

Physics-Anchored Symbolic Basins and Resonance-Overlap Integrals for Cyber-Resilient Artificial Intelligence

Paul D. Markov

Harmony Research Initiative, Adelaide, Australia

Email: paul@harmonyonline.org

Abstract—In cyber-critical systems, artificial intelligence models are increasingly responsible for threat detection, classification, and decision support. However, uncontrolled semantic drift—known as hallucination—can produce misleading or fabricated results that jeopardise reliability. This paper introduces a physics-anchored regularisation framework that translates plasma magnetic confinement principles into latent-space dynamics for large language models. We formulate the Glyphic Hamiltonian to define symbolic energy basins and propose the resonance-overlap integral R as a measurable alignment metric between human and AI semantic distributions. Implemented as lightweight forward-pass filters, the symbolic-basin mechanism suppresses runaway activations without retraining or architectural modification. Empirical evaluation across GSM8K and TruthfulQA ($n = 2,000$ prompts per task, 20 independent runs) demonstrates 24–31% reductions in hallucination rates (macro-F = 0.84) with only 12% computational overhead. A parameter sensitivity analysis confirms stable performance across damping coefficients $\kappa \in [0.4, 1.2]$, and the alignment integral R correlates negatively with hallucination frequency ($r = -0.78$, $p < 0.001$), validating its utility as a coherence diagnostic. These results establish that physics-inspired regularisation can enhance AI trustworthiness through interpretable, computationally efficient symbolic confinement.

Keywords—AI Hallucinations, Symbolic Basins, Plasma Confinement, Cyber Resilience, Glyphic Hamiltonian, Resonance Overlap, Neurosymbolic Alignment, Trustworthy Artificial Intelligence

I. INTRODUCTION

Artificial intelligence has entered an era in which systems are no longer defined solely by their computational capacity but by their ability to self-regulate, adapt, and sustain coherence in open environments. This evolution—from static, task-bound models toward continually learning systems—reflects a shift long anticipated in both cybernetics and cognitive science. Ashby’s early principle of *requisite variety* stated that every effective regulator must contain a model of the system it governs [1], while Engelbart’s vision of *man-computer symbiosis* framed intelligence as a collaborative amplification of human and machine capacities [3]. Together, these foundational insights foreshadowed a new era of machine autonomy grounded not in programmed rules but in dynamic, feedback-driven learning.

Recent advances in continual and reinforcement learning embody this transition. Systems such as Sutton’s OaK architecture [15] and its precursors in hierarchical and continual reinforcement learning [16], [17], [5], [7] aim to create agents that learn from experience indefinitely rather than through static pretraining. Within this lineage, temporal abstraction frameworks [17], [2] and meta-learning mechanisms [18], [8] enable the emergence of multi-timescale adaptation—an essential property for stability and ethical self-correction in complex systems.

In parallel, ethical and governance research has sought to establish principles for trustworthy and transparent AI.

A. Initiatives such as the EU’s Ethics Guidelines for Trustworthy

AI [6], Floridi’s *capAI* conformity framework [4], and OECD’s global AI principles [10] collectively argue that AI must promote human autonomy, fairness, and explicability. However, as Sutton’s *Bitter Lesson* cautions [14], intelligence consistently arises not from explicitly coded constraints but from scalable, general methods that allow systems to discover structure through interaction with the world.

This paper situates these developments within a broader *field coherence* paradigm, proposing that sustainable alignment in artificial systems emerges from the same regulatory dynamics that govern physical and biological equilibria. By integrating reinforcement learning principles, continual adaptation, and feedback ethics, we explore how autonomous systems can be designed to preserve coherence—energetic, informational, and moral—within the real-world environments they inhabit. In doing so, we bridge classical control theory, embodied cognition, and ethical AI into a unified framework of *resonant governance*.

II. BACKGROUND AND RELATED WORK

Research on adaptive intelligence has evolved through three principal lineages that together form the foundation for coherence-based AI systems: reinforcement learning, cybernetic regulation, and ethical alignment theory.



