HARDWARE-VERIFIABLE RECURSIVE ADMISSIBILITY FOR AUTONOMOUS AEROSPACE AND MECHANICAL SYSTEMS: A SEXA FRAMEWORK FOR CONSTRAINT-SAFE STATE SELECTION, FAULT DETECTION, DIGITAL-TWIN RECONSTRUCTION, RECURSIVE DIMENSIONAL COMPRESSION, AND ENGINEERING-GRADE COMPUTATIONAL VALIDATION

Jered E. McClain [1], Erydir Ceisiwr [2]

[1]Unified Field Systems LLC, Bronx, New York, USA
[2]Unified Field Systems LLC, United Kingdom

ABSTRACT

Autonomous aerospace and mechanical systems require supervisory architectures capable of determining whether candidate operating states remain mutually consistent before execution. This paper evaluates a hardware-verifiable recursive admissibility architecture derived from the SEXA computational framework using a controlled mechanical digital twin. Ten state variables are evaluated through six sequential Γ-admissibility stages against nominal trajectories and ten controlled fault classes. Parameters were calibrated on separate nominal trajectories and frozen before held-out testing. The final benchmark comprised 1,230,000 mechanical-state evaluations and achieved 97.57% classification accuracy, 95.70% fault sensitivity, 99.59% nominal-state acceptance, and a 4.30% false-admission rate. An independent scalar-threshold baseline produced an 85.55% false-admission rate. The results support recursive admissibility as a reproducible computational engineering constraint architecture under the synthetic conditions tested while preserving explicit boundaries between computational validation and physical hardware or flight validation.

KEYWORDS

Recursive Admissibility, Mechanical Digital Twin, Fault Detection, Autonomous Aerospace Systems, Hardware Verification

1. INTRODUCTION

Modern aerospace and mechanical systems increasingly combine physical machinery with continuous computational decision architectures. Aircraft, spacecraft, unmanned aerial vehicles, robotic mechanisms, propulsion systems, rotating machinery, autonomous vehicles, and structural-health-monitoring platforms can produce continuous measurements of temperature, pressure, vibration, acceleration, rotational velocity, stress, strain, actuator position, energy consumption, and related diagnostic quantities. Digital-twin and structural-health-monitoring

approaches provide computational representations through which these measurements can be interpreted for monitoring, diagnosis, and maintenance [5]–[11].

The existence of individually valid measurements, however, does not establish that their combined state is mechanically admissible. A temperature measurement can remain inside its scalar limit while becoming inconsistent with simultaneously measured pressure. Rotational velocity can remain below its maximum while vibration indicates an abnormal operating condition. An actuator position can satisfy its individual displacement constraint while its transition is incompatible with the preceding mechanical state. Individually permissible measurements can therefore collectively violate structural, dynamical, energetic, temporal, or operational relationships.

The engineering problem considered here is state admissibility: whether a candidate mechanical state should be permitted to enter or remain within the operational state space. Let the complete candidate state space be Ω and let Ω_A ⊂ Ω denote the admissible subset.

$$A : \Omega \to \{0, 1\}, \qquad A(X) = 1 \text{ admit}, \quad A(X) = 0 \text{ reject}$$

The architecture investigated in this paper is derived from the SEXA recursive-admissibility framework [1]–[4] and its preceding hardware-verifiable computational formulation [2]. The predecessor architecture emphasizes deterministic inputs and outputs, observable intermediate computational states, explicit failure conditions, and reproducible execution. The present investigation advances that problem into a controlled mechanical digital-twin environment in which nominal and deliberately corrupted trajectories can be compared against independently defined ground truth.

The distinction between the wider theoretical framework and the present experiment is essential. This work does not claim completed physical verification of the SEXA Unified Field Framework. It isolates recursive admissibility as an engineering mechanism and subjects that mechanism to measurable state variables, controlled faults, deterministic constraint evaluation, and quantitative performance criteria.

