

Multi-Dimensional Intent Alignment: A Formal Framework for Continuous Coherence Verification

Technical Report v1.0 · Open for Peer Review

Christopher Celaya

Celaya Solutions LLC, El Paso, Texas
hello@celayasolutions.com

April 2026

Technical Report CSR-2026-002

Abstract

This paper formalizes a recurring structural pattern observed across security-critical systems: effective security requires that operator intent be independently specified across multiple dimensions, with system execution permitted only when all dimensions align. I introduce a computable authority function $A(t)$, a non-compensatory coherence composite $C(t)$, and two temporal decay mechanisms. The framework extends prior work by grounding verification in a continuous biological signal, by placing the composite weights under governance rather than fixing them at design time, and by analyzing the game-theoretic attack surface formally. It is implemented as Celaya Chain Protocol, a Layer 3 protocol with live biological attestation via MORTEM. The same mathematical structure appears in cognitive agent prioritization (CLOS), civic accountability infrastructure (EPPE), and legal evidence chain of custody (VERDICT), which suggests a cross-domain coherence principle rather than a domain-specific design choice. Coalition resistance remains an open problem and is stated here formally. Dual-use implications are addressed directly. This report is the framework companion to CSR-2026-001, which develops the same authority function against the trust, reputation, and mechanism design literature.

IEEE Index Terms: intent alignment; multi-dimensional security; coherence verification; biological attestation; blockchain; game theory; temporal decay; proof of coherence

ACM CCS: Security and privacy → Authentication; Computing methodologies → Multi-agent systems; Theory of computation → Algorithmic game theory

1. INTRODUCTION

Security systems fail in predictable ways. Most failures do not occur because a cryptographic primitive was broken. They occur because the system did exactly what it was told while the operator intended something different. The divergence between intent and execution is the fundamental attack surface.

Buterin's analysis of security patterns [1] identifies a structural regularity across effective security systems: they require the operator to specify intent across multiple independent dimensions, and they execute only when those dimensions align. This is not a heuristic. It is a design principle with formal consequences that have not been drawn out.

This paper does three things. It formalizes the multi-dimensional alignment pattern into a computable authority function and a coherence

composite. It extends that framework with temporal decay, biological grounding, and governed parameters. It analyzes the game-theoretic attack surface and identifies coalition resistance as the central open question.

The framework was developed for distributed protocol security, and it turned out to be isomorphic to structures I had already derived independently for cognitive agent prioritization (CLOS), civic accountability (EPPE), and legal evidence integrity (VERDICT). That convergence is itself a finding. It suggests that coherence verification is domain-agnostic, and that the four systems were solving one problem in four vocabularies.

The implementation is live. Eighty three blocks are verified. Biological attestation is streaming. The mathematics is stable. The implications are not, and this paper says so rather than resolving them by assertion.

2. CORE OBSERVATION: SECURITY AS CONTINUOUS DIVERGENCE

2.1 The Fundamental Reframe

Conventional security is binary. Access is granted or denied. This framing obscures the real failure mode. A system can grant access correctly and still execute incorrectly. A system can deny access to a principal holding entirely valid credentials. Binary access control is not security. It is authentication.

Definition 1. *Security is the continuous minimization of divergence between operator intent and system behavior.*

This reframe has immediate consequences. If security is continuous, it cannot be evaluated at a single point in time. It has to be measured across a trajectory. An authority that was valid at authorization time may no longer reflect current

intent. A credential that was earned legitimately may be operating in a context its owner never specified.

2.2 Common Structure Across Security Paradigms

The multi-dimensional alignment pattern appears consistently across independent security domains. Table 1 sets out six instances, including the two systems developed in this work.

System	Axis 1	Axis 2	Alignment condition
Type systems	Code behavior	Type annotations	Must agree at every step
Formal verification	Program code	Mathematical properties	Program satisfies the properties
Transaction simulation	Operator action	Previewed consequences	Operator confirms alignment
Multi-signature recovery	Multiple keys	Multiple authority definitions	All must agree
CLOS priority (this work)	Cognitive axes A, CE, I, CB, S, E	Runtime behavioral state	Geometric mean above threshold
CCP authority (this work)	Intent axes I, R, P, X, B	Temporal continuity	A(t) above the action threshold

Table 1. The multi-dimensional alignment pattern across security paradigms. The generalized form is N independent specifications of intent, with execution permitted only when all of them agree.