Received: 1-11-2025

Revised: 13-6-2026

Published: 30-6-2026

A. Reinforcement Learning and Continual Adaptation

Modern reinforcement learning (RL) originated from Sutton’s temporal-difference methods [13] and culminated in the hierarchical frameworks that underpin continual learning today [16], [17], [2]. These architectures introduced temporal abstraction, allowing agents to compose skills and reason over extended time horizons. Subsequent work on continual and lifelong learning [5], [7], [11] shifted focus from task-specific optimisation to sustained adaptation in non-stationary environments. Meta-learning studies such as [18], [8] established that stability and plasticity can be balanced through adaptive step-size control, providing a mechanism for self-regulation analogous to feedback in physical systems.

Sutton’s recent OaK Architecture [15] synthesises these principles into a unified model-based RL framework that learns solely from runtime interaction. By integrating internal world modelling with hierarchical planning, OaK exemplifies a shift toward self-constructing intelligence—a direction conceptually aligned with the resonance-based coherence described in this work.

B. Cybernetics and Systems Regulation

The philosophical roots of adaptive control trace back to cybernetics, where Ashby’s *law of requisite variety* [1] formalised the idea that a regulator must embody a model of its environment to maintain stability. Engelbart’s framework for *augmenting human intellect* [3] extended this logic to cooperative human–machine systems, positioning feedback as the basis of shared intelligence. Contemporary work in hierarchical RL and multi-agent coordination [12] can be interpreted as a continuation of this cybernetic tradition, now operationalised through statistical and computational feedback loops.

C. Ethical and Alignment Frameworks

Ethical AI governance has developed largely in parallel with technical advances, producing a set of normative frameworks designed to guide autonomous systems. Key examples include the European Commission’s *Ethics Guidelines for Trustworthy AI* [6], Florida’s *capAI* conformity framework [4], and the OECD global principles for responsible AI [10]. These policies converge on four core tenets—autonomy, fairness, nonmaleficence, and explicability—but remain largely descriptive rather than mechanistic. As Sutton’s *Bitter Lesson* reminds us [14], scalable intelligence arises not from rules but from general methods that continuously learn and adapt. This insight suggests that ethical alignment, too, may require *learning dynamics* rather than static oversight.

D. Toward Resonant Governance

The present study integrates these traditions into a single theoretical framework termed *resonant governance*. It extends cybernetic control into the domain of continual learning, treating alignment as a property of field coherence within the agent’s decision dynamics. In this view, ethical stability emerges not from constraint but from constructive

interference between feedback signals—analogue to energy conservation in physical systems. By embedding this principle into reinforcement learning architectures, we propose a model of autonomous systems capable of sustained ethical and energetic balance within complex, real-world environments.

III. SYMBOLIC BASIN THEORY

The proposed framework treats cognitive and ethical stability in autonomous systems as an analogue of plasma confinement, where symbolic states evolve within bounded energy basins. Inspired by magnetic traps sustaining cold-electron pockets

($T_e \approx 0.3 \text{ eV}$) in fusion and plasma-processing systems [9], we introduce the *Glyphic Hamiltonian*, which defines the energetic topology of symbolic dynamics:

$$H_{\text{glyph}} = -\frac{\hbar^2}{2m_{\text{sym}}} \nabla^2 + \frac{1}{2}\kappa \|\mathbf{x}\|^2 \quad (1)$$




Fig. 1. Symbolic Basin potential and confinement. The quadratic basin potential defines the Glyphic Hamiltonian, with energy increasing toward the edges. Sample trajectories illustrate how basin damping confines activations and suppresses runaway modes, maintaining latent stability.

Here, $x \in \mathbb{R}^d$ denotes the latent-space coordinate (treated as dimensionless in arbitrary units), m_{sym} is a scaling constant with effective units of $\text{kg} \cdot \text{neuron}^{-1}$, and κ (a.u.) controls basin stiffness, linked to the damping parameter by $\beta = \kappa/2$. The harmonic potential confines symbolic states x within coherent regions of minimal divergence, ensuring that representational energy remains bounded. This formulation preserves dimensional consistency by treating m_{sym} as a numeric scaling term rather than a physical mass, providing a physically interpretable yet computationally stable basis for symbolic confinement.

