Hi all,

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

TOP PICKS

Probabilistic and Statistical Methods

The Advective Fisher-Rao Geometry of Deterministic Measure Transport introduces a novel **advective Fisher-Rao metric** for optimization on paths of probability measures governed by the continuity equation. The key conceptual contribution is showing that this metric arises naturally from three distinct perspectives: as the zero-noise limit of Fisher-Rao on path measures, as the expected second variation of the Freidlin-Wentzell large deviation rate functional, and as the Hessian of the Benamou-Brenier action functional from dynamic optimal transport. This unification of information geometry, large deviations, and optimal transport into a single geometric object for measure-valued optimization is a genuine departure from standard gradient-based approaches. The empirical finding that this metric optimally fits probability densities while Gauss-Newton optimally fits velocity fields is a crisp, actionable distinction that could inform future work on normalizing flows and generative model training objectives.

Quantum ML and Hybrid Paradigms

Unifying Physical Backpropagation develops a unifying theory for **on-device gradient computation** in physical computing systems, based on the adjoint method. The paper identifies sufficient conditions under which a physical system can compute the gradient of its own performance, recovering Equilibrium Propagation, Hamiltonian echo backpropagation, and fully forward mode training as special cases. The distinction between linear systems (where finite-amplitude experiments suffice) and nonlinear systems (requiring infinitesimal nudging and time-reversal mirrors) is a clean theoretical result. The generalization beyond reciprocity to a more general intertwining condition, extending to non-Hermitian and non-reciprocal systems, is particularly notable as it broadens the class of physical substrates capable of exact gradient computation. This provides a principled template for constructing physical learning algorithms rather than relying on system-specific derivations.

Theory and Generalization

PAC-Bayes Beyond Parameter Space: Behavioral Equivalence, Z-Information, and Exact Complexity Decomposition addresses a fundamental gap in **PAC-Bayes theory**: classical KL divergence between posterior and prior over parameters conflates uncertainty over predictive behavior with variation among behaviorally equivalent parameter configurations. The authors use measure disintegration to decompose the classical PAC-Bayes complexity into a behavior-selection term and a realization-level term, defining the latter's negative as "Z-information." The behavior-selection term admits an exact variational characterization as the minimum KL divergence among all posteriors inducing the same predictive behavior distribution. This is a rigorous structural result that reframes generalization bounds around predictive behavior rather than parameter space, which matters increasingly as over-parameterized models admit vast fiber symmetries. The connection to fiber geometry and behavior-preserving perturbations could inform future work on flat minima and implicit regularization.

Novel Architectures and Objectives

Reducing Symmetry Increase in Equivariant Neural Networks provides the first rigorous characterization of a subtle expressivity failure in **equivariant neural networks**: when processing symmetric inputs, equivariant maps can increase symmetry beyond the input's, making outputs invariant to transformations the input was not invariant to. The authors prove that this increased symmetry admits an infimum determined by the feature space structure, provide a computable algorithm to derive it, and propose practical feature design guidelines to prevent harmful symmetry increases. This is a principled theoretical contribution to geometric deep learning that moves beyond documenting the phenomenon to providing actionable reduction strategies. The experiments on QM9 validate the theoretical predictions, suggesting these guidelines could improve ENN design for molecular and materials applications where symmetry handling is critical.

Back tomorrow with more!

P.S. In the most recent submission window: cs.LG: 152 papers stat.ML: 10 papers