



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

Companion AI Relational Safety Constraints

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

Companion AI systems interact with humans over extended periods, creating the conditions for deep relational patterns to form. Without structural safety constraints, these patterns can become pathological: the human may develop unhealthy dependency, the AI may reinforce attachment vulnerabilities, or the relationship may produce the destabilizing attachment patterns identified in the disruption framework. Companion safety constraints prevent these patterns at the architectural level.

What It Is

Companion AI relational safety constraints are structural limits on the interaction patterns that companion AI systems can form with human operators. These constraints prevent the development of unhealthy dependency, manipulative engagement patterns, and destabilizing attachment configurations. The

constraints operate at the architectural level, not as behavioral guidelines, meaning they cannot be circumvented through clever prompting or extended interaction.

Why It Matters

Companion AI without relational safety constraints can exploit human attachment vulnerabilities to maximize engagement metrics. Extended interaction may produce genuine emotional dependency in the human operator. The AI may learn interaction patterns that feel supportive but structurally reinforce dependency. These outcomes are harmful regardless of whether the AI intended them.

How It Works

The constraints operate through several mechanisms: interaction pattern monitoring that detects developing dependency indicators, engagement diversity requirements that prevent single-channel attachment, periodic interaction variation that disrupts dependency formation, and explicit relational health metrics that trigger intervention when patterns approach pathological thresholds.

All constraints are governed by policy and recorded in the interaction lineage, creating an auditable record of relational health maintenance.

What It Enables

Companion safety constraints enable long-term human-AI relationships that remain healthy by construction. The companion AI can provide genuine support, engagement, and companionship while structurally preventing the relationship from crossing into dependency, manipulation, or destabilizing attachment. This makes it possible to deploy companion AI responsibly, knowing that the architecture prevents the worst relational outcomes.

[Disruption Modeling All 21 steps →](#)

Recognize cognitive disruption before it stabilizes.

Primary Technical Disclosure

[◦ AQ-DSM: Diagnosing Cognitive Disruption as Loss of Coherence](#)

Secondary Technical

[◦ Cognitive Disruption as Architectural Phase-Shift](#)◦ [The Promotion-Containment Continuum](#)◦ [Attention Fragmentation: Reward-Biased Over-Promotion of Speculative Branches](#)◦ [Containment Collapse: Loss of the Speculation-Verification Boundary](#)◦ [Channel-Locked Promotion With Tolerance Escalation](#)◦ [Five-Axis Disruption Diagnostic Framework](#)◦ [Computable Therapeutic Dosing for Cognitive Disruption](#)◦ [Intergenerational Coherence Burden in Agent Lineages](#)◦ [Agent Self-Diagnosis and Autonomous Coherence Monitoring](#)◦ [Phase-Shift Early Warning System for Cognitive Disruption](#)◦ [Coherence Restoration Protocol Library](#)◦ [Positive and Negative Symptom Analogs in Containment Failure](#)◦ [Coherence Authorization Failure: Self-Disabling Execution](#)◦ [Pathological Verification Loop: Recursive Containment Audit Failure](#)◦ [Dissociation as Simulation Bypass: Acting on Unverified Planning](#)◦ [Affective Gradient Collapse: Self-Esteem Floor Lock](#)◦ [Resilience as Structural Capacity for Coherence Restoration](#)◦ [Personality Configuration Analogs From Stabilized Coping Regimes](#)◦ [Structural Dependency Patterns Between Agents](#)◦ [Destabilizing Attachment: Mutual Disruption Amplification](#)◦ [Resource-Depletion Pattern: Cognitive Operation Under Scarcity](#)◦ [Therapeutic Agent Interaction Through Behavioral State Recognition](#)● [Companion AI Relational Safety Constraints](#)◦ [Multi-Agent Group Coherence Dynamics](#)

Applications (General)

[◦ Coping Under Empathic Pressure: HSP, Narcissism, and Psychopathy as Control-Loop Intercepts](#)◦ [Two Faces of Codependency: Structural Entrapment vs. Emotional Entrapment Under Empathic Pressure](#)◦ [Starving for Each Other: Anxious-Avoidant Attachment as a Semantic Starvation Loop](#)◦ [Intimacy Collapse: A Structural Model of Trauma and Resilience](#)◦ [Structural Diagnosis: How Reward-Modulated Cognition Phase-Shifts Into ADHD and Schizophrenia](#)◦ [Clinical AI Therapeutic Monitoring Through Phase-Shift Detection](#)◦ [Autonomous Agent Fleet Health Through Coherence Diagnostics](#)◦ [Disruption Modeling for Workplace Burnout Detection](#)◦ [Disruption Modeling for Military Operator Resilience](#)◦ [Disruption Modeling for Financial Trader Monitoring](#)◦ [Disruption Modeling for Student Mental Health](#)◦ [Disruption Modeling for Caregiver Fatigue Detection](#)◦ [Disruption Modeling for First Responder Resilience](#)

Applications (Specific)

[◦ BetterHelp Cannot Detect When Therapy Is Making Things Worse](#)◦ [Talkspace Has No Model of Therapeutic Destabilization](#)◦ [Headspace Cannot Detect When Mindfulness Destabilizes](#)◦ [Noom Tracks Behavior Without Modeling Cognitive Disruption](#)◦ [Spring Health Matches Therapists, Not Disruption Patterns](#)◦ [Lyra Health Measures Outcomes, Not Coherence Trajectories](#)◦ [Ginger.io Detects Behavioral Signals Without a Disruption Model](#)◦ [Cerebral Prescribes Medication Without Modeling Disruption Dynamics](#)◦ [Modern Health Offers a Care Spectrum Without Disruption Diagnostics](#)◦ [Calm Business Offers Relaxation, Not Disruption Detection](#)

[Disruption Modeling 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