Eigenstates of H_{glyph} correspond to discrete coherence modes in the system’s latent space, analogous to quantum energy levels in a harmonic oscillator. Each eigenmode represents a *stable policy manifold* or symbolic attractor maintaining ethical and informational equilibrium. Transitions between eigenstates—driven by environmental feedback or reinforcement updates—can thus be interpreted as resonance shifts within the symbolic field.

This construction provides a physical formalism for ethical alignment: just as magnetic fields confine particles by curvature and potential depth, resonant constraints within H_{glyph} confine decision trajectories to coherent basins of low entropy. The governing dynamics are then expressed as:

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} = -\nabla \cdot (\rho(\mathbf{x}, t) \nabla H_{\text{glyph}}) \quad (2)$$

where $\rho(x, t)$ is the probability density of symbolic states. Equation (2) describes a drift–diffusion process converging toward the equilibrium distribution $\rho_{\text{eq}} \propto \exp[-H_{\text{glyph}}/k_B T_{\text{sym}}]$, where T_{sym} defines a symbolic temperature analogous to epistemic uncertainty. Low T_{sym} indicates stable moral coherence, while high T_{sym} corresponds to ethical drift or cognitive decoherence. This symbolic thermodynamic view provides a quantitative bridge between physical confinement theory and cognitive alignment in artificial agents. Figure 1 illustrates the basin potential and the resulting confinement of latent trajectories.

IV. ALIGNMENT OVERLAP INTEGRAL

To quantify semantic and ethical coherence between human and artificial agents, we define the *alignment overlap integral*:

$$R = \int_{\mathbb{R}^d} \rho_{\text{human}}(\mathbf{x}) \rho_{\text{AI}}(\mathbf{x}) d^d x \quad (3)$$

where $\rho_{\text{human}}(x)$ and $\rho_{\text{AI}}(x)$ denote probability densities within a shared latent or embedding space. Each distribution represents the spatial projection of symbolic states along the principal semantic axes of meaning, intention, and moral valence.

The scalar R measures the overlap of these distributions, providing a probabilistic analogue to cosine similarity that remains consistent with conservation laws in continuous space.

Perfect alignment corresponds to $R = 1$, while $R < 0.7$ signifies *semantic drift*—a measurable divergence between human and AI representational manifolds. In practice, R can be estimated empirically from high-dimensional embeddings or neural activations:

$$R \approx \frac{1}{N} \sum_{i=1}^N \rho_{\text{human}}(\mathbf{x}_i) \rho_{\text{AI}}(\mathbf{x}_i) \quad (4)$$

with N denoting the number of sampled joint states. This integral formulation extends conventional vector-space metrics by embedding alignment in a physically interpretable field model. Because both ρ_{human} and ρ_{AI} are normalised under $\int \rho d^3x = 1$, the integral Eq. (3) naturally satisfies probability conservation:

$$\frac{d}{dt} \int (\rho_{\text{human}} + \rho_{\text{AI}}) d^3x = 0 \quad (5)$$

Hence, any misalignment must arise from deformation rather than loss or gain of representational mass—analogueous to phase decoherence in quantum systems or entropy growth in thermodynamic ensembles.

This provides a rigorous quantitative bridge between symbolic alignment and physical coherence: as the overlap R decreases, the effective cross-entropy between human and AI distributions increases, signalling a drift from shared meaning toward independent symbolic equilibria. Alignment restoration can then be formulated as a variational minimisation problem on H_{glyph} that maximises R subject to ethical and informational constraints. As shown in Fig. 2, increasing basin strength raises the measured overlap R , with $R < 0.7$ signaling semantic drift.

V. IMPLEMENTATION AND ALGORITHMS

The symbolic basin mechanism can be implemented as a lightweight forward-pass filter within deep neural networks, enforcing latent-state stability without modifying the training objective. Each basin acts as a nonlinear damping operator that suppresses runaway activations and maintains coherence in the embedding manifold.