The generalized form is neither a majority vote nor a weighted average. It is an intersection. The logical structure is conjunction, and that choice is what makes the rest of the framework behave the way it does.

3. THE AUTHORITY EQUATION

3.1 Formal Definition

The authority function A(t) governs whether a principal holds sufficient coherence to execute a given action at time t.

$$A(t) = C(t) \cdot \exp(-\alpha n) \cdot \exp(-\lambda \tau) \quad (1)$$

Here $C(t)$ is the coherence composite at time t , n is the network size in participants, and τ is the time elapsed since the last alignment verification. The rate α is the founder decay parameter and λ is the staleness decay parameter. Both are governed values rather than constants fixed at design time.

Plain language. A principal's authority is their current coherence score, reduced by how far the network has grown past the point where their contribution was proportionally significant, and reduced again by how long they have gone without demonstrating continued alignment.

$A(t)$ is bounded in $[0, 1]$ for all valid inputs. It decays toward zero in two independent ways: as the network scales, and as time passes without continued alignment.

3.2 The Coherence Composite

$C(t)$ is the weighted geometric mean across five independent axes of verified intent.

$$C(t) = (I^\beta \cdot R^\gamma \cdot P^\delta \cdot X^\epsilon \cdot B^\zeta)^{1/(\beta+\gamma+\delta+\epsilon+\zeta)} \quad (2)$$

Plain language. Five separate measurements of whether a principal is who they claim to be, is behaving as they have behaved, is acting within policy, is telling the truth about what they did, and is physically present. The five are multiplied rather than added, so the composite tracks the weakest one.

Axis	Name	Description
I	Identity tier	Verification strength of the principal's identity claim. Increases with hardware attestation, biometric binding, and multi-factor depth.
R	Reputation composite	Multi-axis behavioral history, including peer attestation, action consistency, and long-run pattern stability.
P	Policy compliance	Degree to which the requested action conforms to applicable governance constraints.
X	Audit consistency	Agreement between claimed actions and independently verifiable records.
B	Biological continuity	Continuous verification that the same biological operator is present, implemented through cardiac signature streaming.

Table 2. The five axes of the coherence composite. Each is measured independently and each carries a governed weight.

3.2.1 The Non-Compensatory Property

The geometric mean enforces a constraint that an arithmetic mean cannot: no axis can compensate for a deficiency in another. This follows directly from the behavior of the product as any factor approaches zero.

$$\text{If any axis } x_i \rightarrow 0, \text{ then } C(t) \rightarrow 0, \text{ and therefore } A(t) \rightarrow 0 \quad (3)$$

This is the formal encoding of redundant intent verification. A principal with perfect identity verification, flawless reputation, and full policy compliance still holds zero authority if biological continuity is not established. There is no trading across dimensions.

The non-compensatory property is what separates this from a weighted scoring system. A high score on any dimension cannot rescue a zero on any other. Authority requires coherence, not performance.

3.3 Isomorphism with CLOS

The Cognitive Life Operating System (CLOS) priority equation, developed independently for agent task scheduling [2], has the same mathematical structure.

$$C = A \cdot CE \cdot I \cdot CB \cdot S \cdot E \quad (4)$$

The multiplicative structure enforces the same non-compensatory behavior. An agent with zero available energy has zero cognitive priority regardless of its other scores. An agent with zero coherence cannot execute regardless of how much compute is available to it.

This convergence was not designed. The same principle was derived from two independent problem domains, distributed protocol security and cognitive agent scheduling, before either derivation was compared with the other. The pattern is not domain-specific. It is structural.

4. TEMPORAL DECAY MECHANISMS

4.1 Founder Decay

Early participants in a network accumulate authority under conditions that do not persist. A founder's influence is legitimate while the network is small and their contribution is proportionally significant. As the network scales, that proportionality breaks down, and authority that was earned honestly at one scale becomes unearned at another.

$$founder\ decay = \exp(-\alpha n) \quad (5)$$

As n increases, the founder term drives $A(t)$ toward zero. This enforces a principle: historical authority does not scale with network growth. Founders are not permanent governors. Their authority decays as the system they built becomes independent of them. The rate α is a governed parameter, adjustable through the constitutional weighting system, so a community may choose rapid or slow founder decay based on domain-specific considerations.

4.2 Staleness Decay

Authority earned through past alignment does not remain valid indefinitely. Intent changes. Context changes. A credential verified six months ago does not certify current intent.

$$staleness\ decay = \exp(-\lambda \tau) \quad (6)$$

Here τ is the time elapsed since the last successful alignment verification. A principal who stops demonstrating coherent behavior loses authority at a rate set by λ . There is no coasting on historical reputation.

