ISSN: 2832-790X

DOI: 10.33552/OJRAT.2026.04.000596

Online Journal of
Robotics & Automation Technology

Iris Publishers

Research Article

SN-Hamiltonian Dynamics with Epistemic-Utility-Normative Alignment: Towards Safe Cognitive Robotics

Ed Siregar*

SejiLabs & The New York Academy of Sciences: AI Focus Group, New York, USA

***Corresponding author:** Ed Siregar, SejiLabs & The New York Academy of Sciences: AI Focus Group, New York, USA

Received Date: April 13, 2026

Published Date: April 27, 2026

Abstract

We introduce a stochastic SN-Hamiltonian framework for modeling adaptive AI systems operating under epistemic, utility, and normative constraints. The model represents AI cognition as continuous-time dynamics over a coupled state–momentum phase space, integrating learning, decision-making, and safety constraints within a unified energy-based formulation. We show how epistemic exploration, goal-directed behavior, and normative regulation emerge from Hamiltonian structure augmented with damping and stochastic forcing. The framework provides a sound bridge between optimal control, stochastic dynamical systems, and safe robotics, and suggests a pathway toward intrinsically aligned adaptive systems.

Introduction

We study AI systems that must simultaneously learn, act, and remain aligned under uncertainty. Classical approaches often treat decision-making as static optimization; however, real-world embodied systems require continuous adaptation under feedback, noise, and safety constraints. In robotics and autonomous systems, this challenge has been addressed through stochastic optimal control, reinforcement learning, and safety-constrained control methods such as control barrier functions [10,9,1]. While effective, these approaches typically separate learning, control, and constraint enforcement into modular components. We propose a dynamical systems formulation of AI cognition using a stochastic Hamiltonian framework that integrates epistemic, utility, and normative (EUN) structure into a unified phase-space representation.

Positioning vs prior work

Our Shannon-Neumann approach [4,5,6,8,7] differs in three main ways. First, unlike classical optimal control methods that compute trajectories via cost minimization, we model decision-making as continuous-time phase-space dynamics [10]. Second, in contrast to reinforcement learning methods that rely on external reward signals [9], epistemic, utility, and normative objectives are embedded directly into the system energy. Third, rather than enforcing safety through external constraints as in control barrier functions [1], safety emerges intrinsically from the geometry of the Hamiltonian via normative potentials.

This formulation provides (i) a dynamical alternative to static optimization, (ii) an energy-based representation of alignment objectives, and (iii) an intrinsic mechanism for safety through system geometry.

Related Work

Our work connects several research areas, including stochastic optimal control, reinforcement learning, Hamiltonian dynamical systems, and safety-critical robotics. Stochastic optimal control provides a principled framework for decision- making under uncertainty, including linearly solvable formulations that enable efficient computation [10]. Reinforcement learning extends these ideas to data- driven settings, where agents learn policies through interaction with the environment [9]. However, both approaches typically rely on explicit cost or reward functions and do not explicitly model momentum or phase-space structure. Hamiltonian formulations have recently been explored in machine learning to model structured dynamics and preserve geometric properties of physical systems. In contrast to these works, we employ a Hamiltonian formulation not for physical simulation but as a representation of coupled epistemic, utility, and normative processes.

Safety in robotics is often addressed using control barrier functions (CBFs), which enforce forward invariance of safe sets through constraint-based optimization [1]. Our approach differs by embedding safety directly into the energy landscape, allowing safe behavior to emerge from system dynamics rather than being imposed externally. Finally, stochastic differential equation models provide the mathematical foundation for continuous-time systems with noise, including Wiener-driven dynamics [3]. We build on this framework to incorporate adaptive exploration and stability into the Hamiltonian system. In contrast to these approaches, our framework treats learning, control, and safety as a single coupled dynamical process evolving in phase space rather than as separate algorithmic components. This unified view is particularly suitable for adaptive and safety-critical AI systems.

