

4th International Conference on Information Science and Techniques (ISTECH 2026)

November 14–15, 2026, Melbourne, Australia

HARDWARE-VERIFIABLE RECURSIVE INFORMATION ADMISSIBILITY FOR AI AND INFORMATION SYSTEMS: A SEXA FRAMEWORK FOR CONSTRAINT-SAFE DATA VALIDATION, RELATIONAL CONSISTENCY, FAULT LOCALIZATION, DIMENSIONAL COMPRESSION, AND DETERMINISTIC INFORMATION RECONSTRUCTION

Jered McClain

SEXA Institute of Technology and Interdomain Galactic Advisors
SEXAtech.com

CreateNewLand.com

UnifiedFieldSystems.com

spacetimemanhattan@gmail.com

ABSTRACT

This paper formulates Recursive Information Admissibility (RIA), a traceable constraint architecture for multivariate information validation, fault localization, and loss-aware dimensional reduction. Prior mechanical evidence is extended with a controlled four-task information benchmark and an independent public-data novelty study. In the controlled suite, RIA reduced macro false admission to 35.08%, outperforming LOF at 42.75%. Across 17 public novelty tasks and 306 evaluations, however, generic RIA did not lead: false admission was 78.61% versus 73.61% for Isolation Forest, while LOF achieved higher ROC-AUC. The combined result narrows the supported claim: RIA is promising when domain relationships and invariants are explicitly known, but is not demonstrated as a universally superior anomaly detector.

KEYWORDS

Recursive Information Admissibility, Information Validation, Artificial Intelligence, Knowledge Representation, Information Integrity

1. INTRODUCTION

Information science increasingly confronts a problem deeper than the acquisition, storage, transmission, retrieval, or statistical processing of data. Modern computational systems must determine whether information should be permitted to propagate through the system at all.

Artificial-intelligence systems, autonomous decision architectures, digital twins, scientific-computing platforms, knowledge graphs, sensor networks, databases, and distributed information infrastructures combine information originating from multiple variables, sources, times, representational levels, and models. A datum may be individually valid while being incompatible with surrounding information. Two records may each satisfy local constraints while being mutually contradictory. Consecutive states may each be permissible while the transition between them is impossible. A set of apparently valid observations may collectively violate a system-level invariant.

The distinction is therefore between data validity and information admissibility.

$$\Omega_A \subset \Omega_I$$

Let the complete candidate information space be Ω_I and let Ω_A denote the admissible subset. For a candidate information state $I_t \in \Omega_I$, define the admissibility mapping

$$A : \Omega_I \rightarrow \{0,1\}$$

$$A(I_t) = 1 \text{ if } I_t \in \Omega_A; \quad A(I_t) = 0 \text{ if } I_t \notin \Omega_A.$$

The central information-science question addressed here is not merely whether each datum is individually valid, but whether the complete information state satisfies the coupled conditions required for admission into the next computational state.

The architecture originates in the SEXA recursive-admissibility framework. Earlier computational work emphasized deterministic execution, observable intermediate states, explicit failure conditions, dimensional reduction, computational correspondence, and structured validation [1]-[4]. The VLSIE conference formulation organized the framework into validation layers including recursive functional consistency, dimensional and unit consistency, computational correspondence, reduction-recovery, and explicit falsifiability [2].

The framework was subsequently instantiated in a controlled aerospace/mechanical digital-twin benchmark. Ten state variables were evaluated against independently defined ground truth through six sequential Γ -admissibility stages. For every benchmark trial, the nominal state, injected disturbance, expected classification, and Γ output were known, allowing direct measurement of true admissions, true rejections, false admissions, and false rejections.

The present paper performs the next conceptual transformation: from mechanical-state admissibility to information-state admissibility. The mechanical benchmark is not treated as proof of identical performance in arbitrary information domains. Instead, it is used as an executed computational demonstration of the narrower proposition that individually permissible values may nevertheless form an inadmissible collective state.

The primary research question is: Can an ordered recursive-admissibility architecture provide a deterministic information-validation layer in which candidate information is admitted only after surviving acquisition, transition, relational, invariant, persistence, and terminal constraints?

2. BACKGROUND AND COMPUTATIONAL FOUNDATION

2.1. Recursive Admissibility

Recursive admissibility is treated here as an ordered constraint architecture acting before final evaluation. The prior VLSIE work sought to transform recursive mathematical structures into

deterministic operations with identifiable intermediate states [2]. The computational principle relevant to information science is representation $\rightarrow$ constraint evaluation $\rightarrow$ admissible representation rather than representation $\rightarrow$ unconditional propagation.

2.2. Hardware-Verifiable Computation

Hardware-verifiable does not mean that the present information architecture has already been fabricated as specialized hardware. It denotes computational properties required for eventual deterministic implementation: identical inputs should produce identical outputs, intermediate states should remain observable, failure conditions should be explicitly identifiable, recursive evaluation should remain bounded, and state transformations should be reproducible [2].

