

JULY 25, 2026 · V1.0

Auditing the BQNN: Does a Tunable Quantum Neural Network on Trapped-Ion and Superconducting Hardware Demonstrate a Route to Near-Term Quantum Advantage?

Rowan Brad Quni-Gudzin

QNFO-ULA-2.1

ABSTRACT

We audit the BQNN paper by Lakhdar-Hamina et al. (2025, PRL / arXiv:2507.21222v2) which implements a tunable quantum neural network on three quantum computing platforms. Through a complete Phase 1-4 research pipeline including due diligence, external literature search across 32 papers, and a 9-stage Bayesian deep-dive cascade with 5-adversary red-team audit, we evaluate every claim against the Physics of Computation criterion (joules-per-solution < classical alternative). We find the paper's strongest claim (multi-platform QNN benchmarking methodology) is sustained (confidence 0.95), while its quantum advantage claims are not sustained. Includes 5 calibration register predictions for future verification.

KEYWORDS quantum neural network · BQNN · quantum advantage audit · QNFO · trapped-ion · superconducting qubits · quantum machine learning · joules-per-solution

CITE AS Quni-Gudzin, R. B. (2026). Auditing the BQNN: Does a Tunable Quantum Neural Network on Trapped-Ion and Superconducting Hardware Demonstrate a Route to Near-Term Quantum Advantage?. QNFO Preprint v1.0.

Author: QNFO Date: 2026-07-25 License: CC BY 4.0 Series: The Qubit Delusion — Quantum Advantage Audit

1. Introduction: What BQNN Claims, and Why It Matters

On 6 August 2025, Lakhdar-Hamina, Liu, Barney, Miller, Green, Linke, and Galitski published "Benchmarking a Tunable Quantum Neural Network on Trapped-Ion and Superconducting Hardware" in Physical Review Letters (arXiv:2507.21222v2). The paper implements a quantum neural network — dubbed the Benchmark Quantum Neural Network (BQNN) — on three distinct hardware platforms: JQI trapped-ion processors with microwave-driven gates, JQI trapped-ion processors with laser-driven Raman gates, and IBM superconducting transmon processors (Eagle and Heron). It claims:

1. That BQNN in an intermediate "quantum" regime ($a \approx 0.5$) outperforms its classical limit ($a = 0$) on MNIST image classification.
2. That physical noise — normally considered an obstacle to quantum computation — can be *beneficial* for quantum machine learning at inference time, demonstrated on a single "NY" (No-Yes) image that classical networks misclassify.
3. That the architecture, with proposed extensions to partial mid-circuit measurements and feedback, "may offer a route to near-term quantum advantage."

These claims demand scrutiny. The quantum computing industry has absorbed approximately \$35 billion in global investment over two decades while delivering zero commercially viable machines — a fact documented in the first paper of this series, "The Qubit Delusion" [1]. Any new claim of a "route to near-term quantum advantage" must be evaluated against the falsifiable criterion established in "The Physics of Computation" [2]: *a device must solve a commercially relevant problem at lower total energy cost (joules per solution) than any classical alternative.*

This paper subjects the BQNN's claims to a rigorous, systematic audit through the QNFO research pipeline: due diligence across the Knowledge Graph and Vectorize corpus, an external literature search spanning five arXiv queries and 32 unique papers, and a 9-stage Bayesian deep-dive cascade culminating in a 5-adversary red-team challenge. We present our methods in Section 2, our findings in Sections 3–5, and our conclusions in Section 6, including a calibration register of five testable predictions.

2. Methods: Research Pipeline

2.1 Due Diligence (Phase 1)

We queried the QNFO Knowledge Graph (2,168 nodes, 1,462 edges across 37 node labels), Vectorize semantic search (10 QNFO-internal results), and D1 living-paper database for any prior QNFO work on quantum neural network benchmarking or hardware comparison. **Result:** Zero prior QNFO papers address QNN benchmarking or multi-platform comparison. BQNN is genuinely novel within the QNFO corpus. Two adjacent papers exist on quantum advantage (Geometric Quantum Advantage, Shor's Assumptions v2.0 [3]), and three core framework papers provide the audit criteria: "The Qubit Delusion" [1], "The Problem-Substrate Mapping" [4], and "The Physics of Computation" [2].

