# Reproducible Evidence for Hybrid Mechanical Models

Sean Barnes, Gerald Ross, Carl Henderson*

* Corresponding author: Carl Henderson

Review Article | Frontiers in Integrative Science | PaperSpread Research Publishing | Published 23 September 2026

## ABSTRACT

Reliable deployment in ensembles trained on computational experiments depends on more than obtaining a strong benchmark result. This conceptual analysis uses reproducibility to study how claims travel from data to model output and then to action. Its central thesis is that the unit of assurance must be the digital thread connecting geometry, materials, boundary conditions, solver settings, derived features, and predicted capacity. The reviewed evidence shows recurring risks from hidden distribution change, correlated evaluation error, missing provenance, and optimization objectives that omit downstream costs. In response, the article proposes a layered evaluation program combining controlled perturbations, subgroup and scenario analysis, repeated runs, calibration or selective prediction, and monitoring after release. It also asks who can inspect, override, and learn from failures. By integrating the assigned target papers with established scholarship, the synthesis clarifies which findings transfer across domains and which remain local to a benchmark, dataset, or experimental apparatus. The goal is a testable research program for bounded, traceable, and revisable systems.

## KEYWORDS

reproducible evidence; hybrid mechanical models; benchmark; evaluation; program; reproducible; hybrid

## Introduction

What records allow another team to reconstruct both the solver outputs and the learned surrogate? The difficulty begins with a basic observation: prediction is only one stage of a deployed workflow. Observation, representation, hypothesis retention, computational effort, and action are separate choices, even when an interface makes them appear as one answer. In ensembles trained on computational experiments, these choices interact with institutions and physical processes that the model only partially represents. A high score can coexist with fragile inputs, an evaluator that rewards the wrong surface feature, or an action rule that turns modest uncertainty into a costly decision. For this reason, the synthesis examines reproducibility as an evidence problem. The inquiry focuses on the minimum evidence and comparison set required for a credible claim. For the present focus, reproducibility therefore becomes a criterion for deciding which evidence deserves further scrutiny.

The unit of analysis is deliberately procedural. Its unit is the digital thread connecting geometry, materials, boundary conditions, solver settings, derived features, and predicted capacity. The benefit is that it avoids attributing every problem to the predictor: attributing every failure to model architecture while ignoring data capture, labeling, retrieval, validation, interface design, or organizational incentives. A process-level unit also supports a more precise comparison of the assigned studies. Although their application settings differ, they each make some part of an inference chain more explicit - a reward signal, a graph relation, a physical boundary condition, a confidence gate, or a revision trigger. The comparison examines whether that explicit component remains meaningful when parameterized geometries, finite-element and multiscale outputs, loading histories, material properties, and validation measurements change, and whether the proposed safeguards lead to design screening, uncertainty-aware optimization, targeted simulation, and physical testing rather than merely producing another metric. For the present focus, reproducibility therefore becomes a criterion for deciding which evidence deserves further scrutiny.

## Conceptual Foundations

Reliability analysis must not collapse aleatory variability, epistemic uncertainty, and strategic response into one category. Aleatory variation concerns irreducible differences among cases. Epistemic uncertainty concerns missing knowledge, limited samples, or unsupported regions of the input space. Strategic adaptation occurs when users, adversaries, markets, or institutions respond to the operational tool itself. These sources demand different controls. Repeated measurement can characterize variation; new evidence and conservative extrapolation can reduce epistemic uncertainty; red teaming and feedback-aware testing address adaptation. Combining them into one confidence score hides which intervention is appropriate. A defensible system therefore reports uncertainty in a form tied to an available response rather than presenting confidence as a decorative number. Seen through reproducibility, the issue is not merely technical; it determines which claims remain open to challenge.

The analysis distinguishes evidence, representation, inference, and decision as separate layers. The evidence layer records observations and their provenance. The representation layer turns those observations into features, graphs, tokens, or simulated states. The inference layer produces predictions, explanations, translations, anomaly scores, or candidate actions. The decision layer maps those outputs to design screening, uncertainty-aware optimization, targeted simulation, and physical testing. Reliability can fail at any boundary. Better inference cannot repair evidence that omitted the relevant condition, and a calibrated probability cannot rescue an action rule whose costs were never specified. Conversely, a conservative decision policy may reduce harm even when the underlying model remains imperfect. This layered view makes reproducibility testable because each claim can be assigned to a component and a measurable transition. This point narrows the research task for ensembles trained on computational experiments: compare alternatives under the same information and resource constraints.