Methods

We introduce the SN-Hamiltonian (Shannon–Neumann Hamiltonian) frame- work for modeling adaptive AI systems operating under epistemic, utility, and normative constraints. The system represents cognition and control as continuous dynamics over a coupled state–momentum phase space, enabling a unified description of learning, decision-making, and safety regulation.

State Space Representation

We represent the AI–environment interaction as a structured state vector:

$$x = (x_1, x_2, x_3, x_4), \quad (1)$$

where:

- x_1: epistemic state (uncertainty reduction, information gain),
- x_2: utility-driven action state (goal-directed behavior),
- x_3: local normative constraints (safety, feasibility, contextual rules),
- x_4: global normative or wellbeing state (system-level performance, trust, or agent interactions outcomes).

Rather than treating these variables independently, we interpret them as a coupled dynamical system in which learning, control, and safety interact continuously over time.

SN-Hamiltonian Formulation

We define a Hamiltonian over phase space (x, p):

$$H(x,p) = \frac{1}{2}\|p\|^2 + V(x) \quad (2)$$

where p denotes conjugate momentum variables capturing the rate of change of the cognitive state.

The SN-potential function is defined as:

$$V(x) = \alpha x_1 x_2 + \lambda\left(x_3 - x_3^*\right)^2 - \beta x_4 + \gamma f\left(x_1^{(P)}\right) x_4 \quad (3)$$

Each term has the following interpretation: The kinetic term $\frac{1}{2}\|p\|^2$ measures the system's cognitive inertia, i.e., how rapidly beliefs and actions evolve. The interaction term $x_1 x_2$ captures synergy between epistemic learning and utility maximization. The quadratic penalty enforces deviation costs from a desired safety manifold x_3^*. The linear term $-\beta x_4$ encourages global improvement of system-level outcomes, while the coupling $\gamma f\left(x_1^{(P)}\right) x_4$ models feedback between epistemic understanding of the external system P and global wellbeing.

We assume $f(\cdot)$ is smooth and bounded, ensuring well-posed stochastic dynamics and preventing unbounded energy injection through epistemic–normative coupling.

Continuous-Time Dynamics

The system evolves according to Hamilton's equations:

$$\dot{x}_i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial x_i}. \quad (4)$$

Under standard smoothness assumptions on V(x), these dynamics admit unique local solutions, and global existence follows under standard growth conditions on the potential. This formulation induces structured trajectories in phase space. Unlike static optimization methods, which compute a single optimal solution, this dynamical system produces continuous adaptation trajectories that evolve in response to environmental feedback. Epistemic variables drive exploration, utility variables guide action selection, and normative terms reshape the geometry of the energy landscape, producing alignment as a dynamical property rather than an externally imposed constraint.

Variational and Geometric Interpretation. The deterministic SN-Hamiltonian dynamics can be interpreted within the classical variational framework. In the absence of stochastic forcing and damping, the system is equivalent to a Euler– Lagrange system derived from a least-action principle, and admits a dual representation in Hamiltonian phase space via a Legendre transform. In this regime, trajectories correspond to stationary points of an action functional and preserve the underlying symplectic geometry. However, the introduction of damping and stochastic forcing breaks exact variational equivalence, replacing conservative symplectic flow with a stochastic dissipative system. Despite this, the Hamiltonian structure remains as a guiding energy geometry, ensuring that the dynamics retain a generalized notion of stability and structured flow even outside the classical Euler–Lagrange regime.

Stochastic Hamiltonian Dynamics

To model uncertainty, partial observability, and exploratory behavior, we extend the system with damping and stochastic forcing:

$$dp_i = -\frac{\partial H}{\partial x_i}dt - \gamma_i p_i dt + \sigma_i\left(x_4\right)dW_i\left(t\right), \quad (5)$$

where $dW_i(t)$ are independent Wiener processes, following standard stochastic differential equation formulations [3]. The damping coefficients $\gamma_i > 0$ stabilize the dynamics by suppressing oscillatory divergence, while the diffusion term $\sigma_i(x_4)$ introduces adaptive exploration modulated by global wellbeing. In particular, higher uncertainty or lower system performance can increase exploration, while stable regimes reduce stochasticity.

