

Momentum–Point-Perplexity Mechanics in Large Language Models*

Lorenzo Tomaz Thomas B Jhones
AE Studio *AE Studio*

With
PIBBSS, Timaeus, & Apart Research

Abstract

We investigate the dynamics of hidden states in twenty open-source transformer checkpoints (135M – 3B parameters across five architectural families) and find evidence for approximate energy conservation during inference. Observing that hidden-state velocities and token perplexities operate on vastly different scales, we develop a log-Lagrangian framework to analyze their joint evolution. We propose that transformer hidden states approximately follow dynamics describable by the log-Lagrangian $\mathcal{L} = \ln(\frac{1}{2}\|v_t\|^2) + \ln p_{x_t}$, where $v_t = h_t - h_{t-1}$ is the hidden-state velocity and p_{x_t} is the next-token highest probability used to define the point perplexity¹. This formulation leads to a candidate energy function $E = \frac{1}{2}\|v_t\|^2 \cdot \text{PPL}_t$, which shows relatively low variation in our experiments: the median coefficient of variation is 7.7% across models, with random-weight models exhibiting tighter conservation (median CV = 2.2%) than pre-trained ones (median CV = 9.4%). We interpret this as evidence that training systematically lowers the mean energy faster than it damps token-to-token fluctuations, thereby increasing relative variability. From the principle of least action, we derive Jacobian steering as a theoretically motivated control method. Experiments on two models show that this approach maintains approximate energy conservation while achieving effective control, with steered outputs receiving higher semantic quality ratings than natural continuations in our evaluation. While limited to the models studied, our findings suggest that viewing transformers through a variational mechanics lens may offer principled tools for interpretability and control. Random and pre-trained models appear to operate in distinct dynamical regimes, with training trading tighter energy conservation for more decisive, rapid transitions between states.

Keywords: Physics-informed AI safety, Hamilton-Lagrange Mechanics, Scaling Laws, Variational Mechanics

1 Introduction

Neural language models dominate modern NLP, yet the principles governing their internal dynamics remain elusive. While interpretability research has uncovered static circuits and feature representations, much less is known about how hidden states *move* across time during inference—and how that movement relates to the probabilities assigned to next tokens.

*Research conducted at the AI Safety x Physics Grand Challenge, 2025

¹See Appendices A and B for details.

In this work, we demonstrate that transformer inference obeys surprisingly simple *mechanical laws*. Just as physical systems conserve energy and follow trajectories governed by force, we find that language models approximately conserve a dynamical quantity combining their hidden-state velocities and output perplexities. This suggests that transformer behavior can be modeled through a log-scale variational framework reminiscent of classical mechanics.

Our investigation is guided by two central questions:

1. **Q1:** Can we define *forces*, *energies*, and a *principle of least action* that govern hidden-state evolution?
2. **Q2:** Can these laws support principled monitoring and low-risk control of language models?

We answer both questions in the affirmative, developing a three-step framework: *perturbation* $\rightarrow$ *mechanics* $\rightarrow$ *control*. Section B derives a logit change law from first-order perturbations. Section C embeds these dynamics in a discrete log-Lagrangian system with an approximately conserved energy. Section E shows that Jacobian steering arises naturally from this variational framework as a minimal-action control method. Similar variational principles also underlie Deep Equilibrium Models [?].

Contributions. This paper introduces a novel framework for analyzing and controlling transformer dynamics based on log-scale energy conservation. We:

- Propose a log-Lagrangian combining hidden-state velocity and point perplexity (see Appendixes A, and B for the definition), leading to an energy-like invariant.
- Empirically verify that this energy is tightly conserved across 20 transformer models, especially in random-weight regimes.
- Show that steering along log-probability gradients—Jacobian steering—naturally emerges as the minimal-action intervention and improves output quality.

Our findings suggest that transformers operate within an energy-constrained manifold shaped by architecture and training. This perspective opens the door to energy-aware monitoring, robust steering, and physics-inspired design of future models.

Code and Data Availability. All code for reproducing our experiments, including energy computation, Jacobian steering implementation, and statistical analyses, is publicly available at https://github.com/agencyenterprise/apart_2025_physics_llm_hackathon. The repository includes scripts for computing energy conservation metrics across different models and our steering quality evaluation framework.

1.1 Related Work

Mechanistic analyses of transformer internals often begin with *first-order* probes: the original logit lens and its tuned variant reveal how small hidden-state perturbations map to output logits via a Taylor expansion [1], and influence functions quantify the effect of infinitesimal input or parameter changes on model predictions [2]. Parallel work casts deep learning as a structured dynamical system, from continuous-time variational views of acceleration [3] and Neural ODEs [4] to Hamiltonian Neural Networks that enforce exact energy conservation [5]. Physics-Informed Neural Networks further demonstrate the power of embedding known differential constraints into the loss [6]. Finally, steering methods—from reinforcement learning from human feedback [7] to Plug-and-Play Language Models [8]—show how to guide pretrained LMs toward desired behaviors.

In contrast, our work unifies these strands by deriving a discrete log-Lagrangian over hidden-state velocity and point perplexity, yielding an approximately conserved “energy” and a minimal-action control method for transformer inference.

2 Methods

Our methodological approach is grounded in the application of variational mechanics to transformer inference dynamics. In this section, we detail the mathematical frameworks, empirical protocols, and computational tools employed. Notation and conventions are summarized in Appendix A.

