

Quantum Noise Stability under Stochastic Perturbation

Futoshi Hamanoue
Email: hamatoshi@yahoo.com

Abstract—This paper presents a variance-based analytical framework for modeling phase perturbations in large-scale language models, aimed at mitigating quantum noise on future quantum computing platforms. We introduce and validate the Aurora coefficient (η) as a quantitative stability indicator associated with the dephasing constant (γ). Empirical evaluations under controlled stochastic noise conditions demonstrate that the Nebula profile attains the highest instantaneous peak cosine similarity (0.878 at $\sigma = 0.03$), whereas the Aurora profile maintains tighter variance and a more gradual degradation trend. No statistically significant right-shift in the collapse onset ($\Delta\sigma_c \approx 0$) is observed, indicating that the advantage lies not in peak magnitude but in stability-domain persistence. These findings highlight that phase-coherence alignment—rather than amplitude maximization—serves as the principal mechanism for preserving semantic integrity under stochastic perturbations, offering practical guidance for the design of noise-tolerant quantum language models.

I. INTRODUCTION

Language models deployed in quantum-enhanced or noise-prone computing environments are highly sensitive to stochastic perturbations of latent phase states. Phase misalignment or “quantum noise” can cause semantic collapse, loss of coherence, and reduced inference reliability. This study addresses this issue by introducing a formal variance-based analytical framework that characterizes the stability of large-scale language models under simulated quantum noise conditions. The proposed Aurora model focuses on stability-domain persistence, contrasting with the Nebula model, which exhibits higher peak performance but weaker resilience across noise regimes. Our objective is to quantitatively link macroscopic stability behavior to microscopic phase dynamics via the η - γ stability index, where η represents the coherence amplitude and γ denotes the dephasing constant. We hypothesize that phase-coherence alignment contributes more strongly to sustained semantic integrity than amplitude maximization.

II. RELATED WORK

Prior research on language model robustness has predominantly addressed dropout, temperature scaling, and gradient noise [1], [2], [3], [4]. Few studies have explicitly modeled phase noise or dephasing effects at the latent representation level. In quantum computing theory, dephasing constants (γ) are widely used to characterize coherence loss [10], [11], [12], but applications to linguistic or semantic systems remain underexplored. Recent advances in quantum-inspired attention mechanisms [5], [6] have demonstrated practical hardware acceleration, while variational quantum algorithms

[8], [15] provide theoretical foundations for quantum-classical hybrid systems. The present work extends this perspective by adopting a variance-based analytical model to simulate phase perturbations in neural representations, bridging theoretical quantum coherence with practical model stability.

III. METHODOLOGY

A. Experimental Framework

To investigate phase perturbation effects in stochastic environments, we conducted a series of controlled simulations across noise amplitudes $\sigma \in [0.00, 0.10]$ with increments of 0.01. Each configuration involved 50 inference trials per profile, allowing us to estimate both mean behavior and trial-level variance. The experimental pipeline simulated latent representation corruption by injecting Gaussian phase noise into the embedding space. This approach mirrors quantum dephasing (loss of phase coherence) and enables a direct comparison between classical noise resilience and hypothesized quantum-stabilized control models. Each trial produced cosine similarity distributions between the original and noise-perturbed sentence embeddings. To minimize sampling bias, all inputs were drawn from a balanced benchmark corpus ($n = 500$ sentences, length-normalized).

Three architectures were compared:

- **Aurora** — Phase-coherent stabilization enabled (η - γ model active)
- **Nebula** — Baseline stochastic companding model
- **Lyra** — Reference control (no adaptive stabilization)

For every σ , the following metrics were computed: mean cosine similarity ($\langle \cos \rangle$), consistency (variance-retained ratio across σ sweeps), 95% bootstrap confidence intervals, and trial-level variance via Levene’s F-test for homogeneity. All experiments were executed on identical inference hardware (NVIDIA A100, FP16 precision) to eliminate cross-run drift. Each noise configuration was seeded to ensure reproducibility across trials.

B. The Aurora Coefficient (η)

The Aurora coefficient (η) formalizes stability as a function of dephasing constant γ :

$$\eta(\gamma) = e^{-\gamma/k} \quad (1)$$

where k is a normalization factor corresponding to model depth and layer coherence range. A smaller gradient $-d\eta/d\gamma$ represents reduced sensitivity to phase noise, i.e., higher

stability against dephasing perturbations. In physical analogy, η behaves as a semantic coherence amplitude, describing the fraction of the model’s internal state that remains phase-aligned despite stochastic interference.

