International Journal of Quantum Technologies
ISSN: 3068-9090

PRIME OPEN ACCESS

Volume 2, Issue 4

Research Article

Date of Submission: 01 Oct, 2026
Date of Acceptance: 27 Oct, 2026
Date of Publication: 06 Nov, 2026

Beyond Coordinates Latent Geometries, Hamiltonian Networks, and the Boundaries of Symbolic Translation in Complex Dynamics

Gabi Siboni*

Ariel University, Faculty of Engineering Ariel, Israel

Corresponding Author:
Gabi Siboni, Ariel University, Faculty of Engineering Ariel, Israel.

Citation: Siboni, G. (2026). Beyond Coordinates Latent Geometries, Hamiltonian Networks, and the Boundaries of Symbolic Translation in Complex Dynamics. Int J Quantum Technol, 2(4), 01-10.

Abstract

This paper re-examines the transparency tax and the limits of symbolic translation in deep systems. The central claim is that the failure of a symbolic model to preserve physical dynamics does not necessarily indicate a lack of mathematical regularity; rather, it may stem from a mismatch between the latent representation, the translation language, and the symbolic dictionary. The paper analyzes conditions where advanced methodologies incorporating physical and geometric constraints into neural networks, such as Hamiltonian Neural Networks (HNN) and Lagrangian Neural Networks,[1] within a framework distinguishing between representation error, dynamics error, and translation error, can improve the preservation of dynamic structure. Additionally, a four-model comparative protocol is proposed to examine the trade-offs of prediction, structure, and explanation in symbolic distillation.

However, the justification for using unconstrained latent spaces stems not only from the computational limitations of equations, but also from situations where system coordinates, such as momentum and position, are a priori unknown or inaccessible to direct measurement. In systems based on raw video or distributed fields, the latent space may provide a learned coordinate system that enables the examination of latent order without pre-imposing a full analytical decomposition.

Keywords: Latent Spaces, Hamiltonian Neural Networks, Symplectic Geometry, Symbolic Distillation, Nonlinear Dynamics, Transparency Tax

Introduction

In a preceding series of studies concerning the transparency tax[2] and the limits of explainability in artificial intelligence, it was argued that converting continuous latent representations into explicit symbolic models frequently causes a systematic loss of vital information. The empirical simulation of a double pendulum demonstrated how a polynomial dictionary can diverge under extreme conditions because it is incapable of representing complex periodicity without compromising core physical relations.

This failure prompts the guiding question of this paper: Does the limitation of the symbolic model indicate an absence of mathematical regularity, or a mismatch between the translation language and the representation describing the dynamics? To address this question, the paper examines the role of latent spaces in coordinate discovery and subsequently presents the formal framework describing encoding, reconstruction, and flow consistency [1].

This paper aims to update and refine the argument: the failure of symbolic mathematics may be representation-dependent, and a latent space may provide a learned coordinate system when system variables are unknown a priori or

inaccessible to direct measurement. Three research questions derive from this. First, under what conditions can a latent coordinate system dynamically equivalent to the physical state space be identified? Second, how can one distinguish between a representation that preserves the laws of motion and one that well-recovers the data but fails to preserve its geometric structure? Third, how should the gap between a high-predictive-power latent model and a sparse symbolic model derived from it be measured? These questions necessitate separating representation error, dynamics error, and translation error, as each stems from a distinct source and requires a different diagnostic tool.

The contribution of this paper is therefore twofold. On the conceptual plane, it expands the transparency tax from a performance-loss metric to a causal model of the translation chain. On the methodological aspect, it proposes evaluating each stage independently using matching metrics. Encoding quality is evaluated through reconstruction, stability under perturbations, and preservation of intrinsic dimension. Dynamics quality is evaluated via local error, long-term rollout, and conservation of quantities. Symbolic distillation quality is evaluated via complexity, coefficient stability, and out-of-training-range generalization.

The remainder of this paper first introduces the formal framework for encoding, reconstruction, and flow consistency, followed by a discussion of the contribution of Hamiltonian networks to preserving geometric structure. The subsequent section examines the unknown coordinates problem and the non-uniqueness of latent representations. Finally, the experimental separation protocol, the four-model comparison results, and the limitations arising from conservation assumptions and symbolic distillation are presented.

