This separation is the load-bearing constraint of the mechanism and is enforced structurally rather than by convention. The candidate transition set presented to the weighting step has already passed the anchor's policy-scoped neighborhood publication: every candidate in it is one the user is authorized to traverse. The layer can only reorder that set. It cannot introduce a transition the policy profile excluded, cannot reveal a neighborhood the credential does not authorize, and cannot surface another user's lineage. A user's history can make an admissible result rank first; it can never make an inadmissible result reachable.

Because the boundary is structural, personalization and governance do not compete. A regulated deployment can offer deeply personalized ranking with no risk that personalization becomes a covert access-control bypass, because the governance evaluation that determines admissibility runs before and independently of the weighting overlay. The layer also cannot leak across users: a discovery object references only its own originating user's layer, and the merge of another user's history is impossible by construction rather than prevented by policy.

Operating Parameters

The layer is parameterized so that its behavior is auditable and bounded. Each recorded prior transition carries a recency weight, a repetition count, and an outcome-satisfaction score drawn from whether the traversal it belonged to reached resolution or was abandoned. These combine into a per-candidate bias applied at publication time. A decay half-life, set by policy, attenuates older entries so that a preference the user has moved away from fades rather than persists. A bias ceiling bounds how far the layer may elevate or suppress a candidate, ensuring the overlay shapes ranking without overwhelming the anchor's own relevance and entropy signals. A portability descriptor defines the export format and the scope of what a user may carry between deployments, and a deletion primitive removes the layer and its derived weights on request.

The layer is keyed to a resolved user identity rather than to a session or a device, so it follows the user across sessions and endpoints and is governed by the same biological-identity credential that scopes the user's access. An autonomous agent may maintain a personal lineage layer of its own, keyed to its agent identity, biasing its later choices toward counterparties and neighborhoods its prior delegations resolved successfully.

Alternative Embodiments

In a professional-research embodiment, a user's accumulated medical or legal research lineage biases their later queries toward the neighborhoods their prior work established as authoritative for them, so a clinician and a patient issuing the same query receive the same admissible sources ranked to their respective histories. In a developer embodiment, a user's coding lineage biases later queries toward the libraries and patterns they have repeatedly traversed, without restricting access to alternatives. In an agent-delegation embodiment, an autonomous agent's prior delegation outcomes bias its subsequent counterparty selection toward higher-success partners, a governed preference that improves coordination without hard-coding a partner list.

In a portability embodiment, a user exports their personal lineage layer on leaving a platform and imports it into another deployment of the substrate, retaining the cognitive value of their traversal history across organizational boundaries. In a regulated-scoping embodiment, the layer's contents are themselves subject to retention and residency policy, so that an enterprise can bound where personal lineage is stored and for how long while still offering personalized ranking within those bounds.

Composition

The personal lineage layer composes with the anchor-level neighborhood publication: the layer is applied to the candidate transition set that the publication produces, so it inherits the publication's policy scoping for free and operates strictly inside it. It composes with biological-identity-scoped access: the same credential that determines which neighborhoods a user may see also keys which personal lineage layer applies, binding personalization to a verified identity rather than to a spoofable session. It composes with the aggregate index-evolution signal disclosed for lineage as a whole: collective lineage drives how anchors split, merge, and re-publish for everyone, while the personal layer drives how an already-published neighborhood is ranked for one

user. The two are deliberately separate consumers of traversal history, one shaping the shared structure and one shaping the individual experience, and neither can be used to reconstruct the other.

Prior-Art Distinction

Recommendation systems personalize at the application layer using behavioral matrices and learned embeddings, and the resulting profile is an opaque platform asset the user neither inspects nor carries away. The personal lineage layer operates at the substrate layer as a deterministic weighting over structured candidate transitions, is bounded so that it cannot override governance, and is a user-owned, portable, deletable asset. Personal knowledge graphs and on-device personalization assistants build private user models, but they are typically platform-locked, without explicit lineage continuity, deterministic and bounded weighting, or a portability guarantee. Browser histories and personalization cookies are user-readable but are neither structured as a cognitive asset nor portable across platforms, and they carry no governance boundary distinguishing preference from permission. The distinguishing combination disclosed here is identity-scoped weighting of admissible candidates, bounded so that preference never becomes access, recorded as a portable and deletable user-owned asset, and decaying under a policy-defined half-life.

Disclosure Scope

The policy-scoped, requester-specific anchor neighborhood publication is disclosed in the cognition filing (U.S. Application No. 19/647,395 and its international counterpart) at Section 10.4, the biological-identity scoping of traversal access at Section 10.15, and the use of aggregate traversal lineage as an index-evolution signal at Section 10.20. This article discloses, as an extension of those primitives, the personal lineage layer: a per-user subgraph of accumulated admitted transitions scored by recency, repetition, and outcome satisfaction; the application of that subgraph as a bounded weighting overlay

on the anchor's published candidate transition set; the structural separation between preference modulation and admissibility; the decay half-life, bias ceiling, and portability and deletion primitives that make the layer a user-owned cognitive asset; and the embodiments enumerated above. The scope extends to weighting policies not described whose behavior reduces to bounded reordering of admissible candidates, and to identity bindings, including autonomous-agent identities, for which a portable personal lineage layer is maintained and applied at traversal time.