This mechanism has a specific implication for applied intelligence alignment. A system operating under this framework cannot accumulate authority without continuous biological attestation. Authority cannot be pre-granted and stored.

4.3 Interaction Effects

The two decay mechanisms are multiplicative and independent. A principal can hold a high $C(t)$ and a low $A(t)$ through staleness alone. A principal can hold a recent attestation and a low $C(t)$ through axis deficiency alone. Because $A(t)$ is a product of three terms, all three have to remain non-negligible for meaningful authority to exist.

This creates a clock on authority. Every principal has to keep demonstrating multi-dimensional coherence to maintain operational capacity. The system does not grant and forget. It grants and watches.

5. BIOLOGICAL GROUND

TRUTH: THE B AXIS

5.1 The Credential Transfer Problem

Every digital identity system faces the same fundamental vulnerability. Credentials can be transferred. A private key can be copied. A password can be shared. A hardware token can be stolen or

handed over. The credential proves possession. It does not prove presence.

The B axis addresses this at the protocol level. The biological continuity signal cannot be transferred because it is generated by a specific biological system, continuously, in real time. Presence is the thing being measured, and presence is not a bearer token.

5.2 MORTEM: Cardiac Attestation

The B axis is implemented through MORTEM, a research instrument that streams live pacemaker heartbeat data to a public ledger. The implementation treats a Class III medical device as a non-transferable cryptographic primitive.

Cardiac attestation has four properties that make it usable as ground truth. It is *continuous*: the signal is present for as long as the operator is alive and present, and it cannot be pre-recorded without producing a detectable discontinuity. It is *unique*: cardiac signatures are biometrically distinctive, and the rhythm, interval distribution, and morphology of a pacemaker-regulated signal are tied to that implant and that patient. It is *non-delegable*: a pacemaker cannot be handed to somebody else, because the credential is physically embedded in the operator. And it is *sequentially committed*: each heartbeat is cryptographically chained to the one before it, which makes retroactive forgery detectable rather than merely difficult.

Every other axis can in principle be socially engineered, credential-farmed, or gradually inflated through coordinated behavior. Identity tiers can be manufactured. Reputation can be built by a patient adversary. Policy compliance can be gamed by an actor who reads the policy carefully. Audit consistency can be maintained by an actor who controls the audit surface. B cannot be produced by any of those means. This is not an incremental improvement in identity verification. It is a change of category.

5.3 Implementation Status

The MORTEM biological attestation layer is live on a public devnet with eight autonomous witness agents providing continuous heartbeat attestation. Eighty three blocks are verified. Twelve adversarial threat vectors have been identified with documented recovery paths. Nine property tests pass with deterministic replay confirmed.

Mechanism-level implementation detail is deliberately withheld from this publication, for the dual-use reasons set out in Section 9. The architectural claim is what is offered for review. Researchers interested in a collaborative evaluation of the implementation may contact the author.

MORTEM is not a mortality instrument that happens to interact with the protocol. MORTEM is the axis that makes the whole system non-gameable at the root.

6. GAME-THEORETIC ANALYSIS

6.1 Attack Vector Taxonomy

Three primary attack classes act against C(t). Each targets a different structural assumption of the composite.

6.1.1 Coordinated Attestation Inflation

Mutually reinforcing attestation loops between coordinated principals can inflate the reputation axis artificially. Principal A attests to Principal B, B attests to A, and both accumulate reputation through circular endorsement that no outside party ever validated.

The mitigation is attester credibility weighting. The weight assigned to an attestation is a function of the attester's own C(t), so a circular loop between low-credibility principals produces diminishing marginal reputation gain. This does not resolve the problem. A sufficiently large coalition of high-credibility principals can engage in mutual attestation

that passes the weighting filter. The mitigation reduces the attack surface. It does not eliminate it.

6.1.2 Temporal Gaming

A patient adversary can build authority slowly across all axes and then exploit it during a high-impact window. The adversary accumulates $C(t)$ through genuinely legitimate behavior over months or years, then executes a targeted action before staleness decay becomes effective.

The optimal exploitation window is the interval in which $C(t)$ is high and τ is low. The adversary has to time the action inside that window, and the λ parameter controls how wide it is. A high λ compresses the window and forces the adversary to act faster. A low λ widens it. Selecting λ is a domain-specific calibration problem that requires empirical observation of the threat population's time preferences, and the framework does not specify how to perform that calibration.