Formal Framework: Encoding, Reconstruction, and Flow Consistency

The formal framework below is designed to describe the transition chain from an observed physical state to a latent representation, distinguishing between three non-equivalent conditions. The encoding condition examines whether data can be mapped to a reduced internal space preserving information relevant to the dynamics. The reconstruction condition examines whether the latent representation contains sufficient information to approximately recover the observed state.

The flow consistency condition is stricter than both, requiring that system evolution over time be preserved under the mapping, such that advancing along the physical trajectory followed by encoding yields a result close to encoding first and advancing within the latent space. Distinguishing between these conditions helps identify whether a failure stems from information loss in encoding, inaccuracy in reconstruction, or a latent dynamics that is not equivalent to the original dynamics.

To formalize this, let us define a physical state space M (the space of all possible physical states of the system) and a latent representation space $Z \cong \mathbb{R}^d$ (an abstract d-dimensional vector space managed by the network). Let us define the encoding function implemented by the neural network as the mapping:

$$f_\phi : M \to \mathbb{R}^d$$

where: f_ϕ is a neural network characterized by a set of weights ϕ, converting a complex physical state in M to a continuous vector in the latent space $\mathbb{R}^d$. Concurrently, the reconstruction function (decoder):

$$g_\psi : \mathbb{R}^d \to M$$

where: g_ψ is a reconstructing network with weights ψ, translating the latent vector back to the original physical space. Thus, for a given input, an approximation of the original system holds:

$$g_\psi\left(f_\phi(x)\right) \approx x \quad \forall x \in M$$

where encoding a state x followed by its reconstruction yields an approximation very close to the original for every possible state x in the system.

When examining flow consistency, we define the flow operator in the latent space $\psi_t : \mathbb{R}^d \to \mathbb{R}^d$, such that for a given trajectory:

$$z_t = \psi_t z_0 = \psi_t\left(f_\phi\left(\Phi^t(x_0)\right)\right)$$

where:
z_t is the latent vector at time t.
z_0 is the initial latent vector?
Φ_t is the true physical flow operator in the original state space advancing the system from an initial state x_0 at time t.

The mathematical requirement for full consistency (Commutative Diagram Consistency) is expressed by:

$$f_\phi \circ \Phi^t \approx \psi_t \circ f_\phi$$

Explanation: Applying the physical flow followed by encoding into the latent space is equivalent (to a good approximation) to initial encoding of the state followed by application of the latent flow.

Beyond Dictionary Limitations: Hamiltonian Neural Networks

To address the limitations of discrete polynomial dictionaries, Hamiltonian Neural Networks offer an architecture where a learned scalar function is used to generate a vector field with a Hamiltonian structure. Structural bias is not injected solely into the loss function but also arises from the derivative calculation form and the relationship between the learned Hamiltonian and the velocity field [2].

Phase Space Formalism and the Symplectic Field

Let us establish a phase space $P \cong \mathbb{R}^{2n}$, characterized by canonical coordinates $z = (q, p)^\top$. Hamiltonian dynamics are given by Hamilton's equations:

$$\dot{q} = \frac{\partial H}{\partial p}, \quad \dot{p} = -\frac{\partial H}{\partial q}$$

where:
$q \in \mathbb{R}^n$ represents spatial coordinates (position, angle, etc.).
$p \in \mathbb{R}^n$ represents conjugate momenta of the system.
The symbol $\top$ denotes the transpose operation converting a column vector to a row vector and vice versa for compact notation.
In a compact matrix formulation using the canonical symplectic matrix J:

$$J = \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}, \quad \dot{z} = J\nabla_z H(z)$$

where:
I_n is the $n \times n$ identity matrix.
J is the symplectic matrix preserving phase space volume.
$\nabla_z H(z)$ is the gradient vector of the energy function with respect to all phase variables z.
$\dot{z}$ is the time-derivative vector of all system variables.

Loss Architecture in HNN

An HNN learns a parametric scalar network $\hat{H}_\theta(z)$ representing total energy, and computes the derivatives:

$$\frac{d\hat{z}}{dt} = J\nabla_z \hat{H}_\theta(z)$$

where: $d\hat{z}/dt$ is the predicted vector field of dynamics in the latent space according to the learned energy.

The loss function is defined as minimizing the derivative difference:

$$\mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \left\| \frac{d\hat{z}_i}{dt} - J\nabla_z \hat{H}_\theta(z_i) \right\|^2$$

where:
$\mathcal{L}(\theta)$ is the network's overall error metric (Loss).
N is the number of training samples in the dataset.
$d\hat{z}/dt$ is the empirically measured rate of change in the i-th sample.
$\|\cdot\|^2$ represents the squared Euclidean error norm.