## Relevant Foundations

Several established literatures clarify the design space. Classical fatigue mechanics links cyclic loading, microstructural damage, crack initiation, and crack growth across scales (Suresh 1998). Cumulative-damage rules remain useful baselines precisely because their simplifying assumptions are visible (Miner 1945). Crack-growth laws show how physically meaningful state variables can constrain extrapolation (Paris and Erdogan 1963). Critical-plane analysis demonstrates why multiaxial phase relations cannot be collapsed into a single load magnitude (Fatemi and Socie 1988). Together these findings imply that reproducibility should be evaluated against explicit alternatives and failure costs. In Reproducible Evidence for Hybrid Mechanical Models, this literature supplies a specific test of reproducibility within ensembles trained on computational experiments.

The evidence base offers complementary tests rather than one universal metric. Computational homogenization clarifies which mesoscale details may be summarized and which must remain explicit (Geers, Kouznetsova, and Brekelmans 2010). Model-calibration theory separates parameter uncertainty, observation error, and structural discrepancy (Kennedy and O'Hagan 2001). Global sensitivity analysis measures interactions that one-factor-at-a-time studies can miss (Saltelli et al. 2010). They also caution against interpreting one benchmark, one split, or one explanatory artifact as a complete validation argument. In Reproducible Evidence for Hybrid Mechanical Models, this literature supplies a specific test of reproducibility within ensembles trained on computational experiments.

A credible assurance case can be built from findings that operate at different levels. Random forests offer nonlinear prediction and useful baselines with comparatively modest tuning demands (Breiman 2001). Gradient boosting provides a strong tabular-data benchmark against which more complex ensembles should be justified (Friedman 2001). Additive feature attribution can audit model behavior, although attribution is not a causal explanation (Lundberg and Lee 2017). Their combined value lies in making assumptions visible enough to challenge before deployment and to revise after failure. In Reproducible Evidence for Hybrid Mechanical Models, this literature supplies a specific test of reproducibility within ensembles trained on computational experiments.

## What the Core Studies Establish

MuSK: Multi-Scale Knowledge Learning for Provenance-Graph Anomaly Detection asks how stealthy multi-stage attacks can be detected in heterogeneous system-provenance graphs without abundant attack labels. It uses self-supervised recovery of masked node attributes, meta-path edges, and positional structure, followed by meta-path-guided subgraph verification. The relevant evidence is that evaluation spans nine provenance settings, including DARPA Transparent Computing hosts, and reports high precision, complete recall on those hosts, and efficiency gains from cached incremental expansion. The design joins local attributes, heterogeneous causal paths, and global position while moving decisions from isolated nodes to suspicious subgraphs. In the present review, this work matters because reproducibility cannot be evaluated as an abstract virtue; it must change how evidence is collected and how an action is withheld. Performance may depend on audit completeness, benign workload diversity, attack coverage, and the stability of hand-selected meta-path semantics. (Ma et al. 2026).

A useful test case is BoostAPR: Boosting Automated Program Repair via Execution-Grounded Reinforcement Learning with Dual Reward Models. Rather than treating its output as an isolated score, the study frames how reinforcement learning for program repair can overcome sparse execution feedback and coarse sequence-level credit. Its central mechanism is a three-stage pipeline combining execution-verified supervised traces, sequence- and line-level reward models, and PPO with reward redistribution to critical edit regions, and its evidential basis is that evaluation spans SWE-bench Verified, Defects4J transfer, HumanEval-Java, and QuixBugs, with reported gains in both repository repair and cross-language generalization. Line-level credit assignment makes test execution informative at the granularity where patches actually change behavior. For ensembles trained on computational experiments, the transferable lesson is methodological: a system should preserve the path from signal to decision and expose the conditions under which that path may fail. The boundary is equally important. Benchmark success does not by itself establish maintainability, specification fidelity, or robustness to incomplete and flaky tests. (Li et al. 2026).

Evidence for reproducibility becomes concrete in Synergizing Machine Learning and Multiscale Shakedown Method for Shakedown Loading Capacity Evaluation of Parameterized Lattice Structures. The paper addresses how multiaxial shakedown capacity of parameterized auxetic lattices can be estimated efficiently under uncertain cyclic loading through a multiscale shakedown solver supplies training data to a tuned hybrid ensemble for a peanut-shaped re-entrant lattice, followed by sensitivity analysis. Its evaluation is informative because the reported model reaches low normalized error and high explained variance, while center distance and edge width emerge as influential design variables. Physics-based admissibility and data-driven speed are combined in one design workflow. This does not make the method a universal component; it identifies a design hypothesis that must be re-tested in the article's focal setting. In particular, one lattice family and simulated loading domain do not establish transfer to defects, other topologies, material variability, or experimental fatigue life. The appropriate use of the result is therefore bounded, comparative, and explicit about unresolved deployment conditions (Wang et al. 2025).