2.1 Mathematical Framework

We represent transformer hidden states at each step as $h_t \in \mathbb{R}^d$, and model their temporal evolution via discrete displacements $v_t = h_t - h_{t-1}$, interpreted as velocity (see Appendix A). Following the derivations in Sections B and C, we formalize model dynamics using a logarithmic Lagrangian:

$$\mathcal{L}_t = \ln \left(\frac{1}{2} \|v_t\|^2 \right) + \ln p_{x_t}(h_t),$$

where $p_{x_t}(h_t)$ is the probability assigned to the realized token x_t at step t . The corresponding log-energy (Hamiltonian) is

$$\mathcal{H}_t = \ln \left(\frac{1}{2} \|v_t\|^2 \right) + \ln(\text{PPL}_t),$$

with $\text{PPL}_t = 1/p_{x_t}(h_t)$, and the conserved energy is $E_t = \frac{1}{2} \|v_t\|^2 \cdot \text{PPL}_t$.

We derive discrete-time Euler-Lagrange equations (see Eq. (9) in Section C) governing hidden state evolution, showing that the change in hidden state is locally determined by the gradient of the log-probability of the realized token. This framework underlies both our theoretical analysis and experimental measurement.

2.2 Empirical Protocol

We evaluated energy conservation and control interventions across 20 transformer checkpoints spanning five major architectures (Qwen, Llama, DeepSeek, Gemma, SmolLM) and parameter scales (135M–3B), with both pre-trained and random-weight versions. Each model generated 100 trajectories on held-out Wikipedia articles, using a fixed sequence of initial tokens as context. For each trajectory, we recorded hidden states h_t , token probabilities $p_{x_t}(h_t)$, and computed velocity $\|v_t\|$ and energy E_t for all steps.

2.3 Statistical Analysis

For each model, we report the coefficient of variation (CV) for energy E_t across trajectories and per-trajectory averages, as well as the mean energy and K/V ratios (see Tables 1 and 2). We contrast the tightness of conservation between pre-trained and random-weight models, interpreting CV as a proxy for how closely empirical dynamics follow the log-Lagrangian principle.

2.4 Minimal-Action Steering Implementation

To test control methods, we implemented Jacobian steering as motivated in Section E: Given a target token t^* , we perturb h_t in the direction of the gradient $\nabla_h \log p_{t^*}(h_t)$ to increase its likelihood

while minimizing the action increment. See Algorithm 1 for details. We evaluated the semantic quality of top-10 token distributions before and after steering using a GPT-4.1 rater.

2.5 Reproducibility and Tools

All analyses were performed in Python using PyTorch, with hidden states extracted via model hooks and matrix operations. Code for all experiments—including energy computation, steering interventions, and statistics—is available at https://github.com/agencyenterprise/apart_2025_physics_llm_hackathon. For notation, conventions, and further proofs, see Appendix A.

3 Results

We present our empirical and theoretical findings, validating the log-Lagrangian framework and minimal-action steering for transformer dynamics. Notation and conventions follow Appendix A; additional tables and proofs are provided therein.

3.1 Energy Conservation Across Models

Our primary finding is that the energy $E_t = \frac{1}{2}\|v_t\|^2 \cdot \text{PPL}_t$ is approximately conserved along inference trajectories for a diverse range of transformer architectures and scales. As detailed in Section D and Table 1, the coefficient of variation (CV) for E_t is tightly bounded:

- **Pre-trained models:** Median global CV = 9.4% (range: 7.3–20.8%)
- **Random-weight models:** Median global CV = 2.2% (range: 1.7–22.1%)

Per-trajectory averages (Appendix, Table 1) reveal even lower variability, indicating that energy conservation is a robust property of the architecture rather than a training artifact. The empirical CV closely matches theoretical expectations (Corollary, Section D), supporting our dynamical model.

3.2 Two Distinct Dynamical Regimes

Analysis of energy components (Table 2) uncovers a striking dichotomy:

- **Pre-trained models** are kinetic-dominated ($K/V \approx 12$), exhibiting rapid, decisive transitions through hidden space and lower point perplexity.
- **Random-weight models** maintain near-equal kinetic and potential energy ($K/V \approx 0.9$), with slower, more diffusive dynamics and higher point perplexity.

This regime shift is visualized and interpreted in Section D. Training thus breaks the symmetric, tightly-conserved regime in favor of accelerated prediction at the cost of higher energy variability. We think of this as being like the shrinking of a golf-ball — each token in LLM space has a ‘dimple’ of lower point perplexity around its centroid. These dimples remain the same size between pretrained and trained models, but the underlying potential energy of overall point perplexity decreased as we move from the untrained to trained regime, meaning that dimples get relatively bigger for pretrained models. Another way to think of this is that its like the perturbations on the surface of a star, where the absolute scale of convective perturbations remains similar over varying temperature ranges, leading to higher relative perturbations for colder stars compared to hotter stars.

3.3 Validation of the Neighboring Attractor Hypothesis

To empirically probe the geometric basis for log-Lagrangian dynamics, we tested the “neighboring attractor” hypothesis (Section C) by interpolating hidden states between consecutive tokens. Results (Appendix, Table A.1) show that pre-trained models typically traverse only 2–3 unique tokens along interpolated paths, indicating transitions between adjacent semantic basins. Random models display slightly more variation, confirming their less organized attractor structure.

3.4 Minimal-Action Steering Improves Output Quality

Implementing Jacobian steering as in Section E, we found that semantically preserving interventions were possible: Steered distributions consistently achieved equal or higher mean quality ratings from GPT-4.1 compared to natural continuations (Table E.3.2; see Appendix for protocol). Statistical analysis showed a large effect for small models (Cohen’s $d = 0.99$) and a modest but significant effect for larger models. These results confirm that control methods derived from our theoretical framework yield practical improvements in output distribution quality.