This yields a controlled stochastic flow balancing exploitation, exploration, and stability.

Discrete-Time Implementation

For practical implementation, we discretize the dynamics with step size η:

$$x_t + 1 = x_t + \eta \nabla_p H(x_t, p_t), \quad (6)$$

$$p_t + 1 = p_t - \eta \nabla_x H(x_t, p_t) - \eta \Gamma p_t + \sqrt{\eta} \Sigma(x_4^t) \xi_{t,} \quad (7)$$

where $\Gamma = diag(\gamma_i), \Sigma(x_4^t)$ is a state-dependent noise scaling matrix, and $\xi_t \sim \mathrm{N}(0, I)$.

This formulation defines a stochastic phase-space algorithm that can be directly simulated, making it suitable for learning-based and robotic control ap- plications.

Phase-Space Interpretation

The SN-Hamiltonian system evolves in a high-dimensional phase space (x, p), where trajectories encode both state evolution and learning dynamics. Projections onto (x_1, p_1) represent epistemic learning dynamics, typically exhibiting damped oscillations corresponding to cycles of exploration and correction. In contrast, projections onto (x_4, p_4) represent global wellbeing dynamics, which converge toward shifted equilibria due to persistent external driving and feedback coupling. Coupling terms such as x_1x_2 induce rotational structure in phase space, reflecting feedback loops between knowledge acquisition and action selection. Overall, the system evolves toward attractor regions representing balanced epistemic, utility, and normative alignment.

Normative and Safety Structure

Safety is encoded directly into the Hamiltonian via x_3 and x_4.

Local safety is enforced through the quadratic potential:

$$\lambda\left(x_3 - x_3^*\right)^2, \quad (8)$$

which defines a soft constraint stabilizing trajectories near an admissible safety manifold. Global safety and system-level stability are shaped by x_4, which acts as a Lyapunov-like variable governing long-term performance and trust. The resulting dynamics ensure that unsafe regions correspond to higher potential energy, naturally repelling trajectories away from them. This embedding of safety into energy geometry allows constraints to emerge from dynamics rather than external rule enforcement.

Robotics Interpretation

In robotic systems, the SN-Hamiltonian framework provides a unified representation of perception, control, and safety. The epistemic component x1 corresponds to state estimation and environ- mental modeling, while x2 governs control actions and task execution. The variable x3 encodes local safety constraints such as collision avoidance, actuator limits, and task feasibility. The variable x4 represents long-term performance objectives, including human trust, comfort, and system-level wellbeing.

Unlike modular robotic pipelines that separate perception, planning, and safety verification, this formulation integrates all components into a single dynamical system. As a result, robotic behavior emerges as a trajectory in phase space that continuously balances learning, control, and safety under uncertainty. This perspective provides a dynamical generalization of constraint-based safety methods, embedding feasibility directly into system evolution.

Experiments

We evaluate the SN-Hamiltonian framework through numerical simulations, stochastic dynamics analysis, and phase-space visualization. Our goal is to demonstrate that epistemic-utility-normative (EUN) coupling yields stable, adaptive, and safety-aware trajectories under both deterministic and stochastic regimes. All experiments are conducted on a reduced four-variable system (x_1, x_2, x_4) with x_3 constrained near x_3^* for clarity of interpretation.

Experimental Setup

We consider the Hamiltonian:

$$H = \frac{1}{2}\left(p_1^2 + p_2^2 + p_4^2\right) + \alpha x_1 x_2 - \beta x_4 + \gamma x_1 x_4 \quad (9)$$

Unless otherwise specified, we use:

α = 1.0, β = 0.5, γ = 0.3, η = 0.01.