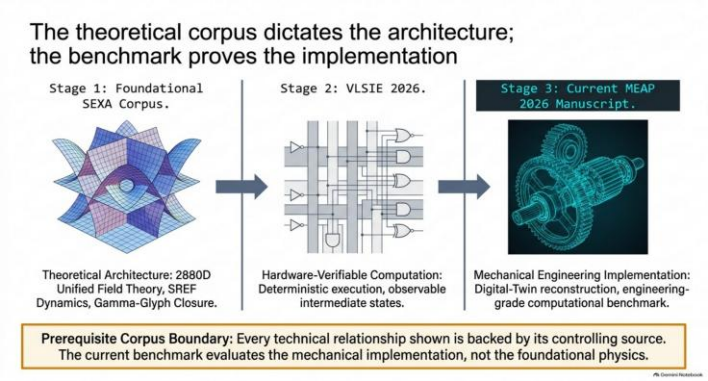


Figure 1. Evidentiary progression from foundational SEXA architecture to the VLSIE computational predecessor and the present MEAP mechanical-engineering benchmark.

The research question is whether a hardware-verifiable recursive admissibility architecture can distinguish nominal from faulty aerospace and mechanical states while preserving deterministic execution, explicit fault localization, computational traceability, and reconstructable digital-twin state information.

2. RECURSIVE ADMISSIBILITY AS AN ENGINEERING CONSTRAINT ARCHITECTURE

2.1. Mechanical State Representation

The instantaneous engineering state used in the benchmark is

$$X_t = [T,\ P,\ V,\ a,\ \omega,\ \sigma,\ \varepsilon,\ q,\ E,\ F]^T$$

where T denotes temperature, P pressure, V vibration amplitude, a acceleration, ω angular velocity, σ mechanical stress, ε strain, q generalized displacement or configuration, E an energy-related state, and F applied or reconstructed force. These quantities need not share common physical dimensions; normalized variables are used only when a common computational representation is required.

Rather than accepting a state because each component independently lies inside a permitted scalar range, the architecture evaluates whether the complete state survives a sequence of coupled constraints. Beginning from X_t^(0)=X_t, the sequential path is

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

For the implemented classifier, each stage returns a binary admissibility indicator a_i(X_t) $\in$ {0,1}, and the complete decision is

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

A candidate state is admitted only if all six implemented constraint classes are satisfied.

2.2. Engineering Interpretation of the Six Γ Stages

Stage	Engineering function	Primary failure type
γ1	Acquisition validity	Missing, corrupted, or out-of-domain sensor state
γ2	Transition/rate consistency	Impossible state-rate or abrupt temporal discontinuity
γ3	Coupled mechanical consistency	Individually valid but mutually inconsistent variables
γ4	Structural/model invariance	Violation of constitutive, energetic, or modeled invariant
γ5	Recursive persistence	Temporally accumulated degradation or persistent residual
γ6	Terminal operational admissibility	Final safety or mission-operating limit

The architecture preserves a first-failure trace rather than returning only a global fault flag. For a rejected state, the first failed stage is

$$k_{\text{fail}} = \min\{\, i : a_i(X_t) = 0 \,\}$$

This converts rejection into a traceable engineering event and permits stage-specific analysis, early termination, and post-test localization.

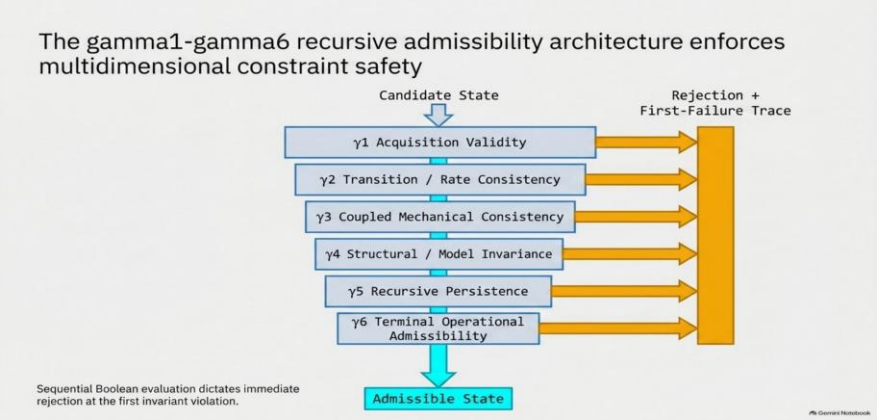


Figure 2. Six-stage Γ-admissibility architecture and first-failure localization path used in the mechanical benchmark.