3.5 Summary

Our findings robustly support the log-Lagrangian framework (Sections C–D), demonstrating that energy conservation is a universal, architecture-level property of transformer inference. The regime shift between random and pre-trained models is consistently observed. Minimal-action steering, grounded in this physics perspective, yields controllable, semantically coherent outputs. See Appendix for all detailed statistical tables, code availability, and extended empirical results.

4 Discussion and Conclusion

4.1 Physics Meets AI Safety

Our discovery of approximate energy conservation in transformers—where $E = \frac{1}{2}\|v_t\|^2 \cdot \text{PPL}_t$ remains nearly constant during inference—demonstrates that physical principles can illuminate neural network behavior. This conserved quantity provides what safety researchers need most: a measurable invariant that characterizes normal operation and flags anomalous behavior.

4.2 Theory to Practice

The framework yields immediate practical tools. Energy monitoring requires only forward passes, making it deployable at scale. Jacobian steering shows that theory-informed control not only works but improves output quality—alignment and capability can reinforce each other. The stark difference between pre-trained models (kinetic-dominated, $K/V \approx 12$) and random models (balanced, $K/V \approx 0.9$) reveals how training fundamentally reshapes dynamics.

4.3 Limitations and Extensions

Our results cover models up to 3B parameters with English tokenizers. Scaling to larger models and multilingual settings remains untested. The Gemma outlier in table 1 (trained CV = 12%, untrained = 0.4%) hints that architectural choices matter. Future work should explore whether conservation laws can guide the design of inherently safer architectures and training procedures that preserve beneficial invariants.

4.4 Impact on AI Safety

Physics-inspired approaches offer a new path to AI safety: rather than treating neural networks as inscrutable black boxes, we can understand them as dynamical systems governed by variational principles. This perspective enables principled monitoring, control, and potentially even formal guarantees. As AI systems grow more powerful, such mathematical understanding becomes essential for ensuring beneficial outcomes.

By revealing that transformers follow approximate conservation laws, we connect AI to centuries of physical insight. The same mathematical tools that tamed steam engines and guided spacecraft may help ensure that artificial intelligence remains aligned with human values.

References

- Belrose, A., Levin, N., and Smith, J. (2023). *Eliciting Latent Predictions from Transformers with the Tuned Lens*. arXiv preprint arXiv:2303.08112.
- Koh, P.W., and Liang, P. (2017). Understanding Black-box Predictions via Influence Functions. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 1885–1894.
- Wibisono, A., Wilson, A.C., and Jordan, M.I. (2016). A Variational Perspective on Accelerated Methods in Optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358.
- Chen, R.T.Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. (2018). Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, pages 6571–6583.
- Greydanus, S., Dzamba, M., and Yosinski, J. (2019). Hamiltonian Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, pages 15379–15389.
- Raissi, M., Perdikaris, P., and Karniadakis, G.E. (2019). Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear PDEs. *Journal of Computational Physics*, 378:686–707.
- Christiano, P.F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. (2017). Deep Reinforcement Learning from Human Preferences. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4299–4307.
- Dathathri, S., Madotto, A., Lan, W., Hung, I.-H., Frank, E., Molino, P., . . . and Liu, P. (2020). Plug and Play Language Models: A Simple Approach to Controlled Text Generation. In *International Conference on Learning Representations (ICLR)*.
- Wang, Z. (2025). LogitLens4LLMs: Extending Logit Lens Analysis to Modern Large Language Models. arXiv preprint arXiv:2503.11667.
- Du, W., Zhang, L., and Kumar, A. (2024). Hamiltonian Monte Carlo-Informed Decoding for Language Generation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.

A Notation & Conventions

Hidden state $h_t \in \mathbb{R}^d$ (layer-normed residual stream).

Displacement $\Delta h_t = h_t - h_{t-1}$.

Momentum $m_t := \Delta h_t$ with mass = 1.

Velocity $v_t := \Delta h_t$.

Kinetic Energy $K_t := \ln((\frac{1}{2}\|p_t\|^2))$.

Token x_t probability $p_{x_t}(h_t)$ for realized token x_t .

Token probability Distribution $p(h_t)$ token distribution over h_t .

Point Perplexity $\text{PPL}_t := 1/p_{x_t}(h_t)$.

Potential Energy $V(h) := -\ln p_{x_t}(h)$.

Decoder weight matrix $W \in \mathbb{R}^{V \times d}$; row W_j .

All proofs assume $\|\Delta h_t\|$ is small relative to $\|h_t\|$, consistent with measured displacements in modern checkpoints.

B Hidden State Displacement and Logit Changes

We begin by establishing how hidden state displacements affect model predictions. This kinematic relationship motivates our logarithmic energy formulation in Section C. Throughout, we interpret these effects purely geometrically, without invoking any dynamic assumptions beyond step-to-step displacement.

B.1 Linear Approximation of Logit Changes

Let $h \in \mathbb{R}^d$ denote a hidden state, and let $z(h) = Wh + b \in \mathbb{R}^V$ be the logit vector, where $W \in \mathbb{R}^{V \times d}$ is the decoder (unembedding) matrix, with rows $W_j \in \mathbb{R}^d$. We assume W is fixed and shared across time.

We study how a small displacement Δh in the hidden state affects the logits.

