

Computing Architectures and Data Systems
REVIEW | COMMISSIONED SYNTHESIS

Reproducible Benchmarking of Machine-Learning Systems

Corey Richardson - Corresponding author

Abstract

This review examines reproducible benchmarking of machine-learning systems. The organizing question is which controls make results portable across hardware, software versions, datasets, and tuning budgets. Ten related scholarly sources are synthesized through a decision-centered framework spanning problem definition, mechanism, measurement, evaluation, implementation, and governance. The review does not invent experiments, pooled estimates, or unreported quantitative results. It instead evaluates the strength and transferability of the available evidence, with particular attention to presenting a narrowly optimized speedup without uncertainty or complete configuration. The resulting framework links technical or empirical performance to explicit use conditions and identifies tests that should precede wider adoption in AI systems and infrastructure evaluation.

Keywords: machine-learning systems, benchmarking, reproducibility, performance, measurement

Synthesis lens	Principal question
Mechanism	the chain connecting evidence inputs, reproducible benchmarking of machine-learning systems, intermediate outputs, and downstream decisions
Evaluation	construct validity, comparative performance, uncertainty, transfer across settings, and decision utility
Primary risk	presenting a narrowly optimized speedup without uncertainty or complete configuration

1. Introduction

Research on reproducible benchmarking of machine-learning systems sits at the boundary between a promising component and a consequential decision. The core question is which controls make results portable across hardware, software versions, datasets, and tuning budgets. Answering it requires more than assembling favorable results. It requires a chain of evidence that makes the problem boundary, mechanism, measurement, comparator, uncertainty, and accountable decision visible to the reader.

This article presents a structured narrative review of ten related publications. Sources were selected for topical relevance, traceable publication metadata, and complementary coverage of technical, empirical, and governance questions. No meta-analytic pooling is attempted because the included studies differ in designs, populations, system boundaries, and outcomes. The purpose is decision-oriented synthesis: to identify what each source can support, where transfer remains uncertain, and which evidence would justify the next stage of use.

The central risk is presenting a narrowly optimized speedup without uncertainty or complete configuration. A top-line metric or broad policy label can appear decisive while concealing the conditions that produced it. Accordingly, the synthesis separates observation from interpretation and implementation from validation. This separation is especially important when the same system creates benefits for one user or institution while transferring cost, risk, or administrative burden to another.

2. Evidence Base and Problem Boundary

A defensible review begins by defining the unit of analysis. For the present topic, the relevant boundary includes inputs, institutional or technical context, the mechanism producing an output, the person or system acting on that output, and the downstream outcome. Evidence falls outside the boundary when it measures only an intermediate proxy yet is used to claim an end-to-end benefit.

Wu et al. (2017) addresses "MoleculeNet: a benchmark for molecular machine learning." In this synthesis, that work is used as a source on problem definition within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Hu et al. (2020) addresses "Open Graph Benchmark: Datasets for Machine Learning on Graphs." In this synthesis, that work is used as a source on conceptual boundary within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Duan et al. (2016) addresses "Benchmarking Deep Reinforcement Learning for Continuous Control." In this synthesis, that work is used as a source on measurement within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Together, these works provide complementary entry points, but their conclusions should not be flattened into a single ranking. Differences in data generation, setting, temporal horizon, and outcome definition may be substantively meaningful. A review should therefore preserve

those differences and state where analogy, rather than direct replication, connects one domain to another.

3. Mechanism and System Design

The operative mechanism is the chain connecting evidence inputs, reproducible benchmarking of machine-learning systems, intermediate outputs, and downstream decisions. This formulation discourages component-only reasoning. A model, platform, policy, assay, or infrastructure service acquires practical meaning only when its interfaces and use conditions are specified. The evidence chain should describe what information is available, what transformation occurs, which assumptions make that transformation credible, and who is authorized to act.

Borlido et al. (2020) addresses "Exchange-correlation functionals for band gaps of solids: benchmark, reparametrization and machine learning." In this synthesis, that work is used as a source on system mechanism within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Purushotham et al. (2018) addresses "Benchmarking deep learning models on large healthcare datasets." In this synthesis, that work is used as a source on representation choice within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Azodi et al. (2019) addresses "Benchmarking Parametric and Machine Learning Models for Genomic Prediction of Complex Traits." In this synthesis, that work is used as a source on failure analysis within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

These studies also show why modular claims require interface tests. A component may perform well while the complete pathway fails because inputs drift, users interpret outputs differently, recovery semantics are ambiguous, or institutional incentives change. Design documentation should therefore include not only the intended pathway but also foreseeable deviations, degraded modes, and conditions under which the system should stop or defer.

4. Evaluation and Uncertainty

Evaluation should center on construct validity, comparative performance, uncertainty, transfer across settings, and decision utility. Average performance remains useful, but it should be accompanied by distributions, error categories, relevant subgroups or operating regimes, and uncertainty at the point where a decision changes. Comparators must isolate the value of the proposed addition rather than compare a mature intervention with an intentionally weak baseline.

Hannousse and Yahiouche (2021) addresses "Towards benchmark datasets for machine learning based website phishing detection: An experimental study." In this synthesis, that work is used as a source on comparative evaluation within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself

establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Wójcikowski et al. (2017) addresses "Performance of machine-learning scoring functions in structure-based virtual screening." In this synthesis, that work is used as a source on uncertainty within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

External validation should introduce meaningful shifts in setting, time, population, workload, market structure, or environmental conditions. A nominally independent sample can still be too similar to test transfer. Conversely, a failed transfer is scientifically informative when the changed conditions are measured and the mechanism of failure is examined rather than hidden in an aggregate score.