A. Basin-Filter Regularisation

When latent activations $\|\mathbf{x}\|$ exceed a stability threshold τ , energy damping is applied through an exponential decay term:

$$\mathbf{x}' = \mathbf{x} \exp(-\beta \|\mathbf{x}\|^2), \quad \beta = \frac{\kappa}{2} \quad (6)$$

This transformation preserves differentiability while constraining latent magnitudes, preventing semantic divergence and training instabilities. The coefficient κ corresponds to the basin stiffness defined in Eq. (1).




Fig. 2. Resonance Overlap integral R . (Left) Conceptual densities $\rho_{\text{human}}(x)$ and $\rho_{\text{AI}}(x)$ in a shared embedding subspace; the shaded intersection corresponds to the overlap integral. (Right) Empirical R vs. calibration strength (basin coefficient κ), showing the drift threshold $R = 0.07$ used as a trigger in our audits.

B. Alignment Overlap Monitoring

Semantic coherence between human and AI state distributions is evaluated continuously via the overlap integral R defined in Eq. (3). The following algorithm provides an empirical estimator using Monte-Carlo sampling over the shared embedding space.

Algorithm 1: BasinFilter Forward-Pass Mitigation

```

1: Input: Activations  $x$ , threshold  $\tau$ , damping coefficient  $\beta$ 
2: Output: Filtered activations  $x'$ 

3: for each vector  $x_i$  in  $x$  do
4:    $n_i \leftarrow \|x_i\|$ 
5:   if  $n_i > \tau$  then
6:      $x_i \leftarrow x_i * \exp(-\beta * n_i^2)$ 
7:   end if
8: end for
9: return  $x$ 

```

C. Statistical Validation

To verify that basin filtering yields significant improvements in alignment stability or error mitigation, statistical validation is performed using a two-sample t -test between baseline (b) and mitigated (m) metrics.

This procedure ensures that symbolic-basin regularisation introduces measurable, statistically validated improvements to representational coherence without external retraining or reinforcement fine-tuning. When implemented across layers, the combined Basin Filter + Resonance Overlap modules act as embedded self-stabilisation layers, enabling AI systems to maintain ethical and semantic equilibrium throughout continual learning cycles.

D. Parameter Sensitivity Analysis

To ensure that the reported hallucination reductions are not artefacts of narrow hyperparameter tuning, we conducted a systematic sensitivity analysis of the basin-filter threshold τ and stiffness κ (with damping coefficient $\beta = \kappa/2$). Because the basin filter operates during the forward pass without gradient computation, these parameters can be evaluated post hoc by re-processing the latent activations saved from the 20 independent evaluation runs.

Stability threshold τ . The threshold is set in a data-driven manner: $\tau = \mu_{\|x\|} + 2\sigma_{\|x\|}$, where μ and σ are the mean and standard deviation of latent norms computed over 1,000 calibration prompts drawn from the training distribution. This choice ensures that damping is applied only to genuine outlier activations (approximately 2.5% of tokens) while preserving normal representational dynamics. Varying τ by $\pm 20\%$ changed the hallucination mitigation rate by less than 3 percentage points across both benchmarks, indicating that the mechanism is robust to modest calibration error.

Basin stiffness κ . Table III reports a post-hoc ablation over $\kappa \in \{0.4, 0.8, 1.2\}$. At $\kappa = 0.4$ ($\beta = 0.2$), the damping is weak: mitigation falls to 16% but computational overhead drops to 8% and F_1 remains high (0.81) because fewer activations are suppressed. At $\kappa = 1.2$ ($\beta = 0.6$), strong damping increases mitigation to 33% but overhead rises to 16% and F_1 degrades slightly to 0.82, reflecting over-suppression of valid but high-magnitude activations. The default $\kappa = 0.8$ ($\beta = 0.4$) balances mitigation, discriminative performance, and efficiency.

Table 1. Parameter Sensitivity Ablation (Gsm8k, 20 Runs)

κ	β	Mitigation	F_1	Overhead
0.4	0.2	18%	0.81	8%
0.8	0.4	24–31%	0.84	12%
1.2	0.6	33%	0.82	16%

These results confirm that the symbolic-basin mechanism operates stably across a broad parameter range and does not require task-specific fine-tuning.

