Exact Predictability and Krylov complexity in large chaotic networks

Vladimir Shaidurov[1], Alkesh Yadav[1], Jonathan Kadmon[1]
[1]Hebrew University Of Jerusalem, Jerusalem, Israel

Random recurrent neural networks are a canonical model for collective dynamics, computation, and chaos in large neural circuits, widely used in theories of both the brain and artificial intelligence[1,2]. In the thermodynamic limit, these disordered networks are typically described by dynamical mean-field theory, which replaces the full deterministic system with a self-consistent effective stochastic process[1]. While this framework successfully captures the statistical behavior of the system, it is generally assumed to lose the information needed to recover single-neuron trajectories. Mean-field theory is therefore usually viewed as a statistical coarse-graining, rather than an exact window into microscopic dynamics.

Here we show that this picture is incomplete. For a broad class of analytic activation functions, the mean-field dynamics are exactly predictable from a single continuous trajectory, despite the apparent stochasticity of the effective description. This property fails sharply for non-smooth nonlinearities, revealing that local regularity determines whether the mean-field description remains effectively deterministic or becomes genuinely stochastic.

In dynamical mean-field theory, the effective dynamics of a typical neuron are described by a self-consistent Gaussian process, and we show that this process is exactly predictable for analytic activation functions. The origin of this result lies in the complex-time analytic continuation of the mean-field covariance kernel, which enforces exponential high-frequency decay of the power spectrum and, by the Paley-Wiener criterion[3], vanishing prediction error. For non-smooth activation functions, by contrast, this analyticity is lost, the spectral tails are no longer exponential, and the mean-field process becomes genuinely stochastic. Dynamical mean-field theory therefore becomes, in the analytic regime, not merely a theory of stationary statistics but a theory of conditional inference: a sufficiently resolved continuous observation of any typical trajectory determines its future exactly.

Viewing dynamical mean-field theory as a theory of conditional inference yields a practical route to prediction. Using a Lanczos recursion, we reformulate the predictable mean-field process as an exact semi-infinite chain, where predictive information propagates through an ordered hierarchy of dynamical modes. In this representation, Krylov complexity—familiar from studies of operator growth in quantum many-body physics[4] —quantifies the intrinsic difficulty of prediction in this classical, dissipative system.

Even when the mean-field dynamics remain exactly predictable, the difficulty of extracting that predictability is not set by the Lyapunov exponent. We find that the Krylov growth rate increases with disorder strength more rapidly than the largest Lyapunov exponent and upper-bounds it, revealing a separation between exact inferability and local dynamical instability. These results suggest a principled route to increasing usable memory capacity in random recurrent networks, with implications for information retention in cortical circuits, recurrent architectures for temporal learning and sequence processing, and the decoding and forecasting of neural dynamics in brain-machine interfaces.

References:

[1] Sompolinsky, et al., PRL. 61, 259 (1988)

[2] Lukovsevivius and Jaeger, Comput. Sci. Rev. 3, 127 (2009).

[3] Dym and McKean, Gaussian Processes, Function Theory, and the Inverse Spectral Problem (Dover, 2008).

[4] Parker et al., Phys. Rev. X 9, 041017 (2019).