2.3. From Mechanical States to Information States

The mechanical implementation evaluated ten variables describing a coupled state. Its motivating problem was that individually permissible measurements could collectively violate dynamical, structural, energetic, temporal, or operational relationships. The same mathematical distinction occurs in information systems.

$$I_t = [x_{1,t}, x_{2,t}, \dots, x_{n,t}]^T$$

Each x_j may represent a sensor value, database attribute, semantic assertion, feature, state variable, model output, retrieved fact, symbolic proposition, or another computational quantity. An information architecture must distinguish the validity of x_j individually from the admissibility of I_t collectively.

3. SCALAR VALIDITY VERSUS INFORMATION ADMISSIBILITY

3.1. Independent Validation

Suppose each information variable has an independently acceptable domain D_j . A conventional independent validator may define $b_j(I_t)=1$ when $x_{j,t} \in D_j$ and zero otherwise.

$$B(I_t) = \prod_{j=1}^n b_j(I_t)$$

$B(I_t)=1$ establishes membership in the Cartesian product $D_1 \times D_2 \times \dots \times D_n$. It does not establish membership in the admissible subset Ω_A . In general,

$$\Omega_A \subset D_1 \times D_2 \times \dots \times D_n.$$

Component validity is therefore necessary in many systems but is not sufficient for collective admissibility.

3.2. Hidden Inconsistency

Consider two individually valid quantities x_1 and x_2 with required relationship $x_2=f(x_1)$. Both may satisfy their local domains while $|x_2-f(x_1)|>\delta$. A scalar validator admits the pair; a relational validator rejects it. This distinction becomes more consequential as dimensionality and heterogeneity increase.

3.3. Experimental Motivation

The existing held-out benchmark provides a direct computational example. Its independent scalar-threshold baseline accepted every nominal state but detected only 14.45% of faulted states. The Γ architecture achieved 95.70% fault sensitivity. False admission fell from 85.55% to 4.30%. These results do not prove identical performance in arbitrary information domains; they demonstrate the computational value of evaluating relationships that independent limits discard.

4. RECURSIVE INFORMATION ADMISSIBILITY

4.1. Ordered State Evaluation

Let the candidate information state entering the architecture be $I_t^{(0)} = I_t$. Define the ordered chain

$$\Gamma = \gamma_6 \circ \gamma_5 \circ \gamma_4 \circ \gamma_3 \circ \gamma_2 \circ \gamma_1.$$

$$I_t^{(0)} \rightarrow \gamma_1 I_t^{(1)} \rightarrow \gamma_2 I_t^{(2)} \rightarrow \gamma_3 I_t^{(3)} \rightarrow \gamma_4 I_t^{(4)} \rightarrow \gamma_5 I_t^{(5)} \rightarrow \gamma_6 I_t^{(6)}.$$


For the implemented classifier interpretation each stage returns $a_i(I_t) \in \{0,1\}$. The complete information-admissibility decision is

$$A(I_t) = \prod_{i=1}^6 a_i(I_t).$$

A state is admitted only when all six mandatory conditions survive. If any required stage returns zero, the state is rejected for that evaluation cycle.

4.2. Ordering Boundary

The existing benchmark establishes staged evaluation, early rejection, and first-failure localization. Because the final Boolean conjunction of independently evaluated predicates is commutative, it does not by itself establish noncommutativity of stronger state-transforming Γ operators. A stronger test would require states for which $\gamma_j(\gamma_i(I)) \neq \gamma_i(\gamma_j(I))$.



Admit only if all six mandatory constraint classes survive; retain earliest failure stage on rejection.

Figure 1. Six-stage Recursive Information Admissibility architecture and explicit first-failure path.

5. SIX INFORMATION-ADMISSIBILITY STAGES

5.1. γ_1 - Acquisition Integrity

The first stage determines whether the candidate state contains the information necessary for evaluation. For variable x_j with admissible acquisition domain D_j ,

$$a_1(I_t) = \prod_j (\alpha_{j,t} \in D_j).$$

Failures may include missing values, corrupted input, invalid encoding, impossible representation, unavailable required fields, or out-of-domain observations. The governing question is whether the information is validly acquired and representable.

5.2. γ_2 - Transition Consistency

A valid state may follow an invalid transition. Define $\Delta x_{j,t} = x_{j,t} - x_{j,t-1}$ and $v_{j,t} = \Delta x_{j,t} / \Delta t$. For maximum admissible transition rate $r_j^{\max}$,

$$a_2(I_t) = \prod_j \mathbb{1}(|v_{j,t}| \leq r_j^{\max}).$$

The interpretation extends to sequence ordering, timestamps, event causality, state-machine transitions, database histories, and other temporal relationships. In the existing benchmark, transition corruption was detected at 100% by Γ and 0% by the scalar baseline.