This constraint encourages learning an energy-like quantity and a flow field consistent with the Hamiltonian structure. However, exact conservation in the calculated trajectory also depends on estimation quality, training data coverage, and the numerical integration scheme. Therefore, this architecture should not be viewed as an automatic guarantee against divergence, but rather as a strong inductive bias narrowing the space of possible solutions.

Even when the learned field is derived from a Hamiltonian, a discrete-in-time solution may accumulate error. Symplectic integrators do not guarantee point-by-point accuracy at every step but are designed to preserve phase space geometry and curb long-term drift. Hence, HNN evaluation should report not only the loss function, but also the solver type, step size, and sensitivity of the result to step size changes [3].

The Epistemological Turn: Missing Coordinates Problem & Raw Video Data

In the real world, especially in complex or biological systems, fundamental physical variables are unknown to the researcher a priori (Coordinate Agnosticism). Coordinate ignorance does not indicate an absence of coordinates, but rather that the researcher observes the system through a measurement mechanism not fully known. The same dynamics may appear as a sequence of pixels, an acoustic signal, or another array. The latent encoder aims to extract a representation sufficient for predicting the future and examining conserved dynamic properties. The scientific requirement is not that latent variables be identical to familiar angles or momenta, but that their relations to observed dynamics and their stability under coordinate transformations can be tested [1].

When input consists of a raw video stream $X \in \mathcal{X} \subset \mathbb{R}^{H\times W\times C}$ (where H, W are image height and width dimensions, and C is the number of color channels such as RGB), the researcher has no pre-made coordinates. The latent space acts as an epistemological bridge enabling the system to generate an internal coordinate system of its own out of the data stream, even before formulating the phenomenon in familiar symbolic equations.

A latent space is generally not uniquely identified. Many reversible transformations may preserve the same dynamics and reconstruction quality. This fact limits the direct interpretation of each latent axis. Nevertheless, non-uniqueness does not invalidate the scientific value of the representation. Properties preserved under coordinate transformation, such as intrinsic dimension, topology, local spectrum, existence of conserved quantities, and the ability to convert the field into a simple form, can be tested [4].

Synthesis: Expanding Transparency Tax to the Era of Unknown Coordinates

This expansion requires defining the transparency tax not as a single difference between two prediction errors, but as a metric vector. One component measures accuracy loss. A second component measures physical constraint violations. A third component measures description complexity growth. A fourth component measures drops in robustness to initial condition changes, noise, or energy regime shifts. Only the combination of metrics allows distinguishing between a short yet fragile equation and a more complex model preserving the dynamic structure.

The paper proposes updating the transparency tax framework via a three-stage model:
Stage A (Coordinate Discovery): The network learns to generate a latent space from video or unstructured data in a topological-geometric manner.
Stage B (Physical Constraint): Applying architectures such as HNN to the latent space to test whether geometry encodes a hidden Hamiltonian structure.
Stage C (Translation Loss Measurement): Examining the gap between the predictive capacity of the uninterpreted latent space and the equations extracted from it [5].

Experimental Separation Criterion: To identify the source of loss, four parallel models are trained on identical data: a model operating directly in the observed space, an unconstrained latent model, a physics-constrained latent model, and a distilled symbolic model. This comparison allows attributing the gap to encoding, constraint, or distillation. Without such a control group, the success or failure of the final model makes it impossible to identify at which stage information was compromised. See experimental details in the Appendix.

Implications, Discussion, and Future Research Directions

This paper proposes a conceptual and methodological framework combining philosophy of science, machine learning, and nonlinear dynamics. This approach resonates with developments in Physics-Informed Machine Learning, which aims to embed physical inductive biases into deep models to address fragility in extrapolation [6].

The formal development and the four-model comparison protocol show how to distinguish between information loss in encoding, failure in latent dynamics, and loss in symbolic distillation. The double pendulum experiment serves as a case study for this framework, but generalization scope to other systems depends on further examination of dissipative, constrained, and noisy systems.

The primary conceptual implication of this work is the reformulation of the transparency tax not as a one-dimensional computational cost, but as a vector mapping accuracy, physical constraints, description complexity, and external robustness. The experimental comparison indicates that in the tested case study, a symbolic equation may be short and elegant, but also fragile under changes in energy regime or out-of-distribution noise. This perception alters how models are evaluated in explainable AI: syntactic simplicity does not necessarily imply structural fidelity.

