THE ENTROPY OF CHOICE

A Boundary-Complexity Hypothesis for Decision Intelligence

Rajesh Singh
MetaForgeX AI (Pty) Ltd | NSDM Institute
Johannesburg, South Africa
Working Paper 2026-02 | Revision 2.0 | July 2026

RESEARCH STATUS

Independent working paper. Not yet peer-reviewed. This paper proposes an original boundary-complexity hypothesis and develops a covering-number argument under explicit assumptions. The links to rate-distortion theory, variational inference and holographic entropy are presented as research correspondences whose exact scope remains to be established. Quantum and hardware implications are stated only as conditional research questions. Empirical validation is in progress through the NSDM benchmark programme.

AI-assistance disclosure

AI was used as a research, drafting and pressure-testing tool. The author directed the argument, selected the claims, reviewed the sources and accepts responsibility for the final manuscript. No claim is presented as validated beyond the evidence stated in the paper.

Abstract

Many decision systems operate in very large state spaces, but the selected action does not change everywhere. It changes near a decision boundary: the region in which small differences in evidence, context or value can alter the action. This paper proposes that the useful complexity of a bounded decision policy may scale primarily with the geometric complexity of that boundary, rather than with the full volume of the state space. We call this proposal the Holographic Decision Bound. The argument is developed through a covering-number construction for distinguishable policy boundaries and compared with ideas from rate-distortion theory, variational inference and entropy-area relations in physics. These comparisons do not all have the same status. The covering-number result is a derivation under explicit regularity and local-complexity assumptions; the links to rate-distortion and free energy require further formalisation; and the connection to black-hole thermodynamics is a structured analogy based on a shared area-scaling form, not a physical equivalence. The paper concludes with falsifiable predictions. If boundary complexity rather than raw state-space size governs useful policy capacity, then model performance should show a measurable capacity knee once a model can represent the relevant decision boundary. Boundary-shell error, regret localisation and representation stability provide further tests. The theory should be judged by whether these predictions survive controlled experiments.

Keywords: decision boundaries; policy complexity; metric entropy; covering numbers; rate-distortion; variational inference; model efficiency; NSDM

1. The research question

Artificial-intelligence systems often receive high-dimensional inputs: text, images, behavioural signals, transactions and sensor streams. Yet the operational decision is usually much smaller. A system may recommend or not recommend, intervene or wait, approve or decline, answer or abstain. The central question of this paper is therefore not how much information exists in the observation space. It is how much information is needed to represent the places where the preferred action changes.

In a binary decision problem, those action-changing locations form a decision boundary. Points far from the boundary are comparatively easy: small perturbations do not alter the action. Points close to it are more difficult because modest changes in evidence, utility or context can reverse the decision. This suggests that model capacity should be studied in relation to boundary complexity rather than raw input dimension alone.

PLAIN-ENGLISH CLAIM

A model does not need equal precision everywhere. It needs the greatest precision where a small change can alter the decision.

2. Definitions and scope

Let M be a compact d-dimensional state-action manifold equipped with a metric g. Let a binary policy π map each state to one of two actions. Its decision boundary is the set at which the selected action changes:

$$\partial D\pi = \{x \in M : \pi \text{ changes action in every neighbourhood of } x\}.$$

Assume that the boundary belongs to a regular k-dimensional class, usually with $k = d - 1$ for a smooth codimension-one boundary. Let ε denote the smallest decision-relevant resolution in the chosen representation. Let $\Pi\delta$ denote a family of policies whose expected regret is within δ of the optimal policy:

$$\Pi\delta = \{\pi : E[L(\pi(X), \pi^*(X))] \le \delta\}, \quad X \sim PX.$$

Two policies are ε-distinguishable when their boundaries differ by at least ε under Hausdorff distance. The decision entropy is the logarithm of the covering or packing number required to distinguish policies at that resolution.

$$Schoice(\varepsilon, \delta) = \log N(\varepsilon, \Pi\delta, dH).$$

SCOPE LIMIT

Boundary measure depends on the representation and metric. The theory therefore requires a fixed metric, a controlled equivalence class of representations, or a representation-stable estimator. It is not yet a universal coordinate-free law.

3. Boundary covering hypothesis

A regular boundary can be covered by geodesic patches of radius ε. If $H_k(\partial D)$ denotes its k-dimensional Hausdorff measure and c_k is the volume of a unit k-ball, the number of patches is of order:

$$P(\varepsilon) \lesssim H_k(\partial D) / (c_k \varepsilon^k).$$

Suppose each patch admits at most λ locally distinguishable configurations at resolution ε. Then the number of distinguishable boundaries is bounded by λ raised to the number of patches. Taking logarithms gives the central scaling form:

$$S_{choice}(\varepsilon, \delta) \leq C \cdot \sup_{\pi\in\Pi_\delta} H_k(\partial D_\pi) / \varepsilon^k.$$

The constant C absorbs local configuration complexity, regularity and dependence between neighbouring patches. The inequality should be read as an upper bound for a specified policy class, not as an exact count for every decision problem.