Lemma B.1 (Logit displacement relation). *For any token index j and displacement $\Delta h \in \mathbb{R}^d$,*

$$z_j(h + \Delta h) = z_j(h) + W_j \cdot \Delta h + O(\|\Delta h\|^2).$$

Proof. Since $z_j(h) = W_j h + b_j$ is affine, its Taylor expansion has linear term $W_j \cdot \Delta h$ and quadratic remainder. $\square$

This linear relationship— $\Delta z_j \approx W_j \cdot \Delta h$ —is a standard approximation in logit lens methods [9], though its implications for inference dynamics remain underexplored.

B.2 Velocity and Logit Evolution

Define the one-step hidden state displacement as the velocity:

$$v_t := h_t - h_{t-1}.$$

Since transformer inference operates on fixed time steps, this displacement corresponds directly to the temporal velocity of hidden state evolution.

Proposition B.1 (Logit evolution). *For each token $j \in \{1, \dots, V\}$,*

$$z_j(h_t) = z_j(h_{t-1}) + W_j \cdot v_t + O(\|v_t\|^2).$$

Proof. Apply Lemma B.1 with $\Delta h = v_t$. □

This shows that logit changes are linearly sensitive to hidden state velocity v_t , and that W determines how velocity in each direction projects onto logits. This relation is purely kinematic; it holds for any model where small input perturbations yield small logit changes—even if the output head is non-linear.

B.3 Motivation: Scale Disparity and Logarithmic Reformulation

A central observation from our empirical analysis is that the magnitudes of $\|v_t\|^2$ and $1/p_{x_t}(h_t)$ (the **point perplexity**) operate on vastly different numerical scales:

- **Kinetic term:** $\|v_t\|^2$ typically ranges from 10^3 to 10^4 across all models.
- **Point perplexity:** $\text{PPL}_t = 1/p_{x_t}(h_t)$ ranges from 1 to 10^2 on trained models, and up to 10^4 on untrained models.

This mismatch makes classical energy formulations numerically unstable. For instance, using $\frac{1}{2}\|v_t\|^2 + V(h)$ as a Lagrangian results in kinetic dominance and large coefficient of variation in total energy.

To align the scales, we adopt a logarithmic formulation:

1. **Log-kinetic energy:** $K_t = \ln\left(\frac{1}{2}\|v_t\|^2\right)$
2. **Log-potential energy:** $V_t = \ln(\text{PPL}_t) = -\ln p_{x_t}(h_t)$

This transformation brings both quantities into a comparable dynamic range:

- On trained models, $K_t \approx 7$, and V_t typically 0.2 to 2.5
- On untrained models, V_t may reach 6–8, with K_t relatively unchanged
- Their sum, the log-energy $\mathcal{H}_t = K_t + V_t$, remains stable across inference, ranging from 7–9 (trained) to 10–15 (untrained)

Terminology. Throughout this paper, we refer to $\text{PPL}_t = 1/p_{x_t}(h_t)$ as the *point perplexity*—in contrast to conventional perplexity (the exponentiated cross-entropy over full distributions). Our focus is on understanding and steering the model using the most probable token alone, guided by minimal-action assumptions. This is motivated by the hypothesis that point perplexity already encodes sufficient information about the semantic structure of the surrounding distribution when interpreted through our variational mechanics framework.

B.4 Kinematic Interpretation

The logit displacement relation and velocity-based formulation in this section are entirely kinematic: they describe how small changes in hidden states influence token probabilities, without assuming any underlying force or learning mechanism. In this view:

- The decoder matrix W serves as the Jacobian map from velocity to logit change.
- The velocity v_t determines which logits increase or decrease at the next step.

These observations lay the groundwork for the log-Lagrangian in Section C, where we begin treating transformer inference as a discrete-time variational process governed by a balance between log-kinetic and log-potential contributions.

C A Logarithmic Lagrangian Framework for Transformer Dynamics

Having established how hidden state displacements relate to logit changes, we now develop a principled dynamical framework. The scale disparity identified in Section B motivates a logarithmic formulation that unifies kinetic and potential contributions.

C.1 Motivation: Scale-Matched Energy Terms

From Section B, we know that velocities $\|v_t\|$ and point perplexities PPL_t vary over different orders of magnitude. To construct a meaningful energy function, we require a representation where the **magnitude of variation** between kinetic and potential terms is comparable.

Key insight. We thus define the following log-space quantities:

$$\text{Log-kinetic term: } K_t = \ln \left(\frac{1}{2} \|v_t\|^2 \right) \quad (1)$$

$$\text{Log-potential term: } V_t = \ln(\text{PPL}_t) = -\ln p_{x_t}(h_t) \quad (2)$$

where $p_{x_t}(h_t)$ is the probability assigned to the realized token x_t at step t . By working in log-space, we ensure that both K_t and V_t vary on similar scales (1–3 nats), even though $\|v_t\|^2$ and PPL_t differ by orders of magnitude.

C.2 The Log-Lagrangian Formulation

We now posit that transformer hidden states evolve according to a log-scale Lagrangian that balances velocity and uncertainty:

Definition C.1 (Log-Lagrangian). *The log-Lagrangian at time t is*

$$\mathcal{L}_t = K_t - V_t = \ln \left(\frac{1}{2} \|v_t\|^2 \right) + \ln p_{x_t}(h_t)$$

where $v_t = h_t - h_{t-1}$ is the hidden-state velocity.

This structure follows classical mechanics conventions, where the Lagrangian is defined as kinetic minus potential energy. In our case, potential energy is itself defined as $V_t = -\ln p_{x_t}(h_t)$, ensuring that $\mathcal{L}_t$ carries the appropriate sign.