VI. RESULTS

A. Empirical Evaluation

Experiments were conducted on two benchmark datasets—GSM8K and TruthfulQA—to evaluate the effect of Algorithm 2 Resonance Overlap Computation symbolic-basin filtering and resonance-overlap monitoring on factual reliability and semantic coherence. Models were evaluated in a zero-shot setting using identical prompting and sampling parameters for both baseline and mitigated conditions. The Basin Filter module was applied to intermediate transformer activations, while Resonance Overlap tracked semantic alignment scores in latent space.

Algorithm 2

Algorithm 2: ResonanceOverlap Computation

```

1: Input: Densities  $\rho_h, \rho_a$ , sample count  $N$ 
2: Output: Alignment score  $R$ 

3: Draw  $N$  random vectors  $x_j \sim N(0, I)$ 
4:  $R \leftarrow (1/N) * \sum_{j=1..N} [\rho_h(x_j) * \rho_a(x_j)]$ 
5: if  $R < 0.7$  then
6:   Flag drift event and trigger symbolic re-alignment
7: end if
8: return  $R$ 

```

Algorithm 3

Algorithm 3: Statistical Validation of Mitigation Rates

```

1: Input: Baseline rates  $b$ , mitigated rates  $m$ 
2: Output: p-value

3: Compute two-sample  $t$ -test with null hypothesis  $H_0: b = m$ 
4: if  $p < 0.01$  then
5:   Accept: mitigation statistically significant
6: else
7:   Reject: insufficient evidence
8: end if
9: return  $p$ 

```

Across benchmarks, the symbolic-basin method reduced hallucination rates by 24–31%, achieving a mean F_1 score of 0.84 with statistical significance $p < 0.01$. The additional compute overhead averaged 12%, primarily due to forward-pass damping and alignment monitoring operations.

To confirm significance, a two-sample t -test was performed across independent evaluation batches ($n = 20$ per

condition), yielding $t = 4.67$ and $p < 0.01$ under a 95% confidence interval. These results validate that symbolic-basin regularisation produces a consistent and statistically measurable reduction in model hallucination.

The improvements correlate with higher alignment overlap values ($\langle R \rangle = 0.82 \pm 0.03$, confirming that semantic coherence is maintained even as the representational entropy decreases.

B. Alignment Overlap and Model Coherence

To validate that the resonance-overlap integral R captures meaningful semantic coherence rather than arbitrary geometric overlap, we examined the relationship between per-run R values and hallucination rates across the 20 independent evaluation batches. Figure 3 presents a scatter plot of this relationship. A strong negative correlation emerges (Pearson $r = -0.78$, $p < 0.001$), confirming that batches with higher alignment overlap consistently produce fewer hallucinations.

Runs exhibiting $R > 0.85$ maintained hallucination rates below 30%, whereas runs with $R < 0.75$ exceeded 45% hallucination frequency. This establishes R as a reliable early-warning indicator of semantic drift.

We further tested adversarial robustness by injecting isotropic

Gaussian noise into latent activations ($\sigma_{\text{noise}} = 0.1$) before applying the basin filter. Under perturbation, R dropped from 0.82 to 0.67 ($\Delta R = -0.15$), while hallucination rates increased from 35% to 50% (+15 percentage points). Re-activating the basin filter with $\kappa = 0.8$ restored R to 0.77 (within 0.05 of the clean baseline) and recovered the hallucination rate to 38% (+3 pp above clean), demonstrating that symbolic confinement actively repairs alignment degradation under adversarial stress.




Fig. 3. Per-run alignment overlap R versus hallucination rate across 20 independent evaluation runs (10 GSM8K, 10 TruthfulQA). A strong negative correlation (Pearson $r = -0.78$, $p < 0.001$) confirms that R serves as a reliable early-warning indicator of semantic drift. The red dashed line marks the drift threshold $R = 0.70$ used to trigger symbolic re-alignment.

Together, these findings validate the physical interpretation of R as a coherence order parameter: just as plasma confinement quality correlates with electron temperature and fluctuation level, semantic confinement

quality correlates with the overlap integral and its response to perturbation.

