Neuron
Article

Probabilistic Population Codes for Bayesian Decision Making

Jeffrey M. Beck,[1,7] Wei Ji Ma,[1,2,7] Roozbeh Kiani,[3] Tim Hanks,[3] Anne K. Churchland,[3] Jamie Roitman,[4] Michael N. Shadlen,[3] Peter E. Latham,[5] and Alexandre Pouget[1,6,*]
[1]Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY 14627, USA
[2]Department of Neuroscience, Baylor College of Medicine, Houston, TX 77030, USA
[3]Howard Hughes Medical Institute and Department of Physiology and Biophysics, University of Washington, Seattle, WA 98195-7330, USA
[4]Department of Psychology, University of Illinois at Chicago, Chicago, IL 60607-7137, USA
[5]Gatsby Computational Neuroscience Unit, London WC1N 3AR, UK
[6]Theoretical Neuroscience Group, Collège de France, Paris 75005, France
[7]These authors contributed equally to this work
*Correspondence: alex@bcs.rochester.edu
DOI 10.1016/j.neuron.2008.09.021

SUMMARY

When making a decision, one must first accumulate evidence, often over time, and then select the appropriate action. Here, we present a neural model of decision making that can perform both evidence accumulation and action selection optimally. More specifically, we show that, given a Poisson-like distribution of spike counts, biological neural networks can accumulate evidence without loss of information through linear integration of neural activity and can select the most likely action through attractor dynamics. This holds for arbitrary correlations, any tuning curves, continuous and discrete variables, and sensory evidence whose reliability varies over time. Our model predicts that the neurons in the lateral intraparietal cortex involved in evidence accumulation encode, on every trial, a probability distribution which predicts the animal's performance. We present experimental evidence consistent with this prediction and discuss other predictions applicable to more general settings.

INTRODUCTION

Decision making affects all aspects of human behavior, on time scales varying from seconds to hours to days. For instance, imagine you are driving your car toward a busy intersection and your brakes fail. Within a few hundred milliseconds, you have to decide where to steer your car. Although this is a task we handle relatively easily, in fact it involves three separate, and nontrivial, stages. First, sensory evidence must be accumulated over time. Here, the sensory evidence consists of the image of cars and people in the intersection. Second, the accumulation must be stopped at some point (waiting too long can have disastrous consequences in this situation). Third, an action must be selected. This task is difficult because the sensory evidence and the response are continuous variables, the reliability of the sensory evidence is a priori unknown, and it can vary greatly over time. For instance, as you get closer to the intersection, your ability to distinguish different objects improves. The reliability of the visual information can also vary from day to day: it is much easier to analyze the scene on a sunny day than on a foggy one.

There is currently no neural model that can deal with this type of decision optimally, where by optimal, we mean that the accumulation of evidence is done without loss of information and that the chosen option is the most likely one given the sensory evidence (we do not address the issue of when to make the decision; see Discussion). Yet, it is essential to understand optimal decision making in the face of multiple choices and unknown and time-varying reliability, since most decisions we make fall into this category. Most models are concerned only with binary decision making, and even with this limitation, cannot deal optimally with sensory evidence of unknown and continuously changing reliability. This problem is conceptual: these models have no clear probabilistic interpretation or, when they do, are limited to situations in which the evidence has a constant and known reliability over time and over trials. As a result, it is unclear how, or even if, they are related to the general case we consider in this paper.

Here, we present the first neural model of decision making that performs sensory evidence accumulation and response selection optimally when there are multiple or a continuum of possible decisions and the reliability of the sensory input varies over time or across trials. This model is built around the observation that spike counts in the brain are close to what we call "Poisson-like" (Ma et al., 2006; Shadlen and Newsome, 1998; Tolhurst et al., 1983). Given this observation, our main contributions are twofold. First, we show that for Poisson-like distributions, optimal evidence accumulation can be performed through simple integration of neural activities, while optimal response selection can be implemented through attractor dynamics. Second, we show (again for Poisson-like distributions of neural activity) that neurons encode the posterior probability distribution over the variables of interest at all times. This latter contribution has far-reaching implications, since it suggests that neurons implicated in simple perceptual decisions represent quantities that are directly relevant to inference, confidence, and belief.