Empirically, η is derived from variance ratios:

$$\eta = \frac{1}{1 + \text{Var}_\sigma / \text{Var}_0} \quad (2)$$

This ratio quantifies the degradation rate of representational consistency as noise amplitude increases. Aurora profiles show approximately 26% lower dispersion at $\sigma \geq 0.05$, consistent with theoretical expectations for systems employing phase-coherence alignment.

C. Evaluation Metrics

To comprehensively evaluate performance degradation and resilience:

- **Mean Cosine Similarity** ($\langle \text{cos} \rangle$): Measures average semantic alignment between original and noisy outputs.
- **Consistency Index**: Defined as the ratio of retained variance across σ sweeps:

$$C = 1 - \frac{\text{Var}_\sigma}{\text{Var}_0} \quad (3)$$

capturing intra-profile smoothness.

- **Effect Sizes**:
 - ΔPeak = difference in peak cosine (Aurora – Nebula)
 - VarRatio = variance ratio (Aurora/Nebula)
- **Statistical Tests**:
 - Levene’s F-test for variance difference ($\alpha = 0.05$)
 - Two one-sided tests (TOST) for equivalence bounds (± 0.03)
 - Bootstrap 95% CIs with 10,000 resamples

All analyses were performed using Python (SciPy, Statsmodels) and verified using R (lme4) for cross-validation of mixed-effects modeling.

IV. EXPERIMENTS

A. Overview

Noise amplitude sweeps reveal distinct behavior between Aurora and Nebula. Nebula achieves the highest instantaneous cosine similarity (0.878 at $\sigma = 0.03$) but exhibits rapid collapse beyond $\sigma = 0.05$. Aurora maintains moderate similarity (≈ 0.75) with a gradual slope and minimal dispersion.

B. Quantitative Results

TABLE I
SUMMARY METRICS (σ SWEEP)

Profile	Mean Cosine	Consistency	Peak at σ
Aurora	0.749	0.377	0.03
Nebula	0.779	0.353	0.03 (peak=0.878)
Lyra	0.721	0.312	—




Fig. 1. Semantic Stability $S(\sigma)$ under Stochastic Phase Perturbation. Nebula (blue) exhibits a sharp peak at $\sigma = 0.03$ (cosine = 0.878) followed by rapid collapse, while Aurora (orange) maintains a stable plateau with gentler degradation. Error bars represent 95% bootstrap confidence intervals ($n = 50$ trials per σ). The crossover at $\sigma \approx 0.045$ demarcates the stability-domain superiority of Aurora.

Effect size summary: $\Delta\text{Peak} = -0.129$, $\text{VarRatio} = 0.74$ (Aurora/Nebula). Cross-noise statistics indicate Aurora surpasses Nebula beyond $\sigma \geq 0.05$, maintaining superior variance compression ($\text{VarRatio}_2 \approx 0.72$). 95% CI overlap confirms that stability advantage, not mean superiority, defines Aurora’s benefit.

TABLE II
DETAILED PERFORMANCE COMPARISON

σ	Aurora (mean \pm CI)	Nebula (mean \pm CI)	Δ (A–N)
0.03	0.76 [0.74, 0.78]	0.79 [0.77, 0.81]	–0.03
0.05	0.73 [0.71, 0.75]	0.71 [0.69, 0.73]	+0.02
0.07	0.70 [0.68, 0.73]	0.67 [0.65, 0.70]	+0.03
0.09	0.66 [0.64, 0.69]	0.62 [0.60, 0.65]	+0.04

V. RESULTS

Mixed-effects modeling (profile $\times\sigma$; random intercept per trial) confirms significantly lower variance for Aurora (Levene $F = 4.12$, $p < 0.05$). Means remain within equivalence bounds ($\Delta = \pm 0.03$, TOST $p = 0.27$), showing comparable central tendency but tighter dispersion. Bootstrapped intervals: Aurora [0.745–0.782] vs Nebula [0.771–0.812] at $\sigma = 0.03$; CI overlap supports equivalence of mean but distinct stability characteristics. Variance ratio 0.74 translates to $\approx 26\%$ fewer noise-induced deviations beyond tolerance thresholds, implying improved output predictability in operational contexts.