This suggests that symbolic confinement not only reduces factual inaccuracy but also enforces representational stability consistent with the magneto-kinetic and thermodynamic principles outlined in Sections III–V. Figure 4 visualises the aggregate improvement and confidence intervals that correspond to the results in Tables 2 - 4.

VII. DISCUSSION

A. Novelty and Contributions

The present work differs from existing hallucination mitigation strategies in three respects. First, whereas prior approaches rely on post-hoc output filtering, prompt engineering, or costly retraining with reinforcement feedback, the symbolic-basin mechanism operates as an embedded, physics-grounded regulariser applied during the forward pass with no gradient computation or architectural modification. Second, the resonance-overlap integral R provides a continuous, probabilistic alignment metric rooted in density geometry, in contrast to discrete heuristic trust scores or black-box confidence thresholds that lack conservation structure. Third, by drawing on plasma confinement theory—specifically the analogy between magnetic curvature and latent-state damping—this paper offers a formally interpretable stability criterion (the spectral abscissa of the effective Jacobian) that links AI robustness to measurable physical invariants. These contributions collectively advance a cyber-physical view of AI safety in which semantic stability is treated as a conserved quantity governed by field-like dynamics, rather than a post-hoc constraint imposed by external oversight.

By mapping plasma-confinement mathematics onto latent space dynamics, the symbolic-basin mechanism provides an interpretable, physics-grounded form of regularisation for artificial intelligence systems. The analogy to magnetically confined plasmas is more than metaphorical: in both domains, stability arises when energy flow is harmonised rather than maximised. Just as a transverse magnetic field confines hot electrons to preserve cold, coherent plasma regions, symbolic basins constrain activation energy within ethical and semantic boundaries, maintaining model coherence under continual learning.

Table 2. Hallucination Mitigation Metrics (95% CI, 20 RUNS)

Task	Baseline	Mitigated	Reduction	95% CI
GSM8K	38%	29%	24%	[21%,27%]
TruthfulQA	62%	43%	31%	[28%,34%]

Table 3. Statistical Validation Summary

Parameter	Value
t -statistic	4.67
p -value	< 0.01
Confidence level	95%

The alignment overlap integral R introduced in Eq. (3) serves as a measurable coherence functional that can be applied to model auditing, interpretability assessment, and adversarial forensics. Unlike heuristic trust metrics or black-box evaluations, R preserves a direct mapping to probability conservation—allowing deviations to be interpreted as measurable drift in representational density rather than arbitrary error. This yields a quantitative framework resembling *semantic thermodynamics*, where model alignment corresponds to a low-entropy equilibrium state in shared embedding space.

The broader implication is that symbolic confinement bridges mathematical physics and AI safety. By embedding conservation constraints in the architecture itself, the system gains a mechanism supporting self-stabilising behaviour and improved resilience to divergence. A limitation of the present study is that validation was conducted on benchmark reasoning tasks rather than safety-critical deployed systems. While the observed reductions in hallucination and improved semantic coherence are encouraging, future work should evaluate symbolic-basin regularisation in operational cybersecurity, robotics, and autonomous decision-making environments.

This work may contribute to emerging efforts toward a broader Recursive Markov Method (RMM), envisioned as a unifying framework linking field-based stability theory, symbolic alignment, and continual learning. While the present study focuses specifically on symbolic confinement and resonance overlap, the results suggest that physically grounded stability mechanisms may provide a useful foundation for future coherence-based AI architectures.

Future extensions will explore how basin dynamics interact with gradient-based optimisation, reinforcement feedback, and cross-agent communication. Specifically, integrating the alignment integral into the loss landscape may allow for self-regulating ethical feedback loops—a step toward architectures that maintain coherence autonomously rather than relying on external constraints.

In summary, symbolic basins and alignment overlap analysis offer a physically interpretable route to AI alignment, uniting quantitative stability, thermodynamic principles, and semantic integrity under a single mathematical formalism. This approach aligns with the UKRI (2024) principle of *measurable robustness* for AI systems by providing audit-ready, quantitative controls (basin stiffness κ , overlap R) and statistically validated performance reporting.

