Hey there,

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

Probabilistic and Statistical Methods

Composing Flow-Matching Energies with Known Physics: Generation, OOD Detection, and Inversion on PDE Fields This paper introduces a method to extract an explicit scalar energy function from flow matching models by using a potential-induced velocity. The energy is obtained purely from the matching regression objective, avoiding variational forms or extra MCMC steps during training. Because this energy composes additively, it can be combined with physics-based energies like PDE residuals for out-of-distribution detection and inverse problems. This is a strong conceptual step toward integrating data-driven priors with known physical structures in a principled, inference-time manner.

Novel Architectures and Objectives

No Gaussian Required: Contrastive Inverse Dynamics for JEPA World Models Joint-Embedding Predictive Architectures (JEPAs) typically rely on prescribed geometry, like an isotropic Gaussian regularizer, to prevent representation collapse. This work replaces that with an action-contrastive inverse-dynamics head (Action-NCE), forcing the latent transition to identify the action that produced it. A collapsed encoder provably fails this discrimination task, providing a distribution-free anti-collapse signal derived directly from transition data. This removes the prescribed geometry bottleneck and eliminates the need for target networks or stop-gradients, which is a notable departure for JEPA objectives.

Recirculation This paper proposes an inference-time architectural enhancement for off-the-shelf foundation models that introduces a specific form of recurrence to track belief states. It bypasses the fundamental limitation that state updates in feedforward transformers are bounded by model depth, without requiring costly retraining of the base model. By leveraging the model itself to inform architectural modifications, adaptive recirculation achieves a 23% reduction in perplexity on the Gemma3 family and a 21% accuracy increase on GSM8k. This suggests a compelling route to architectural evolution guided by a trained network's intrinsic properties.

Quantum ML and Hybrid Paradigms

Iterative tensor network transformations for element-wise evaluation of elementary and filtering functions Tensor networks are powerful for data compression, but their application to general data processing has been limited by the difficulty of performing nonlinear operations. This work introduces iterative tensor network transformations (ITNTs), a framework for evaluating nonlinear filtering functions entirely in the compressed tensor train (TT) domain. The method enables efficient computation on exponentially large datasets while maintaining controlled computational cost, demonstrated by solving Max-SAT instances on spaces up to 2^70 configurations. This provides tensor network methods with a foundational tool for general-purpose data science and large-scale optimization.

Back tomorrow with more!

P.S. In the most recent submission window: cs.LG: 112 papers stat.ML: 21 papers