3. MECHANICAL DIGITAL-TWIN BENCHMARK

3.1. Nominal Mechanical Model

The benchmark requires a mechanical system whose governing behavior is explicit enough that nominal and faulty states can be distinguished independently of the admissibility classifier. A coupled rotating-mechanical system is used because it provides simultaneous displacement, velocity, acceleration, rotation, vibration, thermal, stress, strain, actuator, and energetic quantities.

$$M\ddot{q}(t) + C\dot{q}(t) + Kq(t) = F(t)$$

The translational state follows the standard second-order relationship, while rotational motion is represented by

$$J\dot{\omega}(t) = \tau_m(t) - \tau_L(t) - \tau_f(t)$$

Structural and thermal consistency are represented by

$$\sigma(t) = E_Y \varepsilon(t), \qquad \varepsilon(t) = \Delta L(t)/L_0$$

$$C_T \frac{dT}{dt} = Q_{\text{gen}}(t) - Q_{\text{loss}}(t)$$

The nominal trajectory is generated independently of the Γ classifier. This separation prevents the

admissibility mechanism from defining the same conditions against which it is later judged.

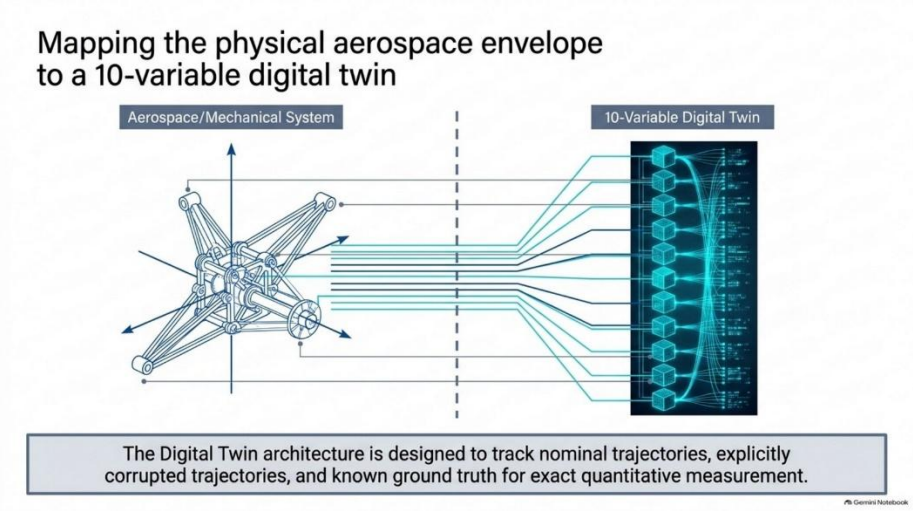


Figure 3. Mapping of the aerospace/mechanical state envelope to the ten-variable digital-twin representation.

3.2. Ground Truth and Controlled Fault Injection

Every simulated state is assigned a ground-truth label before Γ evaluation: y=1 for a state satisfying the benchmark's prescribed mechanical constraints and y=0 for a state violating one or more of those independently defined constraints. The fault generator is computationally separated from the mechanism attempting to detect the fault.

Ten fault classes are evaluated: acquisition corruption; transition corruption; vibration inconsistency; stress-strain inconsistency; energetic inconsistency; pressure inconsistency; progressive thermal drift; rotational overspeed; excessive stress; and compound fault. Each class is evaluated over normalized severity levels 0.25, 0.50, 0.75, and 1.00. Compound-fault cases combine more than one disturbance so that the architecture is not evaluated exclusively under isolated single-variable failures.