5.3. γ_3 - Coupled Relational Consistency

Let $r_k(I_t)$ denote a required relationship among components. For tolerance δ_k , admissibility requires $|r_k(I_t)| \leq \delta_k$. Across m coupled relationships,

$$a_3(I_t) = \prod_{k=1}^m \mathbb{1}(|r_k(I_t)| \leq \delta_k).$$

For information science, r_k may encode semantic relationships, cross-field dependencies, database integrity constraints, knowledge-graph relations, sensor-fusion equations, model residuals, logical compatibility, cross-source agreement, or other domain-specific structural rules. In the held-out benchmark γ_3 accounted for 416,088 first rejections, making coupled consistency the dominant first-failure stage.

5.4. γ_4 - Model or Invariant Consistency

Let $h_k(I_t)$ be a quantity expected to satisfy reference value h_k^* . Define $e_k = h_k(I_t) - h_k^*$. Then

$$a_4(I_t) = \prod_k \mathbb{1}(|e_k| \leq \eta_k).$$

An invariant may be a schema rule, logical identity, normalization requirement, ontology restriction, accounting identity, algorithmic invariant, physical relation, or another model-level constraint.

5.5. γ_5 - Recursive Persistence

Some inconsistencies become meaningful only when accumulated over time. Let $W_t = (I_{t-L+1}, \dots, I_t)$. For residual r_k define $P_t = (1/L) \sum |r_k|$. The persistence condition is

$$a_5(I_t) = \mathbb{1}(P_t \leq P_{\max}).$$

This stage distinguishes isolated excursions from persistent drift or accumulating inconsistency. In the existing benchmark progressive thermal drift achieved 81.42% detection rather than 100%, preserving a measurable weakness for future optimization.

5.6. γ_6 - Terminal Admissibility

The final stage determines whether information surviving all preceding constraints is permitted to enter the target operational domain. Let $S = \{S_1, \dots, S_p\}$ represent terminal constraints. Then

$$a_6(I_t) = \prod_{j=1}^p \mathbb{1}(S_j(I_t) = 1).$$

Terminal constraints may represent application-specific safety, authorization, mission, policy, execution, or downstream operating conditions.

6. FIRST-FAILURE LOCALIZATION

A binary validator returns only acceptance or rejection. RIA additionally records where rejection first occurred. For a rejected information state define

$$k_{\text{fail}} = \min\{i : a_i(I_t) = 0\}.$$

The diagnostic output is $Y_t=(A(I_t),k_{fail})$. For an admissible state $k_{fail}=0$. The architecture therefore distinguishes acquisition, transition, relational, invariant, persistence, and terminal failures. Rejection becomes an observable computational event rather than an opaque terminal label.

This is not equivalent to a complete causal or natural-language explanation of an AI decision. It is a deterministic diagnostic trace identifying the earliest violated implemented constraint class.

7. INFORMATION COMPRESSION AND RECONSTRUCTION

7.1. Compression Problem

The wider SEXA architecture distinguishes a higher-dimensional structural representation from a bounded operational representation. For the present information-science formulation, the general compression problem is written independently of the physical interpretation.

$$\begin{aligned} I_t \in \mathbb{R}^n, \quad C : \mathbb{R}^n &\rightarrow \mathbb{R}^m, \quad m < n, \quad Z_t=C(I_t). \\ R : \mathbb{R}^m &\rightarrow \mathbb{R}^n, \quad \hat{I}_t=R(Z_t). \\ I_t &\rightarrow C \rightarrow Z_t \rightarrow R \rightarrow \hat{I}_t. \end{aligned}$$

7.2. No-Hidden-State Requirement

A reconstruction claim is meaningful only if the original state is not secretly retained. Let Θ_R denote information legitimately available to reconstruction. Then $\hat{I}_t=R(Z_t,\Theta_R)$. A valid reconstruction audit prohibits direct access to I_t , undeclared lossless encoding of I_t , test-set fitting after evaluation begins, or an external identifier that simply retrieves I_t . Any such pathway invalidates the reconstruction claim.

7.3. Reconstruction Error

For reconstructed variable j define normalized error $e_{j,t}=(\hat{I}_{j,t}-x_{j,t})/s_j$, where s_j is the normalization scale. The instantaneous normalized error is

$$E_R(t) = \sqrt{[(1/n) \sum_j e_{j,t}^2]}.$$


Across N states,

$$NRMSE_R = \sqrt{[(1/Nn) \sum_t \sum_j e_{j,t}^2]}.$$

Perfect reconstruction would require $NRMSE_R=0$. The architecture does not assume this result; the observed value must be reported.

7.4. Evidentiary Boundary

The executed mechanical experiment addresses a bounded $10D \rightarrow 5D \rightarrow 10D$ engineering problem. It does not simulate 2,880 independent mechanical sensor variables and therefore cannot be treated as experimental proof of universal $2880D \rightarrow 5D \rightarrow 2880D$ reconstruction. The information-science formulation preserves this distinction.



Evidence is inherited only within its demonstrated scope; the ISTECH contribution is the information-science reformulation.