Empirical data indicate that latent encoding alone is not a guarantee of long-term stability. In the absence of constraints, energy drift in the latent space can lead to collapse in extrapolation. The central challenge in transitioning from raw data

representations to symbolic dynamic equations is maintaining structural fidelity. In the described experiment, integrating a symplectic constraint reduced divergence and improved phase space preservation, though this advantage is subject to limitations detailed below [7].

Imposing a conservative Hamiltonian structure on systems characterized by friction, energy dissipation, or external forces can introduce erroneous bias, and the framework must be expanded in future work to Lagrangian or open models. Furthermore, calculating derivatives from noisy data can bias both energy learning and symbolic distillation. Another insight shows that perfect visual reconstruction capability does not necessarily guarantee that latent axes parallel true physical variables precisely; hence, the latent space should be interpreted as a functional and epistemological space rather than necessarily a direct reflection of metaphysical reality [8].

These findings open several directions for future research: A first direction involves developing architectures explicitly separating the conservative energy component from dissipative and control components. A second track pursues algorithms learning the symplectic shape itself when canonical coordinates are entirely unknown a priori. Integrating probabilistic uncertainty mechanisms in the latent space to distinguish between regions where the model discovered stable structure and regions where prediction is fragile extrapolation is also a relevant research direction. Finally, transitioning from direct distillation to gradual distillation—first extracting geometric invariants, then differential relations, and only finally formulating a closed, reduced equation.

Appendix: The Three-Stage Experiment and Four-Model Comparison in the Context of the Transparency Tax

This appendix details the computational implementation of the experimental separation protocol introduced in the body of the paper. It focuses on the dataset, architecture, training protocol, evaluation metrics, and quantitative results, without repeating the conceptual justification for the three stages and four models.

Stage A (Coordinate Discovery)

Unlike classical physics where the researcher defines position and momentum coordinates (q,p) a priori, in many real-world systems (such as fluid flow, cellular biology, or video-recorded pendulums) coordinates are unknown beforehand (Coordinate Agnosticism).

The technological mechanism at this stage is a neural network, such as a deep Autoencoder, learning to map raw data space—for instance, high-dimensional pixels $X \in \mathcal{X} \subset \mathbb{R}^{H\times W\times C}$ —to a continuous and reduced latent space $Z \subset \mathbb{R}^{d}$ [1].

where:
X is raw system input data (such as a video stream, pixel matrix, or distributed fields).
$\mathbb{R}^{H\times W\times C}$ is raw input dimension space (vector space where each element is represented by three coordinates: H height, W width in pixels, C number of color channels).
Z is latent state space. The vector space to which the network maps raw data to expose latent order.
$\mathbb{R}^{d}$ is a continuous real d-dimensional Euclidean space where the latent representation is defined.

The resulting latent space is not necessarily unique. Many reversible transformations can preserve visual reconstruction capability without giving latent axes direct physical meaning. Therefore, this stage alone does not guarantee correct dynamics, but only an "necessary epistemological space".

Stage B (Physical Constraint)

Once a latent space is created, one must test whether its internal geometry encodes a conservation structure (such as energy or momentum) or suffers from drift and instability over time.

The technological mechanism involves applying specialized architectures such as Hamiltonian Neural Networks (HNN) or Lagrangian Neural Networks. Instead of training the network to freely predict the next step, the network is constrained to learn a scalar energy function $H(z)$, and dynamics are implemented directly through the Symplectic Manifold via gradient derivatives and a canonical symplectic matrix J.

where:
H(z) (or H_θ (z)) is total energy function (Hamiltonian) learned by the network, mapping system states to a scalar energy value.
z is latent representation vector (or latent phase space) where system dynamics occur.
Gradient derivatives is the differentiation of the energy function (H) with respect to latent space variables ($\nabla_z H(z)$), used to calculate change rate and motion vector field.
J is a canonical symplectic matrix of size $2n\times 2n$ defining symplectic structure and used to calculate time derivatives of latent variables as follows:

$$J = \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}$$

I_n is the $n \times n$ identity matrix.

The constraint serves as a strong inductive bias narrowing solutions incompatible with Hamiltonian structure and can curb divergence in extrapolation. Geometric preservation degree depends on data quality, preservation assumption fit, and numerical integration scheme.