5. Translation and Governance

Translation to AI systems and infrastructure evaluation depends on more than technical readiness. Decision rights, documentation, maintenance, monitoring, appeal, and recovery are part of the intervention. The governance requirement is traceable evidence, named responsibility, version control, monitoring, appeal, and explicit revalidation. Each control should be connected to a named failure mode and should identify an owner, evidence source, review interval, and escalation path.

Zech et al. (2018) addresses "Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study." In this synthesis, that work is used as a source on implementation within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

Fu et al. (2022) addresses "Forces are not Enough: Benchmark and Critical Evaluation for Machine Learning Force Fields with Molecular Simulations." In this synthesis, that work is used as a source on governance within reproducible benchmarking of machine-learning systems. The connection is deliberately bounded: the cited study informs the evidence map, but it does not by itself establish readiness for AI systems and infrastructure evaluation. That distinction keeps the review close to the published record and prevents an adjacent result from being converted into a broader causal claim.

A credible implementation plan also defines sunset and revalidation conditions. New data, software updates, institutional restructuring, population change, or shifts in incentives can invalidate previously reasonable assumptions. Monitoring should therefore test the validity of the evidence chain, not merely whether a service remains online or a headline metric remains stable.

6. Research Agenda

Future studies should preregister or otherwise predefine the intended decision, principal outcomes, exclusions, and analysis plan when feasible. They should report missingness, sensitivity analyses, negative or null findings, and the cost of errors to differently situated stakeholders. Reproducible artifacts should include sufficient metadata to reconstruct data provenance, model or analytical versions, parameter choices, and the sequence from evidence to conclusion.

Cross-disciplinary work needs a shared object of explanation rather than a collection of adjacent vocabularies. For reproducible benchmarking of machine-learning systems, that shared object is the complete decision pathway. Technical, clinical, social, environmental, or economic expertise should each change the design or interpretation of the study. When evidence remains incomplete, the useful output is a bounded hypothesis, a transparent uncertainty statement, or a request for additional measurement.

7. Conclusion

The literature supports a disciplined approach to reproducible benchmarking of machine-learning systems: define the decision, preserve system boundaries, test consequential failure modes, connect uncertainty to action, and assign accountable control. The strongest conclusion is not that one method is universally superior. It is that credible use in AI systems and infrastructure evaluation requires an auditable link from source evidence to design choice, evaluation result, and governed decision.

Declarations

Ethics approval: Not applicable. This article synthesizes previously published literature and reports no new research involving humans, animals, human tissue, or identifiable sensitive data.

Consent to participate and consent for publication: Not applicable.

Funding: No external funding is reported for this commissioned synthesis.

Competing interests: The authoring group declares no competing interests for this commissioned synthesis.

AI-assisted writing: Generative AI supported drafting, source organization, and language editing. Accountable authors and editors must verify every claim, citation, and disclosure before publication.

Data, code, and materials availability: No new dataset or code was generated. All evidence is identified in the reference list.

Licence: This manuscript is prepared for publication under the Creative Commons Attribution 4.0 International licence (CC BY 4.0).

References

- Azodi, C. B., Bolger, E., McCarren, A., Roantree, M., Campos, G. D. L., & Shiu, S. H. (2019). Benchmarking Parametric and Machine Learning Models for Genomic Prediction of Complex Traits. *G3 Genes Genomes Genetics*, 9(11), 3691-3702. <https://doi.org/10.1534/g3.119.400498>
- Borlido, P., Schmidt, J., Huran, A. W., Tran, F., Marques, M. A. L., & Botti, S. (2020). Exchange-correlation functionals for band gaps of solids: benchmark, reparametrization and machine learning. *npj Computational Materials*, 6(1). <https://doi.org/10.1038/s41524-020-00360-0>
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., & Abbeel, P. (2016). Benchmarking Deep Reinforcement Learning for Continuous Control. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.1604.06778>
- Fu, X., Wu, Z., Wang, W., Xie, T., Keten, S., Gomez-Bombarelli, R., & Jaakkola, T. (2022). Forces are not Enough: Benchmark and Critical Evaluation for Machine Learning Force Fields with Molecular Simulations. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2210.07237>
- Hannousse, A., & Yahiouche, S. (2021). Towards benchmark datasets for machine learning based website phishing detection: An experimental study. *Engineering Applications of Artificial Intelligence*, 104, 104347. <https://doi.org/10.1016/j.engappai.2021.104347>
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., & Leskovec, J. (2020). Open Graph Benchmark: Datasets for Machine Learning on Graphs. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2005.00687>
- Purushotham, S., Meng, C., Che, Z., & Liu, Y. (2018). Benchmarking deep learning models on large healthcare datasets. *Journal of Biomedical*

- Informatics, 83, 112-134. <https://doi.org/10.1016/j.jbi.2018.04.007>
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., & Pande, V. (2017). MoleculeNet: a benchmark for molecular machine learning. *Chemical Science*, 9(2), 513-530. <https://doi.org/10.1039/c7sc02664a>
- Wójcikowski, M., Ballester, P. J., & Siedlecki, P. (2017). Performance of machine-learning scoring functions in structure-based virtual screening. *Scientific Reports*, 7(1), 46710. <https://doi.org/10.1038/srep46710>
- Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., & Oermann, E. K. (2018). Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLoS Medicine*, 15(11), e1002683. <https://doi.org/10.1371/journal.pmed.1002683>