



[Home](#) [Licensing](#) [Patents](#) [Articles](#)

Slope-Constrained Speculative Simulation

by [Nick Clark](#) | Published March 27, 2026 | [PDF](#)

Planning graph construction constrained by trust slope validation at each speculative step, preventing speculation from diverging beyond trusted behavioral boundaries.

What It Is

Planning graph construction constrained by trust slope validation at each speculative step, preventing speculation from diverging beyond trusted behavioral boundaries.. This mechanism is defined in Chapter 4 of the cognition patent as a structural component of the agent's cognitive architecture, operating through deterministic evaluation rather than heuristic approximation.

Every aspect of this mechanism is specified declaratively in the agent's policy reference, making it auditable, reproducible, and governable without requiring access to the agent's internal decision-making process.

Why It Matters

Without slope-constrained speculative simulation, speculative planning either contaminates verified execution state or is absent entirely. Current systems either allow unrestricted speculation that produces hallucinated commitments, or suppress speculation completely, losing the ability to plan ahead. The structural challenge is maintaining the value of speculation while preventing its risks.

In multi-agent systems, this gap is amplified. Agents that cannot speculate safely cannot coordinate through shared planning. Agents that speculate without containment contaminate the system with unverified commitments that other agents treat as facts. Both outcomes undermine the reliability that autonomous operation requires.

How It Works Structurally

As defined in Chapter 4 of the cognition patent, slope-constrained speculative simulation operates through a deterministic evaluation function embedded within the agent's cognitive architecture. The function receives structured inputs from the agent's canonical fields and produces outputs that govern subsequent processing stages. Every input, computation step, and output is recorded in the agent's lineage, ensuring complete reproducibility.

Planning graph structures are maintained in dedicated memory regions separate from verified execution memory. The containment boundary is enforced at the data structure level, not by convention. Branch lifecycle operations including creation, evaluation, promotion, pruning, and dormancy are governed by policy-defined rules that apply uniformly across all planning operations.

What It Enables

This mechanism enables agents that plan ahead without contaminating their operational state. Speculative exploration produces richer decision-making without the risks that unconstrained speculation introduces. Multi-agent systems gain coordination capabilities through shared planning without requiring centralized schedulers.

Because this mechanism is policy-governed and deterministic, it can be formally analyzed, audited, and certified. Regulatory compliance is demonstrable through structural analysis rather than solely through empirical testing. Different domains can tune the mechanism's parameters through policy configuration without requiring architectural changes, making the same structural capability applicable to autonomous vehicles, companion AI, therapeutic agents, and enterprise systems.

[Forecasting Engine All 21 steps →](#)

Plan before you act. Contain speculation. Promote only what passes.

Primary Technical Disclosure

[◦ Forecasting and Executive Graphs in Autonomous Cognitive Systems](#)

Secondary Technical

[◦ Planning Graphs as First-Class Cognitive Structures](#)◦ [Containment Layer and Delusion Boundary](#)◦ [Branch Classification System](#)◦ [Personality Field as Structural Modifier](#)◦ [Executive Engine Multi-Agent Graph Aggregation](#)◦ [Branch Dormancy and Deferred Promotion](#)◦ [Proactive Speculative Maintenance \(Dream State\)](#)◦ [Planning Graph Archival for Cognitive Forensics](#)◦ [Cross-Agent Planning Graph Visibility](#)● [Slope-Constrained Speculative Simulation](#)◦ [Structural Separation From Verified Memory](#)◦ [Forecasting Engine Architecture](#)◦ [Forecasting Execution Cycle](#)◦ [Emotional Modulation of Planning](#)◦ [Executive Graph Conflict Resolution](#)◦ [Planning Graph Delegation and Forking](#)◦ [Temporal Anchoring and Lifecycle Management](#)◦ [Forecasting as Coordination Primitive](#)◦ [Forecasting-Modulated Discovery Traversal](#)◦ [Forecasting as Confidence Input](#)◦ [Integrity-Constrained Forecasting](#)◦ [Forecasting for Training Curriculum](#)◦ [Biological Signal to Forecasting Coupling](#)◦ [Substrate-Agnostic Forecasting Deployment](#)

Applications (General)

[◦ Surgical Robot Planning Through Governed Speculative Branches](#)◦ [Defense Tactical Planning With Contained Speculation](#)◦ [Forecasting Engine for Logistics Planning](#)◦ [Forecasting Engine for Disaster Response Planning](#)◦ [Forecasting Engine for Financial Portfolio Planning](#)◦ [Forecasting Engine for Construction Project Planning](#)◦ [Forecasting Engine for Epidemic Response Planning](#)◦ [Forecasting Engine for Space Mission Planning](#)

Applications (Specific)

[◦ da Vinci Plans Trajectories, Not Consequences](#)◦ [Anduril's Lattice Plans Missions Without Speculative Containment](#)◦ [Boston Dynamics Plans Motion, Not Missions](#)◦ [Shield AI's Hivemind Cannot Contain Its Own Speculation](#)◦ [MuJoCo Simulates Physics Without Planning Governance](#)◦ [NVIDIA Isaac Sim Renders Worlds Without Governing Plans](#)◦ [Unity ML-Agents Trains Without Governing Speculation](#)◦ [Gazebo Simulates Robots Without Governing Their Plans](#)◦ [Drake Optimizes Trajectories Without Governing Planning Structures](#)◦ [robosuite Benchmarks Manipulation Without Governing Plans](#)

[Forecasting Engine overview →](#)

AQ

deterministic

autonomy

Legal

Subject to one or more pending U.S. and international patent applications, see [Patents](#) for the current list and status. No license, express or implied, is granted. Any use requires a separate written agreement—see [Licensing](#). Patent applications referenced on this site are pending. Claim scope, if any, is subject to examination and may issue in altered form or not at all. See [Legal](#) for terms and conditions.

Adaptive Query™ is a trademark of Nicholas Clark. U.S. federal registration is pending. federal registration. AQ™, AQ Inside™, Adaptive Index™, Adaptive Network™, Semantic Agent™, @AQ™, AQID™, and Adaptive Coin™ are used as trademarks in connection with the Adaptive Query platform and brand. Other names may be trademarks of their respective owners.

Platform operated by Adaptive Query LLC, which provides patent and trademark licensing services. Copyright © 2025-2026 Nicholas Clark. All rights reserved.



-
- [Inventive Steps](#)
- [Licensing](#)
- [Patents](#)
- [Articles](#)
- [Legal](#)
- [Opportunities](#)
- [Sitemap](#)



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



[Invented by Nick Clark](#) | Founding Investors: Devin Wilkie