Latent kinetic Ising models for neural spike trains

Davide Ghio[1]

[1]Aston University, United Kingdom

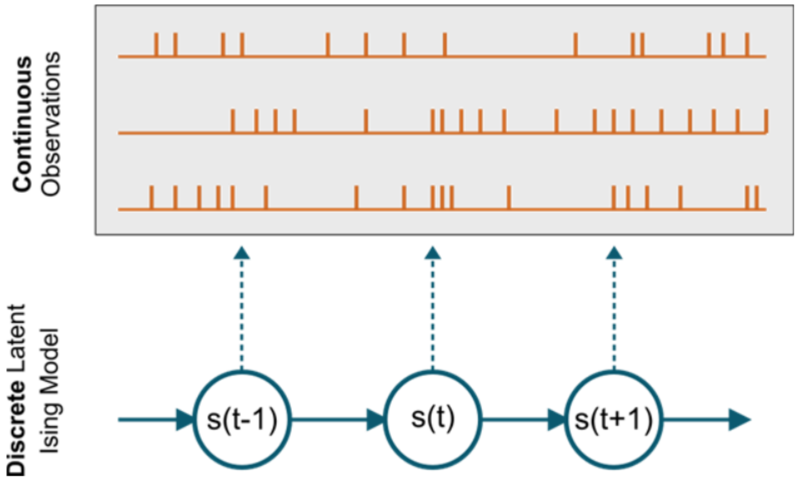

Inferring the underlying functional connectivity of neural circuits from observed spike trains is a central inverse problem in non-equilibrium statistical physics and computational neuroscience. Traditional inference approaches relying on the kinetic Ising model typically require discretizing time into arbitrary bins. This artificially assumes that neural spikes are binary states within fixed windows, thereby failing to capture the continuous nature of neural point processes and inherently blurring fine temporal phenomena like absolute refractoriness.

In this work, we propose a novel analytical framework that strictly separates collective network interactions from single-neuron emission properties. We model the underlying neuronal network as a latent asymmetric kinetic Ising model evolving in discrete time. Crucially, instead of directly equating spins to binned spikes, we treat the observed continuous-time spike trains as realizations of an inhomogeneous Poisson point process conditioned on the latent spin states. The instantaneous conditional intensity of each neuron is modulated both by its latent state and by its own precise spiking history.

To bridge the discrete latent dynamics with the continuous-time observations, we introduce a generalized, analytically tractable emission function that seamlessly captures diverse inter-spike interval (ISI) distributions, from power-law bursting to the exponential tails typical of Leaky Integrate-and-Fire (LIF) neurons. Because the continuous-time likelihood remains fully integrable, we can derive exact analytical gradients for both emission and network parameters without relying on expensive numerical approximations. The inference is driven by a self-consistent mean-field approximation, where the point-process likelihood is naturally incorporated into the mean-field equations through an effective local observation field. To ensure the latent network accurately captures collective interactions rather than being overshadowed by single-neuron emission statistics, we employ a multiphase optimization strategy initialized by empirical ISI tail behaviors.

We validate this theoretical framework on both synthetic models and publicly available multi-electrode recordings. On synthetic data generated from LIF neuronal networks, our latent continuous-time model demonstrates superior structural recovery of the synaptic weight matrix compared to standard discrete-time Ising inference. Traditional inverse approaches suffer from severe estimation errors when fine time-binning leads to highly sparse observation matrices, whereas our formulation naturally leverages continuous spike timings. Furthermore, when applied to multi-electrode datasets, the framework successfully disentangles network-driven correlations from intrinsic single-neuron refractoriness. It accurately reproduces the heterogeneous empirical ISI distributions across different neurons while extracting a robust, biologically plausible functional connectivity graph. By unifying the rich temporal dynamics of point processes with the powerful inverse-problem machinery of the kinetic Ising model, this approach offers a principled, scalable tool for neural data analysis.