Probabilistic Inference Using Stored Examples

Lei Shi, Thomas Griffiths

Helen Wills Neuroscience Institute, Department of Psychology, UC Berkeley

Perception involves inferring the state of the world from sensory evidence, considering a potentially unlimited state space. Human behavior is consistent with the optimal statistical solution to this problem in many tasks, including cue combination [1] and sensorimotor learning [2]. Understanding the neural mechanisms underlying this behavior is of particular importance, since probabilistic computations are notoriously challenging. Here we propose a simple mechanism that can account for such behavior which only involves averaging over a few feature detection neurons which fire at a rate determined by their similarity to a sensory stimulus. We show that this mechanism is equivalent to a Monte Carlo method known as importance sampling, and describe a basic neural circuit that can implement this algorithm.

Specifically, assume that a sensory system consists of a population of neurons, each of which fires in response to a stimulus x at a rate proportional to a tuning curve $f(x, \theta)$, where θ are parameters describing the response properties of that neuron (e.g., the stimulus of maximal response). Then, an average in which each θ receives weight equal to the number of spikes produced by the corresponding neurons will approximate the expectation of θ over the posterior distribution $p(\theta|x)$ produced by applying Bayes' rule with likelihood $p(x|\theta)$ proportional to $f(x, \theta)$ and prior $p(\theta)$ proportional to the number of neurons described by parameters θ.

Simulation of a human sensorimotor learning task [2] shows a population of as few as 50 such neurons accounts well for human behavior. This model can also be easily applied to the case of multiple sensory modalities, such as haptic-visual cue combination [1]. In this case, each neuron combines information from both sensory modalities, in a similar manner to [3]. The weighted sum using the spikes produced by these neurons approximates the expectation of the posterior distribution over the state of the world (such as the length of a block) responsible for producing a stimulus. Simulations show that as few as 50 such neurons can also approximate human haptic-visual cue combination data [1].

The notion of using a population of neurons to represent a probability density is not new [3]. However, we extend this observation in two ways. First, we show that a small subset of neurons, whose tuning curves represent a small set of typical examples from sensory experience, is sufficient for Bayesian inference. Second, our theoretical analysis shows that this mechanism corresponds to a Monte Carlo sampling method. Importance sampling enables one to estimate the expectation of a random variable which has a density function hard to sample from by generating samples from a surrogate distribution and weighting these samples by their "importance". When the surrogate distribution is the prior, importance sampling takes the same form as our model.

References

[1] Marc O Ernst and Martin S Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433, 2002.

[2] K. Körding and D. M. Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427:244–247, 2004.

[3] W. J. Ma, J. Beck, P. Latham, and A. Pouget. Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9:1432–1438, 2006.