Fundamental Finitary Limits to Maximal Mutual Information

A. Chawla
Research Enlightenment And Life Institute, Gurugram, Haryana
editor@realinstitute.live

***Abstract*—In both classical and quantum communication, the idea of perfect noiseless transmission remains an idealization. This paper draws a novel connection between two seemingly unrelated noise floors: the irreducible geometric distortion implied by the manifold hypothesis (MH) in high-dimensional data representation, and the uncertainty-limited precision of quantum nondemolition (QND) measurements using squeezed states. We argue that MH introduces a classical analog of quantum measurement back-action, thereby bounding the capacity of even ideal channels. As a consequence, well-known paradoxes in control theory—such as the insufficiency of infinite-capacity channels for stabilization—can be reinterpreted without invoking hypothetical channels of unbounded Shannon capacity. We propose that this irreducible error, rooted in geometry rather than randomness, constitutes a minimal noise floor that must be engineered around rather than eliminated.**

Key Insight

Even in the absence of stochastic noise, the manifold hypothesis implies a geometric distortion floor, meaning that no communication channel—regardless of Shannon capacity—can be truly noiseless.

The manifold hypothesis, a common assumption in machine learning and signal processing, posits that high-dimensional data often lies on or near a lower-dimensional manifold embedded in ambient space [1]. This geometric constraint has profound implications beyond representation learning; it imposes a structural limitation on the fidelity of encoding and decoding processes. If the data lives on a nonlinear manifold, any attempt to tile the space with uniform quantizers or packing spheres—such as in traditional channel coding—inevitably results in mismatch and distortion. This leads to an in-principle error, not due to stochastic noise but due to geometric incompatibility. Unlike classical additive noise models, this distortion persists even in the limit of perfect hardware and noiseless transmission.

This leads to a surprising insight: Shannon's original sphere-packing bound, while stochastic in formulation, hides an implicit geometric assumption. In traditional coding theory, capacity is defined in terms of achievable rates given additive noise [2]. However, if the codewords are drawn from or intended to represent points on a manifold, then a standard uniform quantizer induces errors orthogonal to the manifold—errors that cannot be made arbitrarily small without infinite resolution adapted to the manifold's geometry [3], [4]. This suggests a classical analog to the Heisenberg uncertainty principle: an irreducible limit to precision arising from the structure of the signal space itself.

This phenomenon mirrors the limits encountered in quantum nondemolition (QND) measurements with squeezed states. In QND measurement, one attempts to measure an observable without perturbing it, which is only possible for a subset of observables that commute with the system's dynamics [5], [6]. Squeezed states allow one to reduce measurement noise in one observable at the cost of increased uncertainty in its conjugate. But crucially, total uncertainty remains bounded by Heisenberg's limit. Similarly, encoding data along a manifold can be made highly precise only at the cost of increased vulnerability to errors in the orthogonal directions. No encoding strategy can eliminate this fundamental trade-off. This analogy draws a conceptual bridge between quantum metrology and classical information theory, both governed by irreducible geometric constraints.

The consequences extend into control theory. In a widely cited argument, Sahai and Mitter demonstrated that certain unstable systems cannot be stabilized over noisy channels—even those with infinite Shannon capacity—due to a lack of anytime reliability [7]. But if our geometric reasoning holds, then truly infinite-capacity channels cannot exist in physical reality, as the manifold-induced error persists even in the absence of stochastic noise. Hence, the foundational paradox motivating the definition of anytime capacity—that infinite capacity can still be insufficient—may be reinterpreted as a geometric constraint. There are no such channels in practice; geometry might forbid them. This strengthens, rather than weakens, the case for anytime communication, as it emerges not from paradox but from necessity.

Furthermore, this geometric noise floor offers a unifying explanation for the empirical success of generative compression and manifold-aware learning models. Neural codecs such as VAEs and diffusion models effectively learn data manifolds and map them to lower-dimensional latent spaces [8], [9]. These systems are not immune to noise, but they encode signal variability in a way that respects manifold constraints, effectively reducing the mismatch error present in traditional quantizers. Similarly, recent work in neural error-correcting codes and learned source-channel coding seeks to align the codebook with the data geometry, implicitly working around the manifold-induced distortion floor [10], [11].

In summary, we propose that the manifold hypothesis imposes a minimal noise floor in communication systems, even in the absence of stochastic channel noise. This parallels the quantum measurement problem, where no measurement can avoid introducing disturbance due to structural constraints. Both settings are governed not by randomness but by irreducible geometry. The task for future systems is not to eliminate this noise, but to engineer around it: by adapting codes to the manifold, using generative models to flatten geometry, or invoking hierarchical representations to progressively reduce error. Just as quantum engineers have learned to live with the Heisenberg limit via squeezing and error correction, classical communication systems must accommodate their own fundamental limits—limits not of bits, but of shape. We invite readers to consider how geometric distortion may reshape bounds across coding, control, and learning.

Acknowledgment

Prepared in collaboration with an LLM.

Appendix

Geometric Illustration of Manifold-Induced Distortion

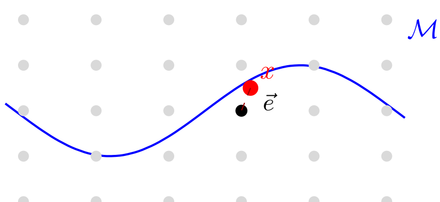

Fig. 1: Illustration of quantization mismatch: a point $x \in \mathcal{M}$ lies off the uniform grid of quantization centers. The projection induces an irreducible geometric error $\vec{e}$ even in the absence of stochastic noise.