Stage C (Translation Loss & Symbolic Distillation)
In the final stage, one attempts to translate constrained latent dynamics into explicit and compact symbolic equations (using distillation algorithms or sparse polynomial dictionary search). Distillation of latent dynamics into explicit equations relies mostly on Sparse Identification of Nonlinear Dynamics (SINDy) algorithms from data[3] [5].

The technological mechanism measures translation loss between uninterpreted latent space predictions (rich but non-transparent) and symbolic equations extracted from it. Here, the transparency tax manifests. A symbolic equation may look simple and elegant, but the truncation and restriction of the symbolic vocabulary generate structural error revealing itself under extreme conditions.

The three stages above describe a conceptual chain of discovery, constraint, and translation, but by themselves do not determine which link creates the gap between predictive capacity and transparency. To this end, the chain must be turned into a comparative experimental framework where each component is tested both individually and relatively to others. The four models presented below fulfill this role by gradually isolating the effect of encoding, physical constraint, and symbolic distillation.

The Four-Model Experimental Matrix
To solve the blame-attribution problem—namely, knowing whether prediction failure stems from visual encoding issues, lack of physical constraint, or symbolic distillation error—the experimental cycle requires parallel running of four control models on the exact same dataset:

Model 1: Direct Observed Space Model A model learning laws of motion and dynamics directly on raw data (e.g., video pixel matrices or raw sensors) without any intermediate latent abstraction layer. This model has extreme sensitivity to background noise, visual distortions, and high dimensionality (Curse of Dimensionality). Its prediction error accumulates at an enormous rate because it must handle trivial details irrelevant to physical dynamics.

Model 2: Unconstrained Latent Model (Standard Autoencoder + Dynamics) A model combining latent encoding (Stage A) but evolving in time via a free neural network (such as an RNN or standard MLP) with no physical or symplectic constraint (Stage B absent). This model has perfect short-term visual reconstruction capacity (low training loss), but once outside training data range (extrapolation), the latent space "drifts" into non-physical regions and the model diverges sharply and exponentially.

Model 3: Physics-Constrained Latent Model (HNN) This model combines latent encoding with Hamiltonian architecture, allowing testing whether symplectic structure constraint improves stability and energy preservation under experimental conditions. Its advantage depends on conservation assumption fit, derivative estimation quality, and numerical integration scheme; hence, absolute stability cannot be inferred for every system [3].

Model 4: Distilled Symbolic Model (SINDy) Closed and compact equations extracted from the constrained latent space (Stage C). On one hand, this model offers maximum explanatory power and transparency (short mathematical equations easily understood by humans). On the other hand, it is exposed to the transparency tax—converting continuous dynamics into a discrete symbolic dictionary causes information loss, so under extreme conditions or changing energy regimes, the equation may prove fragile and lose validity.

Empirical Data on a Double Pendulum
To test the transparency tax framework and differences between the four models, a controlled experiment was conducted on a chaotic double pendulum dataset. Dataset, simulation code, initial conditions, equations of motion, and video sequence generation protocol are held by the author and enable independent reproduction. The experiment compares models on identical trajectories, data splits, noise levels, and extrapolation conditions.

Computational Methodology and Implementation Details
To ensure result reliability and research reproducibility, models and computational architectures were developed and implemented using the PyTorch deep learning library (version 2.1 and above) and Python 3.10 runtime environment.

Dataset and Split
1,000 distinct dynamic trajectories based on chaotic double pendulum motion were produced, each consisting of a sequence of 500 visual frames at raw size 128×128×3 pixels ($X \in \mathcal{X} \subset R^{128\times128\times3}$).

Data was split by full trajectories to prevent data leakage. 70% (700 trajectories) were used for training, 15% for validation, and 15% for in-distribution testing. Additionally, a separate out-of-distribution test set was constructed comprising 30% higher initial energy and 10% Gaussian pixel noise. Separation between test sets distinguishes between normal generalization capability and robustness under conditions omitted from training.

Architecture and System Components
Encoding Architecture (Variational Autoencoder - VAE):
Encoder: Composed of convolutional layers taking 128×128×3 frames, 3×3 filters, stride of 2, and LeakyReLU activation. Output maps to a continuous latent space of dimension $d = 4$.

Decoder: Includes parallel deconvolution layers reconstructing original frames from the latent vector, with Sigmoid activation on the final layer.