6.1.3 Axis Correlation Attacks

The geometric mean in $C(t)$ assumes independence between axes. If the axes are correlated in practice, a capability that allows manipulation of one may enable manipulation of others, and the composite silently becomes a narrower measurement than it appears to be.

Consider an adversary who can forge audit records. That same adversary may also be able to construct misleading policy compliance evidence, because in many deployments audit records and compliance evidence are produced by related systems. Compromising one compromises both, and the product term that was supposed to catch the deficiency never sees it.

This attack requires identifying hidden correlation structure between axes. The defense requires designing axes to be causally independent at the data source level, not merely conceptually independent at the specification level. Empirical measurement of

axis correlation in deployed systems is listed in Section 11 as open work.

6.2 The Coalition Resistance Problem

The central open problem is coalition resistance. The current formulation evaluates $C(t)$ against an absolute threshold. A coordinated coalition that executes mutual attestation and aligned behavior across all axes simultaneously can, in principle, meet that threshold collectively without any individual member deserving it.

Open Problem 1. *Let S be a coalition of principals. Can S coordinate behavior such that every member i has $C_i(t)$ above the threshold, despite no individual member holding genuine multi-dimensional coherence?*

A candidate resolution is relative authority scoring. Instead of evaluating $C(t)$ against an absolute threshold, evaluate it against the distribution of $C(t)$ across all current participants.

$$A_{rel}(t, i) = C_i(t) / E[C(t)] \quad (7)$$

Plain language. A principal is scored against the field rather than against a fixed bar, so inflating everyone's score at once moves nobody up.

Under relative scoring, coordinated coalition inflation becomes a treadmill. If all coalition members inflate together, their relative scores are unchanged. A new participant entering the network with genuine coherence achieves a competitive score without coordinating with anybody. The coalition gains nothing from mutual inflation, which is the same payoff restructuring that the quadratic voting literature applies to Sybil resistance [3].

Three questions remain open. Does relative scoring create perverse incentives during network formation, when the reference distribution is thin and a small number of participants set the mean? How should relative scoring interact with the temporal

decay terms, which are absolute quantities? And what is the equilibrium strategy for a rational coalition under relative scoring, given that the coalition can also choose to shrink the reference population rather than inflate against it?

7. APPLICATION: CONTINUOUS BIOLOGICAL ALIGNMENT

7.1 The Limits of Training-Time Alignment

Current alignment approaches for machine systems operate primarily at training time. Reinforcement learning from human feedback, constitutional methods, and preference modeling all attempt to encode human intent into model weights before deployment [6], [7]. The model then acts on that pre-encoded intent without ongoing human authorization.

This architecture has a structural weakness. It cannot respond to intent that changes after training. Human values drift. Context changes. The specific action a human would authorize in a given situation cannot be fully specified in advance, and a system whose only representation of intent is a set of weights has no channel through which the update could arrive.

7.2 Runtime Intent Attestation

The framework provides an alternative: continuous biological attestation as a runtime authority signal. A system governed by $A(t)$ can only execute actions while the human operator's biological state is continuously attesting to presence and intent. Table 3 places this against existing approaches.

Approach	Intent specification	Ongoing verification
Learning from human feedback	Training time	None
Constitutional methods	Training time	None
Human in the loop	Per-action approval	Per action
Coherence authority (this work)	Multi-axis, continuous	Biological attestation

Table 3. Where intent is specified and whether it is verified again at runtime. The distinguishing property of the present framework is the second column, not the first.

7.3 Properties of Biologically Grounded Alignment

An autonomous system operating under this authority function with MORTEM attestation has four properties that follow from the mathematics rather than from policy.

Liveness-conditional authority. The system cannot take consequential actions if the human operator is dead, unconscious, or absent. Authority does not persist beyond biological presence, because the B axis goes to zero and the composite follows it.

Non-coercible design. Because the B axis requires a continuous voluntary biological signal, coerced operation is in principle detectable. Stress physiological signatures differ from baseline. This is a property of the signal, and its reliability in adversarial conditions is not yet established empirically.

No authority storage. Authority cannot be pre-granted and cached. Each action requires a fresh $A(t)$ evaluation against a current B axis reading, which means an agent cannot bank authorization during a period of presence and spend it later.