Semantic Discovery (</semantic-discovery>)

[All 36 steps → \(/inventive-steps\)](/inventive-steps)

Search, inference, and execution as one governed step.

PRIMARY TECHNICAL DISCLOSURE

- [Governed Semantic Discovery: Search, Inference, and Execution Through Adaptive Traversal \(/articles/governed-semantic-discovery-search-inference-and-execution-through-adaptive-traversal\)](/articles/governed-semantic-discovery-search-inference-and-execution-through-adaptive-traversal)

SECONDARY TECHNICAL

- [The Adaptive Index as Unified Search-Inference-Execution Substrate \(/articles/semantic-discovery/unified-substrate\)](/articles/semantic-discovery/unified-substrate)
- [Three-in-One Traversal: Search, Inference, and Execution in a Single Step \(/articles/semantic-discovery/three-in-one-traversal\)](/articles/semantic-discovery/three-in-one-traversal)
- [The Discovery Object: A Traversal-Native Semantic Agent \(/articles/semantic-discovery/discovery-object\)](/articles/semantic-discovery/discovery-object)
- [Post-PageRank Semantic Ranking: Relevance Through Governed Traversal \(/articles/semantic-discovery/post-pagerank\)](/articles/semantic-discovery/post-pagerank)
- [Persistent Semantic State: Eliminating Prompt Reconstruction \(/articles/semantic-discovery/persistent-state\)](/articles/semantic-discovery/persistent-state)
- [Traversal Lineage as Index Evolution Signal \(/articles/semantic-discovery/traversal-lineage\)](/articles/semantic-discovery/traversal-lineage)
- [Anchor Semantic Neighborhood Publication \(/articles/semantic-discovery/semantic-neighborhoods\)](/articles/semantic-discovery/semantic-neighborhoods)
- [Inference-Time Execution Control as Traversal Primitive \(/articles/semantic-discovery/inference-governance\)](/articles/semantic-discovery/inference-governance)

- [Anchor Self-Organization Under Entropy and Load Pressure \(/articles/semantic-discovery/anchor-self-organization\)](/articles/semantic-discovery/anchor-self-organization).
- [Alias Resolution as Navigational Traversal \(/articles/semantic-discovery/alias-resolution\)](/articles/semantic-discovery/alias-resolution).
- [Three Discovery Operating Modes: Human Search, Agent Reasoning, Answer Synthesis \(/articles/semantic-discovery/operating-modes\)](/articles/semantic-discovery/operating-modes)
- [Model-Agnostic Semantic Discovery \(/articles/semantic-discovery/model-agnostic\)](/articles/semantic-discovery/model-agnostic)
- [Affect-Modulated Discovery Traversal \(/articles/semantic-discovery/affect-modulated-traversal\)](/articles/semantic-discovery/affect-modulated-traversal)
- [Confidence-Gated Discovery Traversal \(/articles/semantic-discovery/confidence-gated-traversal\)](/articles/semantic-discovery/confidence-gated-traversal)
- [Integrity-Tracked Traversal Drift Detection \(/articles/semantic-discovery/integrity-tracked-drift\)](/articles/semantic-discovery/integrity-tracked-drift).
- [Biological Identity-Scoped Access During Discovery \(/articles/semantic-discovery/biological-access\)](/articles/semantic-discovery/biological-access)
- [Rights-Grade Anchor Governance for Content Discovery \(/articles/semantic-discovery/rights-grade-anchors\)](/articles/semantic-discovery/rights-grade-anchors).
- [Forecasting-Shaped Discovery Traversal \(/articles/semantic-discovery/forecasting-shaped\)](/articles/semantic-discovery/forecasting-shaped).
- [Capability-Constrained Anchor Accessibility \(/articles/semantic-discovery/capability-constrained\)](/articles/semantic-discovery/capability-constrained).
- [Collaborative Multi-Object Discovery Traversal \(/articles/semantic-discovery/collaborative-traversal\)](/articles/semantic-discovery/collaborative-traversal).
- [Discovery-Driven Sensor Invocation Closed Loop \(/articles/semantic-discovery/sensor-invocation-loop\)](/articles/semantic-discovery/sensor-invocation-loop).
- [Cross-Platform Credentialed Reader Activation \(/articles/semantic-discovery/credentialed-reader-activation\)](/articles/semantic-discovery/credentialed-reader-activation).
- [LLM-as-Bootstrap: Why Anchor Inference Engines Shrink as the Lineage Matures \(/articles/semantic-discovery/maturation-engine-substitution\)](/articles/semantic-discovery/maturation-engine-substitution)
- **[Personal Cognitive Asset: How Per-User Lineage Re-Weights the Same Substrate \(/articles/semantic-discovery/personal-lineage-layer\)](/articles/semantic-discovery/personal-lineage-layer)**
- [Loki, the Dog, and the Symbol Grounding Problem \(/articles/semantic-discovery/hybrid-symbol-grounding\)](/articles/semantic-discovery/hybrid-symbol-grounding).