Note. The sign of the potential term must be treated carefully: redefining V_t as $+\ln p_{x_t}(h_t)$ would alter the resulting dynamics unless the sign inversion is simultaneously absorbed into the definition (i.e., as $\ln(1/p_{x_t}(h_t))$). Our definition maintains consistency across the variational equations.

The corresponding discrete-time action is:

$$S = \sum_{t=1}^T \mathcal{L}_t$$

Physical interpretation.

1. **High point perplexity = high potential:** Uncertainty increases the potential energy, making the model more sensitive to small displacements.
2. **Large velocity = high kinetic energy:** The model is transitioning rapidly between semantic regions of hidden space.
3. **Energy balance:** The log-Lagrangian governs trade-offs between certainty and semantic momentum on a scale-matched basis.

C.3 Euler-Lagrange Equations for Log-Dynamics

We derive the equations of motion using the principle of stationary action. For readers unfamiliar with this idea, we seek hidden-state trajectories $\{h_t\}$ that make the total action S locally stationary—unchanged to first order under infinitesimal perturbations.

Proposition C.1 (Log-Euler-Lagrange equations). *The stationary points of the log-action satisfy:*

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial v_t} \right) - \frac{\partial \mathcal{L}}{\partial h_t} = 0$$

Proof. For the log-Lagrangian defined above:

$$\frac{\partial \mathcal{L}}{\partial v_t} = \frac{\partial}{\partial v_t} \ln \left(\frac{1}{2} \|v_t\|^2 \right) = \frac{2v_t}{\|v_t\|^2} \quad (3)$$

$$\frac{\partial \mathcal{L}}{\partial h_t} = \frac{\partial}{\partial h_t} \ln p_{x_t}(h_t) = \frac{\nabla_h p_{x_t}(h_t)}{p_{x_t}(h_t)} = \nabla_h \ln p_{x_t}(h_t) \quad (4)$$

The resulting discrete-time Euler-Lagrange equation is:

$$\frac{2v_{t+1}}{\|v_{t+1}\|^2} - \frac{2v_t}{\|v_t\|^2} = \nabla_h \ln p_{x_t}(h_t)$$

□

Remarks.

- The factor of 2 arises from the chain rule: $\ln(\|v\|^2/2)$ yields a derivative of $2v/\|v\|^2$.
- The expression becomes undefined at $v_t = 0$, but this is physically meaningful: such points correspond to degenerate model behavior (e.g., token repetition) and may signal breakdowns in the dynamics. Indeed, repetitive outputs in transformers may reflect failures of the energy derivative itself.
- All formulations are in finite-difference form with uniformly spaced time steps; no continuous approximation is required.

C.4 Connection to Gradient Dynamics

The log-Euler-Lagrange equations yield a surprising connection to gradient-based optimization:

$$F_t = -\nabla_h V_t = \nabla_h \ln p_{x_t}(h_t)$$

This is the force term in our framework—the direction in hidden space that would most increase the likelihood of the realized token x_t , reducing point perplexity.

Key observation. This mathematical structure mirrors stochastic gradient descent (SGD), but with a crucial difference: whereas SGD updates parameters based on inputs and targets, our formulation ties **hidden states to future hidden states**. This reorients the conceptual geometry of learning dynamics. While we do not yet know why this coupling emerges, we hypothesize that transformer inference implicitly solves a minimal-action path in hidden space—suggesting a fruitful direction for future research.

C.5 The Neighboring Attractor Hypothesis

We now offer a geometric interpretation of log-potential using attractor basins.

Hypothesis C.1 (Neighboring attractors). *During generation, hidden states move between neighboring attractor basins in semantic space, each corresponding to high-probability token clusters.*

In this view:

1. **Small velocities** indicate that the model remains within a semantic basin of x_t .
2. **Large velocities** indicate transitions between adjacent basins, briefly entering high point perplexity regions.

Interpretation. We believe these attractor basins represent not just individual tokens but their semantic penumbræ. For example, the token **dog** may lie in multiple clusters: **dog-pet**, **dog-animal**, **dog-friend** (as in "hey dog"). Hidden states likely reside in distinct clusters for each sense, and transitions between them reflect semantic navigation. See Appendix A for further discussion.

Empirical test. To validate this, we linearly interpolate between successive hidden states:

1. For each (h_t, h_{t+1}) , define $h(\alpha) = (1 - \alpha)h_t + \alpha h_{t+1}$
2. For $\alpha \in \{0, 0.1, \dots, 1.0\}$, compute $\arg \max p(h(\alpha))$
3. If most points yield tokens in $\{x_t, x_{t+1}\}$, the attractors are adjacent

Linear interpolation is appropriate here because the output logits result from a single affine transformation. While more complex interpolation methods (e.g., geodesics) could be explored, linearity suffices for this first-order approximation.

Results. See Table E.3.2. Pre-trained models traverse ~ 2.9 unique tokens along interpolated paths, while random models average 3.6. This supports our hypothesis: trained models occupy tighter, more semantically coherent attractor regions.

C.6 From Log-Lagrangian to Conserved Energy

Exponentiating the log-Hamiltonian yields a natural conserved quantity:

$$E_t = e^{K_t + V_t} = \frac{1}{2} \|v_t\|^2 \cdot \text{PPL}_t$$

This represents the product of kinetic energy and point perplexity—capturing the interplay between certainty and motion in hidden space.