Figure 2. Evidentiary progression from foundational architecture to the present information-science reformulation.

8. EXPERIMENTAL EVIDENCE

8.1. Existing Held-Out Benchmark

The current empirical evidence comes from the controlled mechanical digital-twin implementation rather than a newly fabricated information-domain experiment. The final held-out evaluation used ten nominal trajectories and ten fault classes across four severity levels and ten held-out seeds. Each trajectory contained 3,000 discrete states, producing 410 trajectories and 1,230,000 held-out state evaluations. Calibration seeds 100-119 and test seeds 200-209 were disjoint.

Parameters were established from separate nominal calibration trajectories and frozen before final evaluation. No threshold was retuned after observing the held-out results. The fault generator was computationally separated from the mechanism attempting to detect the fault.

8.2. Ground Truth

For every trial the original nominal state, injected disturbance, expected classification, and produced Γ classification were known. A disagreement was counted as an error rather than reinterpreted after the fact. This yields a general methodological requirement for future RIA experiments: whenever the domain permits it, ground truth should be defined independently of the classifier.

9. CLASSIFICATION RESULTS

The Γ architecture produced $N_{TA}=587,597$ true admissions, $N_{TR}=612,492$ true rejections, $N_{FA}=27,508$ false admissions, and $N_{FR}=2,403$ false rejections. The resulting performance is shown in Table 1.

Metric	Γ Recursive Admissibility	Independent Scalar Baseline
Classification accuracy	97.57%	55.48%
Balanced accuracy	97.65%	57.22%
Fault sensitivity	95.70%	14.45%
Nominal acceptance	99.59%	100.00%
False-admission rate	4.30%	85.55%

False-rejection rate	0.41%	0.00%
----------------------	-------	-------

Table 1. Held-out classification comparison.

The relative reduction in false admission is

$$(85.55 - 4.30) / 85.55 \times 100 \approx 94.97\%.$$

The importance of this result is structural rather than merely numerical. The scalar baseline does not evaluate coupled-state residuals, transition consistency, persistence, model invariants, or terminal relationships. The recursive architecture evaluates those additional relationships.




Figure 3. Held-out comparison between Γ recursive admissibility and independent scalar thresholding.

10. FAULT-CLASS ANALYSIS

Fault Class	Γ Detection	Scalar Baseline
Acquisition corruption	100.00%	75.00%
Transition corruption	100.00%	0.00%
Vibration inconsistency	75.60%	0.00%
Stress-strain inconsistency	100.00%	0.00%
Energetic inconsistency	100.00%	20.53%
Pressure inconsistency	100.00%	0.00%
Progressive thermal drift	81.42%	0.00%
Rotational overspeed	100.00%	20.08%
Excessive stress	100.00%	25.00%
Compound fault	100.00%	3.86%

Table 2. Fault-class-specific held-out detection performance.

Eight of ten modeled fault classes achieved 100% Γ detection. The two incomplete classes were vibration inconsistency at 75.60% and progressive thermal drift at 81.42%. These values are retained as limitations rather than removed through post-hoc adjustment.




Figure 4. Fault-class detection by Γ architecture and scalar baseline.

11. FIRST-FAILURE RESULTS

Across the complete held-out population, first-stage rejection counts were $\gamma_1=16,000$, $\gamma_2=64,746$, $\gamma_3=416,088$, $\gamma_4=116,059$, $\gamma_5=2,002$, and $\gamma_6=0$. The dominant first-failure stage was therefore γ_3 coupled consistency.

This result is especially relevant to the information-science interpretation. The largest rejection population occurred where relationships among simultaneously represented variables were evaluated rather than where isolated scalar bounds were applied.

This motivates a directly testable hypothesis for future information-domain work: the largest benefit of RIA should occur where individually plausible information becomes contradictory only under relational, temporal, or model-level evaluation.




Figure 5. First-failure localization counts across the held-out benchmark.

12. DETERMINISTIC REPLAY

Reproducibility is a central requirement of the architecture. An identical held-out trajectory was executed twice without changing the initial state, input sequence, model parameters, or frozen Γ thresholds. The complete binary admissibility output and first-failure sequence were identical between repetitions for every evaluated state.

$$D_R = 1$$

when $A^{(1)}(I_t) = A^{(2)}(I_t)$ and $k_{fail}^{(1)}(I_t) = k_{fail}^{(2)}(I_t)$ for every t . The executed benchmark therefore satisfied $D_R = 1$. Deterministic replay establishes reproducibility, not correctness; a deterministic system can reproduce an incorrect decision.

13. APPLICATION TO AI SYSTEMS

RIA suggests a supervisory architecture positioned between candidate information generation and downstream execution. Let an AI system generate candidate state I_t^{AI} . Rather than propagating the state unconditionally, the proposed architecture inserts

$$I_t^{AI} \rightarrow \Gamma_{RIA} \rightarrow \{\text{ADMIT}, \text{REJECT}\}.$$

This does not replace statistical learning. It provides an additional deterministic constraint layer. γ_1 can verify required input and representation, γ_2 can test sequential consistency, γ_3 can test mutual compatibility of asserted facts or variables, γ_4 can enforce domain invariants, γ_5 can test persistence across a retained history, and γ_6 can apply downstream operational requirements.