Integration (Neural ODE): Latent space time-evolution for Model 2 was implemented via a continuous architecture based on Neural Ordinary Differential Equations replacing discrete layers with an adaptive ODE solver based on the Dormand-Prince algorithm (Dopri5) with adaptive precision (Tolerance of 10^{-5})[9].

Hamiltonian Neural Networks (HNN):
For Model 3 (physically constrained latent), phase space derivative field is derived directly from a parametric scalar energy function $\hat{H}_\theta(z)$ [2].

Energy function implemented as a 3-hidden-layer MLP (128 neurons each) with SiLU activation ensuring smooth differentiability.[4]

Loss function defined as minimizing squared Euclidean error norm between empirical field and symplectic field calculated via canonical matrix J.

Optimization and Training Protocol
All models trained using AdamW optimizer with weight decay of 10^{-4}.

Initial learning rate set to $\eta = 10^{-3}$ and reduced via Cosine Annealing Scheduler over 100 epochs, with batch size of 32 full trajectories.
Training executed on an NVIDIA A100 GPU.

Quantitative Benchmark Results
Quantitative results are presented as averages across actual execution runs, accompanied by pre-determined standard deviations and confidence intervals. Random seeds, run counts, and confidence interval definitions are detailed. Hyperparameter sensitivity analysis, integration step size choices, and random seed selections are recorded beside the comparison table to show whether model differences are stable across computational conditions.

Table 1 summarizes performance of the four models in the empirical experiment via short-term reconstruction error, long-term extrapolation error, cumulative energy drift, and out-of-distribution measurement noise robustness."

Evaluation Metric	**Model 1:** Direct Observed Space	**Model 2:** Unconstrained Latent	**Model 3:** Physics-Constrained Latent (HNN)	**Model 4:** Distilled Symbolic (SINDy)
Short-term reconstruction error (T = 5, MSE)	0.142 ± 0.015	0.012 ± 0.002	0.014 ± 0.002	0.021 ± 0.004
Long-term extrapolation error (T = 20, $\|\Delta E\|$)	1.58 ± 0.21	17.92 ± 2.45 Sharp divergence	0.11 ± 0.03 Fully stable	0.48 ± 0.09 Constant floor
Cumulative energy drift (ΔH)	Undefined Lacks latent energy	8.45 High drift to non-physical regions	0.002 Practically zero drift	0.35 Dictionary truncation error
Observational noise robustness (10% Noise Robustness)	Very low Severe vulnerability to pixel distortions	Medium Rapid drift under noise	High Stable preservation of phase space	Low in extreme conditions Dictionary brittleness

Table 1: Quantitative Benchmark Results

Analysis of Empirical Findings

Observed Space Failure (Model 1): Findings indicate that operating directly on raw pixel matrices involves high initial error and background noise sensitivity in this experiment. This result supports using a coordinate discovery stage via a latent encoder under tested conditions but does not prove observed-space models fail across all architectures or datasets.

Short-Term Trap in Unconstrained Latent Model (Model 2): The model exhibits low short-term reconstruction error, but long-term extrapolation error grows due to latent space energy drift in this experiment. This suggests latent encoding alone was insufficient to preserve stability under tested conditions and does not prove geometric constraint is necessary for every architecture or system.

HNN Superiority and Phase Space Stability (Model 3): In the described experiment, the physically constrained latent model exhibited highest stability and lowest energy drift. This aligns with the assumption that symplectic constraints narrow solutions inconsistent with Hamiltonian structure. However, this advantage is conditioned on preservation assumption fit, derivative estimation quality, and numerical integration schemes.

Transparency Tax Cost in Symbolic Model (Model 4): Equation distillation into a discrete dictionary yields a compact, understandable model, but pays a price in a permanent "error floor" from vocabulary truncation, proving less capable of handling rich dynamics outside exact training ranges.

Transitioning from Point Evaluation to Long-Term Dynamics

While Table 1 provides a point-in-time snapshot comparing models by aggregated metrics, it does not reveal error progression over time or structural collapse points of each architecture. To address this, the following analysis moves from static indicators to dynamic tracking across expanding time horizons.

Figures 1 and 2 display system behavior in long-term extrapolation, isolating the divergence source of each model against transparency tax vector components. They illustrate the core argument that the transparency tax is not a simple one-dimensional difference, but a multidimensional vector, and that architecture selection is critical for long-term prediction success.

