Hey there,

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

TOP PICKS

Theory and Generalization

Neural Quadratic Forms: A Unified Minimal Model for Sudden Learning and Scaling Laws -- This paper derives a universal leading-order expansion for neural network training dynamics from a symmetry argument: because network units are interchangeable, smoothness plus vanishing-gradient-at-origin forces a quadratic form $\Tr[WW^\top A(x)]$ where all architectural detail collapses into a single structure matrix $A(x)$. Perceptrons, attention layers, mixtures of experts, and convolutions all reduce to this one model at different A, and the training dynamics close on the order parameter $M = WW^\top$, yielding Lotka-Volterra equations whose modes switch on sequentially. The theory explains both sudden learning (plateau-then-drop) and smooth power-law scaling as two regimes of the same flow, with the power-law exponent predicted analytically. Co-authored by Tomaso Poggio and Isaac Chuang (MIT), this is a genuine conceptual unification rather than an incremental empirical study.

Quantum ML and Hybrid Paradigms

Exponential Quantum Advantage for Learning Signals with a Single Qubit -- This work shows that coupling a single controllable qubit to a classical sensor can exponentially reduce the number of measurements needed to learn Fourier coefficients, temporal correlations, and observable transformations, with an experimental demonstration achieving 10^7-fold measurement reductions on a superconducting cavity-qubit architecture. The theoretical backbone is Quantum Phase-Space Inference (QΨ), a framework that converts experimental objectives and constraints into tight lower bounds and optimal learning algorithms while producing a certificate of quantum advantage. QΨ extends beyond the regimes captured by quantum Fisher information, which is notable because it provides a systematic method for identifying rigorous quantum advantages in practical tasks rather than abstract oracle models. The authors demonstrate applicability to dark matter detection and wireless communication simulations, grounding the claims in realistic settings.

Probabilistic and Statistical Methods

The Data Geometry of Masking Diffusion: Certified-Optimal Schedules via Unmasking Growth Complexity -- Martin Wainwright introduces a path-resolved measure called unmasking growth complexity (UGC) whose local increments directly control KL discretization error in masking diffusion for discrete sampling. The key result is a set of certified-optimal samplers that achieve a prescribed KL error with high probability and iteration complexity within a constant factor of the oracle procedure, with UGC increments estimable from samples via coupled reveal trajectories. In the fine-partition limit, the squared integral of the square-root UGC density determines the sharp leading-order optimal Euler discretization error, and examples show $\tilde\Omega(\sqrt{d})$ improvements over coarse schedules with a constant number of adaptively placed blocks. This is a rigorous information-geometric treatment of discrete diffusion scheduling that moves well beyond heuristic schedule design.

Novel Architectures and Objectives / Transferable Methods from Physics

Equivariant Learning of a Transferable Three-Dimensional Classical Density Functional -- Bingqing Cheng demonstrates that the excess free-energy functional of classical density functional theory can be learned directly from 3D equilibrium density fields while preserving spatial symmetry and variational consistency, without any free-energy or chemical-potential labels. A single learned functional transfers across temperatures, system sizes, and statistical ensembles, recovering structure factors, equations of state, liquid-vapor coexistence, and interfacial broadening -- none of which are training targets. The method predicts non-trivial phenomena in complex 3D geometries such as solvent-depleted bridge forces between colloids and adsorption in gyroid pores. This is a principled equivariant architecture grounded in thermodynamic variational structure rather than empirical loss engineering, and it represents a genuine transfer of physics methodology into a reusable ML primitive.

Until the next announcement!

P.S. In the most recent submission window: cs.LG: 133 papers stat.ML: 26 papers