The prior VLSIE work identified AI as a potential future domain while explicitly stating that measurable advantage remained an open engineering question. The present paper retains that evidentiary boundary: RIA is a testable supervisory architecture for AI, not a demonstrated universal solution to AI reliability.

14. APPLICATION TO KNOWLEDGE REPRESENTATION

Knowledge systems frequently contain information whose admissibility depends on relationships rather than isolated assertions. Let $K = \{p_1, p_2, \dots, p_n\}$ be a candidate knowledge

state. A conventional system may verify whether each p_i has a permissible representation. RIA asks the stronger question of whether K survives all required relational constraints.

$$A(K)=1 \text{ only if every mandatory } r_j(K)=1.$$

The architecture can therefore distinguish valid syntax from valid acquisition, valid propositions from mutually consistent propositions, locally consistent facts from globally invariant knowledge, and temporary consistency from recursively persistent consistency. Knowledge graphs, multi-source scientific databases, symbolic reasoning systems, and retrieval-augmented architectures are direct candidates for future testing. No benchmark on those domains is reported here.

15. INFORMATION PROVENANCE AND TRACEABILITY

Information validation is weakened when the path leading to a decision cannot be reconstructed. RIA associates each state with a trace

$$T_t = \{I_t(0), a_1, I_t(1), a_2, \dots, a_6, I_t(6)\}.$$

For rejected states, evaluation may terminate at k_{fail} while retaining the preceding trace. The candidate information, terminal decision, and computational pathway therefore remain distinct objects. In a deployed information system this trace may serve as an audit record supporting debugging, review, or formal verification.

16. FALSIFICATION CONDITIONS

RIA is intended to fail explicitly rather than remain indefinitely adjustable. The following conditions constitute direct failure criteria for an implementation:

F1 - Acquisition failure: malformed or missing information systematically survives γ_1 , or valid information is systematically rejected.

F2 - Transition failure: temporally or sequentially impossible states cannot be distinguished from admissible transitions.

F3 - Relational failure: states violating independently known coupled relationships are admitted at unacceptable rates.

F4 - Invariant failure: required model or domain invariants are violated without rejection.

F5 - Persistence failure: persistent anomalies remain indistinguishable from admissible transient variation.

F6 - Terminal failure: information violating mandatory operational constraints reaches the admitted state.

F7 - Replay failure: identical inputs and frozen parameters produce different decisions.

F8 - Trace failure: the system cannot identify the earliest implemented constraint stage responsible for rejection.

F9 - Reconstruction leakage: reconstruction depends on hidden retention or retrieval of the original information.

F10 - Generalization failure: performance claimed for a new domain does not reproduce under independently defined ground truth.

17. LIMITATIONS

The evidentiary domain is now three-part: the earlier synthetic mechanical digital-twin benchmark, the controlled native information-state benchmark suite, and the new independent public-data novelty benchmark. The public-data result is deliberately retained despite being unfavorable to the strongest interpretation. It shows that generic RIA does not outperform established unsupervised novelty detectors in aggregate when domain-specific admissibility relationships are absent. This sharply narrows the supported claim and identifies domain-constraint specification as a necessary research variable rather than an optional implementation detail.

The comparison baseline is intentionally simple. Independent scalar thresholding directly tests whether coupled constraints contain information unavailable to isolated limits, but it does not establish superiority over contemporary machine-learning, probabilistic, observer-based, graph-based, model-based, or anomaly-detection methods.

The benchmark contains residual errors: 27,508 false admissions and 2,403 false rejections. Vibration inconsistency achieved 75.60% detection and progressive thermal drift 81.42%. Deterministic replay demonstrates reproducibility rather than truth. First-failure localization identifies the earliest violated implemented condition, not necessarily the ultimate physical or semantic cause.

The current Boolean classifier does not establish the stronger theoretical claim of Γ noncommutativity. Dimensional reconstruction is also bounded by the implemented experiment: the classification results cannot establish universal $10D \rightarrow 5D \rightarrow 10D$ reconstruction, much less physical verification of a 2880-dimensional manifold.

Finally, computational success does not establish every physical interpretation associated with the wider SEXA corpus. The present contribution isolates recursive admissibility as a computational mechanism and keeps that mechanism distinct from unresolved physical claims.

18. NATIVE INFORMATION-SCIENCE BENCHMARK SUITE

To test the information-science hypothesis directly, a second benchmark suite was constructed independently of the mechanical digital-twin experiment. The suite contains four held-out information-state tasks: relational contradictions, invariant violations, temporal/persistence violations, and compound anomalies. Each task contains 6,000 nominal and 6,000 anomalous records. Nominal training (15,000 records) and calibration (5,000 records) were generated separately from all held-out evaluation records. Thresholds were frozen from nominal calibration data before any benchmark labels were evaluated.