Auditable continuity. Every action is linked to a specific heartbeat attestation on the ledger. The biological state at the moment of authorization is permanently recorded, which is what makes the first

three properties checkable after the fact rather than merely asserted.

CLOS, the 37-agent cognitive operating system, is the natural integration target. CLOS agents operating under this authority function would execute only while the biological attestation layer is active, which produces a system that cannot act autonomously in any consequential sense while the human is absent.

8. CROSS-INSTRUMENT INTEGRATION

8.1 EPPE: Civic Accountability

The El Paso Proof Engine anchors civic records, water system data, and grid performance to an immutable chain. Raw anchoring without a trust layer produces timestamped data and nothing more. The authority function gives the engine a credibility weighting for its attestors.

A city councilmember's attestation is not equal to a random wallet. A civil engineer's attestation on infrastructure data is not equal to an anonymous claim. C(t) makes that distinction computable and transparent, because the authority of a claim becomes a function of the measurable coherence of the claimant rather than their self-reported credentials. That turns a record system into a veracity layer, where public records carry not only timestamps but authority scores that can be inspected and disputed.

8.2 VERDICT: Evidence Chain of Custody

VERDICT is a four-agent legal intelligence pipeline for evidence processing and timeline reconstruction. A document that arrives with an authority score attached to its originating agent is a materially different object from an unsigned claim.

The chain-of-custody problem in legal contexts maps directly onto the B axis: proving that the person who created a document was who they claimed to be, was present at the time of creation, and has not had their credentials transferred or spoofed. Attestation at document creation time

provides a biometrically grounded proof of authorship for which current legal process has no direct equivalent. Whether such a proof should be admissible, and under what governance, is a question for that field rather than this one.

8.3 The Coherence Principle

Four systems, formulated for four unrelated problems, converged on the same mathematical structure. That is the finding this section reports.

Coherence verification is the common structure underlying effective security, cognitive scheduling, civic accountability, and legal evidence integrity. The domain-specific implementations differ. The mathematical form does not.

This has architectural consequences beyond the observation itself. Systems built on the principle should be able to share verification infrastructure, exchange authority signals, and compose coherence claims across domains. A principal verified at the biological layer in one context carries that verification into every context that shares the framework, which is an interoperability property rather than a convenience.

9. ETHICAL CONSIDERATIONS

9.1 The Dual-Use Problem

This framework is domain-agnostic, and that is the problem. The same mechanism that enforces security in distributed systems can enforce compliance in population management systems. The same biological attestation that proves a researcher's presence can track a worker's every movement. The same coherence scoring that weights civic attestation can weight social credit.

This is not a theoretical concern held at arm's length. The implementation is complete. The mathematics works. The question of whether it should be deployed is not answerable by the

mathematics, and nothing in the framework answers it.

9.2 The Unchallengeable System Problem

A system that cannot be gamed may become a system that cannot be challenged. If coherence verification is the basis for authority, and authority is required to modify the coherence verification system, then the system is recursively self-protecting against exactly the challenges that might reveal its failure modes.

This is not a defect that can be patched. It is a structural property of any sufficiently robust authority system. The constitutional governance layer, with its adjustable weights and its governing process, is the designed mitigation. A constitutional process governed by coherence scores is itself subject to the critique, and I do not claim to have escaped it.

9.3 Biological Data as an Authority Signal

Using continuous cardiac data as a protocol primitive raises questions that existing frameworks for biometric data governance do not address. The data is not held in a database that can be deleted. It is permanently committed to a ledger. A heartbeat attestation recorded in 2026 is recoverable in 2040.

That permanence is architecturally valuable and ethically significant at the same time, and for the same reason. The property that makes the attestation a reliable ground truth is the property that makes it an irrevocable disclosure. A principal who participates in a biologically grounded protocol cannot later withdraw their biological data from the record. Voluntary, informed consent is a necessary precondition for any deployment, and the architecture does not enforce it. Governance does.

9.4 Disclosure Commitment

This paper is published with full disclosure of these concerns rather than an appendix that manages them.

The framework is presented for technical validation, for game-theoretic analysis, and for ethical evaluation on equal footing. I do not present it as unambiguously beneficial. I present it as technically novel with implications that are not fully understood.

10. RELATED WORK

The framework draws on and extends several established lines of research. Table 4 states what each contributes and where the present work departs from it.