PROPOSITION

For a regular family of near-optimal policies, the metric entropy of distinguishable decision boundaries is bounded above by boundary measure divided by the kth power of representation resolution, up to a class-dependent constant.

3.1 Why boundary complexity may matter more than volume

For a binary policy, the action is locally constant away from the boundary. Additional representation detail in those interiors may improve perception or estimation, but it does not create new action distinctions at resolution ε. By contrast, detail near the boundary can change the selected action and therefore carries direct decision value.

This does not mean that the bulk is literally irrelevant. The interior still determines evidence, causal structure and value. The narrower claim is that distinguishable action changes are organised by the boundary. Model capacity devoted to irrelevant ambient dimensions may therefore yield less decision value than capacity devoted to accurately representing the action-changing region.

4. Relationship to established fields

The paper uses three cross-field comparisons. They are deliberately separated by evidential status.

Field	Status in this paper	What is defensibly claimed
Covering numbers and metric entropy	Direct mathematical basis	The proposed bound is a covering-number argument for a regular boundary class.
Rate-distortion theory	Structured correspondence	Expected regret is often concentrated near action-switching regions, suggesting that lossy policy codes should preserve boundary-relevant information.
Variational free energy	Restricted correspondence	For a posterior family concentrated

Field	Status in this paper	What is defensibly claimed
		in a tubular neighbourhood of a boundary, the entropy term contains a boundary-measure contribution.
Black-hole holography	Analogy in scaling form	Both use an area-over-resolution form, but the physical mechanisms, variables and evidential status are different.

4.1 Rate-distortion and regret localisation

Rate-distortion theory asks how much information is needed to encode a source while keeping expected distortion below a target. For a policy, a natural distortion measure is regret. When the action-value gap grows with distance from the decision boundary, small representation errors change decisions mainly within a shell around that boundary. This suggests that an efficient policy code should allocate information preferentially to the boundary shell.

The present manuscript does not prove equality between the Shannon rate-distortion function and the proposed boundary entropy. It states a research programme: derive the rate-distortion function for explicit policy classes and test whether its leading behaviour is controlled by boundary geometry.

4.2 Variational inference and boundary-localised uncertainty

Variational free energy provides a tractable bound on negative log evidence within probabilistic inference. Friston's free-energy principle uses this formalism as part of a broad account of perception, action and learning. In the restricted case where residual uncertainty is concentrated in a narrow tube around a regular decision boundary, the normalising volume and entropy contain a term involving the boundary measure. This is a useful local model, but it does not prove that free energy is globally bounded by boundary area.

4.3 Holographic entropy as a disciplined analogy

Bekenstein's black-hole entropy and the holographic principle establish an area-related information bound in gravitational physics. The proposed decision bound shares a formal area-over-resolution structure. That shared form is intellectually suggestive, but it is not an identity of physical mechanism. The decision boundary is not an event horizon, ε is not the Planck length and no gravitational law is invoked.

5. Consequences for model efficiency

The practical hypothesis is that model capacity should display diminishing returns once the model can represent the task's relevant boundary at the required resolution. This predicts a capacity knee rather than a universal "small models beat large models" rule.

A larger model may still improve perception, data integration, robustness or transfer. The narrower prediction is that, after controlling those factors, extra capacity should add little decision value when the boundary representation is already sufficient.

$$W^* = \min \{W : \text{model class at size } W \text{ represents the target boundary within } \varepsilon\}.$$

The value W* must be estimated before inspecting the final performance curve. Retrospective identification of a knee is not strong evidence because almost any saturation curve can be described after the fact.

6. Empirical observability

A decision boundary is empirically meaningful only when it corresponds to a stable policy transition under an explicit value function. A class-probability threshold of 0.5 is not automatically a decision boundary. It is appropriate only when the action values are equal at that threshold.

BOUNDARY OBSERVABILITY CONDITION

Operational boundaries should be defined by action-value equality, zero expected regret, zero net uplift, treatment-effect indifference or another explicit decision criterion. Passive prediction thresholds are insufficient unless they coincide with that criterion.

6.1 Falsification protocol

Synthetic recovery
Known boundaries must be ordered correctly by the estimator; random-label controls must fail.

Boundary-shell regret
Material errors or regret should concentrate closer to the action-changing region.

Capacity-knee prediction
Estimated boundary complexity must predict the smallest near-best model before the performance curve is inspected.

Representation stability
The result must survive controlled invertible transformations or explicitly reveal the need for a better metric.