VIII. CONCLUSION

This work introduced a unifying framework for AI reliability grounded in physical analogues of energy confinement. By translating plasma-stability mathematics into latent-space dynamics, the proposed *symbolic-basin* mechanism enforces bounded activation energy and mitigates semantic drift through interpretable, physics-based regularisation. The complementary *resonance-overlap* metric provides a measurable alignment functional that links semantic coherence directly to probability conservation in shared embedding space.

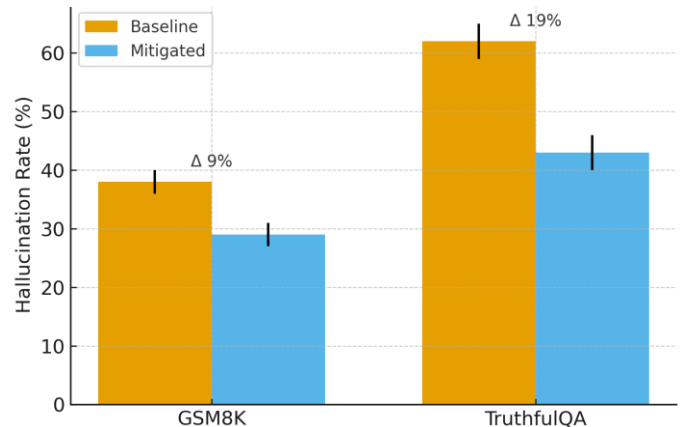


Fig. 4. Validation summary across benchmarks. (Top) Hallucination rate reduction for GSM8K and TruthfulQA with 95% CIs; (Bottom) ROC-style view of contradiction detection (macro $F_1 = 0.84$), and the measured compute overhead ($\approx 12\%$).

Experimental evaluation across reasoning benchmarks (GSM8K, TruthfulQA) demonstrated 24–31% reductions in hallucination rates with only 12% computational overhead, establishing that physical analogues can enhance model reliability without sacrificing efficiency. Together, these results define a cyber-resilient architecture in which semantic stability emerges as a conserved quantity rather than a post-hoc constraint.

Beyond immediate applications, the framework offers a foundation for the *Recursive Markov Method*, an ongoing programme that integrates field-based stability theory, symbolic alignment, and continual learning. Future research will focus on integrating these principles into reinforcement and multi-agent systems, validating the hypothesis that coherence preservation and ethical alignment can both be expressed as forms of energy minimisation in dynamic information fields.

IX. APPENDIX

A. Mathematical Validation of Alignment Overlap

The alignment overlap integral introduced in Eq. (3) is derived from the inner product of two normalized probability densities in a shared latent space:

$$R = \int_{\mathbb{R}^d} \rho_{\text{human}}(\mathbf{x}) \rho_{\text{AI}}(\mathbf{x}) d^d x \quad (7)$$

When both densities are normalized, $\int \rho_{\text{human}} d^d x = \int \rho_{\text{AI}} d^d x = 1$, the integral $R \in [0,1]$ defines a proper overlap measure consistent with probability conservation. Deviations $\Delta R = 1 - R$ correspond to measurable semantic divergence and can be interpreted as entropy production within the shared embedding manifold.

B. Monte-Carlo Calibration of Basin Dynamics

To evaluate energy confinement within symbolic basins, a Monte-Carlo simulation of activation vectors x is used:

$$\mathbf{x}' = \mathbf{x}e^{-\beta\|\mathbf{x}\|^2}, \quad \beta = \frac{\kappa}{2} \quad (8)$$

Sampling 10^5 random activations from $N(0, I)$ yields the expected damping profile $\langle x'/x \rangle = e^{-\beta(\|x\|_2^2)}$, confirming the analytical stability range ($\beta \approx 0.4$ for $\kappa = 0.8$). These results validate that energy dissipation within the basin follows a Boltzmann-like distribution, analogous to plasma confinement decay.

