Memory Is Distributed Across Time in Chaotic Networks

Alkesh Yadav[1], Vladimir Shaidurov[1], Jonathan Kadmon[1]

[1]The Hebrew University Of Jerusalem, Jerusalem, Israel

Large random recurrent neural networks are canonical disordered dynamical systems in both neuroscience and machine learning [1,2]. In the thermodynamic limit, they undergo a transition from stable dynamics to deterministic chaos, and the vicinity of this transition is widely believed to be computationally advantageous because long temporal correlations enhance memory of past inputs [3,4]. Dynamical mean-field theory provides the standard description of this regime by replacing quenched disorder with an effective stochastic process. Within this framework, memory capacity has been studied extensively and is typically found to peak at the edge of chaos [3,4]. However, this perspective is built largely on instantaneous statistics and places external noise and the apparent randomness generated by chaotic recurrence on nearly the same footing, even though the underlying network dynamics are fully deterministic.

Here we show that this neglected deterministic structure leaves a clear imprint on the mean-field dynamics. Even in the presence of external white noise, recurrent interactions generate strong non-Markovian temporal correlations, so that memory is distributed across the trajectory rather than localized in the instantaneous state. We capture this residual structure through a hierarchical decomposition into effective Markovian modes, which reveals how smooth dynamical components preserve information about the past despite stochastic drive.

This residual temporal structure reshapes how memory should be quantified in noisy recurrent networks. Standard analyses treat memory capacity as a property of the network state at a single instant. We instead ask how much information about past inputs can be recovered from multiple observations along the trajectory. This reveals how retained information depends on the number and placement of temporal observations, and which sampling times maximize recovery. By exploiting these temporal correlations, such temporally extended readouts recover substantially more information than instantaneous sampling, offsetting losses due to finite spatial sampling.

This analysis also reshapes the usual interpretation of the edge of chaos. In the standard picture, memory is maximized near criticality because the relevant observables are taken to be equal-time states. Once information distributed across the trajectory is included, the edge of chaos remains the optimal regime, but its advantage becomes less sharply confined: useful memory traces persist over a broader neighborhood of the transition. The edge of chaos is thus not simply a singular operating point, but the center of a wider regime whose temporal correlations can be harvested by appropriate readout. This broadening reduces the need for precise fine-tuning to criticality and helps explain how biological and artificial recurrent systems can achieve strong memory performance without operating exactly at the transition.

More broadly, our results show that memory in noisy recurrent networks resides not only in instantaneous activity, but in the temporal correlations generated by recurrence, disorder, and chaos.

References:

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

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

[3] Schuecker et al., Phys. Rev. X 8, 041029 (2018)

[4] Toyoizumi and Abbott, Phys. Rev. E 84, 051908 (2011)