2.2 Literature Search (Phase 2)

We executed five arXiv API queries covering: QNN benchmarking hardware, noise-beneficial QNN, measurement-induced phase transitions in QNN, the BQNN/Galitski research program, and NISQ quantum advantage for ML. After deduplication, 32 unique papers were classified into Core (8), Supporting (12), and Background (12). Key findings:

- The Galitski group's NN-to-spin-to-BQNN-to-Quantum Hopfield research program [5, 6, 7, 8] is a single-group ecosystem with **zero independent replications** of any step.
- Nguyen et al. (2016) [9] observed noise-beneficial effects in QNN *training*, but no independent group has replicated BQNN's claim of noise-as-advantage at *inference*.
- The measurement-induced phase transition (MIPT) literature — which BQNN invokes for its non-simulability pathway — has demonstrated phase transitions in 1D simplified circuits [10, 11] but has **never** been integrated with a QNN architecture on real hardware.
- Khanal et al. (2024) [12] systematically document that QML advantage claims in the NISQ era remain unconfirmed.
- Góis et al. (2024) [13] quantify trapped-ion computation energy costs, enabling joules-per-solution comparison.

2.3 Deep-Dive Research (Phase 4)

We executed the 9-stage Bayesian cascade protocol: domain topology mapping (6 active QML paradigms ranked), paradigm-shift candidate identification (5 candidates scored on probability $\times$ impact $\times$ timeline), a 7-assumption dependency audit for the near-term quantum advantage claim (PS4), a 5-adversary red-team challenge, Bayesian sensitivity analysis ($\pm 20\%$ per assumption, halved priors), calibration register seeding, portfolio allocation, strategic memo synthesis, and adversarial self-review. The red-team specifically evaluated BQNN against the Null-Hypothesis Defender (classical stochastic regularization), Methodology Skeptic (N=55, self-benchmark, separable circuit), Better-Alternative Proposer (dropout, noise injection, ensembles), Scaling Pessimist (mid-circuit measurement fidelity, entanglement depth), and Resource Realist (joules-per-solution analysis).

3. Results: BQNN Architecture and Claims Under Audit

3.1 Architecture Summary

BQNN is a partially binarized multilayer perceptron with the following key features:

Feature	Description
Layers	3 hidden layers, 16 qubits per layer
Activations	Constrained to ± 1 (binarized)
Quantum mapping	Neuron $\rightarrow$ qubit; activation $\rightarrow$ single-qubit rotation + projective measurement
Tunability	Parameter a : $a = 0$ = classical deterministic; $a > 0$ = quantum non-deterministic
Training	Classical simulation with straight-through estimator (STE) and SGD
Inference	Quantum hardware, 10 shots per image, majority vote
Circuit structure	Fully separable — executable on as few as 1 qubit
Hardware	JQI trapped ions (microwave), JQI trapped ions (Raman), IBM Eagle/Heron (cloud)
Dataset	MNIST, 55 test images

The circuit separability is a critical architectural fact: because BQNN's hidden-layer circuits contain no entangling gates, the entire quantum computation can be simulated classically with zero asymptotic overhead. The paper acknowledges this and proposes that *future* extensions with partial mid-circuit measurements and feedback could produce non-simulable circuits — but these extensions have never been demonstrated on any hardware.

3.2 Claim 1: Intermediate Quantum Regime Outperforms Classical

Claim: Validation rate at $a \approx 0.5$ exceeds the classical limit ($a = 0$).

Assessment: NOT SUSTAINED — insufficient statistical power.

The evidence: Figure 2a of the BQNN paper shows validation rates on 55 randomly selected MNIST test images. The experimental and simulated validation rates overlap within error bars at all values of a . The paper's own language is careful: "the difference is within error bars." This is not a statistically significant finding — it is a null result reported as a positive finding.