APPLICATIONS · GENERAL

- [Enterprise Knowledge Management Through Governed Traversal \(/articles/semantic-discovery/enterprise-knowledge-management\)](/articles/semantic-discovery/enterprise-knowledge-management).
- [AI-Native Search That Replaces PageRank With Contextual Relevance \(/articles/semantic-discovery/ai-native-search\)](/articles/semantic-discovery/ai-native-search).
- [Semantic Discovery for Scientific Research \(/articles/semantic-discovery/scientific-research-discovery\)](/articles/semantic-discovery/scientific-research-discovery)

- [Semantic Discovery for Legal Case Research \(/articles/semantic-discovery/legal-case-research\)](/articles/semantic-discovery/legal-case-research)
- [Semantic Discovery for Patent Landscape Analysis \(/articles/semantic-discovery/patent-landscape-analysis\)](/articles/semantic-discovery/patent-landscape-analysis)
- [Semantic Discovery for Medical Literature Search \(/articles/semantic-discovery/medical-literature-search\)](/articles/semantic-discovery/medical-literature-search)
- [Semantic Discovery for Competitive Intelligence \(/articles/semantic-discovery/competitive-intelligence\)](/articles/semantic-discovery/competitive-intelligence)
- [Semantic Discovery for Regulatory Compliance Search \(/articles/semantic-discovery/regulatory-compliance-search\)](/articles/semantic-discovery/regulatory-compliance-search)
- [Discovery-Coordinated Multi-Sensor Perception \(/articles/semantic-discovery/coordinated-perception\)](/articles/semantic-discovery/coordinated-perception)
- [Post-AirTag Cross-Platform Object Tracking \(/articles/semantic-discovery/post-airtag-tracking\)](/articles/semantic-discovery/post-airtag-tracking)
- [Use the World as Memory: The Brain Strategy for AI \(/articles/semantic-discovery/world-as-memory\)](/articles/semantic-discovery/world-as-memory)

APPLICATIONS · SPECIFIC

- [Google Search Retrieves Results, Not Understanding \(/articles/semantic-discovery/google-search\)](/articles/semantic-discovery/google-search)
- [Perplexity Answers Questions Without Discovery State \(/articles/semantic-discovery/perplexity\)](/articles/semantic-discovery/perplexity)
- [Elasticsearch Indexes Documents, Not Discovery \(/articles/semantic-discovery/elasticsearch\)](/articles/semantic-discovery/elasticsearch)
- [Algolia Optimizes Relevance Without Discovery State \(/articles/semantic-discovery/algolia\)](/articles/semantic-discovery/algolia)
- [Pinecone Finds Vectors, Not Understanding \(/articles/semantic-discovery/pinecone\)](/articles/semantic-discovery/pinecone)
- [Weaviate Stores Semantics Without Discovery Governance \(/articles/semantic-discovery/weaviate\)](/articles/semantic-discovery/weaviate)
- [You.com Answers Questions but Does Not Govern Discovery \(/articles/semantic-discovery/you-com\)](/articles/semantic-discovery/you-com)
- [Brave Search Built an Independent Index Without Governed Traversal \(/articles/semantic-discovery/brave-search\)](/articles/semantic-discovery/brave-search)
- [Kagi Charges for Better Results, Not Governed Discovery \(/articles/semantic-discovery/kagi\)](/articles/semantic-discovery/kagi)
- [Metaphor Systems Predicts Links but Does Not Govern Traversal \(/articles/semantic-discovery/metaphor-systems\)](/articles/semantic-discovery/metaphor-systems)
- [Glean Indexes Enterprise Knowledge Without Governing Its Discovery \(/articles/semantic-discovery/glean\)](/articles/semantic-discovery/glean)
- [Coveo Personalizes Retrieval, Not Discovery Governance \(/articles/semantic-discovery/coveo\)](/articles/semantic-discovery/coveo)
- [Apple Find My Lacks Cross-Authority Reader Activation \(/articles/semantic-discovery/apple-find-my\)](/articles/semantic-discovery/apple-find-my)

- [Google Find My Network Needs Credentialed Cross-Activation \(/articles/semantic-discovery/google-find-my\)](/articles/semantic-discovery/google-find-my).
- [IETF DULT Specifies Behavior, Not Architecture \(/articles/semantic-discovery/ietf-dult\)](/articles/semantic-discovery/ietf-dult).
- [Glean Enterprise Search and Work AI \(/articles/semantic-discovery/glean-enterprise-search\)](/articles/semantic-discovery/glean-enterprise-search).
- [GraphRAG, but with Governance: Where Microsoft's Architecture Stops Short \(/articles/semantic-discovery/microsoft-graphrag\)](/articles/semantic-discovery/microsoft-graphrag).
- [Memory Layers for Agents: Why Mem0, Zep, and Letta Get Close \(/articles/semantic-discovery/memory-for-agents\)](/articles/semantic-discovery/memory-for-agents).

[Semantic Discovery overview → \(/semantic-discovery\)](/semantic-discovery)