## Testing Claims Rather Than Scores

The first validation artifact should list the claims the operational tool is expected to support. Each intended claim - such as improved robustness, safer abstention, faster updating, or better structural localization - needs a target population, comparator, outcome, and boundary condition. Random splits are insufficient when related samples share a patient, repository, instrument run, season, organization, or simulated design family. Grouped and temporal splits better approximate the novelty encountered after field use. Stress tests should vary one factor at a time when attribution is important, then combine factors to measure interaction. Repeated runs are necessary whenever decoding, optimization, retrieval, or numerical kernels can vary. Reporting only the best seed or a pooled mean makes operational instability disappear. In ensembles trained on computational experiments, this distinction should be tested before it changes an operational decision.

Metric choice must be derived from the decision. Discrimination measures whether cases are ranked correctly; calibration asks whether confidence corresponds to observed frequency; selective risk asks what error remains after deferral; robustness measures performance across defined perturbations; and utility assigns costs to actions. These are not interchangeable. For reproducibility, the minimum report should include overall performance, slice-level results, uncertainty intervals, coverage-risk curves, stability across runs, and failure examples linked back to parameterized geometries, finite-element and multiscale outputs, loading histories, material properties, and validation measurements. If the application updates incrementally, the evaluation must also test forgetting, stale evidence, and rollback. If an automatic judge supplies

labels, a sample needs independent human review and content-preserving perturbations that reveal whether the judge is grading substance. Seen through reproducibility, the issue is not merely technical; it determines which claims remain open to challenge.

## A Traceable System Architecture

A workable architecture for ensembles trained on computational experiments begins with typed evidence. Each record should identify its source, time, collection conditions, transformations, and relation to other records. Graphs are useful when relations carry meaning, but graph construction is itself a hypothesis. An edge may denote temporal adjacency, physical causation, code dependency, shared infrastructure, or analyst assertion; treating these as interchangeable creates false certainty. The architecture should therefore preserve edge types and confidence, retain alternative links when evidence is ambiguous, and version both the data schema and the rules that construct the graph. A model may aggregate this structure, but the original evidence must remain recoverable for audit. Seen through reproducibility, the issue is not merely technical; it determines which claims remain open to challenge.

A uniform inference budget is rarely the best operational policy. Easy, well-supported cases can take a fast path. Shifted, ambiguous, or high-cost cases should activate additional precision, retrieval, alternative drafts, simulation, or expert review. This conditional design is preferable to applying maximum computation everywhere because resource cost is part of deployment quality. It is also preferable to an unconditional fast path because efficiency can amplify a systematic error at scale. The trigger must be evaluated as carefully as the expensive model: coverage, false reassurance, subgroup effects, latency, and the consequences of abstention all matter. The output is not simply a prediction but a package containing evidence links, uncertainty, the action taken, and the conditions that caused escalation. The implication is especially consequential in ensembles trained on computational experiments, where a small modeling choice can redirect attention and resources.

## Deployment Controls

Operational rules are the governance expression of the evidence. A field use specification should name allowed uses, prohibited uses, escalation thresholds, data-retention rules, and the person or role that can halt the operational tool. Monitoring should track input shift, missingness, abstention, disagreement, latency, overrides, and downstream outcomes. A rising override rate may indicate model drift, a changed operational sequence, or new user behavior; treating it only as operator noncompliance wastes evidence. Incident review should reconstruct the entire chain, including the estimator and the surrounding interface. Versioned models without versioned prompts, retrieval indexes, rules, and datasets do not support reliable reconstruction. In ensembles trained on computational experiments, this distinction should be tested before it changes an operational decision.

A credible human-review process must be specified and tested. Reviewers need enough context to challenge the output, but not so much generated rationale that weak evidence is obscured by volume. Escalation queues require prioritization and workload limits. A reject option that sends nearly everything to experts is safe only on paper; a reject option that rarely fires may be equally ineffective. For ensembles trained on computational experiments, oversight should therefore be evaluated with realistic staffing, time pressure, and disagreement. The system should record the final action and its reason without turning that record into surveillance unrelated to the stated purpose. Contestability, privacy, and security are properties of this institutional process as much as of the predictive component. Seen through reproducibility, the issue is not merely technical; it determines which claims remain open to challenge.