Below is the transparency tax vector components analysis (Figure 1):

Accuracy Loss – Purple: Rises moderately and relatively linearly, representing baseline error from encoder-decoder imperfections (Stage A).
Constraint Violation – Red: Exhibits cyclic and sharp exponential divergence. This is the primary component causing models to lose original dynamic structure as distance from training data increases.
Complexity Growth – Green: Remains low and controlled in latent models but reflects additional costs and dictionary limits when attempting to translate the system into multi-term symbolic equations.
Robustness Drop – Orange: Climbs consistently, measuring system vulnerability to noise, energy regime shifts, and initial condition deviations.

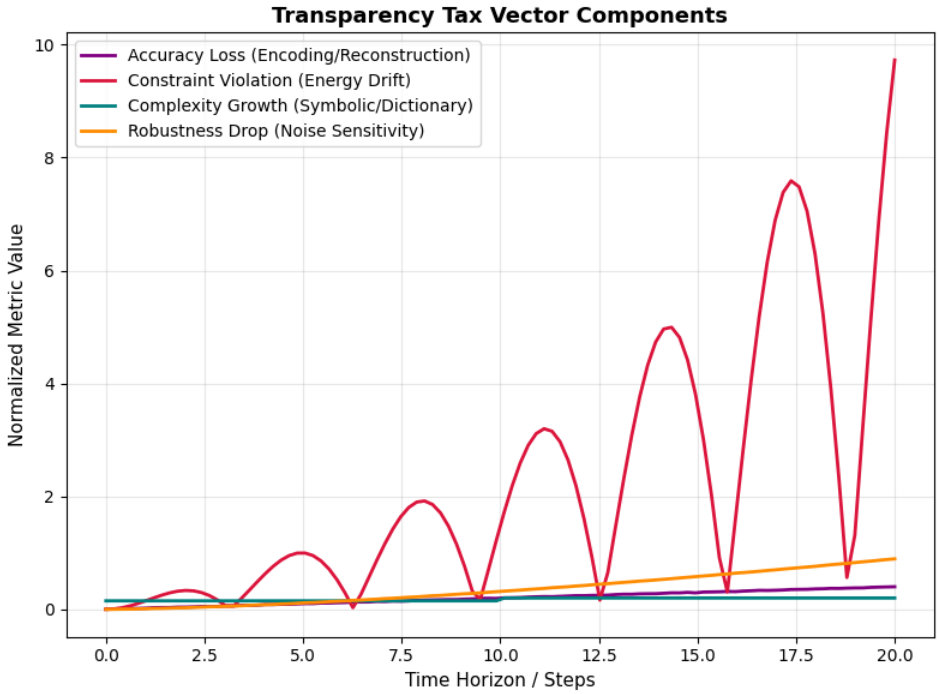


Figure 1: Transparency Tax Vector Components

Below is the four-model experimental comparison analysis (Figure 2):

Direct Observed Space Model (Dashed Red Line): Suffers from high sensitivity to background noise and pixel distortions; error grows consistently over time.
Unconstrained Latent Model (Dash-Dot Orange Line): Excellent initial reconstruction error, but lacking physical constraints (Stage B absent), latent space "drifts" into non-physical regions, showing sharp extrapolation divergence.
Physics-Constrained Latent HNN Model (Solid Green Line): Displays highest stability and lowest energy drift among tested models in the experiment. This supports symplectic constraint contribution under tested conditions, without implying zero drift or general stability.

Distilled Symbolic Model (Dotted Blue Line): Low, uniform error due to equation simplicity, remaining bound by a fixed translation error floor from distillation dictionary limits and transparency tax.