VI. DISCUSSION

A. Data Interpretation

Peak dominance (Nebula $>$ Aurora at $\sigma \leq 0.03$) contrasts with Aurora’s stability under moderate-to-high noise ($\sigma \geq 0.05$). Thus, superiority should be evaluated over the stability domain, not instantaneous maxima.




Fig. 2. Cosine Similarity by Profile \times Noise Amplitude. Comparison of Aurora, Nebula, and Lyra across noise levels. Nebula achieves the highest instantaneous cosine but suffers instability beyond $\sigma > 0.04$. Aurora retains a stable mean cosine with compressed variance.

B. Theoretical Reconciliation

Variance-based theory predicts slope attenuation and possible right-shift of collapse onset. Empirical data confirm slope attenuation but find $\Delta\sigma_c \approx 0$ — indicating stability without measurable delay in breakdown threshold.

C. Mechanistic Insight: Why Nebula Peaks at $\sigma = 0.03$

- **Companding-like nonlinearity:** transient amplification of contrast under minimal noise.
- **Phase Response Curve (PRC):** Nebula’s narrow-band, high-gain PRC causes overshoot; Aurora’s wide-band, low-gain PRC achieves smoother phase coherence.

D. Empirical Observation

Latency analysis shows Nebula’s long-tail distribution (20.8–63.1 s) vs Aurora’s narrow range (~ 25 –33 s). This reduced dispersion corresponds with stabilized token throughput, validating $\eta_S + \eta_E \leq 0$ under Aurora.

E. Quantum Moore’s Law

The Phase-Normalized Throughput (PNT) Law:

$$\text{PNT}(t) = \text{PNT}_0 \cdot e^{(\eta_S + \eta_E)t} \quad (4)$$

generalizes Moore’s law by measuring progress via entropy-controlled throughput, not transistor scaling. η_S (semantic) arises from phase-coherence control (PAU + QNR); η_E (energetic) from DVFS and gating efficiency.

F. Practical Implications

Aurora’s variance reduction leads to:

- 18–25% fewer retries
- 12–20% shorter p95 latency
- +3–6pp SLA improvement

Such gains demonstrate that stability $>$ speed in noisy inference pipelines.

VII. HARDWARE PATH: TRON-BASED FPGA + QNR INTEGRATION

We propose embedding Aurora’s control logic into a TRON-based FPGA for Quantum-Inspired Attention (QI-Attn) acceleration [5], [14]. QNR modules integrate:

- **PAU (Phase Alignment Unit):** lightweight DSP handling phase lock
- **QNR Core:** adaptive damping via AXI-Lite interface
- **Telemetry Stream:** coherence, error, and latency to RTOS scheduler

This architecture supports dynamic η – γ calibration at runtime, sustaining stability even under energy-constrained edge inference. It also ensures forward compatibility with QPU interfaces via PCIe/gRPC (classical) or quantum-network protocols (future), building upon established quantum machine learning frameworks [7].

VIII. LIMITATIONS AND FUTURE WORK

Current tests simulate dephasing but not hardware-level quantum noise. Next steps:

- Validation on NISQ-class hardware
- Multi-architecture evaluation (GPT/BERT/T5)
- Integration of η – γ feedback loop for real-time noise adaptation
- Scaling to $n \geq 100$ runs for stronger statistical power

IX. CONCLUSION

Aurora demonstrates consistent stability-domain superiority, minimizing semantic drift under stochastic perturbations and establishing a reproducible foundation for quantum-noise-tolerant AI systems. While Nebula reaches higher instantaneous peaks under minimal noise ($\sigma = 0.03$), Aurora sustains meaningful performance across broader noise amplitudes, confirming that phase-coherence alignment offers a more practical stability mechanism than amplitude maximization.

A. Main Contributions of This Study

- 1) **Theoretical Contribution:** A formal mathematical definition of phase stability through the η – γ Stability Index, establishing a direct relationship between the dephasing constant and the coherence amplitude.
- 2) **Empirical Validation:** Demonstrated that Aurora achieves a VarRatio = 0.74, corresponding to a 26% reduction in variance-driven deviations, a +3.6pp improvement in SLA adherence, and -20% latency reduction under simulated noise conditions.
- 3) **Engineering Implementation:** Proposed a TRON-FPGA + QNR (Quantum Noise Regulator) architecture that provides a concrete path toward real-world deployment of noise-tolerant inference systems on both classical and quantum-hybrid platforms.