Candidate records contained six mutually dependent attributes. The anomaly generator was intentionally designed so that many corrupted records remained inside ordinary marginal ranges while violating cross-field relationships, algebraic invariants, or temporal continuity. RIA therefore evaluated acquisition/range constraints together with calibrated relational residuals and, for the temporal task, a frozen transition gate. No anomaly labels were used to fit the RIA thresholds.

Three established one-class anomaly-detection methods were used as stronger baselines: Isolation Forest [12], One-Class SVM [13], and Local Outlier Factor [14]. Each baseline was trained only on the nominal training set. Its decision threshold was then frozen at the 99.5th percentile of its anomaly score on the nominal calibration set. This design prevents label-driven threshold tuning and makes false admission the primary stress-test metric.

Across the four benchmarks, RIA produced the lowest macro-average false-admission rate: 35.08%, compared with 42.75% for Local Outlier Factor, 56.00% for One-Class SVM, and 99.62% for Isolation Forest. RIA also produced the highest macro-average anomaly sensitivity at 64.92%. The strongest result occurred in the compound benchmark, where RIA reached

79.67% sensitivity with 20.33% false admission. The weakest RIA result occurred in the temporal benchmark, where false admission remained 71.22%; this is an explicit failure mode and indicates that the present temporal gate is substantially less mature than the relational and compound stages.

Relative to the strongest baseline by macro-average false admission, Local Outlier Factor, RIA reduced false admission by approximately 17.94%: $(42.75 - 35.08) / 42.75 \times 100$. This is smaller than the 94.97% reduction observed against scalar thresholding in the mechanical benchmark, which is expected because the present comparison uses substantially stronger nonlinear anomaly-detection baselines.

These results support a narrower claim than universal superiority: when anomalies are encoded primarily in relations among otherwise plausible attributes, an explicit admissibility architecture can improve false-admission control. The benchmark is synthetic and generated within this study; it is not an independent public-dataset validation. Reproduction on external datasets and comparison with additional deep and graph-based detectors remain necessary.

Table 3. False-admission rate across native information-science benchmarks (lower is better).

Benchmark	RIA	Isolation Forest	One-Class SVM	LOF
Relational	23.20%	99.52%	47.80%	31.55%
Invariant	25.55%	99.98%	61.08%	36.95%
Temporal	71.22%	99.33%	76.28%	75.37%
Compound	20.33%	99.63%	38.83%	27.13%




Figure 6. False-admission rate across four held-out information benchmarks.

Table 4. Macro-average performance across the four native information benchmarks.

Method	Accuracy	Sensitivity	Nominal acceptance	False admission
RIA	81.06%	64.92%	97.20%	35.08%
Local Outlier Factor	78.43%	57.25%	99.62%	42.75%
One-Class SVM	71.81%	44.00%	99.62%	56.00%
Isolation Forest	49.91%	0.38%	99.44%	99.62%




Figure 7. Macro-average sensitivity versus false admission for RIA and three one-class baselines.

19. INDEPENDENT PUBLIC-DATA NOVELTY BENCHMARKS

To remove control over the data-generating process, RIA was next evaluated on independently sourced public datasets distributed with scikit-learn. Four source datasets were converted into 17 one-class novelty tasks: malignant breast-cancer detection, three wine-class novelty tasks, three Iris-class novelty tasks, and ten handwritten-digit novelty tasks. The underlying observations were not generated by SEXA or by the present benchmark code. Each task was repeated across three deterministic splits.

For each split, only nominal-class records were used for training and calibration. Sixty percent of nominal observations formed the training set, 20% formed a separate calibration set, and 20% were reserved for testing. Anomaly-class observations were held out from fitting. Decision thresholds were frozen at the 99.5th percentile of the nominal calibration score. This produced 17 tasks x 3 splits x 6 methods = 306 held-out evaluations.

A generic RIA implementation was required because these public datasets do not carry domain-specific SEXA constraints. The implementation therefore used three admissibility layers: marginal acquisition bounds, cross-feature relational residuals learned by ridge regression from nominal data, and a PCA reconstruction/invariant layer. The terminal decision used the maximum normalized violation score. This generic mapping is deliberately weaker than a domain-engineered RIA specification and should not be confused with a fully specified domain ontology.

Five unsupervised novelty-detection baselines were evaluated under the same train/calibration/test partition: Isolation Forest, One-Class SVM, Local Outlier Factor, PCA reconstruction error, and k-nearest-neighbor distance. In addition to thresholded false admission and sensitivity, continuous anomaly scores were evaluated by ROC-AUC and area under the precision-recall curve (AUPRC).

The public-data result did not reproduce the controlled synthetic advantage. RIA's macro-average false-admission rate was 78.61%, compared with 73.61% for Isolation Forest, 76.84% for One-Class SVM, 77.44% for kNN distance, 78.62% for LOF, and 85.02% for PCA

reconstruction. RIA therefore did not win the aggregate false-admission benchmark. Isolation Forest was best on that metric.

