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 / Transferable Methods from Adjacent Fields

Stochastic Counterdiabatic Driving via Biorthogonal Liouvillian Eigenmodes by Sandeep Suresh Cranganore, Sebastian Lehner, Johannes Brandstetter, and Max Welling. This paper constructs a numerical framework for eliminating non-adiabatic lag in stochastic thermodynamic systems using gauge-type transforms rather than generalized coordinate transforms. The key innovation is a biorthogonal spectral decomposition of the time-dependent Fokker-Planck generator that yields an exact counterdiabatic correction, directly analogous to Berry's transitionless driving in quantum systems. The results are striking: non-adiabatic lag is suppressed by roughly twelve orders of magnitude in total variation distance, with vanishing dissipated work across all protocol speeds. This matters because it provides a principled, physics-grounded alternative to learned diffeomorphisms for escorted free energy simulations, and the quantum-to-stochastic analogy could inspire new objective functions in generative modeling and sampling. Max Welling is a co-author.

Theory and Generalization

Beyond ICA: Identifiability by Symmetry Breaking by Pengzhou Wu. This paper proves identifiability of deep generative models with piecewise-affine decoders and Gaussian mixture priors in a purely unsupervised setting, replacing the usual continuity assumptions with three algebraic contrast principles: domain contrast, mechanism contrast, and interaction contrast. The approach is the first to make algebraic symmetry-breaking the engine of nonlinear identifiability, the first to admit discontinuous decoders, and the first to handle fully non-injective decoders where every observation admits multiple latent codes. The hierarchy from law identifiability through map identifiability to posterior and pointwise identifiability is clean and constructive. This matters for causal representation learning because it decouples structural identification from pointwise inversion, and the algebraic conditions are on the data-generating process rather than on learning methods, making them fundamentally different from existing identifiability results.

Probabilistic and Statistical Methods / Novel Architectures and Objectives

All in One: Generative Modeling as Mean-Field Game Design by Kun Zhao and Xu Chen. This work frames twelve prominent continuous-time generative models -- including CNFs, OT-Flow, score-based models, and Schrodinger Bridges -- as special cases of a single mean-field game variational problem specified by four composable cost functions. The authors exploit two previously unexplored dimensions: the interaction term (set to zero in most existing models) and the application of MFG solvers to generative modeling. They also propose DI-Flow, which uses a differentiable entropy functional to encourage mode coverage, and demonstrate that learning-based MFG solvers substantially outperform standard neural training on stochastic dynamics. The unification is not merely taxonomic -- the shared API is shown to be lossless relative to hand-coded implementations, and the cost-tuple parameterization opens a systematic design space for new generative objectives derived from game-theoretic rather than purely probabilistic principles.

Theory and Generalization

Grokking on the Weight-Decay Clock: A Rate Hierarchy from Softly Broken Symmetries by Taeyoung Kim. This paper identifies an exactly solvable late-time relaxation mechanism for grokking in linear models trained with full-batch heavy-ball optimization and weight decay, with a locally quadratic extension to nonlinear networks. The core insight is a "grokking subspace" -- a population-active component of the empirical null space where training predictions remain unchanged, leaving weight decay as the sole restoring force and producing slow dissipative relaxation governed by an exact discrete-time and continuous-time law. The theory yields explicit iteration-scale predictions for grokking time, recovers the familiar $(1-\beta)/(\eta\lambda)$ scaling, and distinguishes coupled L_2 regularization from decoupled weight decay. All theoretical identities are verified without fitted parameters in a synthetic model, and the predicted scaling holds for genuine delayed generalization in modular addition. This matters because it provides a mechanistic, parameter-free account of a phenomenon that has resisted principled explanation, and the null-space decomposition could generalize to other implicit regularization phenomena.

Back tomorrow with more!

P.S. In the most recent submission window: cs.LG: 250 papers stat.ML: 41 papers