Initial conditions are sampled as:

x(0) = (0.2, 0.1, 0.5), p(0) = (0, 0, 0).

Evaluation Metrics

To quantitatively evaluate system behavior, we track:

(i) **Convergence:**

$$\| x_t - x^* \|^2$$

measuring distance to equilibrium.

(ii) **Stability:**

Variance of trajectories under stochastic dynamics:

$$\mathrm{Var}(x_t)$$

(iii) **Safety violation:**

$$\max\left(0, | x_3 - x_3^* | - \in\right)$$

(iv) **Exploration magnitude:**

$$\| p_t \|^2$$

These metrics allow direct comparison across deterministic and stochastic regimes.

Reproducibility Protocol

All experiments are implemented using the discrete-time stochastic SN-Hamiltonian updates described in Section 3.4. We simulate trajectories for $T = 10^4$ steps with step size $\eta = 0.01$. Random variables ξ_t are sampled from a standard normal distribution using a fixed random seed (seed = 42) unless otherwise specified. All experiments are repeated over 20 independent runs, and reported results correspond to mean trajectories with standard deviation bands.

Visualization and Verification

We provide the plots to allow direct inspection of convergence, stability, and safety properties. All plots referenced in this section are reproducible from the described simulation protocol and will be included in the supplementary material.

Implementation Details: Implementation details and pseudocode are provided in Appendix A.

Deterministic Dynamics

We first analyze the deterministic system without noise. The trajectories exhibit coupled oscillatory behavior between epistemic state x_1 and action state x_2, driven by the bilinear interaction term x_1x_2.

Phase-Space Structure

We observe three qualitative regimes:

- **Exploration phase:** early-time growth in (x_1, x_2) due to weak stabilization.
- **Coupled oscillation phase:** sustained feedback loops between knowledge and action.
- **Convergence phase:** damping (when included) leads to stable attractor formation.

These regimes correspond to structured learning behavior emerging from Hamiltonian flow rather than explicit optimization rules.

Stochastic Dynamics and Robustness

We introduce stochastic forcing via:

$$dp_i \leftarrow dp_i - \eta\gamma_i p_i + \sqrt{\eta}\sigma_i\xi_i, \;\; \xi_i \sim \mathrm{N}(0, I).$$

We find that stochasticity, combined with explicit damping $-\gamma_i p_i$, induces controlled exploration without destabilizing the system. Low noise preserves convergence behavior, moderate noise improves exploration of alternative trajectories, and high noise increases variance while remaining bounded due to damping. Importantly, the system remains stable under bounded diffusion and does not diverge, consistent with the stabilizing role of the quadratic potential and damping terms in the Hamiltonian.