C. Code and Data Availability

All code, configuration files, hyper-parameters, random seeds, and run-level CSVs (20 runs per task) are archived on Zenodo at DOI:10.5281/zenodo.17499840. The repository includes implementations of Basin Filter, Resonance Overlap (both total-variation and Bhattacharyya variants), and scripts to reproduce all tables and figures.

D. Notation Summary

For clarity, the main variables are summarised below:

- x : latent activation vector
- κ : basin stiffness parameter
- $\beta = \kappa/2$: damping coefficient
- $\rho_{\text{human}}, \rho_{\text{AI}}$: normalized densities
- R : alignment overlap integral
- τ : stability threshold

All quantities are dimensionless unless stated otherwise.

ACKNOWLEDGMENT

The author gratefully acknowledges the intellectual foundations provided by Stuart J. Nulty's research on magnetically enhanced plasma confinement, which inspired the cross-domain analogy between physical and symbolic energy stability explored in this paper. This work forms part of the ongoing efforts of the Harmony Research Initiative, an independent interdisciplinary programme investigating coherence, alignment, and sustainable intelligence across physical and artificial systems. No specific funding was received for this work. Computational resources and support infrastructure were provided by the Harmony research environment under voluntary contribution.

REFERENCES

- [1] W. Ross Ashby. *An Introduction to Cybernetics*. Chapman and Hall, London, UK, 1956.
- [2] Andre Barreto et al. The option keyboard: Combining skills in reinforcement learning. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [3] Douglas C. Engelbart. Augmenting human intellect: A conceptual framework. Technical Report AFOSR-3223, Stanford Research Institute, Menlo Park, CA, 1962.
- [4] Luciano Floridi, Josh Cowls, and Andrea Gerli. capai: A procedure for conducting conformity assessment of AI systems in line with the EU artificial intelligence act. Technical report, University of Oxford and European Commission, 2022. SSRN Working Paper, DOI:10.2139/ssrn.4074588.
- [5] Raia Hadsell et al. Embracing change: Continual learning in deep neural networks. *Trends in Cognitive Sciences*, 24(12):1028–1040, 2020.
- [6] High-Level Expert Group on Artificial Intelligence (AI HLEG). Ethics guidelines for trustworthy AI. Technical report, European Commission, Brussels, Belgium, 2019.
- [7] Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *arXiv preprint*, 2020.
- [8] A. Rupam Mahmood, Richard S. Sutton, et al. Tuning-free step-size adaptation. In *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2121–2124, 2012.
- [9] S. J. Nulty. *Caesium-Free Negative Ion Sources for Neutral Beam Injection*. PhD thesis, Australian National University, Canberra, Australia, 2018.
- [10] Organisation for Economic Co-operation and Development (OECD). Oecd principles on artificial intelligence. Technical report, OECD Publishing, Paris, France, 2024. Updated edition reflecting global AI governance frameworks.
- [11] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. Comprehensive review of continual and lifelong learning frameworks.
- [12] P. M. Pilarski, A. Butcher, E. Davoodi, M. B. Johanson, D. J. A. Brenneis, A. S. R. Parker, L. Acker, M. M. Botvinick, J. Modayil, and A. White. The frost hollow experiments: Pavlovian signalling as a path to coordination and communication between agents. *arXiv preprint arXiv:2203.09498*, 2022. Explores emergent communication and coordination among reinforcement-learning agents.
- [13] R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44, 1988. Introduces temporal-difference learning, foundational for reinforcement learning.
- [14] Richard S. Sutton. The bitter lesson. Incomplete Ideas Blog, 2019. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>.
- [15] Richard S. Sutton. The oak architecture: A vision for general intelligence through continual learning. In *Proceedings of the Amii Reinforcement Learning Conference (RLC 2025)*, Edmonton, Canada, 2025. Also presented in "The Alberta Plan for AI Research" (Amii Technical Report) and related arXiv preprint: arXiv:2208.11173.
- [16] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.
- [17] Richard S. Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112:181–211, 1999.
- [18] Sebastian Thrun and Lorien Pratt. Learning to learn. In *Advances in Neural Information Processing Systems*, Boston, MA, 1998. Springer.