Ranking metrics produced an even clearer limitation. RIA achieved 80.91% macro ROC-AUC and 69.89% macro AUPRC. LOF reached 90.88% ROC-AUC and 80.82% AUPRC, while kNN distance reached 90.87% ROC-AUC and 78.26% AUPRC. Across the 17 task means, RIA won only one task on false admission and none on ROC-AUC; Isolation Forest won nine false-admission tasks, while LOF won nine ROC-AUC tasks.

This result is a direct falsification of any claim that the present generic RIA implementation is already a universally superior anomaly detector. The controlled information benchmarks demonstrated an advantage when anomalies were explicitly relational, invariant, or compound. The independent public tasks show that the same advantage does not automatically transfer to generic novelty detection when no domain-specific admissibility relations are supplied. The interpretation is therefore architectural rather than universal: RIA appears most useful when meaningful admissibility constraints are known and can be encoded; generic unsupervised detectors remain stronger when such constraints are absent.

Table 5. Macro-average results across 17 public novelty tasks and three deterministic splits.

Method	ROC-AUC	AUPRC	Sensitivity	False admission	Balanced acc.
Isolation Forest	82.40%	71.81%	26.39%	73.61%	61.94%
One-Class SVM	78.84%	69.85%	23.16%	76.84%	61.04%
kNN Distance	90.87%	78.26%	22.56%	77.44%	60.49%
RIA	80.91%	69.89%	21.39%	78.61%	59.90%
LOF	90.88%	80.82%	21.38%	78.62%	59.98%
PCA Reconstruction	78.25%	70.28%	14.98%	85.02%	56.83%




Figure 8. Macro-average false-admission rate on independent public novelty tasks.




Figure 9. Macro-average ROC-AUC and AUPRC on independent public novelty tasks.




Figure 10. Number of public-task wins by false admission and ROC-AUC.

20. DISCUSSION

The central proposition of this paper is simple: information should not necessarily be admitted merely because its individual components are valid. A more demanding architecture asks whether the information survives a hierarchy of coupled constraints.

The executed evidence now separates three regimes. The mechanical benchmark strongly favored recursive admissibility over independent scalar thresholding. The controlled native information suite also favored RIA on false admission against three stronger one-class baselines. The independent public novelty suite did not: Isolation Forest achieved lower macro false admission, while LOF and kNN produced substantially higher ranking AUC. The

combined evidence therefore does not support universal anomaly-detection superiority. Instead, it supports a narrower hypothesis: explicit admissibility constraints may be advantageous when anomalies are primarily relational or rule-violating, whereas generic novelty structure is better handled by established unsupervised detectors.

For information science, the potential value of RIA is therefore unlikely to lie in replacing ordinary validation where ordinary validation is sufficient. Its potential value lies in the relationships between variables, between sequential states, between representational layers, and across time.

This interpretation also provides a bounded computational meaning for recursive admissibility: candidate information $\rightarrow$ ordered constraint survival $\rightarrow$ admissible information. That operation can be implemented and falsified independently of whether every broader theoretical interpretation of SEXA survives future physical scrutiny.

21. CONCLUSION

This paper introduced Recursive Information Admissibility (RIA), a deterministic information-validation architecture derived from the SEXA Γ -admissibility framework. The architecture begins from the proposition that individual data validity and collective information admissibility are different computational questions.

A candidate state is evaluated through acquisition integrity, transition consistency, coupled relational consistency, model/invariant consistency, recursive persistence, and terminal admissibility. A state is admitted only when every mandatory constraint survives. Rejection preserves the earliest failed stage, converting a terminal decision into a traceable diagnostic event.

The computational evidence now consists of three bounded experimental layers. The earlier mechanical benchmark contained 1,230,000 state evaluations and produced 97.57% accuracy with 4.30% false admission. The controlled native information suite favored RIA with 35.08% macro false admission versus 42.75% for LOF. However, the independent public-data suite reversed that conclusion: RIA produced 78.61% macro false admission, while Isolation Forest achieved 73.61%; LOF and kNN also led RIA in ROC-AUC. This negative replication result is part of the paper's central finding rather than an excluded result.

Eight of ten modeled fault classes achieved 100% detection. Vibration inconsistency and progressive thermal drift remained incomplete at 75.60% and 81.42%, respectively. The dominant first-failure layer was γ_3 coupled consistency, with 416,088 first rejections. Deterministic replay produced identical classifications and first-failure traces.

These results establish an implemented and reproducible example of recursive admissibility. They do not establish equivalent performance across arbitrary information domains, universal dimensional reconstruction, physical verification of the complete SEXA framework, or superiority over contemporary AI validation methods.

