Neural Networks 21 (2008) 1247–1260

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

# Integrated Bayesian models of learning and decision making for saccadic eye movements☆

Kay H. Brodersen[a,b,*], Will D. Penny[a], Lee M. Harrison[a], Jean Daunizeau[a], Christian C. Ruff[c], Emrah Duzel[c], Karl J. Friston[a], Klaas E. Stephan[a,d]

[a] Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK
[b] Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB), John Radcliffe Hospital, University of Oxford, Oxford OX3 9DU, UK
[c] Institute of Cognitive Neuroscience, University College London, 17 Queen Square, London WC1N 3AR, UK
[d] Branco-Weiss-Laboratory, Institute for Empirical Research in Economics, University of Zurich, Switzerland

ARTICLE INFO

*Article history:*
Received 7 September 2007
Received in revised form
29 August 2008
Accepted 31 August 2008

*Keywords:*
Saccades
Decision making
Reaction time
Bayesian learning
Model comparison

ABSTRACT

The neurophysiology of eye movements has been studied extensively, and several computational models have been proposed for decision-making processes that underlie the generation of eye movements towards a visual stimulus in a situation of uncertainty. One class of models, known as linear rise-to-threshold models, provides an economical, yet broadly applicable, explanation for the observed variability in the latency between the onset of a peripheral visual target and the saccade towards it. So far, however, these models do not account for the dynamics of learning across a sequence of stimuli, and they do not apply to situations in which subjects are exposed to events with conditional probabilities. In this methodological paper, we extend the class of linear rise-to-threshold models to address these limitations. Specifically, we reformulate previous models in terms of a generative, hierarchical model, by combining two separate sub-models that account for the interplay between learning of target locations across trials and the decision-making process within trials. We derive a maximum-likelihood scheme for parameter estimation as well as model comparison on the basis of log likelihood ratios. The utility of the integrated model is demonstrated by applying it to empirical saccade data acquired from three healthy subjects. Model comparison is used (i) to show that eye movements do not only reflect marginal but also conditional probabilities of target locations, and (ii) to reveal subject-specific learning profiles over trials. These individual learning profiles are sufficiently distinct that test samples can be successfully mapped onto the correct subject by a naïve Bayes classifier. Altogether, our approach extends the class of linear rise-to-threshold models of saccadic decision making, overcomes some of their previous limitations, and enables statistical inference both about learning of target locations across trials and the decision-making process within trials.

© 2008 Elsevier Ltd. All rights reserved.

## 1. Introduction

In order to survive in a competitive, dynamic environment, animals must be able to integrate past experience with sensory evidence to infer the current state of the world and execute a behavioural response. Marked progress in our understanding of the neural basis of decision making has been achieved by focusing on sensory-driven decisions, such as the simple question of where to look next. Studying decision making in sensorimotor systems like the oculomotor system has the advantage that one can exploit a large body of neuroanatomical and neurophysiological knowledge that has been accumulated over the past decades. It seems conceivable that studying the neuronal mechanisms of visual-saccadic decision making could provide us with a blueprint of how the brain implements other sensorimotor decisions, or even deliver "a model for understanding decision making in general" (Glimcher, 2003).

The decision processes that underlie rapid eye movements towards a target have been studied in a variety of experimental paradigms. One seminal series of studies is based on the *random dot-motion* task designed by Newsome and colleagues (Newsome & Pare, 1988). In an initial fixed-duration version of this task, monkeys were trained to discriminate the motion direction of a set of moving dots with varying degrees of coherence, and indicate the perceived motion by a leftward or rightward

☆ This research has been funded by the Wellcome Trust (VS/06/UCL/A18), the German Academic Exchange Service (DAAD, D/06/49008), and the Stiftung Familie Klee (Frankfurt/Main).

* Corresponding author at: Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK.
*E-mail address:* kay.brodersen@gmx.net (K.H. Brodersen).

0893-6080/$ – see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2008.08.007

saccade (Newsome, 1997; Newsome, Britten, & Movshon, 1989; Newsome, Britten, Salzman, & Movshon, 1990; Salzman, Britten, & Newsome, 1990). Subsequently, Shadlen, Britten, Newsome, and Movshon (1996) suggested a computational explanation of the neuronal mechanisms producing the resulting saccade and provided experimental verification of its key assumptions (Gold & Shadlen, 2000; Kim & Shadlen, 1999; Shadlen et al., 1996; Shadlen & Newsome, 2001). In particular, they identified a gradual rise of spiking activity in the lateral intraparietal (LIP) area integrating motion direction-specific signals from the middle temporal (MT) area (Shadlen & Newsome, 1996, 2001).

Based on a reaction time version of the same task (Roitman & Shadlen, 2002), Shadlen and colleagues advanced the hypothesis that rising activity before a saccade, which had also been observed in the frontal eye fields (FEF), represented the ratio of the log likelihoods that the two possible eye movements would be executed (Gold & Shadlen, 2000, 2001). Based on their decision-theoretic analysis, they suggested that log likelihood ratios might be used as "a natural currency for trading off sensory information, prior probability and expected value to form a perceptual decision" (Gold & Shadlen, 2001).