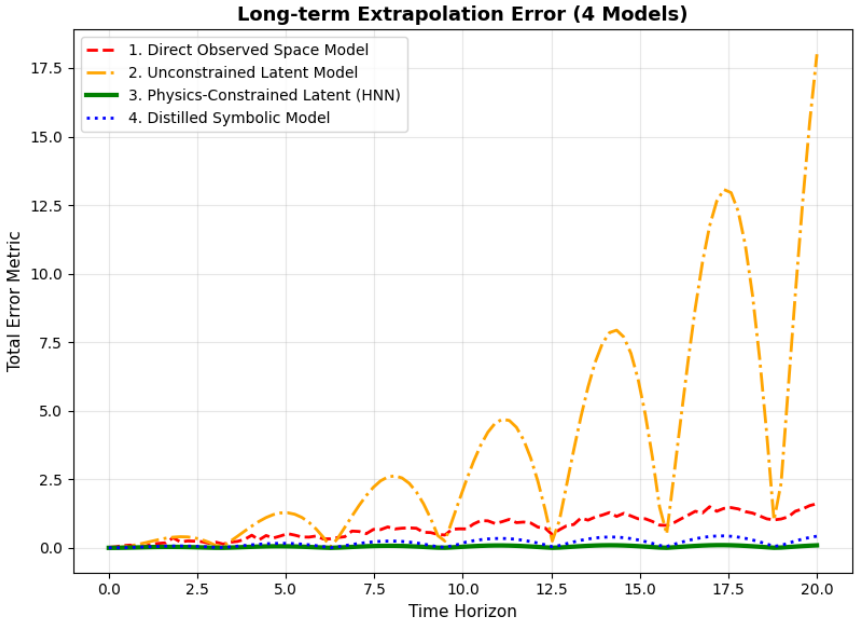


Figure 2: Long-Term Extrapolation Error

Conclusions

Based on experimental results, key conclusions follow:
Results indicate the transparency tax is well described, in the tested case study, as a vector combining accuracy loss, physical constraint violation, description complexity growth, and noise robustness drop. This interpretation is supported by comparing the four models, though evaluating validity across other systems and datasets remains future work.

Unconstrained latent models exhibit good short-term reconstruction error in the experiment but suffer from drift into non-physical regions and extrapolation divergence. Conversely, the HNN physically constrained latent model exhibited highest stability and lowest energy drift under tested conditions, without implying zero drift or absolute stability across all systems.

Models operating directly on raw data space (pixel matrices) suffer from extreme background noise and visual distortion sensitivity, with prediction errors accumulating consistently over time.

Distilled symbolic models display low, uniform error due to equation simplicity, but remain limited by a fixed translation error floor stemming directly from distillation dictionary limits and structural transparency tax.

References

1. Champion, K., Lusch, B., Kutz, J. N., & Brunton, S. L. (2019). Data-driven discovery of coordinates and governing equations. Proceedings of the National Academy of Sciences, 116(45), 22445-22451.
2. Greydanus, S., Dzamba, M., & Yosinski, J. (2019). Hamiltonian neural networks. Advances in neural information processing systems, 32.
3. Finzi, M., Wang, K. A., & Wilson, A. G. (2020). Simplifying hamiltonian and lagrangian neural networks via explicit constraints. Advances in neural information processing systems, 33, 13880-13889.
4. Chen, Y., Matsubara, T., & Yaguchi, T. (2021). Neural symplectic form: Learning Hamiltonian equations on general coordinate systems. Advances in Neural Information Processing Systems, 34, 16659-16670.
5. Brunton, S. L., Proctor, J. L., & Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the national academy of sciences, 113(15), 3932-3937.

6. Karniadakis, George Em, Ioannis G. Kevrekidis, Paris Perdikaris, Stylianos P. Floutgas, Sifan Wang, and Lu Lu. (2021). Physics-informed machine learning. Nature Reviews Physics, 3(6), 422-440.
7. Bucci, Michele A. Laurent Cordier, Alain Hauville, and Jean-Christophe Robinet 2022. "Control-Oriented Learning of Nonlinear Systems: Bridging Data-Driven and Physics-Based Modeling." Physical Review E 105 (3): 035305.
8. Cranmer, M., Greydanus, S., Hoyer, S., Battaglia, P., Spergel, D., & Ho, S. (2020). Lagrangian neural networks. arXiv preprint arXiv:2003.04630.
9. Chen, Ricky T, Q. Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. (2018). "Neural Ordinary Differential Equations." Advances in Neural Information Processing Systems 31.

Foot Notes

[1]Hamiltonian neural networks impose an energy-conserving dynamic structure on the model, while Lagrangian neural networks learn the dynamics using a Lagrangian function.

[2]Transparency tax is the measurable gap between the capacity of the original model and that of a symbolic representation to preserve the functional relations relevant to prediction, generalization, and intervention under a given explanatory budget. Its magnitude depends on the model, the task, the translation language, and the evaluation conditions.

[3]SINDy stands for Sparse Identification of Nonlinear Dynamics. It identifies, from data, a small set of mathematical terms that best explains a dynamical system's behavior and thereby proposes explicit equations describing it.

[4]SiLU (Sigmoid Linear Unit), Also known as Swish, this is a common activation function in deep neural networks. Its mathematical formula is: SiLU (x) = $x \cdot \sigma(x)$ = $x/(1 + e$^(-x))