Hey there,

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 / Probabilistic and Statistical Methods

Nuclear Quantum Effects as a Denoising Problem (arXiv:2607.19680)

This paper establishes a striking equivalence: a denoiser trained purely on classical Boltzmann statistics, when composed at sampling time with an analytic Gaussian component encoding quantum context (mass, temperature, dissipation, boundary conditions), exactly reproduces the quantum Boltzmann distribution of nuclei. The key insight is that the noise structure of generative modeling and quantum fluctuations share the same quadratic action, so transfer across quantum contexts is exact and requires no retraining as long as the training noise does not exceed the intrinsic quantum uncertainty. This is a genuine conceptual departure -- it reframes quantum simulation as a composition problem rather than a training problem, and the invariance across isotopic mass, temperature, and path boundary conditions (including bosonic exchange) is proven, not just empirically suggested. The work bridges **transferable methods from physics** and **probabilistic methods** in a way that could reshape how generative models are applied to quantum mechanical systems.

Theory and Generalization / Emergent Behavior and Agentic AI

Active Inference as a Convex Markov Decision Process (arXiv:2607.20152)

The authors recast Active Inference's expected free energy (EFE) minimization as policy optimization in a convex MDP, showing that the pragmatic component is linear in predictive state marginals (equivalent to standard reward maximization) while the epistemic value introduces a nonlinear, policy-dependent -- i.e., performative -- reward. This decomposition is important because it cleanly separates what is genuinely novel about AIF from what reduces to classical RL, and it yields a mirror descent algorithm with local linearization compatible with actor-critic methods. The framing of epistemic drive as performative reward, coupled with convergence analysis and policy improvement guarantees, provides a long-sought theoretical grounding for AIF within modern optimization theory. The connection to performative prediction is non-trivial and opens a route for principled analysis of world-model learning coupled with policy optimization.

Probabilistic and Statistical Methods / Novel Architectures and Objectives

Boltzmann-Expected Molecular Design with Decoupled Annealing Flows (arXiv:2607.19519)

DECAF reframes 3D molecular design from single-conformer optimization to **Boltzmann-expected design**, where the target is an ensemble statistic rather than a property of one structure. The architecture factorizes the joint distribution into a graph-conditioned flow acting as a Boltzmann emulator and a coordinate-conditioned flow proposing graphs from 3D information, alternated with a simulated-annealing acceptance rule. The most novel contribution is higher-moment design -- jointly optimizing variance and skewness of ensemble properties -- which no existing 3D generative model does, and the authors verify these designs with all-atom MD simulations. The no-retraining-to-change-objectives property is a practical advantage, but the conceptual contribution is the shift from point estimation to ensemble-aware optimization in generative molecular design.

Probabilistic and Statistical Methods / Theory and Generalization

Analytic Distribution of Classifier-Free Guidance for Schedule Design (arXiv:2607.19725)

This work derives exact analytic path-integral representations of the distributions induced by classifier-free guidance (CFG) through the probability flow ODE, exposing that the standard product-distribution heuristic $p_0^\omega q_0^{1-\omega}$ is wrong and that CFG introduces an exponential path-integral correction modulated by the schedule $\omega(t)-1$. The characterization of how score discrepancies accumulate along sampling trajectories is new and provides a principled basis for schedule design. The proposed DG-CFG schedule balances timestep contributions while accounting for low-noise score-error amplification, yielding concrete improvements on Stable Diffusion 1.5. This matters because CFG is the default conditional generation mechanism in diffusion models, yet its theoretical properties have been poorly understood -- this is a rare case of exact analysis replacing heuristics in a widely deployed technique.

Novel Architectures and Objectives / Theory and Generalization

Koopman Dreamer: Spectrally Constrained Latent Dynamics for Stable World-Model Imagination (arXiv:2607.19719)

Koopman Dreamer replaces the usual neural transition function in latent world models with a spectrally constrained deterministic core built from 2D rotation-scaling blocks with bounded radii, directly representing damping, rotation, and near-periodic modes. The architectural choice is grounded in Koopman operator theory, and the authors derive a multi-step rollout-error bound that cleanly separates spectral amplification from stochastic-state mismatch, providing a mechanistic explanation for long-horizon instability in Dreamer-style models. The combination of posterior-conditioned EMA teacher targets with one-step consistency and open-loop prediction objectives addresses the train-imagination gap in a principled way. This is not just another world-model variant -- the spectral backbone gives direct control over modal persistence, which is the root cause of error accumulation in long rollouts.

BY CATEGORY

cs.LG

stat.ML

- Asymptotically Optimal Regret for Reinforcement Learning without Horizon Dependence
- RELTA-SGLD: Relative-Growth Localized Taming for Nonconvex Stochastic-Gradient Langevin Learning
- Strong Gravitational Lensing Posterior Sampling in Pixel-Space Using Diffusion Models and Recurrent Inference Machines
- Total Variation Distance Estimation in Autoregressive Models
- Adaptive Bayesian Online Learning via Expert Aggregation
- Optimal Recalibration of an Online Predictor
- Statistical Inference for Rank Allocation in Low-Rank Adaptation
- Deep Shape Regression for Planar Curves with Multimodal Covariates
- Data-Poisoning Audits for Causal Effect Estimation
- Adaptive deep nonparametric regression from dependent data under covariate shift
- Directional Kernel Mean Difference: A Fast Signed Statistic for Univariate Distribution Comparison
- Quantum Kernels and the Cross-Section of Stock Returns: Anatomy of a Vanishing Advantage
- A Bayesian Framework for Built-in Input Dimension Reduction for Gaussian Process Modeling

---

Back tomorrow with more!

P.S. In the most recent submission window: cs.LG: 115 papers stat.ML: 17 papers