Another key series of studies was carried out by Hanes, Schall, and colleagues, who investigated an *oddball* task (as well as the *countermanding* paradigm; Hanes and Carpenter (1999)) to study how neural signals in the FEFs would finally trigger the initiation of saccades (Hanes & Schall, 1996; Hanes, Thompson, & Schall, 1995; Schall & Thompson, 1999; Thompson, Bichot, & Schall, 1997; Thompson, Hanes, Bichot, & Schall, 1996). In their oddball task, monkeys were trained to indicate, by an eye movement, the location of the oddball within a circular arrangement of visual stimuli around a central fixation dot. They showed that FEF activity was consistent with psychophysical models about oddball reaction time tasks (Luce, 1986; Ratcliff, 1978; Sternberg, 1969a, 1969b). Specifically, their findings supported the notion that the saccadic decision would be made as soon as gradually increasing neural activity in the FEFs had crossed a biophysical threshold (Hanes, Patterson, & Schall, 1998; Schall & Thompson, 1999).

Motivated by the question of why saccadic latencies displayed large variance in all of the above tasks, an even simpler reaction time paradigm was investigated by Carpenter and colleagues (Carpenter & Williams, 1995; Reddi & Carpenter, 2000). In their *saccade-to-target* reaction time task, human subjects were asked to shift their gaze from a central fixation stimulus to an eccentric target as soon as it appeared on the screen. The critical manipulation was to vary the uncertainty about where the target would appear (Basso & Wurtz, 1997, 1998). It was found that saccade latencies became shorter with increasing prior probability of the corresponding target location. Specifically, response speed was found to be proportional to the log prior probability of target location (Basso & Wurtz, 1997, 1998; Carpenter & Williams, 1995).

The behavioural and electrophysiological findings from all three paradigms described above are consistent with the notion of a saccade being elicited once some gradually rising neuronal activity crosses a biophysical threshold. This idea has been formalized in terms of various mechanisms known as *rise-to-threshold* accumulator models. These models aim to provide a computational abstraction of a biophysically conceivable mechanism that explains saccade latencies and their variability across trials (for reviews see Glimcher (2001, 2003), Gold and Shadlen (2001), Platt (2002), Ratcliff and Smith (2004), Schall (2001, 2003), Smith and Ratcliff (2004) and Usher and McClelland (2001)).

In the context of saccadic decision making with a fixed set of potential target locations, rise-to-threshold models assume that subjects maintain a set of hypotheses each of which corresponds to one such location (Carpenter & Williams, 1995; Gold & Shadlen, 2002; McMillen & Holmes, 2006; Shadlen & Gold, 2004). As the stimulus appears, a measure of evidence for each of these hypotheses is continuously refined, implemented as a competition between alternative decision signals in the brain. At any given point in post-stimulus time, these decision signals might, for example, represent the posterior probabilities of the target hypotheses, as derived from the subject's prior (Basso & Wurtz, 1997, 1998; Platt & Glimcher, 1999) and the sensory evidence (i.e., the likelihood of the data) collected up to that point in time (Carpenter, 2004; Carpenter & Williams, 1995). As soon as one such signal reaches a preset threshold, a saccade is elicited towards the corresponding target. Depending on the way in which information is assumed to be accumulated over time, two specific types of rise-to-threshold model are often distinguished: random-walk models and linear rise-to-threshold models.

*Random-walk* or *diffusion* models are fundamentally based on a sequential probability ratio test that is being carried out continually (Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Ratcliff, Zandt, & McKoon, 1999; Wald, 1945). In these models, each new incoming piece of sensory evidence either increases or decreases a single decision variable until it has drifted beyond a threshold associated with the saccadic movement towards a particular target. The decision variable represents the relative evidence for the two alternatives (Ratcliff & Rouder, 1998). However, in the case of a simple saccade-to-target task in a high-contrast setting with highly salient targets, it has been questioned whether a random-walk process for target detection provides a sufficient explanation for the large variability in latencies (Carpenter, 2004; Carpenter & Reddi, 2001; Reddi, 2001).

In *linear* rise-to-threshold models, randomness is introduced as trial-by-trial changes in the otherwise constant rate of rise of the decision signal. This notion has been formalized by Carpenter in a model termed 'LATER' (linear approach to threshold with ergodic rate; Carpenter and Williams (1995), Leach and Carpenter (2001), Reddi, Asrress, and Carpenter (2003)). Like other rise-to-threshold models, LATER proposes that a saccade towards a target is elicited as soon as a neural decision signal has reached a particular threshold. But unlike other rise-to-threshold models (e. g., Grice (1968) and Nazir and Jacobs (1991)), it assumes a fixed threshold and a linear increase whose rate is subject to variation *across* trials, yet fixed *within* a given trial (for a debate on the relationship between the two approaches see Carpenter and Reddi (2001), Ratcliff (2001), Usher and McClelland (2001)). The neurophysiological recordings by Schall and colleagues (Hanes & Schall, 1996; Schall & Thompson, 1999) are consistent with these key assumptions of the LATER model: they had observed that the threshold for saccade release seemed to be constant, whereas the slope of the rise in activity varied considerably across trials (see Fig. 2a).

In their experiments on the saccade-to-target task, Carpenter and colleagues found that the observed saccadic latency was a function of the log probability of the corresponding target location: the more likely the target location, the shorter the latency (Carpenter & Williams, 1995). LATER accounts for this relationship by assuming that the learned a priori target probabilities determine the baseline levels of the decision signals, but not their rates of rise (cf. biased choice theory by Luce (1963)). Carpenter and colleagues used LATER to produce remarkably accurate predictions of human latency distributions in the saccade-to-target task as well as variations of it (Asrress & Carpenter, 2001; Carpenter & Williams, 1995; Leach & Carpenter, 2001; Reddi et al., 2003; Reddi & Carpenter, 2000).

Download English Version:

https://daneshyari.com/en/article/10326518

Download Persian Version:

https://daneshyari.com/article/10326518

Daneshyari.com