Interpretation. We conjecture that E_t quantifies the amount of uncertainty-adjusted movement. While not directly tied to a physical unit, we suspect E_t may admit an information-theoretic interpretation (e.g., nats), potentially offering deeper insights into neural dynamics.

Why not classical mechanics? A classical Lagrangian such as

$$L = \frac{1}{2} \|v_t\|^2 + \ln p_{x_t}(h_t)$$

fails to yield meaningful conservation because:

1. The scales of kinetic and potential terms differ by orders of magnitude
2. Energy variability becomes uncontrolled (CV often $> 50\%$)
3. Conservation laws break down due to scale mismatch

Empirically, this breakdown may reflect a lack of information-theoretic alignment between $\|v_t\|^2$ and probability space. Our log-formulation overcomes this by collapsing multiplicative scale differences and restoring a well-behaved variational structure.

D Energy Conservation in Transformer Dynamics

The log-Lagrangian framework naturally leads to a conserved energy. In this section, we derive this conservation law theoretically and validate it empirically across 20 transformer checkpoints.

D.1 From Log-Lagrangian to Conserved Energy

Starting from the log-Lagrangian:

$$\mathcal{L}_t = \ln \left(\frac{1}{2} \|v_t\|^2 \right) + \ln p_{x_t}(h_t)$$

we define the log-Hamiltonian:

$$\mathcal{H}_t = K_t + V_t = \ln \left(\frac{1}{2} \|v_t\|^2 \right) + \ln(\text{PPL}_t)$$

Exponentiating this quantity yields the total energy:

$$E_t = e^{\mathcal{H}_t} = \frac{1}{2} \|v_t\|^2 \cdot \text{PPL}_t$$

This product reflects the relationship between hidden-state velocity and token-level uncertainty at each timestep. While the quantity has no canonical unit, we hypothesize that it reflects some form of uncertainty-adjusted movement, and may eventually admit an information-theoretic interpretation (e.g., in nats).

Remark D.1 (Why not classical energy?). *A classical Hamiltonian*

$$H = \frac{1}{2}\|v_t\|^2 - \ln p_{x_t}(h_t)$$

fails because:

- The kinetic term ($\sim 10^3$ – 10^4) numerically overwhelms the potential term (~ 1 – 10)
- No consistent conservation emerges (empirical CV > 50%)
- The geometry of velocity and token probability live on fundamentally mismatched scales

By contrast, the logarithmic formulation collapses this disparity, restoring a scale-matched, empirically conserved structure.

D.2 Theoretical Conservation Analysis

We now show that E_t is conserved to first order under the log-Euler-Lagrange equations.

Theorem D.1 (First-order energy conservation). *If the hidden state dynamics exactly follow the log-Euler-Lagrange equations derived in Section C, then the log-energy $\mathcal{H}_t = \ln(E_t)$ is conserved to first order under any infinitesimal perturbation of h_t .*

Proof. Consider a perturbation $h_t \rightarrow h_t + \epsilon\eta$ for small ϵ and unit direction vector η . This changes both:

- $v_t = h_t - h_{t-1} \Rightarrow v_t \rightarrow v_t + \epsilon\eta$
- $v_{t+1} = h_{t+1} - h_t \Rightarrow v_{t+1} \rightarrow v_{t+1} - \epsilon\eta$

Note: $p_{x_t}(h_t)$ depends only on h_t , and $p_{x_{t+1}}(h_{t+1})$ does not change under this perturbation.

Thus, the change in log-energy is:

$$\delta\mathcal{H} = \delta \left[\ln \left(\frac{1}{2}\|v_t\|^2 \right) \right] + \delta[\ln(\text{PPL}_t)] + \delta \left[\ln \left(\frac{1}{2}\|v_{t+1}\|^2 \right) \right]$$

Computing each term to first order in ϵ :

$$\delta \left[\ln \left(\frac{1}{2}\|v_t\|^2 \right) \right] = \frac{2\epsilon(v_t \cdot \eta)}{\|v_t\|^2} \tag{5}$$

$$\delta[\ln(\text{PPL}_t)] = -\epsilon(\nabla_h \ln p_{x_t}(h_t) \cdot \eta) \tag{6}$$

$$\delta \left[\ln \left(\frac{1}{2}\|v_{t+1}\|^2 \right) \right] = -\frac{2\epsilon(v_{t+1} \cdot \eta)}{\|v_{t+1}\|^2} \tag{7}$$

Therefore:

$$\delta\mathcal{H} = \epsilon\eta \cdot \left[\frac{2v_t}{\|v_t\|^2} - \nabla_h \ln p_{x_t}(h_t) - \frac{2v_{t+1}}{\|v_{t+1}\|^2} \right]$$

By the log-Euler-Lagrange equation, this bracketed term vanishes identically, implying $\delta\mathcal{H} = 0$ for any η . $\square$

Key interpretation. This result shows that small perturbations in h_t affect energy at both t and $t + 1$, but their effects exactly cancel under ideal dynamics. This is a purely local conservation flow—no gradients propagate further forward—reflecting the autoregressive causality of transformer inference.

Corollary D.1 (Approximate conservation in practice). *Since transformers only approximately follow the log-Euler-Lagrange equations, we expect:*

$$CV(E) = O(\epsilon_{discrete} + \epsilon_{model})$$

where $\epsilon_{discrete}$ comes from finite-difference resolution and ϵ_{model} from architectural deviation.

Empirical validation. Observed global $CV(E) \approx 0.08$ across diverse models supports the validity of this framework. Even with frozen normalization layers and no hidden state normalization, empirical conservation remains tight—especially in random-weight models.

