A dynamical system perspective on the mean-field limit of spatially structured recurrent neural networks

Louis Pezon | Swiss Federal Institute of Technology in Lausanne

Title:

A dynamical system perspective on the mean-field limit of spatially structured recurrent neural networks

Authors: L. Pezon (1,*), W. Gerstner (1).

(1) Brain Mind Institute, École Polytechnique Fédérale de Lausanne, Switzerland.

(*) Presenting author.

Abstract:

A wide class of computational neuroscience models involves large spatially structured recurrent neural networks, where interactions between neurons are described by a kernel function of individual "locations" of the neurons. These locations can represent spatial position on the cortical sheet, or neuronal tuning to task or stimulus variables [1-5]. Mean-field (i.e. large network) limits of such spatially structured networks can be represented as neural fields — an approach consistently used since the 1970s [1-4], that conforms naturally to the spatial structure of the connectivity. This approach has been justified using heuristic coarse-graining arguments [1,4]. Yet, the question of how well the dynamics of the limit neural field captures that of the finite-size network, and how it depends on various parameters of the model, remains incompletely understood.

We study a fully connected and spatially structured recurrent network of linear-nonlinear-Poisson spiking neurons. Neurons are assigned i.i.d. locations in an underlying space (called the "similarity space"), and synaptic weights are given by a kernel function of neuronal locations (as in [5,6]). Recently, the convergence of the empirical measure of neuronal trajectories to the solution of a neural field equation was proven rigorously [6] (see also [7] for networks of integrate-and-fire neurons). Yet, this approach has key limitations. First, it

requires the convergence of the empirical measure of neuronal locations, which depends critically on the dimension of the similarity space and is not needed for the convergence of a "typical" trajectory. Second, focusing on the convergence of trajectories typically yields bounds (using Grönwall-like inequalities) that grow exponentially with time, revealing little about the system's dynamical structure, such as its potential fixed points and their stability properties.

To address these concerns, we adopt a dynamical system perspective: we express the dynamics of both the finite-size network and the limit neural field as the flow of a vector field over a Hilbert space [5]. The finite-size network dynamics is characterised by a random vector field that depends on the empirical measure of neuronal locations. We obtain concentration inequalities for this vector field and its spatial derivatives, at every point of its domain. This allows us to quantify how potential fixed points of the finite-size dynamics, and their associated eigenvalues, differ from those of the limit dynamics. Critically, the obtained bounds depend on the spectral properties of the connectivity kernel and on the nonlinearity of the neuron model, but not on the dimension of the underlying similarity space.

Our approach offers a new perspective on studying the mean-field limit of large structured neural networks, shifting focus from the convergence of individual trajectories to the underlying dynamical structure. This approach could further enable obtaining quantitative guarantees on the similarity between the finite-size network dynamics and its limit, and aligns with the modern paradigm shift to population dynamics in systems neuroscience [8].

References:

[1] H.R. Wilson, J.D. Cowan. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. Kybernetik (1973).
DOI: 10.1007/BF00288786

[2] K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. The Journal of Neuroscience (1996).
DOI: 10.1523/JNEUROSCI.16-06-02112.1996

[3] R. Ben-Yishai, R.L. Bar-Or, H. Sompolinsky. Theory of orientation tuning in visual cortex. Proceedings of the National Academy of Sciences (1995).
DOI: 10.1073/pnas.92.9.3844

[4] W. Gerstner. Time structure of the activity in neural network models. Physical Review E (1995).
DOI: 10.1103/PhysRevE.51.738

[5] L. Pezon, V. Schmutz, W. Gerstner. Linking neural manifolds to circuit structure in recurrent networks. Preprint, bioArxiv (2024).
DOI: 10.1101/2024.02.28.582565

[6] M. Chevallier, et al. Mean field limits for nonlinear spatially extended Hawkes processes with exponential memory kernels. Stochastic Processes and their Applications (2019).
DOI: 10.1016/j.spa.2018.02.007

[7] P.-E. Jabin, V. Schmutz, D. Zhou. Non-exchangeable networks of integrate-and-fire neurons: spatially-extended mean-field limit of the empirical measure. Preprint, arXiv (2024).
DOI: 10.48550/arXiv.2409.06325

[8] S. Vyas, et al. Computation through Neural Population Dynamics. Annual Review of Neuroscience (2020).
DOI: 10.1146/annurev-neuro-092619-094115