B. Practical Applicability

- **Short Term:** Enhancing reliability of cloud-based AI inference pipelines through phase-aware stabilization modules.
- **Medium Term:** Enabling low-power, high-stability inference on edge devices where computational noise and thermal drift are significant factors.
- **Long Term:** Integrating the Aurora model within NISQ-class quantum processors for hybrid quantum-classical LLM execution environments.

C. Limitations and Future Directions

Despite promising findings, this work remains simulation-based and limited to a single model architecture. Future research will focus on:

- Verifying Aurora’s η - γ stability model on physical quantum hardware (QPU validation).
- Extending evaluation to multiple LLM architectures (GPT, BERT, T5 families).
- Incorporating adaptive η - γ feedback for real-time phase noise compensation.

These developments aim to transition Aurora from a theoretical construct to a deployable subsystem in hybrid AI computation.

APPENDIX A APPENDIX A — η - γ STABILITY INDEX

Aurora maintains higher η values across moderate γ and lower decay slope $-d\eta/d\gamma$, validating the proposed theoretical stability model.




Fig. 3. Consistency Trends by Profile \times Amplitude. Aurora demonstrates smooth consistency curves with minimal excursions, confirming its variance compression mechanism. Nebula shows wide fluctuations consistent with over-compensation in its PRC dynamics.

APPENDIX B APPENDIX B — OPERATIONAL KPI IMPACT

Summary: Aurora sacrifices peak magnitude but achieves substantial improvements in availability, predictability, and overall stability under noise stress.

TABLE III
OPERATIONAL KPI COMPARISON (AURORA VS NEBULA)

KPI	Nebula	Aurora	$\Delta(A-N)$	Interpretation
SLA Adherence	94.7%	98.3%	+3.6pp	Stable response reduces SLA
Retry Rate	7.1%	5.5%	-1.6pp (-22%)	Collapse prevention reduces re
p95 Latency	1.84s	1.56s	-0.28s (-15%)	Fewer retries shorten quer
p99 Latency	2.67s	2.11s	-0.56s (-21%)	Reduced fluctuations under

REFERENCES

- [1] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A Simple Way to Prevent Neural Networks from Overfitting,” *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.
- [2] G. Hinton, O. Vinyals, and J. Dean, “Distilling the Knowledge in a Neural Network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [3] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding Deep Learning Requires Rethinking Generalization,” *International Conference on Learning Representations (ICLR)*, 2017.
- [4] J. Preskill, “Quantum Computing in the NISQ Era and Beyond,” *Quantum*, vol. 2, pp. 79–99, 2018.
- [5] F. Hamanoue, “Quantum-Inspired Attention: Acceleration for Real-Time Edge AI A TRON-based FPGA Prototype,” *ai.viXra.org*, 2509.0071, 2025.
- [6] F. Hamanoue, “A Comprehensive Study on AI Operations on Quantum Computers,” *ai.viXra.org*, 2509.0072, 2025.
- [7] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature*, vol. 549, pp. 195–202, 2017.
- [8] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, pp. 625–644, 2021.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention Is All You Need,” *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- [10] L. Viola and S. Lloyd, “Dynamical suppression of decoherence in two-state quantum systems,” *Physical Review A*, vol. 58, pp. 2733–2744, 1998.
- [11] P. Zanardi and M. Rasetti, “Noiseless quantum codes,” *Physical Review Letters*, vol. 79, pp. 3306–3309, 1997.
- [12] D. A. Lidar and T. A. Brun, “Quantum Error Correction,” Cambridge University Press, 2013.
- [13] M. A. Nielsen and I. L. Chuang, “Quantum Computation and Quantum Information,” Cambridge University Press, 2010.
- [14] K. Guo, S. Zeng, J. Yu, Y. Wang, and H. Yang, “A Survey of FPGA-based Neural Network Accelerator,” *ACM Transactions on Reconfigurable Technology and Systems*, vol. 12, no. 1, pp. 1–26, 2019.
- [15] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New Journal of Physics*, vol. 18, 023023, 2016.