ELSEVIER

Neurocomputing 38–40 (2001) 1511–1517

NEUROCOMPUTING

www.elsevier.com/locate/neucom

Stochastic optimal control and the human oculomotor system[☆]

Liam Paninski*[,1], Michael J. Hawken

Center for Neural Science, New York University, 4 Washington Place, New York, NY 10003, USA

Abstract

Neuroscientists have long been interested in how efficiently we solve probabilistic sensory problems. In order to explore analogous questions in the motor domain, we observed the eye movements of human subjects attempting to track a visual target which moved stochastically across a computer screen. The subjects' behavior was then compared to a mathematically-derived bound on the best performance possible in such a task. The subjects were able to perform surprisingly near the optimum under the conditions examined. These results constitute an important step in determining the efficiency of the nervous system in the context of ongoing behavior.

Keywords: Optimal control; Oculomotor; Psychophysics; Statistical modeling; Prediction

1. Introduction

Much of our fascination with the nervous system is based on the feeling that brains seem to efficiently solve difficult statistical problems in real time. How good, in fact, are we? Quantitative answers to this simple question, even in the most specialized of experimental preparations, can lead to deep insight into the mechanisms underlying

[☆]This work was supported by NIH Grant R01 EY 08300. LP is supported by a predoctoral fellowship from the Howard Hughes Medical Institute. We thank S. Geman, N. Hatsopoulos, and P. Glimcher for many thought-provoking discussions.

[1] http://www.cns.nyu.edu/ ~ liam.

* Corresponding author.

E-mail address: liam@cns.nyu.edu (L. Paninski).

0925-2312/01/$ - see front matter
PII: S0925-2312(01)00541-0

the behavior of nervous systems in general (see e.g. [1]). This idea has been pursued quite successfully in the sensory domain, but has received less attention from the motor psychophysics community. This is unfortunate, as motor systems face a subclass of problems with some fascinating special characteristics; these "motor problems" are inherently sequential, involving feedback with the environment. What we decide to do now often affects what we will be able to do (and, often, what we will perceive) in the future, and this constant interaction with the external world has important consequences for the design of behavioral systems.

Thus, we were led to study a well-understood, experimentally accessible test system—the human oculomotor system—in the context of a simple probabilistic visual tracking task. This experimental setting has obvious ethological relevance and permits a rigorous definition of the upper bounds on the performance we can expect of the system. In other words, we can provide a well-defined answer to the question of how well this system can perform this task, and this, we can hope, will lead to a deeper understanding of the computational efficiency of motor systems in general. While the oculomotor system has been modeled from an optimal control point of view before (see [2] for a review), the approach of rigorously comparing the performance of the system to a well-defined, statistically-derived upper bound, as in the perceptual setting, seems to be new.

2. Experimental methods

Two subjects were instructed to track a visual stimulus in such a way that an average error was minimized (see below). The stimulus consisted of a small black dot on a uniform gray field, presented on a television monitor under computer control. The horizontal position of the target as a function of time was given by a Markov chain on position and velocity, while the vertical position was constant. In some experimental conditions, the target moved in a continuous fashion, perhaps with a few jumps in any given trial; in others, the motion of the target was of a completely discrete nature. (The data shown here was taken from the former "smooth" experimental condition.) As the subject tracked the target, we recorded the horizontal position of one eye using a dual-Purkinje-image infrared eyetracker [3].

An instantaneous error signal was defined by comparing this eye position signal to the simultaneous position of the visual target; the total error over the whole trial was defined as a weighted average of this instantaneous error. We delivered auditory feedback (a tone whose instantaneous pitch was proportional to the error computed up to that time step) in real time, as well as visual feedback (a numerical representation of the error) after every trial, to allow subjects to purposefully adjust their tracking strategy. To give the subject ample opportunity to learn the error metric and the statistical properties of the stimulus for a given experimental condition, we presented the subjects with at least 200 practice trials (each about three seconds long) before any data were analyzed.

3. Analysis and results

The goal of the series of experiments described here is to compare human behavior to a well-defined, theoretically derived optimum. The initial task, then, is to find this optimum, to find the best possible tracking strategy, if it exists, and to compute its average error; later we will discuss in more depth exactly how the comparison to the optimum was made. Any optimal tracking strategy must depend on three items: first, the (Markov) probability distribution from which we draw the dot paths; second, the error metric used for the given experiment; and third, our model of the biological constraints on the oculomotor system. Of these items, only the third requires further discussion. For our purposes, the eye moves in two ways: in "pursuit" and "saccade" mode [4]. Pursuit behavior is characterized by smooth motion, with speeds not exceeding about 100°/s; saccades, on the other hand, are large, brief jumps of gaze direction. The important point is that the eye cannot move or change speed too quickly without making a saccade, and the oculomotor system is not typically capable of making more than about six saccades per second. This effectively places an upper bound on the average velocity of the eye over the course of a behavioral trial. Moreover, the eye does not respond to visual stimuli instantaneously but rather lags by some positive reaction time. These constraints, coupled with the probabilistic nature of the stimulus, bound the best achievable average tracking error away from zero in general, and thus provide us with the upper bound on performance we are looking for.