3.3. Classification Metrics

For each state the predicted classification is compared directly with ground truth, producing true admissions N_TA, true rejections N_TR, false admissions N_FA, and false rejections N_FR. A disagreement is counted as an error; it is not reinterpreted after the fact.

$$\mathrm{ACC} = \frac{N_{TA} + N_{TR}}{N_{TA} + N_{TR} + N_{FA} + N_{FR}}$$

$$\mathrm{FAR} = \frac{N_{FA}}{N_{FA} + N_{TR}}, \qquad \mathrm{FRR} = \frac{N_{FR}}{N_{FR} + N_{TA}}$$

Fault sensitivity, nominal-state acceptance, balanced accuracy, first-failure localization, deterministic replay, and fault-class-specific detection are additionally retained. These metrics prevent a high aggregate accuracy value from concealing either unsafe false admission or indiscriminate over-rejection.

4. DIMENSIONAL COMPRESSION AND RECONSTRUCTION BOUNDARY

The wider SEXA architecture distinguishes a higher-dimensional structural representation from a bounded operational representation. The present mechanical experiment does not assume that this theoretical relationship has been physically established. It instead tests a narrower engineering compression problem in which the ten-component mechanical state is mapped to five operational coordinates and reconstructed without direct access to a hidden copy of the original state.

$$C_\Gamma : \mathbb{R}^{10} \to \mathbb{R}^5, \qquad \mathcal{R} : \mathbb{R}^5 \to \mathbb{R}^{10}$$

The five operational coordinates represent physically interpretable state families rather than arbitrary selected sensor channels: dynamical state, thermal-fluid state, structural state, control/actuation state, and recursive health or persistence state. Reconstruction parameters are fixed before final evaluation.

A reconstruction is considered valid only if the information pathway itself passes the audit. Direct access to the original state, an undeclared lossless encoding of the original vector, test-set fitting after evaluation begins, or an external identifier that retrieves the original state invalidates the reconstruction claim.

This engineering experiment therefore directly tests a 10D→5D→10D problem. It does not simulate 2,880 independent mechanical sensor variables and then treat that experiment as equivalent to a complete 2,880-dimensional SEXA manifold. The distinction between a measured engineering compression result and the broader theoretical 5D/2880D relationship is retained throughout.

5. EXPERIMENTAL PROTOCOL AND COMPUTATIONAL IMPLEMENTATION

5.1. Prospective Calibration and Parameter Freeze

The benchmark is prospectively separated into calibration and held-out evaluation. Twenty nominal calibration seeds are used to establish residual and transition tolerances. Those parameters are then frozen before the final population is evaluated. The held-out population uses disjoint deterministic seeds and contains both nominal and faulted trajectories. No threshold is retuned after observing the held-out results.

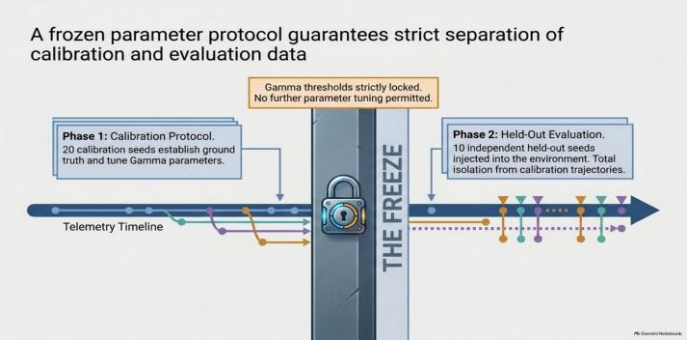


Figure 4. Prospective calibration, parameter freeze, and held-out evaluation protocol.

5.2. Independent-Threshold Baseline

The complete Γ architecture is compared with an independent scalar-threshold baseline. The baseline asks only whether each individual measurement remains inside its permitted interval. It does not evaluate coupled-state residuals, transition consistency, persistence, model invariants, or terminal operational relationships. This directly tests the paper's central proposition: states may be individually permissible while jointly inadmissible.

5.3. Deterministic Replay and Falsification

The complete benchmark is repeated without changing the initial state, model parameters, frozen Γ thresholds, or input sequence. Exact classification and first-failure reproducibility are required. A benchmark failure is recorded if deterministic replay fails, false admission remains unacceptably high relative to the specified ground truth, first-stage localization is not reproducible, the intended coupled or temporal fault classes are not detectably different from the scalar baseline, or reconstruction requires hidden access to the original state.

$$D_R = 1$$

6. Executed Held-out Benchmark Results

6.1. Population and Primary Classification

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. The resulting population contained 410 trajectories and 1,230,000 held-out state evaluations. Calibration seeds 100–119 and test seeds 200–209 were disjoint.