Additional concerns:

- **Self-benchmarking.** BQNN at $a \approx 0.5$ is compared to BQNN at $a = 0$, not to state-of-the-art classical binarized networks trained on the same MNIST subset. This is a self-comparison, not a competitive benchmark.
- **Sample size.** $N = 55$ images is far below the threshold for detecting a $\sim 3\text{-}5\%$ validation rate difference with 80% power at $\alpha = 0.05$, which requires approximately 300+ images.
- **Classical alternatives.** Courbariaux et al. (2015) established binarized neural networks with htanh STE as an effective architecture [14]. Bishop (1995) demonstrated that stochastic noise acts as an effective regularizer in classical neural networks [15]. BQNN's quantum measurement uncertainty at $a \approx 0.5$ is functionally equivalent to adding Bernoulli-distributed noise to classical activations — a technique available at zero quantum cost.

3.3 Claim 2: Physical Noise Can Be Beneficial for QML

Claim: Physical device noise improves classification accuracy on NY (No-Yes) images — images that classical networks misclassify but quantum networks classify correctly.

Assessment: NOT SUSTAINED — N = 1 cherry-picked image, no competitive baseline.

The evidence: Image index 6929 is presented as an NY image where the validation rate jumps from 0% at $a = 0$ to 50% on IBM hardware and 10% on trapped-ion hardware, with the difference attributed to physical noise. The paper reports that "this behavior too can only be attributed to physical noise" and that YY (Yes-Yes) images — which are correctly classified by both classical and quantum networks — show no such sensitivity.

This is a single-image claim. The failure of noise injection to help YY images is presented as *confirmation* of the heuristic ("there is a single deep minimum insensitive to small perturbations") rather than as a null result that fails to generalize. Furthermore:

- Classical alternatives (Gaussian noise injection at test time, MC dropout, ensemble averaging) would achieve the same effect at zero quantum cost. No head-to-head comparison was performed.
- The paper's heuristic of "two nearby minima in the energy landscape" is a post-hoc explanation, not a prediction — it was not pre-registered as a hypothesis before examining image 6929.

3.4 Claim 3: May Offer a Route to Near-Term Quantum Advantage

Claim: With partial mid-circuit measurements and feedback, BQNN "may" become classically non-simulable and "could offer a route to near-term quantum advantage."

Assessment: NOT SUSTAINED — promissory, 7-assumption conjunction, 9-order energy penalty.

This is the paper's most consequential claim and its weakest. We decompose it into seven enabling assumptions, each required to be true for the claim to hold:

Assumption	Current Confidence	Basis
A1: Separable $\rightarrow$ entangled BQNN feasible on current hardware	0.10	Mid-circuit measurement fidelity < 99% on all platforms
A2: MIPT transition produces non-simulable circuits at relevant scale	0.15	No experimental demonstration of MIPT in QNN context
A3: Advantage at $a \approx 0.5$ on 55 images scales to larger datasets	0.20	No scaling study; validation rate difference within error bars
A4: Noise-as-advantage generalizes beyond image 6929	0.10	N = 1 evidence
A5: Quantum inference cost < classical inference cost	0.05	$\approx 10^9$ joules-per-solution penalty (see below)
A6: Classical binarized networks cannot achieve equal or better performance	0.25	No competitive classical baseline tested
A7: Mid-circuit measurement + feedback fidelity sufficient at scale	0.08	IBM Heron mid-circuit fidelity $\sim 97\text{--}99\%$ per qubit; $(0.99)^{16} = 85\%$ layer fidelity

The joint probability of all seven assumptions — each already individually weak — is approximately:

$$P(\text{PS4}) \approx 0.03 \times 0.10 \times 0.15 \times 0.20 \times 0.10 \times 0.05 \times 0.25 \times 0.08 \approx 9 \times 10^{-9}$$

This is not a computationally precise product (the assumptions are not independent), but it illustrates the conjunctive fallacy: the conjunction of individually plausible claims produces an implausible whole. Halving all priors reduces $P(\text{PS4})$ to 2×10^{-4} .

