Jakob Macke: Correlations and signatures of criticality in neural population models

Large-scale recording methods make it possible to measure the statistics of neural population activity, and thereby to gain insights into the principles that govern the collective activity of neural ensembles. One hypothesis that has emerged from this approach is that neural populations are poised at a 'thermodynamic critical point', and that this has important functional consequences (Tkacik et al 2014). Support for this hypothesis has come from studies that computed the specific heat, a measure of global population statistics, for groups of neurons subsampled from population recordings. These studies have found two effects which—in physical systems—indicate a critical point: First, specific heat diverges with population size N. Second, when manipulating population statistics by introducing a 'temperature' in analogy to statistical mechanics, the maximum heat moves towards unit-temperature for large populations.

What mechanisms can explain these observations? We show that both effects arise in a simple simulation of retinal population activity. They robustly appear across a range of parameters including biologically implausible ones, and can be understood analytically in simple models. The specific heat grows with N whenever the (average) correlation is independent of N, which is always true when uniformly subsampling a large, correlated population. For weakly correlated populations, the rate of divergence of the specific heat is proportional to the correlation strength. Thus, if retinal population codes were optimized to maximize specific heat, then this would predict that they seek to increase correlations. This is incongruent with theories of efficient coding that make the opposite prediction. We find criticality in a simple and parsimonious model of retinal processing, and without the need for fine-tuning or adaptation. This suggests that signatures of criticality might not require an optimized coding strategy, but rather arise as consequence of sub-sampling a stimulus-driven neural population (Aitchison et al 2014).