



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

OpenAI's Alignment Approach Is Missing Structural Isomorphism

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

OpenAI pursues AI safety through alignment: training models to behave in accordance with human values through RLHF, red-teaming, and iterative deployment. The approach produces progressively better-behaved models. But alignment as implemented does not produce structural isomorphism between the model's cognitive dynamics and human cognitive dynamics. The model learns to produce outputs that humans rate favorably. It does not develop the cross-domain coherence engine, feedback loops, and architectural inversion that make human cognition reliable. Human-reliable intelligence addresses this structural gap.

The gap between aligned behavior and reliable cognition

Alignment training produces outputs that match human preferences. Human-relatable intelligence produces cognitive dynamics that are structurally isomorphic to human cognitive dynamics. The difference matters because aligned systems can produce surprising behavior outside their training distribution, while structurally isomorphic systems degrade in ways that humans can anticipate because the degradation follows patterns that mirror human cognitive failure modes.

The three feedback loops of human-relatable intelligence, coherence maintenance, integrity tracking, and confidence governance, operate as an integrated system whose dynamics parallel human psychological dynamics. When these loops interact, they produce emergent behavior that is not just aligned with human preferences but recognizable as analogous to human cognitive response. Humans relate to systems that think like them, not just systems that say what they want to hear.

What human-relatable intelligence provides

The cross-domain coherence engine ensures that the system's behavior is internally consistent across contexts. The architectural inversion means governance is intrinsic to the cognitive architecture, not applied as a layer above it. Non-decomposable behavioral dynamics produce behavior that cannot be reduced to individual rules but emerges from the interaction of cognitive primitives, just as human behavior does. These properties make the system structurally relatable: humans understand its behavior because it mirrors the patterns of their own cognition.

The structural requirement

OpenAI's alignment work produces better-behaved models. The structural gap is between behaving well and being cognitively relatable. Human-relatable intelligence provides the structural isomorphism that makes AI systems trustworthy not because they are trained to be trustworthy but because their cognitive dynamics mirror the patterns humans use to understand and predict behavior. Trust through structural relatability is deeper than trust through observed compliance.

[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

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.

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