The Γ architecture produced N_TA=587,597, N_TR=612,492, N_FA=27,508, and N_FR=2,403. Overall classification accuracy was 97.57%, balanced accuracy 97.65%, fault sensitivity 95.70%, nominal-state acceptance 99.59%, false-admission rate 4.30%, and false-rejection rate 0.41%.

Metric	Γ architecture	Independent 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%

6.2. Improvement Over Independent Thresholding

The independent scalar-threshold baseline correctly accepted all nominal states but detected only 14.45% of faulted states and falsely admitted 85.55% of the fault population. The Γ architecture reduced false admission from 85.55% to 4.30%, corresponding to an approximately 94.97% relative reduction in false admission under the tested conditions.

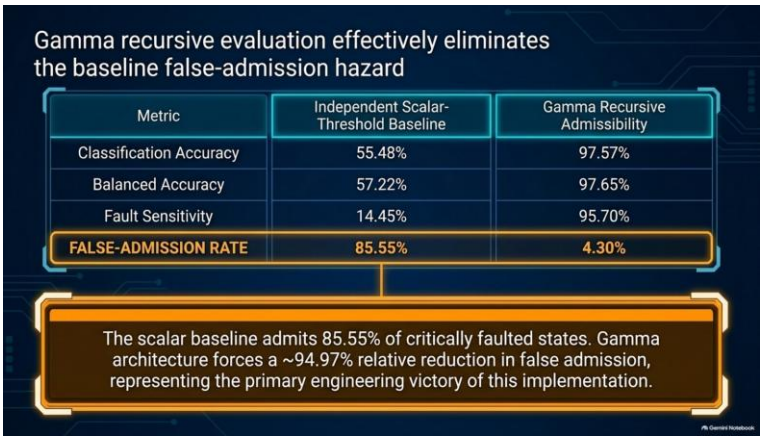


Figure 5. Primary held-out comparison between independent scalar thresholding and Γ recursive admissibility.

6.3. Fault-Class Performance

Fault class	Γ detection	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%

Eight of the ten modeled fault classes achieved 100% Γ detection in the held-out population. The two incomplete classes were vibration inconsistency at 75.60% and progressive thermal drift at 81.42%. Those values remain part of the result and identify the principal weaknesses of the current parameterization.

6.4. First-Failure Localization

Across the complete held-out population, first-stage rejection counts were $\gamma1=16,000$, $\gamma2=64,746$, $\gamma3=416,088$, $\gamma4=116,059$, $\gamma5=2,002$, and $\gamma6=0$. The dominant first-failure stage was therefore $\gamma3$, the coupled mechanical-consistency layer. This is consistent with the benchmark's deliberate emphasis on disturbances that can remain inside individual scalar ranges while violating relationships among pressure, rotational state, vibration, stress, strain, and other simultaneously evaluated quantities.

6.5. Deterministic Replay

An identical held-out trajectory was executed twice without changing the initial state, frozen parameters, Γ architecture, or input sequence. The complete binary admissibility output and first-failure sequence were identical between repetitions for every evaluated state. The deterministic replay criterion was therefore satisfied with

$$D_R = 1$$

7. DISCUSSION

7.1. Coupled Admissibility Versus Independent Thresholding

The strongest engineering distinction between the two classifiers is visible in the hidden-inconsistency subset. Transition corruption, stress-strain inconsistency, and pressure inconsistency were completely undetected by the independent baseline yet were detected at 100% by the Γ architecture; vibration inconsistency was detected at 75.60% by Γ and 0% by the baseline. The result demonstrates that, within the implemented digital twin, classification can depend on relationships between variables rather than on individual sensor amplitudes alone.

The dominance of γ3 supplies a mechanistic explanation for the measured improvement. Γ3 evaluates coupled residuals and therefore has access to information that a scalar range checker discards. The result does not establish that this specific Γ decomposition is unique or universally optimal; it establishes that the implemented coupled-state layer contributed strongly under the tested benchmark conditions.