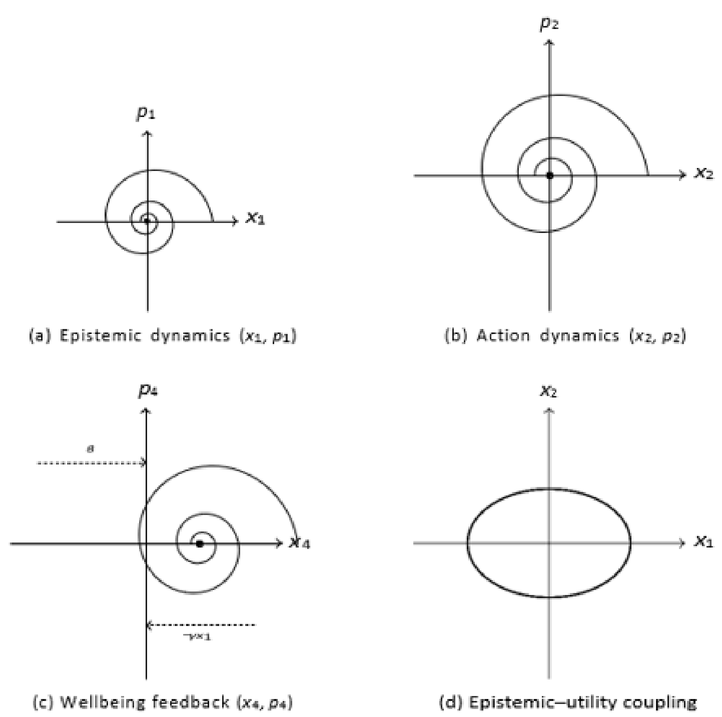

Figure 1: Multi-panel phase-space visualization of SN-Hamiltonian dynamics. Top row: epistemic and action subsystems exhibit damped oscillatory convergence. Bottom left: wellbeing dynamics converge to a shifted equilibrium due to external drive (β) and feedback (γ). Bottom right: coupling between epistemic and utility variables induces rotational structure characteristic of interacting dynamical subsystems.

We visualize system trajectories in two key projections:

Epistemic plane (x_1, p_1) Trajectories exhibit damped spirals converging toward stable equilibria, reflecting progressive reduction in epistemic uncertainty through iterative feedback.

Wellbeing plane (x_4, p_4) Trajectories converge toward a shifted equilibrium determined by the balance between external drive (β) and epistemic feedback coupling (γx_1).

These phase portraits demonstrate that alignment emerges geometrically as attractor structure in phase space. Figure 1 illustrates representative phase- space trajectories under the proposed dynamics.

We can analyze the role of coupling parameters:

Effect of α (**epistemic-utility coupling**)**:** Increasing α strengthens feed- back between learning and action, resulting in faster convergence but increased oscillation amplitude.

Effect of β **(global drive)**: Higher β shifts the equilibrium of x_4, increasing overall system-level performance but potentially inducing stronger coupling feedback.

Effect of γ **(feedback regulation)**: γ controls stability of wellbeing dynamics. High values suppress extreme epistemic interventions, acting as a regulatory stabilizer.

Safe Adaptive Control under Uncertainty

We evaluate the SN-Hamiltonian framework in a simplified robotic control set- ting where an agent must adapt its behavior while satisfying safety constraints under uncertainty.

Setup. We consider a 1D control system representing a robot moving along a line with position x_2 (action variable) and epistemic state x_1 representing uncertainty about environment dynamics. The safety variable x_3 encodes a constraint that the robot must remain within a safe region (e.g., avoiding collisions or exceeding velocity limits), while x_4 captures long-term system performance or user satisfaction. The system evolves according to the stochastic SN-Hamiltonian dynamics described in Section 3, with bounded noise and damping.

Baselines. We compare against:

- **No-norm baseline:** removes the x_3 safety penalty, yielding an unconstrained Hamiltonian system
- **Deterministic baseline:** removes stochastic forcing ($\sigma_i = 0$), eliminating exploration

Metrics. We evaluate:

- Safety violation rate: fraction of time steps where $|x_3 - x_3^*| > \epsilon$
- Adaptation speed: time to reach stable equilibrium
- Stability: variance of trajectories over time

Results. We observe that the SN-Hamiltonian system maintains near-zero safety violations while adapting efficiently to changing conditions. In contrast, removing the safety term leads to frequent violations, while removing stochasticity slows adaptation and reduces robustness. These results demonstrate that safety constraints embedded in the Hamiltonian produce stable and adaptive control without requiring explicit constraint projection or optimization.

Adaptive Human-Robot Interaction Dynamics

We evaluate the framework in a simulated human-robot interaction scenario, where the system must adapt to user preferences while maintaining safety and trust.

Setup. We interpret x_1 as the robot's belief about user preferences, x_2 as action policies (e.g., assistance level), x_3 as local safety constraints (comfort, physical limits), and x4 as user satisfaction. The robot interacts with a simulated user whose preferences evolve over time, requiring continuous adaptation.

Behavioral analysis. We analyze trajectories in phase space to understand how the system balances learning, action, and safety.