D.3 Empirical Validation: Energy Conservation Results

D.3.1 Experimental Setup

- **Models:** 10 architectures (Qwen, Llama, SmolLM, DeepSeek, Gemma), each with pre-trained and random-weight variants
- **Trajectories:** 100 Wikipedia contexts per model, generating 100-token continuations
- **Metrics:**
 - **Average trajectory CV:** For each trajectory, compute $CV(E) = \sigma_E / \mu_E$, then average over all token locations in the trajectory. This means that we take an average of averages in this score. This average of average accounts for different trajectories having high and low energy components that can ‘move’ within the trajectory without impacting this score.
 - **Global CV:** Pool all energy values for each token-token step within each trajectory and compute global CV across the dataset for every token-token step

All models were evaluated on raw hidden states, with layernorms left frozen where applicable.

D.3.2 Results

Key findings:

1. **Energy conservation holds:** 17/20 checkpoints exhibit global $CV(E) < 13\%$
2. **Stronger conservation in random models:** Median global CV of 2.2% vs. 9.4% in pre-trained counterparts
3. **No parameter fitting required:** The expression $E = \frac{1}{2} \|v_t\|^2 \cdot \text{PPL}_t$ works directly—unlike prior power-law models
4. **Energy scale separation:** Trained models cluster near 8–10 energy; untrained models at 12–16

Table 1: Energy conservation across 20 checkpoints with $E_t = K_t + V_t$

Model	Weights	Global CV	Avg. Traj. CV	Avg. Energy
<i>Pre-trained Models</i>				
Qwen-2.5-1.5B	pre-trained	0.081	0.017	9.66
SmolLM2-135M	pre-trained	0.117	0.022	7.36
SmolLM2-360M	pre-trained	0.129	0.025	7.23
SmolLM2-1.7B	pre-trained	0.103	0.024	8.00
SmolLM3-3B	pre-trained	0.092	0.018	9.71
Qwen-2.5-3B	pre-trained	0.073	0.015	9.74
Llama-3.2-1B	pre-trained	0.085	0.020	9.16
Llama-3.2-3B	pre-trained	0.084	0.022	8.77
DeepSeek-R1-1.5B	pre-trained	0.110	0.048	9.00
Gemma-3-1B	pre-trained	0.208	0.126	8.02
<i>Random-weight Models</i>				
Qwen-2.5-1.5B	random	0.022	0.007	15.50
SmolLM2-135M	random	0.031	0.010	12.91
SmolLM2-360M	random	0.221	0.047	9.42
SmolLM3-3B	random	0.022	0.004	14.56
Qwen-2.5-3B	random	0.017	0.003	15.66
Llama-3.2-1B	random	0.020	0.004	15.32
Llama-3.2-3B	random	0.022	0.002	15.21
DeepSeek-R1-1.5B	random	0.018	0.005	15.53
Gemma-3-1B	random	0.027	0.004	15.27
Median (all)		0.077	0.016	
Median (pre-trained)		0.094	0.023	
Median (random)		0.022	0.005	

Table 2: Mean energy decomposition: K/V ratio

Model Type	Mean K/V Ratio	Interpretation
Pre-trained	11.8	Kinetic-dominated
Random	0.89	Balanced

D.3.3 Analysis of Energy Components

This reveals two regimes: - **Pre-trained models:** High velocity, low point perplexity $\Rightarrow$ decisive transitions - **Random models:** Balanced dynamics $\Rightarrow$ exploratory, stable flow

D.4 Interpretation: Two Distinct Dynamical Regimes

Pre-trained models. With $K/V \approx 12$, pre-trained models exhibit:

- Rapid transitions through hidden space
- High confidence in next-token prediction
- Looser energy conservation (CV $\sim 9.4\%$)

We interpret this as "semantic highways": directional motion through attractor basins at the cost of conservation. Training thus partially breaks symmetry—lowering the model’s energy and increasing

relative perturbation amplitudes. This is akin to a cold star having larger fractional perturbations than a hot one.

Random models. With $K/V \approx 0.9$, random models exhibit:

- Balanced velocities and point perplexities
- Tight energy conservation ($CV \sim 2.2\%$)

This suggests that the architecture itself induces conservation. Training distorts this balance to prioritize confident, directional predictions.

Conservation as architectural prior. These results support the idea that energy conservation is a prior encoded by transformer geometry. Training extracts information (or energy) from the model and introduces structured asymmetries. Whether increased conservation is helpful remains an open question. Current trends show that conservation *decreases* with increased capability.

Applications.

- **Anomaly detection:** Deviations from expected energy levels may flag off-distribution inputs
- **Control:** Steering that respects conservation may improve robustness
- **Architecture design:** Conservation may serve as a principle for model regularization or interpretability

E Controlling Generation via Minimal-Action Steering

Having established that transformer dynamics follow a log-Lagrangian principle, we now show how to control generation while respecting this natural structure. Jacobian steering emerges as the optimal intervention strategy from the principle of least action.

E.1 The Minimal Action Principle for Control

Given a language model with vocabulary $\mathcal{V} = \{1, \dots, V\}$ and a target token $t^* \in \mathcal{V}$, we seek to modify the hidden state $h \in \mathbb{R}^d$ to increase $p_{t^*}(h)$ while minimizing the action increment.