Framework	Structural similarity	Key difference
Quadratic voting [3]	Non-linear authority scaling	No biological grounding, no temporal decay
Conviction voting [4]	Time-weighted authority accumulation	No multi-axis coherence composite
EigenTrust [5]	Reputation through network propagation	No biological axis, no founder decay
Learning from human feedback [6]	Human intent specification	Training time only, no runtime attestation
Constitutional methods [7]	Principle-based constraint	No biological verification, no temporal decay
Multi-signature schemes	Multi-dimensional authorization	Binary, no continuous decay, no biological axis

Table 4. Structural relationship to prior work. In each case the difference is either the multi-axis non-compensatory composite, the temporal decay terms, or the biological axis.

The geometric mean structure itself has precedent in multi-criteria decision analysis, where it is chosen precisely because it is non-compensatory. Its application as a security primitive, with biological grounding and temporal decay attached, is to the best of my knowledge novel. CSR-2026-001 [10] develops that positioning against the trust and

reputation literature in detail; this report is the framework-side companion to it.

11. CONCLUSION

This paper formalizes a recurring structural pattern in security-critical systems into a computable authority function with three properties: non-compensatory multi-dimensional coherence, continuous temporal decay, and biological ground truth verification.

The framework is live. The mathematics is stable. The game-theoretic properties are partially analyzed, with one open problem, coalition resistance, identified as requiring further work and stated formally in Section 6.2.

The cross-domain convergence on this structure, from four independent derivations in protocol security, cognitive agent scheduling, civic accountability, and legal evidence integrity, suggests a principle that generalizes past any single application domain. That is a claim about structure, not about merit.

Four things remain open. Coalition resistance requires a formal proof or disproof of relative scoring as a solution. Axis independence requires empirical measurement of the correlation between the identity, reputation, policy, and audit axes in deployed systems. Alignment integration requires a formal specification of CLOS operating under this authority function. And the ethical framework requires governance standards for biometric data committed to permanent public ledgers, which do not currently exist.

The framework can enforce alignment. It can also enable social control. A system that cannot be gamed may become unchallengeable. These are not edge cases. They are design properties. This work is published because the implications require more minds than one lab.

REFERENCES

- [1] V. Buterin, "How I think about security," personal blog, 2024. Available: vitalik.eth.limo
- [2] C. Celaya, "CLOS: Cognitive Life Operating System, Priority Equation and Agent Architecture," Celaya Solutions LLC, internal technical specification, El Paso, TX, 2025.
- [3] E. G. Weyl and E. A. Posner, Radical Markets: Uprooting Capitalism and Democracy for a Just Society. Princeton, NJ: Princeton University Press, 2018.
- [4] J. Emmett, "Conviction Voting: A Novel Continuous Decision Making Alternative to Governance," Medium / 1Hive, 2019.
- [5] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The EigenTrust Algorithm for Reputation Management in P2P Networks," in Proc. 12th Int. Conf. World Wide Web (WWW '03), pp. 640–651, ACM, 2003.
- [6] P. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep Reinforcement Learning from Human Preferences," in Proc. NeurIPS, 2017.
- [7] Y. Bai, A. Jones, et al., "Constitutional applied intelligence: Harmlessness from applied intelligence Feedback," Anthropic technical report, arXiv:2212.08073, 2022.
- [8] C. Celaya, "MORTEM: Biological Attestation via Cardiac Signature Streaming," Celaya Solutions LLC, technical report, El Paso, TX, 2026.
- [9] C. Celaya, "Celaya Chain Protocol: Proof of Coherence Consensus," Celaya Solutions LLC, technical report, El Paso, TX, 2026. Available: celayasolutions.com/research/CCP-Scholar-v1
- [10] C. Celaya, "Coherence as a Security Primitive: A Multi-Dimensional Intent Alignment Framework with Biological Ground Truth," Celaya Solutions LLC, technical report CSR-2026-001, El Paso, TX, 2026. Available: celayasolutions.com/research/CSR-2026-001-Scholar-v1

Technical Report CSR-2026-002 | Celaya Solutions Research | El Paso, Texas | April 2026 | Open for peer review. Celaya Solutions LLC is an independent research laboratory building instruments that make coherence visible and actionable across cognitive, civic, and distributed systems. Contact hello@celayasolutions.com.

HOW TO CITE

Celaya, C. (2026). Multi-Dimensional Intent Alignment: A Formal Framework for Continuous Coherence Verification (Technical Report CSR-2026-002). Celaya Solutions Research Lab.

<https://celayasolutions.com/research/CSR-2026-002-Scholar-v1>

COPY CITATION

License: [CC BY 4.0](#)