Hi all,

Here is your disruptive ML/AI digest for the day.

Here's what you folks might find interesting today.

TOP PICKS

Theory and Generalization

All Routes Lead to Collapse

This paper looks like one of the sharper conceptual pieces in today's batch. The core claim is that attention sinks, representation collapse, and related pathologies are not quirks of transformers, but generic consequences of *content-based routing under a mismatched similarity metric*. What is novel here is the unifying reframing: softmax attention is cast as Boltzmann-weighted aggregation over Euclidean distances with a missing norm term, and that omission is used to explain why routing mechanisms compensate by concentrating mass and collapsing representations. The authors then test the mechanism across very different architectures, including transformers, graph attention, selective state-space models, recurrent mixers, and residual routing over depth, which makes the argument much broader than a transformer pathology note. If you care about inductive biases, training dynamics, or mechanistic explanations of failure modes, this is exactly the kind of cross-architecture diagnosis worth reading. I would flag it as especially relevant because it proposes a general principle about routing, not just an empirical observation about one model family.

Provable Benefits of RLVR over SFT for Reasoning Models: Learning to Backtrack Efficiently

This is a clean theory paper aimed at a very current question: why reinforcement-style post-training can improve reasoning beyond supervised fine-tuning. The interesting move is to model chain-of-thought as graph pathfinding and then prove that SFT on gold trajectories does not teach efficient backtracking, while RL with verifiable rewards can. The claimed result is an exponential separation in inference-time compute, which is a strong and concrete statement rather than vague intuition about "exploration" or "search". That makes it relevant both for emergent reasoning behavior and for formal analyses of when outcome-based training changes the learned algorithm. I also like that the paper goes one step further and argues the RLVR traces can be distilled back into a base model, which connects theory to practical training pipelines. If you are tracking mechanistic accounts of reasoning or principled differences between post-training objectives, this is one of the strongest fits today.

Convergence of Gradient Descent for General Neural Network Architectures Beyond the NTK Regime

This one is squarely in the "broad theory with modern architectures" bucket. The main contribution is a convergence framework for gradient descent that operates at the level of network blocks and is claimed to cover architectures including pre-normalized transformers, while explicitly going beyond the NTK regime. The technical novelty seems to be the combination of an iterate-dependent PL-type inequality with generalized smoothness and relaxed dissipativity arguments, which is more structurally ambitious than the usual narrow-architecture convergence result. The practical angle is also useful: the paper interprets the theorem under Xavier initialization and argues that stable learning-rate scale depends on depth and bottleneck dimensions rather than just maximum width. That kind of statement could matter if you are interested in how architectural structure shapes optimization behavior. It is not flashy, but it looks like a serious attempt to say something general and nontrivial about training dynamics in contemporary networks.

Novel Architectures and Objectives

[Error Highways: Scaling Predictive Coding to Very Deep Networks]

This is one of the more genuinely off-mainstream architecture/objective papers in the list. The authors tackle a longstanding limitation of predictive coding networks by modifying the free-energy formulation so that selected hidden layers receive direct linear couplings to the clamped output error. The key idea, highway error propagation, is meant to preserve local predictive-coding updates while avoiding the depth-wise attenuation that normally kills learning in deep PCNs. That matters because predictive coding is often discussed as a biologically plausible alternative to backprop, but usually fails to scale beyond shallow settings; here they claim robustness up to 128-layer MLPs. The experiments are still on MNIST and Fashion-MNIST, so I would not overread the empirical scope, but the conceptual contribution is clear: a structural fix for deep local-learning systems rather than another backprop variant. If you are interested in alternatives to gradient backpropagation or in learning rules with different computational primitives, this is worth a close look.

[A Markov Chain Approach to Preference Alignment]

This paper stands out because it reframes preference alignment as a *Markovian* process over outputs rather than as scalar reward optimization or minimax game solving. The proposed MCHF procedure uses pairwise preferences directly to define a transition kernel, then studies convergence to the stationary distribution, with the rate controlled by a seminorm that measures non-transitivity in the preference utility. That is a more structural treatment of preference data than standard RLHF reductions, and the paper also ties MCHF, NLHF, and RLHF together through perturbation analysis around the RLHF solution. The relevance here is less about immediate deployment and more about introducing a different mathematical object for alignment: a stochastic process over generations rather than reward fitting. For anyone interested in principled objectives, information-theoretic or probabilistic views of alignment, or alternatives to standard RLHF formalism, this is one of the more original papers in today's set. It also looks like the kind of work that could open up new algorithmic directions if the formulation proves useful.

BY CATEGORY

Novel Architectures and Objectives

[Scalable Physics-Inspired Transformers for Spin Glasses]

[Prime Fourier Embeddings: A Principled Basis for Modular Arithmetic]

[Learning Graphs through Continuous Information Entropy Fields]

[SkyJEPA: Learning Long-Horizon World Models for Zero-Shot Sim-to-Real Control of Quadrotors]

Theory and Generalization

[Topological Out-of-Domain Generalization in Dynamical Systems Reconstruction]

Probabilistic and Statistical Methods

[Flow Annealing Posterior Sampling for Function-Space Regression and Inverse Problems]

[Neural Operator Processes for Probabilistic Operator Learning under Partial Observations]

[Robust Diffusion Models via Divergence-Induced Weighted Denoising]

[Generative Robust Optimisation]

Optimization and Learning Dynamics

Escaping the Variance Trap: Jacobian-Free Dynamics for Root-Finding Bilevel Optimization

Back tomorrow with more!

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