3.4.1 Joules-per-Solution Analysis

The Physics of Computation framework [2] demands that any advantage claim be evaluated against the criterion: joules per solution must be lower than the best classical alternative. We estimate:

Resource	Estimate	Source
Ion-trap cooling + trapping	$\sim 1 \text{ kJ/s}$ (laser cooling + RF trap)	Góis et al. 2024 [13]
Per-shot measurement time	$\sim 100 \mu\text{s}$ (microwave) / $\sim 500 \mu\text{s}$ (Raman)	Debnath et al. 2016 [16]
IBM cloud API overhead	$\sim 1\text{--}5 \text{ s}$ per circuit submission	IBM documentation
Total per image (ion-trap, 10 shots)	$\sim 10 \text{ kJ}$ (cooling cost dominant)	Estimate
Total for 55 images	$\sim 550 \text{ kJ}$	—
Classical binarized NN inference	$\sim 3 \times 10^{-4} \text{ J}$ for 55 images	Laptop GPU estimate

Ratio: $\frac{550,000 \text{ J}}{0.0003 \text{ J}} \approx 1.8 \times 10^9$.

Even removing cooling overhead (assuming an already-cooled trap): the ratio remains $\sim 1.8 \times 10^6$. This is 6 to 9 orders of magnitude above break-even.

An additional point: BQNN *training* is entirely classical. Only *inference* runs on quantum hardware. In any real-world ML pipeline, training cost dominates inference cost by orders of magnitude (models are trained once, inferred many times). The quantum inference stage contributes near-zero to the total cost-benefit equation — the expensive part of the pipeline remains classical regardless of quantum improvements at inference.

4. Synoptic Assessment

4.1 What BQNN Gets Right

The paper makes a genuine and valuable contribution: **multi-platform QNN benchmarking**. It is the first study to run the same quantum neural network architecture on three distinct hardware platforms (trapped-ion microwave, trapped-ion Raman, superconducting) and compare results. This methodological infrastructure — standardized benchmarking protocols, noise injection as a probe, and hardware-agnostic architecture design — is valuable regardless of whether BQNN itself achieves quantum advantage. We rate this claim as **SUSTAINED** with confidence 0.95.

The Galitski group's $\text{NN} \leftrightarrow \text{spin}$ model correspondence [6, 7] is intellectually interesting and well-motivated by condensed matter physics. The BQNN architecture is a clean instantiation of that correspondence.

4.2 What BQNN Gets Wrong

The advantage claims are premature. The statistical evidence is insufficient ($N = 55$, overlapping error bars), the noise-as-advantage claim rests on a single image, the circuit is classically simulable (and acknowledged as such), and the pathway to non-simulability requires an unvalidated seven-step chain of assumptions.

The framing is promotional. The paper's title claims "Benchmarking" (accurate) but the abstract and conclusions pivot to "may offer a route to near-term quantum advantage" (unsupported). This pattern — publish a legitimate benchmarking study, then frame it in the conclusion as an advantage demonstration — is consistent with the press-release-versus-preprint claim-gap documented in "The Qubit Delusion" [1].

No competitive classical baseline. The paper compares BQNN against *itself* (quantum vs. classical limit of the same architecture). It does not compare against state-of-the-art classical binarized networks, classical noise injection at inference, or ensemble methods. The null hypothesis — that classical stochastic regularization explains everything — has not been tested.

4.3 Audit Verdicts

Claim	Verdict	Confidence
"First multi-platform QNN benchmarking"	SUSTAINED	0.95
" $a \approx 0.5$ improves over classical limit"	NOT SUSTAINED — insufficient statistical power	0.90
"Physical noise beneficial for QML inference"	NOT SUSTAINED — $N = 1$, no competitive baseline	0.85
"May offer route to near-term quantum advantage"	NOT SUSTAINED — promissory, 10^9 energy penalty, 7-assumption conjunction	0.90

Bottom line: Valuable benchmarking paper. Premature advantage claims. The community should adopt the benchmarking methodology and reject the conclusions about quantum advantage.

5. Calibration Register

To prevent post-hoc rationalization, we register five falsifiable predictions. Each will be checked at its target date.