Mathematical Footnote.: Let $\mathcal{M}$ be a smooth d-dimensional Riemannian manifold embedded in $\mathbb{R}^n$, with curvature bounded above by $\kappa > 0$. Suppose quantization centers form a grid $\mathcal{C}_\delta$ of spacing δ in the ambient space $\mathbb{R}^n$. Then the minimum distortion due to off-manifold projection satisfies:

$$\min_{c \in \mathcal{C}_\delta} \|x - c\| \geq \delta \cdot \sin\theta(x) \quad \text{for } x \in \mathcal{M}, \tag{1}$$

where $\theta(x)$ is the angle between the tangent space $T_x\mathcal{M}$ and the vector to the nearest center. Since $T_x\mathcal{M}$ varies continuously and curvature is nonzero, there exists $\theta_{\min} > 0$ such that $\sin\theta(x) \geq \sin\theta_{\min} > 0$ locally. Thus, no finite-resolution uniform grid avoids this geometric mismatch.

Conclusion.: This geometric distortion is irreducible and structural: it arises not from probabilistic noise but from curvature and dimensionality reduction. As grid resolution $\delta \to 0$, $\|e\| \to 0$ only if the grid adapts to the manifold's tangent geometry, such as via learned or geodesic tiling—supporting our thesis that there exists a minimal classical noise floor imposed by geometry.

Appendix

On Epistemic vs. Representational Limits

To preempt potential conflation, we offer a brief clarification on the nature of the limitations discussed, particularly in distinguishing quantum back-action from classical geometric distortion.

The representational limit imposed by the manifold hypothesis is structural: data constrained to a nonlinear manifold cannot be faithfully encoded using uniform quantization schemes in the ambient space without incurring error orthogonal to the manifold. This distortion arises from the mismatch between the intrinsic geometry of the data and the geometry of the encoding space, and persists regardless of observer knowledge or probabilistic uncertainty. It is, in that sense, a representational or ontological limit.

In contrast, quantum back-action stems from the dynamical (cf. the Heisenberg equations of motion) non-commutativity of observables in quantum mechanics. Specifically, let $\hat{x}$ and $\hat{p}$ be canonically conjugate observables satisfying $[\hat{x}, \hat{p}] = i\hbar$. Any attempt to reduce the uncertainty in $\hat{x}$ necessarily increases uncertainty in $\hat{p}$, not due to lack of knowledge, but because the Heisenberg equations of motion ensure that the act of measuring $\hat{x}$ perturbs $\hat{p}$:

$$\frac{d\hat{p}}{dt} = -\frac{\partial \hat{H}}{\partial \hat{x}} \quad \text{and} \quad \frac{d\hat{x}}{dt} = \frac{\partial \hat{H}}{\partial \hat{p}}, \tag{2}$$

where $\hat{H}$ is the system's Hamiltonian. In a quantum nondemolition (QND) measurement, one attempts to circumvent this by choosing observables that commute with the Hamiltonian, thereby limiting the induced back-action. But even then, the Heisenberg uncertainty principle bounds the joint precision with which certain observables can be known or preserved under measurement.

This stands in contrast to classical statistical uncertainty, which reflects epistemic limitations—the observer's lack of complete knowledge about the system state. For example, additive white Gaussian noise in a channel is an epistemic model of environmental variability; it represents what we do not know or control. In contrast, both quantum back-action and manifold-induced distortion are not reducible to better modeling; they persist even with perfect information and ideal observers.

We thus reframe our central analogy: the manifold-induced distortion in classical systems is not an epistemic uncertainty, but a representational limit similar in spirit—but not in mechanism—to quantum measurement back-action. Both are governed by structural constraints: one by geometry, the other by algebraic commutation relations. Each imposes a minimal distortion that cannot be eliminated, only managed.

References

[1] C. Fefferman, S. Mitter, and H. Narayanan, "Testing the manifold hypothesis," *Journal of the American Mathematical Society*, vol. 29, no. 4, pp. 983–1049, 2016.

[2] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.

[3] V. Shah and S. Verdu, "The geometry of mutual information," *IEEE Transactions on Information Theory*, vol. 66, no. 1, pp. 402–432, 2020.
[4] M. Belkin, D. Hsu, and P. Mitra, "Reconciling modern machine-learning practice and the bias–variance trade-off," *PNAS*, vol. 116, no. 32, pp. 15849–15854, 2019.
[5] V. B. Braginsky and F. Y. Khalili, *Quantum Measurement*. Cambridge University Press, 1996.
[6] C. M. Caves, "Quantum-mechanical noise in an interferometer," *Physical Review D*, vol. 23, no. 8, p. 1693, 1981.
[7] A. Sahai and S. K. Mitter, "The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link," *IEEE Transactions on Information Theory*, vol. 52, no. 8, pp. 3369–3395, 2006.
[8] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," *arXiv preprint arXiv:1312.6114*, 2013.
[9] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *CVPR*, 2022.
[10] J. Choi and Y. C. Eldar, "Neural layered decoding for error correction," *IEEE Transactions on Neural Networks and Learning Systems*, 2019.
[11] E. Bourtsoulatze, D. B. Kurka, and D. Gesbert, "Deep joint source-channel coding for wireless image transmission," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 3, pp. 567–579, 2019.