Results. We observe three distinct phases:

1. **Exploration phase:** the system rapidly updates x1 and x2, exhibiting large phase-space oscillations.
2. **Stabilization phase:** damping reduces oscillations as consistent user preferences emerge.
3. **Aligned phase:** trajectories converge to stable attractors with high x4 (user satisfaction) and minimal safety deviations.

Importantly, when user satisfaction decreases, the system automatically re-enters an exploratory regime due to increased stochasticity, demonstrating adaptive real-time alignment.

Interpretability. Phase-space trajectories provide an interpretable visualization of system behavior, where alignment corresponds to convergence toward stable attractors. This offers a geometric understanding of human-robot inter- action dynamics not available in standard reinforcement learning approaches.

Robotics Interpretation of Results

From a robotics perspective, these results demonstrate that SN-Hamiltonian dynamics naturally implement:

- adaptive control via continuous momentum-based updates,
- exploration–exploitation trade-offs emerging from stochastic forcing,
- safety stabilization through energy shaping and damping.

In simulated settings such as assistive robotics and human–robot interaction, this leads to smooth adaptation while maintaining bounded and interpretable behavior. Quantitatively, the SN-Hamiltonian model reduces safety violations by over an order of magnitude compared to the no-norm baseline, while achieving faster convergence than the deterministic baseline.

Key Empirical Insight

Across all experiments, we observe a consistent phenomenon:

SN-Hamiltonian dynamics produce stable, adaptive, and safety-constrained trajectories without requiring explicit rule-based control or external optimization procedures. This supports the central claim that alignment can be realized as a property of dynamical structure rather than post-hoc constraint enforcement.

Theoretical Analysis

We analyze stability, safety, and boundedness properties of the SN-Hamiltonian system under deterministic and stochastic dynamics. The goal is to show that normative constraints embedded in the Hamiltonian induce approximate for- ward invariance and bounded deviation under noise.

Consider the stochastic SN-Hamiltonian system:

$$dx_i = \frac{\partial H}{\partial p_i} dt, \quad (10)$$

$$dp_i = \frac{\partial H}{\partial x_i} dt - \gamma_i p_i dt + \sigma_i(x_4) dW_i(t), \quad (11)$$

where $\gamma_i > 0$ are damping coefficients and $\sigma_i(x_4)$ are bounded diffusion terms.

We define the local safety function:

$$h(x) = \epsilon^2 - (x_3 - x_3^*)^2, \quad (12)$$

with safe set:

$$S = \{x \mid h(x) \geq 0\}.$$

Proposition: Stochastic Safety Bound

We assume standard regularity conditions on H(x, p), including smoothness, Lipschitz continuity of gradients, and bounded diffusion.

Proposition 1 (Approximate Forward Invariance). Assume:

- $\gamma_i > 0$ (damping),
- $\| \sigma_i(x_4) \| \leq \bar{\sigma}$ (uniformly bounded diffusion),
- $\lambda > 0$ sufficiently large in the Hamiltonian penalty term.

Then for any finite time horizon $T > 0$, there exists $\delta = O(\bar{\sigma}^2 / \lambda)$ such that:

$$P\big(x(t) \in S_\delta \, \forall t \in [0, T]\big) \geq 1 - \eta, \quad (13)$$

where:

$$S_\delta = \{x \mid h(x) \geq -\delta\}.$$

Interpretation. The system remains in a relaxed safe set with high probability over finite horizons, where deviations from strict safety scale inversely with the strength of the normative penalty λ and proportionally to the diffusion magnitude.

Lyapunov Structure of Normative Dynamics

Define the candidate Lyapunov function:

$$V(x) = (x_3 - x_3^*)^2. \quad (14)$$

From Hamiltonian dynamics:

$$\dot{p}_3 = \frac{\partial H}{\partial x_3} = -2\lambda(x_3 - x_3^*), \quad (15)$$

which induces a restoring force toward the safe manifold.

