Hey all,

Here is your disruptive ML/AI digest for the day. Here's what you folks might find interesting today.

TOP PICKS

Novel Architectures and Objectives

[Quantum Port-Hamiltonian Neural Networks: Learning Conservative and Dissipative Dynamics via Measurement-Induced Nonlinearity]

This paper introduces Q-pHNNs, parameterised quantum circuits that learn classical dynamics while preserving Hamiltonian structure by construction. The key conceptual move is the Isomorphic Hamiltonian Mapping: the skew-symmetric interconnection matrix J maps to unitary gate evolution, while the positive-semidefinite dissipation matrix R maps to Measurement-Induced NonLinearity (MINL) via mid-circuit measurement and classical feedforward. This means conservation and passivity are enforced architecturally rather than through penalty terms, which is a genuine departure from standard neural ODE approaches. The authors demonstrate four instantiations including a topology-entangled quantum GNN for coupled-phasor networks, and show 100% energy monotonicity for the MINL circuit on damped oscillator experiments. The framework sits squarely at the intersection of **quantum ML** and **geometric/physical architectures**, and the idea that dissipation can be realised through measurement-induced nonlinearity is a fresh primitive worth watching.

[Learning Forced Multibody Dynamics on Lie Groups]

The authors propose learning dynamics directly on manifold-valued configuration spaces by formulating the architecture around discrete forced Euler-Lagrange equations on Lie groups, using only position data. This is notable because most learned dynamics models operate in Euclidean space and require velocity measurements, which are often noisy or unavailable. By respecting the geometric structure of mechanical systems, the method preserves geometric invariants and conservation laws as a natural consequence of the formulation rather than through auxiliary losses. The extension to multibody systems and external control inputs broadens applicability considerably. This is a clean example of how **equivariant/geometric architectures** can be derived from variational mechanics rather than designed empirically, and the position-only requirement is practically important for robotics and physical simulation.

[What Does Goodness Measure? A Likelihood-Ratio Account of Forward-Forward Learning]

Hinton's Forward-Forward algorithm treats its scalar goodness function and inter-layer normalization as heuristics. This paper shows they are not heuristic at all: under an explicit generative model, squared goodness is the sufficient statistic of a likelihood-ratio test between two zero-mean populations differing in scale, and the FF threshold is the decision boundary. The generalization is elegant -- anisotropic populations yield a Mahalanobis goodness, heavy-tailed populations yield a saturating statistic whose slope is posterior precision (i.e., divisive normalization). The paper also explains why inter-layer normalization must remove length while preserving per-coordinate energy, and identifies a scale-inflation shortcut in the pairwise objective that a whitened goodness eliminates. This is a **novel objective** grounded in statistical decision theory, and it gives FF a principled foundation that could guide future variants beyond ad hoc design choices.

Theory and Generalization

[Gradient Flow Dynamics and Implicit Bias of Diagonal Linear Networks under Infinitesimal Initialization]

This work extends the implicit bias analysis of diagonal linear networks to both deep architectures and a broader class of two-layer variants under infinitesimal initialization. The authors prove that training trajectories converge to a modified l1 norm minimizer, generalizing prior results from Pesme and Flammarion. The identification of the Structural Invariant Manifold (SIM) as the geometric object shaping learning dynamics is the key theoretical contribution -- it provides a mechanistic explanation for why the implicit bias takes this specific form. This matters for **theory and generalization** because it sharpens our understanding of how depth and architecture interact with initialization scale to determine which solution gradient descent selects, and the SIM framework may extend to other architectures beyond diagonal networks.

Probabilistic and Statistical Methods

[LatentFlow: A General Framework for Conditioning Stochastic Processes]

LatentFlow reduces the hard problem of conditioning stochastic processes to latent-space inference by writing any process as the deterministic image of a tractable innovation, then pulling the likelihood back through that map and sampling with a guided probability flow. The construction is provably exact at the level of the target law, with approximation entering only through explicitly reducible sources (finite terminal noising, Monte Carlo guidance, time discretisation). Crucially, it requires no training -- conditioning reduces to solving a single reverse-time SDE, enabling conditional sampling in seconds on a CPU across model classes that have never shared a scalable method: stochastic PDEs, heavy-tailed processes, point processes, and simulator-defined models. This is a significant contribution to **simulation-based inference** and **probabilistic methods** because it unifies conditioning across process types that previously required bespoke constructions, and the training-free property makes it immediately deployable.

BY CATEGORY

cs.LG

- Accelerated Mixing Time of Randomized Hamiltonian Monte Carlo
- The Computational Basis of Confidence in Large Language Models
- Ensemble Controlled-Flow Filtering for Implicit Data Assimilation
- The Geometry of Memorization: Finite-Time Spectral Sensitivity as a Diagnostic for Flow Matching Models
- AVQ-Attention: Adaptive Vector-Quantized Attention
- Learning the Graphical Nature of Symmetries
- SinAE: A Single-Architecture Flow-Matching Autoencoder for Cross-Domain Atomic Systems
- What Makes a Representational Prior Work? Feature Families, Label-Free Invariances, and Critical Windows in Grokking
- Fisher Rank Inflation: A Spectral Signature of Memorization under Label Noise
- A Shortcut to Statistically Steady-State Turbulence with Flow Matching

stat.ML

- Contrast-Free ICA and Causal Inference via Wasserstein Distances to the Gaussian
- Causal Graphs, Markov Properties and Do-calculus for Stochastic Differential Equations
- Cluster-Weighted EDMD
- Sharp Optimal Algorithm for Derivative-Free Stochastic Convex Optimization in One Dimension
- PolarBM: Complex-valued Boltzmann Machine for Modeling Audio Signals in Polar and Log-polar Coordinates
- Falsifying Causal Graphs With Outlier Events
- MixCIT: A Kernel Based Local-Polynomial Debiased Test for Conditional Independence on Mixed-Type Data
- Thompson Sampling Is 2-Competitive for Mistakes
- Dynamic Online Processor-Native Inference for State Estimation

Until the next announcement!

P.S. In the most recent submission window: cs.LG: 100 papers stat.ML: 18 papers