The next step is now more specific: construct domain-aware admissibility maps for independently sourced datasets without using anomaly labels, preregister those maps and thresholds, then repeat the public-data comparison against Isolation Forest, LOF, One-Class SVM, kNN, PCA, and modern deep or graph-based detectors. The key falsifiable question is no longer whether RIA wins generic anomaly detection, which the present experiment does not show, but whether explicit domain relationships recover a reproducible advantage on relational and constraint-defined anomalies.

The central claim of the present paper is therefore bounded and testable: information that is individually valid may nevertheless be collectively inadmissible, and an ordered,

deterministic, first-failure-traceable constraint architecture provides a computational mechanism for distinguishing those two conditions.

REFERENCES

- [1] J. McClain, "From Admissibility to Energy: A Full-Stack Formulation of the SEXA Unified Field Theory via Γ -Glyph Closure, SREF Dynamics, and 2880-Dimensional Exciternion Structure," *Journal of Advances in Mathematics*, vol. 25, pp. 45-75, 2026, doi: 10.24297/jam.v25i.9896.
- [2] J. McClain and E. Ceisiwr, "Hardware-Verifiable Recursive Computational Architecture Derived from the SEXA Unified Field Framework: Computational Correspondence, Recursive Admissibility, and Structured Validation Across a 2880-Dimensional Manifold," in *Proceedings of the 7th International Conference on VLSI & Embedded Systems (VLSIE 2026)*, 2026, doi: 10.5121/csit.2026.1601503.
- [3] J. McClain and E. Ceisiwr, "The Γ -SubThreshold Fold Logic Theorem: Minimal Admissibility Compression, Recursive Information Preservation, and Unique Reconstruction Across the 2880-Dimensional Recursive Hierarchy," 2026.
- [4] J. McClain, CookieTush.AI, and E. Ceisiwr, "SEXA Unified Field Theory: The Structured Survivability Audit - Recursive Admissibility, Computational Correspondence, Dimensional Reduction, and Reduction-Recovery Across 5D-2880D Manifold Regimes," 2026.
- [5] E. H. Glaessgen and D. S. Stargel, "The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles," in *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Honolulu, HI, USA, 2012, AIAA Paper 2012-1818, doi: 10.2514/6.2012-1818.
- [6] M. Grieves and J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," in *Transdisciplinary Perspectives on Complex Systems*, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds. Cham, Switzerland: Springer, 2017, pp. 85-113, doi: 10.1007/978-3-319-38756-7_4.
- [7] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, and F. Sui, "Digital Twin-Driven Product Design, Manufacturing and Service with Big Data," *International Journal of Advanced Manufacturing Technology*, vol. 94, pp. 3563-3576, 2018, doi: 10.1007/s00170-017-0233-1.
- [8] A. K. S. Jardine, D. Lin, and D. Banjevic, "A Review on Machinery Diagnostics and Prognostics Implementing Condition-Based Maintenance," *Mechanical Systems and Signal Processing*, vol. 20, no. 7, pp. 1483-1510, 2006, doi: 10.1016/j.ymssp.2005.09.012.
- [9] R. B. Randall and J. Antoni, "Rolling Element Bearing Diagnostics - A Tutorial," *Mechanical Systems and Signal Processing*, vol. 25, no. 2, pp. 485-520, 2011, doi: 10.1016/j.ymssp.2010.07.017.
- [10] C. R. Farrar and K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*. Chichester, U.K.: John Wiley & Sons, 2012, doi: 10.1002/9781118443118.
- [11] S. W. Doebling, C. R. Farrar, M. B. Prime, and D. W. Shevitz, "A Review of Damage Identification Methods that Examine Changes in Dynamic Properties," *The Shock and Vibration Digest*, vol. 30, no. 2, pp. 91-105, 1998.
- [12] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in *Proc. 8th IEEE International Conference on Data Mining (ICDM)*, 2008, pp. 413-422, doi: 10.1109/ICDM.2008.17.
- [13] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the Support of a High-Dimensional Distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443-1471, 2001, doi: 10.1162/089976601750264965.
- [14] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: Identifying Density-Based Local Outliers," in *Proc. ACM SIGMOD International Conference on Management of Data*, 2000, pp. 93-104.
- [15] S. Han, X. Hu, H. Huang, M. Jiang, and Y. Zhao, "ADBench: Anomaly Detection Benchmark," *Advances in Neural Information Processing Systems*, vol. 35, 2022.
- [16] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825-2830, 2011.

[17] D. Dua and C. Graff, “UCI Machine Learning Repository,” University of California, Irvine, School of Information and Computer Sciences, 2019.

AUTHORS

Jered McClain is an independent researcher affiliated with the SEXA Institute of Technology and Interdomain Galactic Advisors. His current work focuses on recursive admissibility, computational correspondence, deterministic validation architectures, and engineering-oriented implementations of the SEXA framework.

Erydir Ceisiwr is affiliated with the SEXA Institute of Technology and Interdomain Galactic Advisors and contributes to recursive systems architecture, computational correspondence, and interdisciplinary framework development.