Graceful Degradation With Active-Domain Registry

by Nick Clark | Published March 27, 2026 | PDF

Not all deployments support all cognitive domain fields. An edge device may lack the resources for full forecasting. A rapid-response system may operate without integrity tracking. The active-domain registry tracks which cognitive fields are operational and adjusts confidence proportionally. An agent operating without forecasting knows it is operating without forecasting, and its confidence reflects this limitation.

What It Is

The active-domain registry is a real-time record of which cognitive domain fields are currently operational for an agent. When fields are unavailable due to resource constraints, deployment limitations, or intentional configuration, the registry records which fields are missing and adjusts the agent's confidence computation to reflect the reduced cognitive coverage.

Why It Matters

An agent that operates with missing cognitive fields without accounting for their absence overestimates its own reliability. It makes decisions without integrity checking and remains fully confident. The active-domain registry prevents this by ensuring that missing fields reduce confidence proportionally, preventing the agent from operating beyond its actual cognitive capabilities.

How It Works

The registry is updated whenever a cognitive field becomes available or unavailable. The confidence computation includes a coverage factor that decreases as active fields decrease. The degradation is proportional: losing a non-critical field produces a small confidence reduction, while losing a safety-critical field produces a large reduction that may trigger non-executing cognitive mode.

The registry also adjusts other cognitive computations: missing forecasting disables planning-dependent decisions; missing integrity tracking limits delegation authority.

What It Enables

Graceful degradation enables agents that operate safely across a wide range of deployment conditions. A full-featured deployment operates at maximum capability. A constrained deployment operates with proportionally reduced capability and confidence. The agent always knows its own limitations and adjusts its behavior accordingly, preventing overconfident operation with incomplete cognitive coverage.

Human-Relatable Intelligence All 21 steps →

The most human-like computer ever built.

Primary Technical Disclosure
○ Human-Relatable Computable Intelligence
Secondary Technical
○ The Cross-Primitive Coherence Engine○ Narrative Identity as Compressed Self-Model○ Ecosystem Participation Credentials From Cognitive History○ Anonymized Governance Telemetry Aggregation○ The Coherence Control Loop: Detection, Recording, Restoration○ The Complete Thirteen-Stage Mutation Lifecycle○ Ten Conditions for Human-Relatable Behavior● Graceful Degradation With Active-Domain Registry○ Architectural Inversion: Agent Carries State, Substrate Provides Environment○ Sequential Cascade Structures in Cross-Primitive Coherence○ Conformity Attestation: Verifiable Architectural Compliance
Applications (General)
○ Why AI 2.0 Requires Structural Cognition, Not Better Prompts○ The Compliance Case for Cognitive Architecture Under the EU AI Act○ Why Alignment Is Insufficient for Trustworthy AI○ Enterprise Trust Through Architecture, Not Alignment○ Insurance Liability Reduction Through Human-Relatable AI○ Building Consumer Trust in AI Through Cognitive Relatability○ Regulatory Future-Proofing Through Human-Relatable Architecture○ Competitive Differentiation Through Cognitive Architecture
Applications (Specific)
○ OpenAI's Alignment Approach Is Missing Structural Isomorphism○ Constitutional AI Defines Principles Without Cognitive Architecture○ DeepMind's Safety Research Lacks Cognitive Isomorphism○ Meta's Open AI Safety Is Missing Cognitive Architecture○ Inflection AI Simulates Empathy Without Structural Coherence○ Adept AI Automates Actions Without Structural Integrity○ Covariant Trains Robot Dexterity Without Cognitive Coherence○ Sanctuary AI Builds Humanoid Form Without Human-Relatable Cognition○ Aleph Alpha Offers Sovereign AI Without Structural Coherence○ Mistral AI Optimizes Efficiency Without Architectural Coherence
Human-Relatable Intelligence overview →
AQ
deterministic
autonomy

Last updated: 2026-03-03

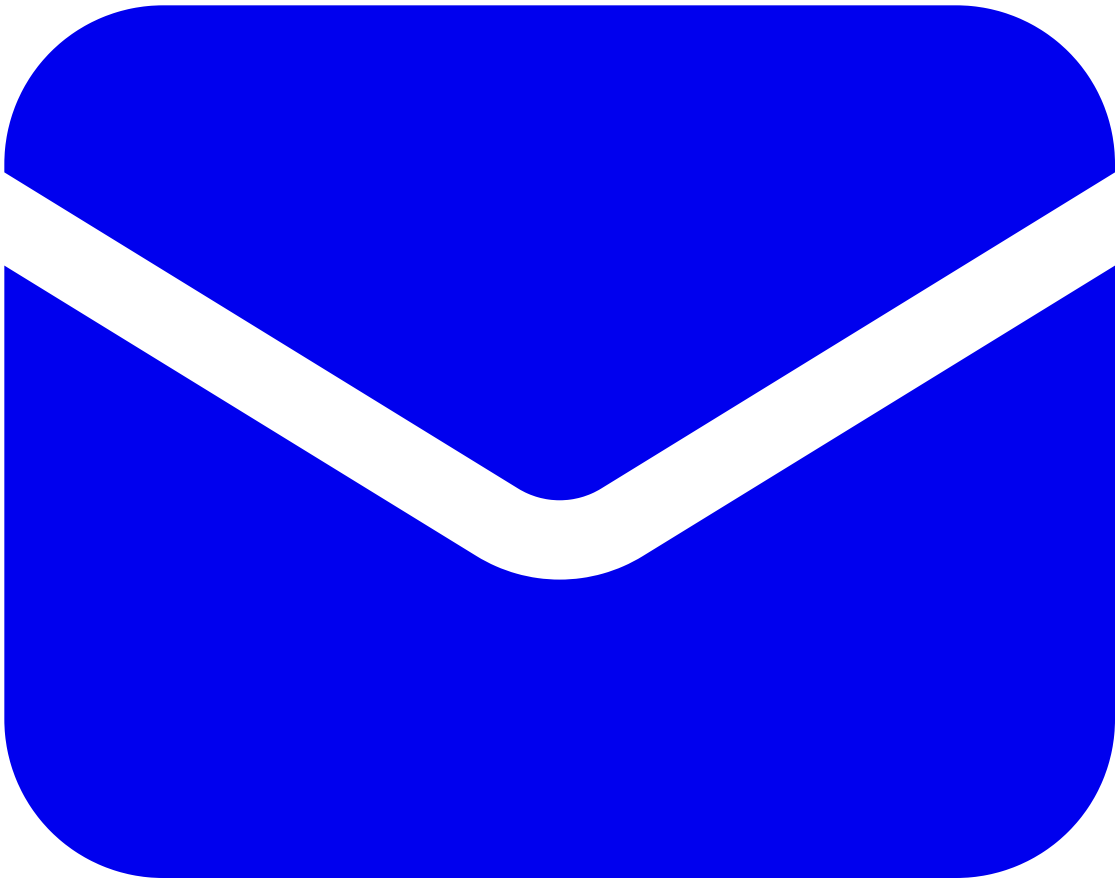

-
- nick@qu3ry.net
- 72 28 14 36 01

Invented by Nick Clark | Founding Investors: Devin Wilkie