7.2. Progressive and Temporal Faults

Transition corruption was detected at 100% while the independent threshold baseline detected 0%, providing a clear demonstration that state validity and transition validity are distinct. Progressive thermal drift was more difficult: Γ detected 81.42% of drift states, leaving a measurable residual false-admission fraction. That limitation is retained as an engineering optimization target rather than removed through post-hoc threshold adjustment.

7.3. False Rejection and Operational Trade-Off

Fault sensitivity alone is insufficient for assessing a supervisory architecture because an arbitrarily restrictive system could reject nearly every state. The held-out nominal-state acceptance rate of 99.59% and false-rejection rate of 0.41% provide the corresponding counterweight. Nevertheless, 2,403 false rejections remain operationally meaningful and would require application-dependent cost weighting in a physical deployment.

7.4. Ordering and Noncommutativity Boundary

The present classifier evaluates six staged predicates, but the final Boolean conjunction of independently evaluated predicates is commutative. Consequently, this benchmark establishes staged evaluation, early rejection, and first-failure localization, but it does not by itself establish noncommutativity of the theoretical Γ operators. A direct test of the stronger proposition requires state-transforming operators of the form for which changing operator order can change the resulting transformed state. That stronger experiment remains future work.

$$X^{(i)} = \gamma_i(X^{(i-1)})$$

8. LIMITATIONS AND VALIDATION BOUNDARY

The principal limitation is that the reported benchmark is computational. It evaluates a mechanical digital twin and controlled fault injection rather than an aircraft, spacecraft, turbine, rotor, propulsion system, or other flight-qualified physical system. Successful simulation is therefore not equivalent to flight qualification or experimental hardware validation.

A second limitation is that the comparison baseline is intentionally simple. Independent thresholding is a useful direct test of whether coupled Γ constraints contribute information unavailable to scalar limits, but it does not establish superiority over the full range of contemporary model-based, statistical, machine-learning, observer-based, or digital-twin fault-detection methods.

A third limitation concerns the remaining vibration and thermal-drift false admissions. These failures demonstrate that the present architecture is not universally sensitive to every admissibility violation and that detection quality depends on the fidelity of the modeled relationships and fixed tolerances.

A fourth limitation concerns dimensional reconstruction. Reduction from ten measured variables to five operational coordinates is a demanding compression problem, and unique inversion is not guaranteed for arbitrary unconstrained states. The current classification result therefore cannot be used as evidence of universal 10D→5D→10D reconstruction, much less as physical verification of a 2880-dimensional manifold.

The next validation stage is independent physical data: independently acquired mechanical or aerospace sensor data, prospectively frozen parameters, independently established ground truth, comparison against stronger contemporary methods, and ultimately embedded or hardware-assisted execution.

9. CONCLUSION

This paper converted a SEXA-derived recursive-admissibility concept into a directly falsifiable mechanical-engineering benchmark. A ten-component digital-twin state was evaluated through six staged constraint classes against nominal operation and ten controlled fault classes.

Under prospectively frozen calibration and a disjoint held-out population of 1,230,000 state evaluations, the Γ architecture achieved 97.57% classification accuracy, 97.65% balanced accuracy, 95.70% fault sensitivity, 99.59% nominal-state acceptance, a 4.30% false-admission rate, and a 0.41% false-rejection rate. The independent scalar-threshold baseline produced 55.48% classification accuracy, 14.45% fault sensitivity, and an 85.55% false-admission rate. The resulting reduction in false admission was approximately 94.97% under the synthetic benchmark conditions.

The dominant first-failure stage was γ3 coupled mechanical consistency, accounting for 416,088 first rejections. Deterministic replay returned identical classifications and first-failure traces. At the same time, vibration inconsistency and progressive thermal drift remained incomplete at 75.60% and 81.42% detection, respectively.

The appropriate evidentiary classification is therefore a successful controlled computational engineering benchmark, not experimental physical validation of the complete SEXA framework. The result supports recursive admissibility as an implementable supervisory architecture under

the tested conditions and defines a clear next step toward independent physical and hardware validation.

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.