My Research: The stochastic brain

Susanne Ditlevsen

August 2, 2017

Randomness is ubiquitous in all biological systems. The nervous system is hugely affected by stochasticity, and noise and variability are fundamental of brain function. Information processing is faced with the challenge of acting in this unpredictable and random world, but even if cellular and molecular processes of life itself are noisy and variability in perception and behavior is observed for equal sensory inputs, the brain displays a remarkable precision, essential for survival. Is this ever-present noise a caprice of nature, which evolution has taught us to deal with the best we can? Or does stochasticity play a constructive role, increasing reliability and robustness, and the brain is a probabilistic device because this makes us more fit to survive? Is the noise beneficial or detrimental?

I study stochastic models of biological systems, and develop statistical tools for inferring biological relevant quantities from data. I am particularly interested in how the brain process sensory input. Many probabilistic models show some surprising and different dynamics not present in the deterministic models. These effects might be important for understanding the biological system, as well as essential to incorporate into a statistical analysis to get reliable inferences. The standard approach to deal with the noisy and highly variable data is to average over trials, to presumably get a more reliable output for further analysis. This might blur or entirely remove essential characteristics and mechanisms, which are fundamental for understanding the underlying function of the system under study. More advanced statistical methods and stochastic models are paramount to disentangle the finer mechanisms, because single trials carry information, which is not maintained in the average behavior. For example, a fundamental question concerning representation of the visual world in our brain is how a cortical cell responds when presented with more than a single stimulus. Recently, we found that cells in primate visual cortex presented with a pair of stimuli respond predominantly to one stimulus at a time in a probabilistic fashion, rather than a weighted average response. This provides a mechanism by which the representational identity of multiple stimuli in complex visual scenes can be maintained despite the large receptive fields. Another example is estimation of inhibitory and excitatory conductances from membrane potential measurements. The synaptic input is stochastic and nonreproducible so using repetitions will yield erroneous results.

Let $Z(t)$ be some d-dimensional process of interest. It could for example be the time evolution of a neuron, where one coordinate represents the membrane potential, and two coordinates represent excitatory and inhibitory synaptic input. Typically, only the membrane potential can be measured, and the synaptic input has to be inferred from the effects it has on the membrane potential dynamics.

In general, say we observe $X(t)$ at time points $t_0, t_1, \ldots, t_n$, whereas $Y(t)$ is unobserved, and $Z(t) = (X(t), Y(t))$. Then our data are $x_0, x_1, \ldots, x_n$, where x_j is the observation of $X(t_j)$. In the example above, $X(t)$ is the membrane potential at time t, and $Y(t)$ is 2-dimensional, being the excitatory and the inhibitory synaptic input at time t. The model for $Z(t)$ can for example be given as the solution to a stochastic differential equation. Examples of questions I have looked at in my research are the following:

- How to find a reasonable model for Z, such that the data are well described by the model, and the model can help us answer fundamental questions about the biology behind it.
- What are the properties of the model?
- What role does the probabilistic behavior play?
- How can we infer parameters or other quantities of interest from data?
- How do we estimate the most likely time course of the unobserved coordinates $Y(t)$ from observations of $X(t)$?
- Say we can manipulate $X(t)$, for example by injecting current into the neuron. What is the optimal perturbation course in order for $X(t)$ to be as close as possible to a predefined target? This has applications in brain-machine interfaces.
- Another application is in experimental design. What is the optimal perturbation course in order for us to learn as much as possible about the parameters?
- Assume $X(t)$ is multidimensional, and each coordinate represents a stochastic oscillator. An important question is if the oscillators are coupled, so the interest lies in recovering the functional graph behind the dynamics.
- Assume $X(t)$ is only observed when it crosses a certain threshold, after which it is reset to some predefined value. Then our data are the crossing times $t_1, \ldots, t_n$. An example is extracellular recordings of neurons, such that only spike times are observed, but not the underlying membrane potential evolution. These spike trains are naturally modeled by point processes. Can we estimate parameters of the underlying process from such spike trains?
- From spike trains measured in the visual cortex in monkeys while they are presented with various visual stimuli on a computer screen, can we infer how visual information is processed in the brain?

The mathematical tools applied to answer these questions are many, such as maximum likelihood theory, solutions of partial differential equations, numerical approximation schemes, Hawkes processes, sequential Monte Carlo methods, stochastic approximation expectation maximization algorithm, optimal stopping, mean field limits, interacting particle systems, stochastic optimal control, maximization of mutual information between posterior of unknown parameters and the distribution of the hitting times, stochastic integration theory, co-integration theory, dynamical systems theory, etc.