Hello folks,

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

TOP PICKS

Novel Architectures and Objectives

Emergent Models: Intelligence from Tiny Substrates

This paper proposes a genuine conceptual departure from the differentiable feed-forward paradigm: modeling is treated not as learning a closed-form input-output map but as the emergence of computational behaviors within simple open-ended dynamical systems such as cellular automata. The substrates iterate a fixed local rule over a latent space for an adaptive number of steps, trained by evolutionary search rather than gradient descent. The authors prove that some instances are latent-universal (can realize any partial computable function by varying only the initial condition), and empirically show that tens to hundreds of parameters can extrapolate exactly on arithmetic functions and support online adaptation. While not competitive as an architecture yet, this widens the design space of ML in a way that is not reducible to known Transformer/CNN/GNN variants, making it a foundational contribution to alternative computational primitives.

Variation Brownian Kernel Ladders

This work introduces a path-atomic function-space framework that separates nonlinear recursive dictionary construction from linear variation superposition, built on Brownian RKHS geometry rather than standard neural network parameterizations. Each recursive dictionary atom composes unit-ball profiles from the Brownian RKHS, and the full space is a signed-measure variation hull of the completed dictionary. The authors establish variation-controlled Holder regularity, compactness, strict growth with depth under a local non-degeneracy condition, and derive Rademacher and generalization bounds through Brownian quadratic chaos and VC entropy. The two-stage approximant analysis yields an $M^{-1/2} + m^{-1/2}$ error bound with at most $2M$ active outer-profile basis contributions. This is a principled alternative to depth-benefit arguments that typically rely on informal complexity claims, grounding the architecture in rigorous function-space theory.

Theory and Generalization

Non-Shattering at and Above the Dynamical Temperature in the Spherical Pure p-Spin Model

This paper tackles the geometry of spherical pure p-spin glass landscapes, which serve as statistical-mechanical models for understanding optimization landscapes and training dynamics in high-dimensional learning. The author proves that shattering (the fragmentation of the solution space into exponentially many clusters) does not occur below certain overlap thresholds for $0 < \beta \leq \beta_{\mathrm{sh}}(p)$, combining deterministic $N+1$ bounds for disjoint bands with a general-p sign law showing subdominant free energy. For $p=3$, the criteria exhaust every fixed overlap, meaning the landscape is not shattered at any $T \geq T_{\mathrm{sh}}$. This partially resolves a conjecture by Ben Arous and Jagannath and introduces new proof techniques that could transfer to analyzing loss landscape geometry in neural networks, where shattering-like phenomena are invoked to explain generalization and optimization.

Quantum ML and Hybrid Paradigms

Classical Limits of Spectral Filtering in Quantum Generative Models

This paper rigorously examines whether spectral filtering in quantum circuit Born machines can produce a genuine quantum-classical separation for generative modeling. The author derives necessary and sufficient conditions under which the filtered quantum output can be exactly matched by classical post-processing of samples from the unfiltered model at matched sampling cost. The key result is a dichotomy for magnitude (attenuating) filters: either the filtered output is a constant-size Fourier object with an efficient classical sampler, or the passband must widen until no fixed frequency is attenuated and the filter no longer smooths. In neither case does the filter create a separation; whatever separation survives is inherited from the spectral phase of the input state, which is invisible to Born-rule training loss and set by initialization. This is a sharp negative result that clarifies the boundaries of quantum advantage in generative modeling and narrows the space of potentially useful quantum operations to pure phase filters.

Back tomorrow with more!

P.S. In the most recent submission window: cs.LG: 115 papers stat.ML: 17 papers