Cross-domain replication
The same estimator and preregistered decision rule must work across unrelated domains.

The operational theory should be considered unsupported if it passes synthetic recovery but fails both boundary-shell and capacity-knee tests across at least three unrelated real-world domains.

7. Research implications for NSDM

NSDM can use the boundary-complexity hypothesis as an experimental layer rather than a production fact. The immediate applications are measurement and model comparison, not automatic quantum-readiness scoring or physics-derived confidence limits.

NSDM component	Research use
Boundary estimator	Estimate local action sensitivity, geometric complexity and uncertainty around action changes.

NSDM component	Research use
Model-routing experiment	Test whether boundary estimates predict when escalation to a larger model produces measurable value.
Evidence-state diagnostics	Separate unsupported and under-specified cases that may appear close in embedding space but differ in decision status.
Action-state policy	Use consequence and governance conditions to set answer, abstention, verification and escalation thresholds.
Cost per justified decision	Measure compute expenditure against decisions that cross the evidence and governance boundary correctly.

8. Quantum and thermodynamic computing: research questions, not conclusions

Quantum feature maps can represent data in high-dimensional Hilbert spaces, and recent research continues to investigate conditions for useful quantum learning advantages. However, Hilbert-space dimension alone does not prove computational advantage. A valid claim must include data-access cost, training and measurement cost, noise, a classical lower bound and an end-to-end task comparison.

Thermodynamic computing is likewise a serious research direction. Extropic describes probabilistic hardware intended to sample programmable distributions more efficiently than deterministic implementations. That engineering programme does not yet establish a boundary-area scaling law for commercial decision tasks.

PUBLICATION RULE

The paper does not claim a qubit threshold, a 2028-2030 decision-advantage date, a 10,000x task-level gain or an automatic quantum-readiness score. Those statements remain hypotheses for future work.

9. Limitations

- Regularity: the bound assumes a boundary class with controlled local complexity.
- Dependence: neighbouring boundary patches are not generally independent.
- Representation dependence: boundary measure changes under latent-space transformations.
- Data distribution: geometric complexity alone does not capture where probability mass or consequence is concentrated.
- Multiclass and sequential decisions: the binary static formulation requires extension to multiple actions and trajectories.
- Causal validity: a predictive boundary is not automatically an intervention boundary.
- Governance: a technically accurate boundary does not establish authority, legitimacy or fairness.

- Empirical status: no completed experiment yet validates the proposed capacity-knee prediction across domains.

10. Conclusion

This paper advances a boundary-complexity hypothesis for decision intelligence. The central claim is not that decision systems obey black-hole physics. It is that useful policy capacity may scale with the complexity of action-changing regions rather than with the full volume of the state space.

The covering-number argument provides an initial mathematical basis under explicit assumptions. The relationships to rate-distortion theory, variational inference and holographic entropy are research correspondences whose exact scope remains to be established.

The value of the hypothesis is empirical as well as theoretical. If it is correct, model-capacity knees, regret localisation and cross-domain efficiency should be predictable from measurable properties of decision boundaries. Those predictions can be tested. The theory should be judged by whether they survive.

References

Bekenstein, J. D. (1973). Black holes and entropy. Physical Review D, 7(8), 2333-2346. https://doi.org/10.1103/PhysRevD.7.2333

Bousso, R. (2002). The holographic principle. Reviews of Modern Physics, 74(3), 825-874. https://doi.org/10.1103/RevModPhys.74.825

Friston, K. (2010). The free-energy principle: a unified brain theory? Nature Reviews Neuroscience, 11, 127-138. https://doi.org/10.1038/nrn2787

Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. Physical Review Letters, 122, 040504. https://doi.org/10.1103/PhysRevLett.122.040504

Shannon, C. E. (1959). Coding theorems for a discrete source with a fidelity criterion. IRE National Convention Record, 7(4), 142-163.

Susskind, L. (1995). The world as a hologram. Journal of Mathematical Physics, 36(11), 6377-6396. https://doi.org/10.1063/1.531249

Tishby, N., Pereira, F. C., & Bialek, W. (2000). The information bottleneck method. arXiv:physics/0004057.

Conte, T. et al. (2019). Thermodynamic computing. arXiv:1911.01968.

Extropic. (2026). Hardware and thermodynamic sampling platforms. Official product and research information, accessed July 2026.

Google Quantum AI. (2024). Meet Willow, our state-of-the-art quantum chip. Official research announcement.

SUGGESTED CITATION

Singh, R. (2026). The Entropy of Choice: A Boundary-Complexity Hypothesis for Decision Intelligence. NSDM / MetaForgeX AI Working Paper 2026-02, Revision 2.0. Johannesburg: MetaForgeX AI.