## Experiments the Field Still Needs

A stronger evidence base requires test sets designed around operational failure modes. Cases should represent unsupported regions, ambiguous evidence, conflicting modalities, rare but costly conditions, and realistic adversarial transformations. Each case needs provenance and a reason for inclusion. The second priority is intervention testing: when uncertainty is detected, compare additional computation, retrieval, alternative reasoning paths, model ensembles, and human escalation under the same resource budget. This reveals whether a proposed safeguard actually changes the decision or only changes the explanation. The third priority is transfer. A method should be re-evaluated across sites, devices, languages, organizations, or physical regimes before broad claims are made. In ensembles trained on computational experiments, this distinction should be tested before it changes an operational decision.

A complementary research program should improve how evidence accumulates over time. Benchmarks should maintain negative results, failed branches, and changed definitions so later work does not learn only from successful demonstrations. Shared reporting should include data lineage, versioned testing code, confidence intervals, and enough error material to reproduce major conclusions. Prospective studies should predefine monitoring and stopping rules, then measure how people adapt to the operational tool. Finally, researchers should compare simple baselines with complex architectures under equivalent information. Complexity is justified when it adds a measurable capability - for example, structural localization, calibrated deferral, or incremental updating - not merely when it creates a richer narrative around the same errors. Seen through reproducibility, the issue is not merely technical; it determines which claims remain open to challenge.

## Conclusion

Reproducibility is credible only when it is attached to an clearly stated evidence chain and decision policy. The focal literature contributes several ways to improve the chain: structured exploration, typed graphs, conditional revision, adaptive computation, physical constraints, and repeated testing. A reported benchmark gain still leaves the boundary of safe transfer to be established. For ensembles trained on computational experiments, the practical standard should be a traceable workflow that measures shift and instability, supports design screening, uncertainty-aware optimization, targeted simulation, and physical testing, and preserves enough evidence for an independent reviewer to reconstruct failure. Future work should compare safeguards under realistic resource and staffing limits, publish negative and subgroup results, and treat monitors and automated judges as components requiring their own validation. A dependable system need not answer every case; it must connect each action to supporting evidence, respond appropriately to uncertainty, and leave an auditable record. The implication is especially consequential in ensembles trained on computational experiments, where a small modeling choice can redirect attention and resources.

## References

1. Ma, Mingjun, et al. "MuSK: Multi-Scale Knowledge Learning for Provenance-Graph Anomaly Detection." Computer Networks 289 (2026): 112728.

2. Li, Yuanhao, et al. "BoostAPR: Boosting Automated Program Repair via Execution-Grounded Reinforcement Learning with Dual Reward Models." arXiv preprint arXiv:2605.09134 (2026).

3. Wang, Lizhe, et al. "Synergizing Machine Learning and Multiscale Shakedown Method for Shakedown Loading Capacity Evaluation of Parameterized Lattice Structures." Extreme Mechanics Letters 75 (2025): 102297.

4. Suresh, S. Fatigue of Materials. 2nd ed., Cambridge University Press, 1998.

5. Miner, M. A. "Cumulative Damage in Fatigue." Journal of Applied Mechanics, vol. 12, 1945, pp. A159-A164.

6. Paris, P., and F. Erdogan. "A Critical Analysis of Crack Propagation Laws." Journal of Basic Engineering, vol. 85, no. 4, 1963, pp. 528-533.

7. Fatemi, Ali, and Darrell F. Socie. "A Critical Plane Approach to Multiaxial Fatigue Damage Including Out-of-Phase Loading." Fatigue & Fracture of Engineering Materials & Structures, vol. 11, no. 3, 1988, pp. 149-165.

8. Geers, M. G. D., V. G. Kouznetsova, and W. A. M. Brekelmans. "Multi-Scale Computational Homogenization: Trends and Challenges." Journal of Computational and Applied Mathematics, vol. 234, no. 7, 2010, pp. 2175-2182.

9. Kennedy, Marc C., and Anthony O'Hagan. "Bayesian Calibration of Computer Models." Journal of the Royal Statistical Society: Series B, vol. 63, no. 3, 2001, pp. 425-464.

10. Saltelli, Andrea, et al. "Variance Based Sensitivity Analysis of Model Output: Design and Estimator for the Total Sensitivity Index." Computer Physics Communications, vol. 181, no. 2, 2010, pp. 259-270.

11. Breiman, Leo. "Random Forests." Machine Learning, vol. 45, 2001, pp. 5-32.

12. Friedman, Jerome H. "Greedy Function Approximation: A Gradient Boosting Machine." Annals of Statistics, vol. 29, no. 5, 2001, pp. 1189-1232.

13. Lundberg, Scott M., and Su-In Lee. "A Unified Approach to Interpreting Model Predictions." Advances in Neural Information Processing Systems, vol. 30, 2017.