Hi all,

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

Probabilistic and Statistical Methods

Information Geometry of Message Passing This paper derives a natural-gradient message passing (NGMP) rule from the Bethe free energy by constraining edge marginals to an exponential family. The key novelty is that the update rule is local and retains the part of the exact message that the receiving family can represent, rather than averaging the factor under neighboring beliefs. This matters because it provides a geometrically grounded variational inference scheme that improves uncertainty calibration in non-conjugate settings, such as partially observed latent chains, where standard variational message passing degrades.

Theory and Generalization

Sparse Prototype Code Underlies Classification and Prediction Across Modalities The authors derive an analytical mean-field theory for neural representations, demonstrating that classification in deep networks is governed by a sparse, centroid-aligned structure embedded within high-dimensional space. The novelty lies in showing that within-class variability is not random but strongly correlated with the true class and rival class centroids, and that accurate prediction requires only a small set of these coordinates. This matters because it provides a parsimonious, mechanistic theory of representational geometry that connects empirical deep learning observations to sparse-feature extraction approaches like sparse autoencoders.

Improving the matrix multiplication exponent with modern optimization and AlphaEvolve This work improves the upper bound on the matrix multiplication exponent omega to less than 2.371177 by reformulating the optimization problem at the core of combination loss analysis and applying a new ML-based optimization algorithm refined by AlphaEvolve. The conceptual novelty here is the successful application of evolutionary ML techniques to discover novel mathematical constructions that push theoretical computer science forward. Coming from DeepMind, this paper highlights a disruptive paradigm where AI acts as a primary driver of fundamental mathematical discovery rather than just an empirical tool.

Novel Architectures and Objectives

MiNO: Cotangent-bundle propagator learning for PDEs Instead of learning solution fields or solution maps, this paper introduces the microlocal neural operator (MiNO), which learns the propagator itself: the phase and amplitude in phase space. The inductive bias relies on the fact that the rule moving a discontinuity can be far smoother than the field it generates, using the eikonal equation for phase and the transport equation for amplitude. This matters because it allows sharp fronts and caustics to be handled as propagation geometry rather than pointwise fitted fields, yielding significantly better accuracy with smaller models on discontinuous advection benchmarks.

Until the next announcement!

P.S. In the most recent submission window: cs.LG: 250 papers stat.ML: 39 papers