

TIER-2 COUNSEL VALIDATION & IP SUFFICIENCY **DISCLOSURE**

Sovereign Principal, Chief Engineer, & Counsel-Only (Non-Enabling | Representation Evaluation)

Disclosing Party:

Tony Green

Founder & Sole Inventor

C.O.R.E. Values-Based Roundtables & Consulting LLC, d/b/a C.O.R.E. AxisAI

1. Purpose of Disclosure

This disclosure is provided solely to allow qualified legal counsel to assess:

- **Patent strength and scope**
- **Trade-secret posture and enforceability**
- **IP defense and misappropriation risk**
- **Transaction readiness for acquisition or licensing**

This Tier-2 Legal Disclosure is provided to prospective acquirers or licensors solely for the purpose of assessing legal sufficiency and determining whether to proceed to **paid Tier-3 diligence**; it shall not constitute buyer diligence, technical disclosure and/or enablement, or transaction negotiation.

Any use of Tier-2 information to influence internal development, governance design, or research direction constitutes prohibited use.

2. About C.O.R.E.

C.O.R.E. is a **consent-based, longitudinal human self-governance telemetry system that produces governance-grade signals** usable by machines, institutions, and societies operating under pressure. C.O.R.E. is not a governance product. It is a **choke point** in the evolution of **autonomous intelligence**. Alignment failure has proven structural not merely technical.

What follows is not a product description, implementation outline, or technical disclosure. It is a constrained view of the implication surface created when such a signal exists.

As AI systems progress toward greater **autonomy, adaptive learning, and agentic behavior**, the primary challenge is no longer intelligence, scale, or optimization. The unresolved **constraint is governance under uncertainty**—how systems regulate themselves when speed exceeds oversight and prediction fails.

No amount of RLHF, constitutional rules, synthetic scenarios, or internal oversight can ever observe value formation, only value expression. **If this data existed, someone would already be using it.** If such a signal existed at scale, it would represent a structurally different governance substrate than any approach currently in use.

C.O.R.E. addresses this constraint externally. It does not modify model internals, inject beliefs, or rely on preference optimization. Instead, it applies patterns derived from observed human self-regulation over time as a live governance layer that persists across model upgrades, architectures, and deployment contexts.

This allows governance to survive:

- **continual learning**
- **agent swarms**
- **recursive improvement**
- **cross-model deployment**

without retraining or architectural dependency.

The existence of this signal creates implications across multiple domains that are currently being approached independently, expensively, and without a shared reference standard.

3. Asset Overview

C.O.R.E. AxisAI is a Human-Values Operating System (HVOS) designed to generate consent-based, structured human moral-reasoning telemetry and apply it to AI alignment and governance through a model-agnostic control layer.

The asset consists of:

- **A patent-pending system (non-publication requested)**
- **A trade-secret protected methodology**
- **A portfolio of 32 copyrighted IP verticals**
- **A repeatable behavioral data instrument**
- **A governance wrapper architecture (black-box disclosed only)**

No implementation details are disclosed at Tier-2.

Implication classes include—but are not limited to:

- **AI governance and alignment beyond preference learning**
- **Autonomous systems restraint under adversarial pressure**
- **Regulatory auditability where intent must be demonstrated**
- **Institutional decision governance under incentive distortion**
- **Long-horizon risk identification and drift detection before failure manifests**

These domains are listed here to convey breadth, not completeness.

A critical implication of C.O.R.E. is early detection of governance deviation—identifying when decision patterns drift outside acceptable human self-regulation envelopes before harm occurs.

This is not failure detection. It is pre-failure governance deviation awareness.

Such a capability repositions governance from reactive control to anticipatory infrastructure.

4. Patent Posture

U.S. Utility Patent Application: (Non-Provisional)

Application No.: 19/403,575

Filed: November 28, 2025

Title: AI-Driven Systems and Methods for Behavioral Telemetry, Moral Reasoning Alignment, and Scalable Values-Based Roundtables

Claims: 48 total (7 independent)

Priority: Provisional No. 63/830,316 (June 25, 2025)

Status: Non-publication requested

High-level claim scope (non-enabling):

- **Structured capture of human values-based decision telemetry**
- **Longitudinal self-judgment and delta tracking**
- **Use of telemetry for AI governance, alignment, and auditability**
- **External, model-agnostic governance enforcement layer**

Claims, embodiments, figures, and prosecution strategy are not disclosed at this tier.

5. Trade-Secret Position

Trade secrets include:

- **Behavioral telemetry capture constraints**
- **Integrity confidence and anti-gaming mechanisms**
- **Delta normalization logic**
- **Governance enforcement logic**
- **Longitudinal weighting concepts**
- **Distribution, consent, and cadence controls**

No formulas, thresholds, weights, schemas, prompts, or code are disclosed.

Trade-secret protection:

- **Tiered disclosure controls**
- **NDA-gated access**

- **Counsel-supervised disclosure**
- **Non-invertible aggregation**
- **Parameter vaulting**

6. Intellectual Property Portfolio

32 Independent Copyrighted Verticals:

- **Each vertical targets a distinct population and pressure environment**
- **Content is not the asset; signal generation is the asset**
- **Vertical descriptions**
- **Operational mechanics remain confidential**

No curricula, session flows, or facilitator mechanics are disclosed at Tier-2.

7. Transaction Posture

The company is structured for:

- **Full acquisition of assets and IP**
- **Potential exclusive licensing under strict control**
- **Founder-controlled transaction process**

Tier-2 disclosure does not include valuation, revenue streams, or buyers list.

8. What Is Explicitly Excluded from Tier-2

The Following Remains Undisclosed:

- **the internal signal structure**
- **dependency ordering**
- **scaling constraints**
- **benchmark construction**
- **replication boundaries**
- **patent claims text or prosecution history**
- **technical architectures or data schemas**
- **algorithms, models, or scoring systems**
- **training methods or datasets**
- **calibration, weights, or thresholds**
- **buyer diligence materials**
- **Tier-3 vault contents**

At this stage, you have visibility into the **implication surface**, not the internal structure.

9. Disclosure Statement

Tier-3 exists to determine whether the internal dependencies actually support these implications, and whether the non-replicability claims hold under scrutiny.

Because confirmation alone changes strategic posture, Tier-3 access is gated, time bound, and conducted as paid diligence.

Tier-3 validates:

- **That the signal category exists as claimed**
- **That it cannot be generated synthetically, internally, or through scaling**
- **That the dependency graph renders non-ownership strategically unsafe**

It does not disclose full mechanics, code, or transferable implementation.

Tier-3 validation is offered only through a paid diligence engagement in the amount of **Two Million-Five Hundred Thousand U.S. Dollars (USD \$2,500,000.00)**.

This gate exists because confirmation of the internal dependency structure alone materially alters strategic posture and creates asymmetric risk if accessed without commitment. The diligence fee functions as a seriousness filter, an alignment mechanism, and a control on information exposure at a point where understanding precedes ownership.

At this stage, the only decision required is whether confirmation of this dependency structure would materially alter your organization's strategic posture.