Given these three ingredients, an "optimal control" is defined as a function, a rule that will look at the past behavior of the dot and of the eye and decide where to put the eye on the next time step (time is measured in steps of 14 ms, at the frame update rate of the monitor). Eye paths are then generated by applying this rule recursively, updating our representation of the past behavior of the dot and the eye on each time step. We say this rule is optimal only if it minimizes, on average, the error function we defined above. Choosing this best rule is a nontrivial task; to optimally decide what to do at any time 0, we theoretically have to look many time steps ahead, since our actions at the present time step determine what we'll be able to do in the future. Obviously, we cannot look infinitely far into the future to decide what to do now, so we have to introduce a "horizon," or "depth" parameter—the number of time steps we'll allow ourselves to look ahead as we decide what to do at the present step. When we choose a fixed, finite value for this foresight parameter, we finally arrive at a well-defined optimization problem whose solutions are our optimal controls. Due to space limitations, we cannot describe here how we actually computed these optimal rules: full details are available at http://www.cns.nyu.edu/∼liam.

In general, the optimal control is not uniquely defined; this means that we cannot directly compare what the subjects did to what the optimal control would have done in the same situation on a trial-by-trial basis, because "what the optimal control would have done" is not necessarily well-defined. However, we can always compare the subjects' performance to the optimum "in the mean"; that is, we'll compare the distribution of errors for an optimal control and subjects over a large ensemble of trials. At the same time, we have a natural null hypothesis to test against, corresponding to

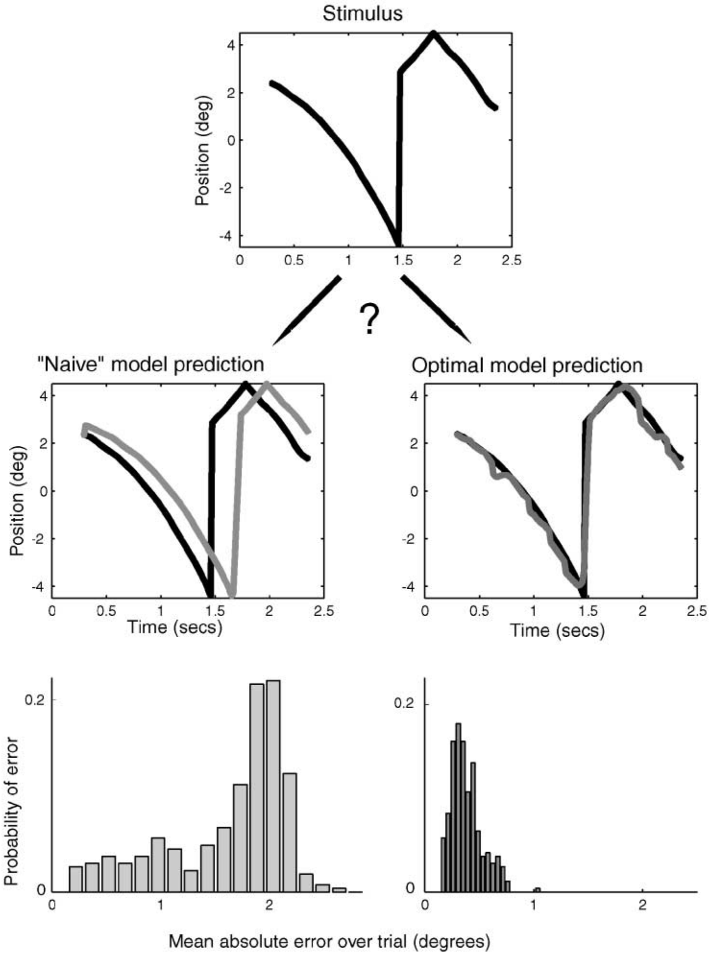

Fig. 1. General experimental strategy. Top: an example stimulus trace (horizontal stimulus target position vs. time). We compared our subjects' behavior to that predicted by the two models described in the text: "naive" and optimal. The comparison was done qualitatively by examining sample paths (middle row), or more quantitatively, by examining the distributions of error incurred by the subjects and the models over several hundred trials (bottom row).

what we may call the "naive" strategy: the eye merely follows the stimulus, going wherever the target was observed last. Between the upper bounds provided by the optimal model and the null point provided by this "naive" model, we have a well-defined scale by which we can judge the performance of our subjects. (See Figs. 1 and 3.)