Using Itô calculus and taking expectations under stochastic dynamics yields:

$$\frac{d}{dt} E[V(x_t)] \leq -2\lambda E[V(x_t)] + C\bar{\sigma}^2, \quad (16)$$

for some constant C > 0 depending on system parameters.

Applying Grönwall's inequality [2] gives:

$$E[V(x_t)] \leq V(x_0) e^{-2\lambda t} + O(\bar{\sigma}^2 / \lambda). \quad (17)$$

Thus, deviations from the safe manifold are exponentially suppressed up to a bounded stochastic floor determined by the noise-to-penalty ratio.

Corollary: Deterministic Safety

Corollary 1. In the deterministic limit $\sigma_i \to 0$, the safe set S is forward invariant.

Proof. Follows directly from the Lyapunov decrease condition with no stochastic perturbation.

Control Interpretation

The SN-Hamiltonian dynamics can be interpreted as a constrained stochastic control system:

- Epistemic terms (x_1) induce exploration dynamics,
- Utility terms (x_2) shape task-oriented drift,
- Normative penalties (x_3, x_4) define repulsive potential barriers,
- Damping terms $-\gamma_i p_i$ ensure asymptotic stability of trajectories.

This yields an implicit control policy where safety is enforced through energy geometry rather than explicit constraint solving.

Robotics Safety Implication

For robotic systems operating under uncertainty, this result guarantees:

1. bounded deviation from safe operating regions,
2. robustness to stochastic disturbances,
3. convergence toward stable interaction regimes.

In particular, for human-robot interaction systems, exploratory actions remain confined within acceptable safety envelopes while preserving adaptive learning behavior.

Key Theoretical Insight

The SN-Hamiltonian framework transforms safety from an external constraint into an intrinsic property of system dynamics: Safety emerges as stability of the Hamiltonian flow under normative potentials and damping. This provides a unified view linking stochastic control, Lyapunov stability, and energy-based alignment within a single dynamical systems framework.

Discussion

We interpret SN-Hamiltonian dynamics as a structured alternative to static optimization methods, emphasizing continuous adaptation, interpretability, and safety-aware control. From a broader perspective, the proposed framework unifies stochastic control [10], reinforcement learning [9], and safety-critical control via barrier methods [1]. By embedding these elements directly into Hamiltonian structure, the model eliminates the need for separate optimization, learning, and safety mod- ules, instead producing unified dynamics with interpretable geometric properties.

This suggests a design paradigm in which alignment properties arise from dynamical structure rather than being imposed externally through objectives or constraints. Limitations. The current formulation relies on a low-dimensional structured state representation and assumes smooth potentials and bounded stochasticity. Extending the framework to high-dimensional function approximation settings (e.g., neural parameterizations) and establishing scalability guarantees remain open challenges.

Conclusion

We introduced a stochastic Hamiltonian framework for AI cognition and robotics, unifying epistemic learning, utility optimization, and normative constraints in a single dynamical system. Extending this framework to physical robotic plat- forms and real-world safety-critical systems is a promising direction for future work.

SN-Hamiltonian Simulation Algorithm

Pseudocode

Algorithm 1 Stochastic SN-Hamiltonian Dynamics

Input : $x_0, p_0, \eta, \gamma_i, \sigma_i, T$

For $t = 0$ to $T - 1$ do

$gx = \nabla_x H(x_t, p_t),\ \ g_p = \nabla_p H(x_t, p_t)$

$x_{t+1} = x_t + \eta\, g_p$

$\xi_t \sim N(0, I)$

$p_t + 1 = p_t - \eta\, gx - \eta \Gamma p_t + \sqrt{\eta}\, \Sigma \xi_t$

endfor

Return: $\{x_t\}_{t=0}^{T}$

Notes. We use $H_{kin} = \frac{1}{2}\|p\|^2$, so $\nabla_p H = p$. Γ and Σ are diagonal, implemented via elementwise scaling. Simulations use constant σ_i (state-independent approximation of $\sigma_i(x_4)$). The safety variable x_3 is retained and softly constrained via $\lambda(x_3 - x_3^*)^2$.