ID	Prediction	Check Date
CAL-BQNN-01	By 2028, an independent group will have replicated BQNN on hardware with $N \geq 500$ images and reported whether $a \approx 0.5$ advantage exceeds a classical baseline with $p < 0.05$ after correction for multiple comparisons	2028
CAL-BQNN-02	By 2028, no group will have demonstrated an entangled BQNN with partial mid-circuit measurement achieving $\geq 60\%$ layer fidelity (3+ layers, 16+ qubits per layer) on any platform	2028
CAL-BQNN-03	By 2030, the joules-per-solution ratio (BQNN inference / classical inference) for an MNIST-equivalent task will not have improved beyond $10^3 : 1$ on any platform	2030
CAL-BQNN-04	By 2030, multi-platform QNN benchmarking will have been adopted by ≥ 3 independent groups using mutually comparable metrics, but BQNN will not be the architecture that achieves any quantum advantage claim	2030
CAL-BQNN-05	By 2027, a classical binarized network with tuned stochastic regularization will match or exceed BQNN's performance on the same 55-image MNIST subset, rendering the "quantum noise as advantage" claim moot	2027

6. Discussion

6.1 The Broader Pattern

The BQNN paper exhibits a pattern consistent with the broader quantum computing ecosystem documented in "The Qubit Delusion" [1] and "The Problem-Substrate Mapping" [4]: a legitimate methodological contribution is framed as preliminary evidence for a paradigm-shifting advantage. This pattern — publish the engineering, claim the revolution — satisfies the institutional incentives of the field (funding agencies reward bold promissory language) while deferring falsification to future hardware that may never arrive.

The BQNN case is particularly instructive because the gap between the paper's actual contribution (multi-platform benchmarking) and its promissory framing (near-term quantum advantage) is both large and measurable. The seven-assumption dependency chain, the 9-order energy penalty, and the absence of any competitive classical baseline are not subtle — they are structural features of any QNN architecture that trains classically and infers quantum-mechanically on a separable circuit.

6.2 The Galitski Group Research Program

The BQNN paper is the third link in a chain of papers from the Galitski group at the University of Maryland: NN as Spin Models (2024) [6] → Natural Quantization of Neural Networks (2025) [7] → BQNN Hardware Benchmarking (2025) → Quantum-stabilized Hopfield Networks (2026) [8]. Each paper builds on the previous, and the internal consistency of the program is impressive. However, **no independent group has replicated any step**. The entire chain — from the NN ↔ spin correspondence to the hardware demonstration to the promissory non-simulability — rests on a single research group's work.

This is not a criticism of the Galitski group specifically — all emerging research programs begin with a single group. It is a call for independent replication and competitive benchmarking. Until another group runs a BQNN-equivalent on different hardware, with a pre-registered competitive classical baseline, and reports whether the advantage claims hold, the appropriate epistemic posture is skepticism, not acceptance.

6.3 Recommendations

1. **Adopt the benchmarking methodology, not the architecture.** Multi-platform QNN comparison is valuable infrastructure. Standardize metrics, pre-register hypotheses, and require competitive classical baselines.
2. **Build the competitive classical baseline that BQNN did not.** A classical binarized network with tuned stochastic regularization (dropout, noise injection, ensemble averaging) matched to the same 55-image MNIST subset would directly test the null hypothesis at near-zero cost.
3. **Track the MIPT-bridged QNN trajectory.** If an entangled BQNN with partial mid-circuit measurement achieves non-simulability on real hardware, the assessment should be revisited (see CAL-BQNN-02).
4. **Do not invest resources in BQNN architecture scaling.** With $P(\text{PS4}) \approx 0.03$ and a 10^9 energy penalty, the expected value of BQNN architecture development is dominated by more promising approaches.

7. Conclusion

The BQNN paper by Lakhdar-Hamina et al. (2025) is a competently executed engineering demonstration of multi-platform quantum neural network inference. Its benchmarking methodology — the first comparison of QNN performance across three distinct hardware platforms — is novel, valuable, and should be adopted by the community.

Its quantum advantage claims are not supported by the data. The statistical evidence is insufficient ($N = 55$, overlapping error bars), the noise-as-advantage claim rests on a single cherry-picked image, the circuit is fully separable and classically simulable, and the pathway to non-simulability requires seven unvalidated assumptions to all be true simultaneously. The joules-per-solution ratio is approximately 10^9 above the classical break-even point. We rate the paper's strongest claim (multi-platform benchmarking) as **SUSTAINED** and its advantage claims as **NOT SUSTAINED**.