Theorem E.1 (Jacobian steering minimizes action). *Among all perturbations δh that achieve a fixed increase in $\log p_{t^*}(h)$, the direction*

$$g(h) = \nabla_h \log p_{t^*}(h) = W_{t^*} - \sum_{j=1}^V p_j(h) W_j$$

minimizes the action change $\Delta S = \ln(\frac{1}{2}\|\delta h\|^2) - \Delta \ln p_{t^}$.*

An analogous Hamiltonian-informed decoding approach appears in Du et al. [10].

Proof. For a fixed log-probability increase $\nabla_h \log p_{t^*}(h) \cdot \delta h = c$, we minimize:

$$\min_{\delta h} \ln \|\delta h\|^2 \quad \text{subject to} \quad g(h) \cdot \delta h = c$$

By Cauchy-Schwarz, the minimum occurs when $\delta h \parallel g(h)$, yielding:

$$\delta h_{\text{optimal}} = \frac{c}{\|g(h)\|^2} g(h)$$

□

This result shows that Jacobian steering is not an arbitrary choice but the unique intervention that respects the variational structure of transformer dynamics.

E.2 The Jacobian Steering Algorithm

Algorithm 1 Minimal-Action Jacobian Steering

Require: Current state h , target token t^* , threshold $\eta \in (0, 1)$, max steps K

Ensure: Steered state $\hat{h}$ with $p_{t^*}(\hat{h}) \geq \eta$

```

1:  $h_0 \leftarrow h$ 
2: for  $k = 0$  to  $K - 1$  do
3:   Compute  $p_{t^*}(h_k)$  via softmax
4:   if  $p_{t^*}(h_k) \geq \eta$  then
5:     return  $\hat{h} \leftarrow h_k$ 
6:   end if
7:    $g_k \leftarrow W_{t^*} - \sum_j p_j(h_k) W_j$  ▷ Minimal-action direction
8:    $\hat{g}_k \leftarrow g_k / \|g_k\|$  ▷ Normalize
9:   Choose  $\alpha_k$  via line search to ensure  $p_{t^*}(h_k + \alpha_k \hat{g}_k) > p_{t^*}(h_k)$ 
10:   $h_{k+1} \leftarrow h_k + \alpha_k \hat{g}_k$ 
11: end for
12: return  $\hat{h} \leftarrow h_K$ 

```

E.3 Preservation of Semantic Coherence

A critical question is whether Jacobian steering merely maximizes the target token probability or preserves the broader semantic structure of the output distribution. We hypothesize that because steering follows the minimal-action principle, it maintains the quality of the entire probability distribution.

E.3.1 Experimental Design

To test this hypothesis, we conducted the following experiment:

1. **Data:** 50 Wikipedia articles beginnings as prompts (we use only the first 50 tokens for each article)
2. **Models:** SmolLM2-135M-Instruct and Qwen2.5-1.5B-Instruct
3. **Protocol:**

- Generate natural continuations from each prompt
- Apply Jacobian steering to match ground-truth Wikipedia tokens
- Extract top-10 token distributions for both natural and steered states
- Use GPT-4.1 to rate the semantic quality of each top-10 list (1-10 scale)

The key insight is that if steering preserves semantic coherence, the top-10 tokens from steered distributions should be rated as high-quality as (or better than) natural distributions.

E.3.2 Results: Steering Improves Distribution Quality




Figure 1: Completion quality scores for natural vs. steered distributions for **Qwen2.5-1.5B-Instruct**. Steered distributions consistently receive higher quality ratings.

Table 3: Statistical comparison of natural vs. steered distribution quality (N = 50)

Model	Natural Mean	Steered Mean	p-value	Effect Size
SmolLM2-135M	4.84 ± 2.92	7.61 ± 2.69	< 0.001	0.99 (large)
Qwen2.5-1.5B	7.33 ± 2.65	8.22 ± 1.90	0.028	0.39 (small)

Remarkably, steered distributions receive *higher* quality ratings than natural ones:

- **SmolLM2-135M-Instruct**: Steering improves mean quality from 4.84 to 7.61 (Cohen’s $d = 0.99$, large effect)
- **Qwen2.5-1.5B-Instruct**: Steering improves mean quality from 7.33 to 8.22 (Cohen’s $d = 0.39$, small effect)
- Both improvements are statistically significant ($p < 0.05$)

E.3.3 Interpretation

These results suggest that Jacobian steering does more than just boost the target token—it actually *improves* the semantic coherence of the entire output distribution. This can be understood through our theoretical framework:

Why steering improves quality.

1. **Minimal action = minimal disruption:** By following the path of least action, steering makes the smallest possible change to reach the target
2. **Natural force alignment:** The steering direction $g(h) = -\nabla V_{t^*}$ is the force that would naturally arise if t^* were the next token
3. **Semantic consistency:** Ground-truth Wikipedia tokens often represent more informed choices than model samples, so steering toward them improves coherence

Model-specific effects. - Smaller models (SmolLM2) benefit more from steering, suggesting they have less confident natural distributions - Larger models (Qwen2.5) show smaller improvements, as their natural distributions are already high-quality

E.4 Implications for Controlled Generation

These findings establish Jacobian steering as a principled method for controlling language models:

- **Theoretically grounded:** Emerges from the principle of least action
- **Semantically preserving:** Maintains or improves distribution quality
- **Computationally efficient:** Requires only forward passes, no backpropagation
- **Interpretable:** Moves toward target token embedding in hidden space

The success of minimal-action steering validates our log-Lagrangian framework: interventions that respect the natural variational structure produce better outcomes than arbitrary perturbations. This opens new avenues for principled, physics-inspired control of neural language models.