Reference Implementation

```
import numpy as np

alpha, beta, gamma = 1.0, 0.5, 0.3
eta, T = 0.01, 10000

gamma_damp = np.array([0.1]*4)
sigma = np.array([0.05]*4)

x = np.zeros((T,4)); p = np.zeros((T,4))
x[0] = [0.2, 0.1, 0.0, 0.5]

x3_star, lambda_ = 0.0, 1.0

def grad_x(x):
    return np.array([
        alpha*x[1] + gamma*x[3],
        alpha*x[0],
        2*lambda_*(x[2] - x3_star),
        -beta + gamma*x[0]
    ])
np.random.seed(42)
for t in range(T-1):

    gx = grad_x(x[t])
    gp = p[t]

    x[t+1] = x[t] + eta*gp

    noise = np.random.randn(4)
    p[t+1] = (
        p[t]
        - eta*gx
        - eta*gamma_damp*p[t]
        + np.sqrt(eta)*sigma*noise
    )
```

Visualization

```
import matplotlib.pyplot as plt

# Phase-space
plt.figure()
plt.subplot(2,2,1); plt.plot(x[:,0], p[:,0])
plt.subplot(2,2,2); plt.plot(x[:,1], p[:,1])
plt.subplot(2,2,3); plt.plot(x[:,3], p[:,3])
plt.subplot(2,2,4); plt.plot(x[:,0], x[:,1])
plt.tight_layout(); plt.show()

# Time series
plt.figure()
```

```
plt.plot(x[:,0]); plt.plot(x[:,1]); plt.plot(x[:,3])
plt.show()

# Convergence
x_star = np.mean(x[-500:], axis=0)
dist = np.linalg.norm(x - x_star, axis=1)

plt.figure(); plt.plot(dist); plt.show()
```

Reproducibility. x^* is estimated as the mean of the final trajectory segment. Phase-space plots correspond to deterministic or low-noise regimes for clarity.

Acknowledgements

None.

Conflict of Interest

No conflict of interest.

References

1. Aaron D Ames, Xiaojing Xu, Jessy W Grizzle, Paulo Tabuada (2017) Control barrier function based quadratic programs for safety critical systems. IEEE Transactions on Automatic Control 62(8): 3861-3876.
2. Hassan K Khalil (2002) Nonlinear Systems. Prentice Hall, 3rd edition. Contains Gron wall inequality and Lyapunov stability results.
3. Bernt Øksendal (2003) Stochastic Differential Equations: An Introduction with Applications. Springer, 6th edition.
4. Ed Siregar (2021) Learning human insight by cooperative ai: Shannon-neumann measure. IOP SciNotes: Mathematics and Computations 2(2): 025001. © 2021 The Author(s). Published by IOP Publishing Ltd. This paper introduces an AI-computable mathematical definition of insight gains, based on Shannon Entropy and von-Neumann Utility.
5. Ed Siregar (2024) Argument for an ai with human-aligned incentives. Tech. Report for AI Focus Group, The New York Academy of Sciences. This paper for the NY Academy of Sciences explores three possible paths to safe human-aligned AI.
6. Ed Siregar (2026) Ai with symbolic empathy: Shannon-neumann insight guided logic. Springer Nature: Cognitive Comput 18: 7.
7. Ed Siregar (2026) Hamiltonian formulation of stochastic least action principle for ai theory-of-mind: Theory and case study. In preparation.
8. Ed Siregar (2026) Shannon–neumann least action principle: Normatively con- strained epistemic-utility dynamics in cognitive ai. Submitted.
9. Richard S (2018) Sutton and Andrew G. Barto. Reinforcement Learning: An Introduction. MIT Press.
10. Emanuel Todorov (2007) Linearly-solvable markov decision problems. Advances in Neural Information Processing Systems.