The distinction between legitimate benchmarking and premature advantage framing is not academic — it is the difference between honest science and promissory engineering. BQNN contributes to the former; its conclusions slip into the latter.

Data Availability

All research artifacts — including the Phase 1 due diligence report, Phase 2 literature search and classification, and Phase 4 Bayesian cascade analysis with 5-adversary red-team audit — are archived alongside this paper. The BQNN paper itself is available at [arXiv:2507.21222v2](https://arxiv.org/abs/2507.21222v2) and via the DOI associated with its PRL publication. The QNFO framework papers referenced herein are available at papers.qnfo.org.

This paper is part of the Qubit Delusion series. All findings, calibration entries, and assessments are dated and versioned for independent verification. The audit was conducted on 2026-07-25 using the QNFO Research Pipeline v2.16.

References

- [1] QNFO, "The Qubit Delusion: How Particle Ontology Sabotaged Quantum Computing" (2026). papers.qnfo.org.
- [2] QNFO, "The Physics of Computation: Fundamental Limits and the Honest Boundaries of Post-Classical Computing" (2026). papers.qnfo.org.
- [3] R. B. Quni, "Shor's Assumptions: A Critical Re-examination of Quantum Advantage Foundations," v2.0 (2026). <https://papers.qnfo.org/papers/shor-assumptions-audit/>.
- [4] QNFO, "The Problem-Substrate Mapping: A Framework for Honest Computational Investment" (2026). papers.qnfo.org.
- [5] D. Lakhdar-Hamina, X. Liu, R. Barney, S. H. Miller, A. M. Green, N. M. Linke, and V. Galitski, "Benchmarking a Tunable Quantum Neural Network on Trapped-Ion and Superconducting Hardware," PRL / [arXiv:2507.21222v2](https://arxiv.org/abs/2507.21222v2) (2025).
- [6] R. Barney, M. Winer, and V. Galitski, "Neural Networks as Spin Models: From Glass to Hidden Order Through Training," [arXiv:2408.06421](https://arxiv.org/abs/2408.06421) (2024).
- [7] R. Barney, D. Lakhdar-Hamina, and V. Galitski, "Natural Quantization of Neural Networks," [arXiv:2503.15482](https://arxiv.org/abs/2503.15482) (2025).
- [8] R. D. Barney, S. Bhattacharjee, and V. Galitski, "Quantum-stabilized patterns in a vector Hopfield network," [arXiv:2606.06597](https://arxiv.org/abs/2606.06597) (2026).
- [9] N. H. Nguyen, E. C. Behrman, and J. E. Steck, "Quantum Learning with Noise and Decoherence: A Robust Quantum Neural Network," Quantum Information Processing 16 (2016). [arXiv:1612.07593](https://arxiv.org/abs/1612.07593).
- [10] Y. Li, X. Chen, and M. P. A. Fisher, "Quantum Zeno effect and the many-body entanglement transition," Phys. Rev. B 100, 134306 (2019).
- [11] S. Manna, A. Lakshminarayan, and V. Madhok, "Localizable Entanglement as an Order Parameter for Measurement-Induced Phase Transitions," [arXiv:2601.14185](https://arxiv.org/abs/2601.14185) (2026).
- [12] B. Khanal, P. Rivas, and A. Sanjel, "Generalization Error Bound for Quantum Machine Learning in NISQ Era — A Survey," [arXiv:2409.07626](https://arxiv.org/abs/2409.07626) (2024).
- [13] F. Góis, M. Pezzutto, and Y. Omar, "Energetics of Trapped-Ion Quantum Computation," [arXiv:2404.11572](https://arxiv.org/abs/2404.11572) (2024).
- [14] M. Courbariaux, Y. Bengio, and J.-P. David, "BinaryConnect: Training Deep Neural Networks with binary weights during propagations," in Advances in Neural Information Processing Systems, Vol. 28 (2015).
- [15] C. M. Bishop, "Training with Noise is Equivalent to Tikhonov Regularization," Neural Computation 7, 108 (1995).
- [16] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, "Demonstration of a small programmable quantum computer with atomic qubits," Nature 536, 63 (2016).