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 / Transferable Methods from Physics

1. Deep neural networks as lattice gauge theories (Jefferson & Ramakrishnan)

This paper extends the NN/QFT duality into a genuine (0+1)-dimensional lattice gauge theory where network layers act as lattice sites and weight matrices as gauge fields on links. The authors compute the tree-level neuron-neuron propagator describing layer variance evolution and build out Feynman diagram machinery for perturbative corrections in 1/N, including infinitely-many loop diagrams mediating inter-layer interactions. This is a real conceptual departure: it provides a field-theoretic framework for studying higher-point correlations and information propagation in deep networks, rather than the usual empirical or mean-field analyses. The connection between permutation symmetry of layers and gauge structure is the kind of structural insight that could open new analytical tools for understanding training dynamics and representation flow.

1. Learning Deterministic and Stochastic Forced Hamiltonian Systems (Brantner & Tyranowski)

The authors develop a geometric framework for learning forced Hamiltonian systems using structure-preserving neural networks derived from the Lagrange-d'Alembert principle. The key novelty is the introduction of a Lagrange-d'Alembert map with a proven C^r convergence theorem, leading to Generalized Forced Hamiltonian Neural Networks (GFHNNs) that concatenate Lagrange-d'Alembert-Euler maps. They extend this to stochastic forced Hamiltonian systems by interpreting Stratonovich integrals as parameters, and prove a universal approximation theorem for the architecture. The reported long-time stability improvements over non-geometric ResNets, with substantially less training data, matter because they demonstrate that the right inductive bias from symplectic geometry is not just theoretically elegant but practically consequential.

Theory and Generalization

1. Exact Algebraic Computation of Learning Coefficients for Two-Dimensional Singular Models (Sergeant-Perthuis, Tsigaridas, Tsukahara)

This paper presents the first deterministic algorithm for exactly computing local Real Log Canonical Thresholds (RLCTs) for any two-dimensional model whose KL divergence is contact-equivalent to a polynomial. This is significant because the WBIC and singular learning theory more broadly have relied on sampling-based estimators for these learning coefficients, which are the correct asymptotic objects for model selection in singular models like deep neural networks. The authors derive complexity bounds and demonstrate effectiveness on polynomial neural networks. Beyond providing ground truth for calibrating MCMC-based estimators, exact computation reveals algebraic structure in learning coefficients that sampling methods cannot detect, which could deepen our understanding of how model complexity manifests in singular statistical manifolds.

Quantum ML and Hybrid Paradigms

1. An Irreducible Quantum Advantage in Aligning World Models with Reality (Lumbreras, Ma, Thompson, Gu)

The authors construct classical worlds for which every finite classical world model provably fails along a specific trajectory: it either loses the ability to distinguish actions or repeatedly assigns highest expected reward to suboptimal actions, with nonvanishing average error. Crucially, each such world admits a quantum world model using a single qutrit that reproduces it exactly, maintaining perfect policy alignment. This is not a quantum speedup claim but a provable representational separation: classical memory cannot suffice even when the true world is itself classical. The result is a clean, constructive demonstration that quantum structure in world models provides an irreducible advantage for agent alignment, which is directly relevant to the foundations of model-based reinforcement learning and simulation.

Until the next announcement!

P.S. In the most recent submission window: cs.LG: 111 papers stat.ML: 16 papers