After all this discussion, what did we find? The data shows a clear trend away from the null hypothesis represented by the "naive" strategy. Subjects performed about as

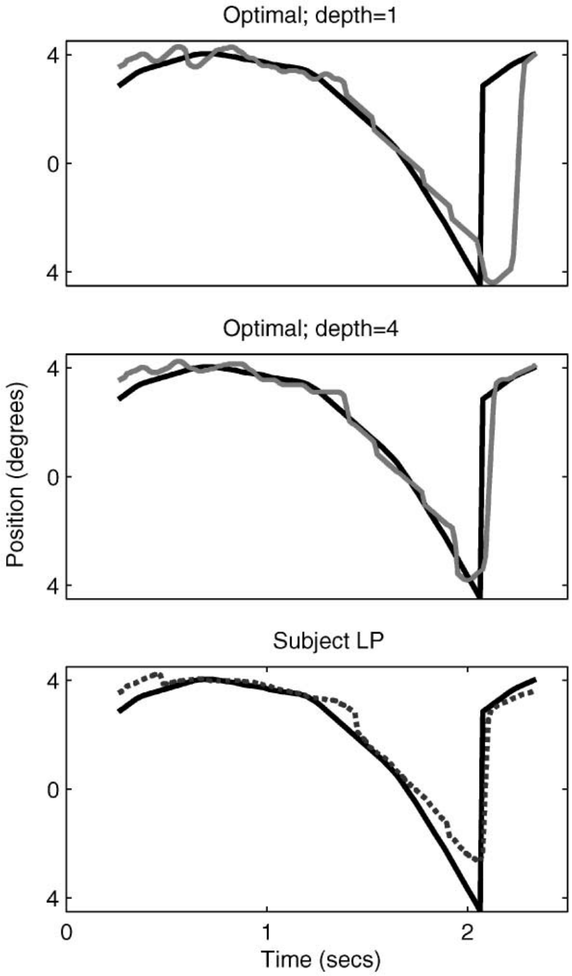

Fig. 2. Example data trace (subject LP) and corresponding optimal model predictions. (Dark line: stimulus; lighter line: action of subject or model.) The top path was generated with the depth, or horizon, variable for details) set to 1, while the middle path had a depth of 4; subject LP's eye path appears in the bottom panel. Note the greater "foresight" displayed by the depth = 4 model and subject LP at the large jump in the middle of the trial; the subject appears to anticipate the jump and plan for it accordingly.

well as the optimal (depth = 1) model predicted. Eye position traces consistently showed predictive behavior of the type seen in Fig. 2.

4. Discussion

The "ideal observer" idea is therefore useful in examining time-varying behavior; we believe that the optimal stochastic control approach presented unifies previous work

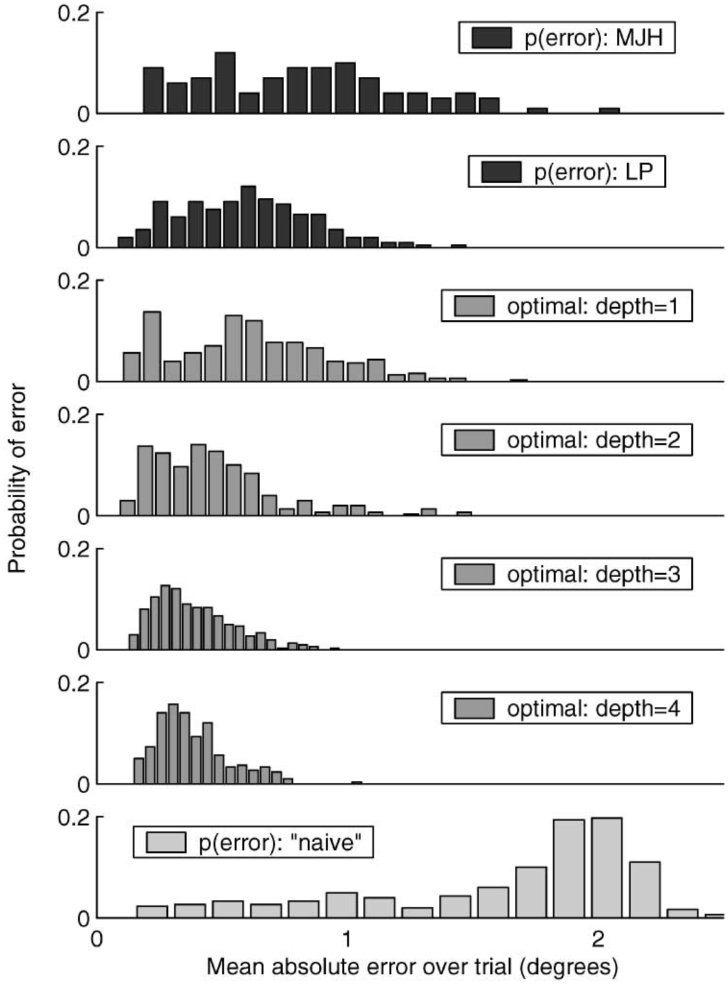

Fig. 3. Summary error distributions for the experiment shown in Fig. 2. Here we can perform the comparison of the subjects' performance to that predicted by the optimal and "naive" models. It is clearly easy to reject the "naive" null hypothesis (bottom histogram); the data (top two histograms) seem to agree best with the optimal (depth = 1) model.

examining predictive vs. reactive components of the oculomotor system (and indeed, of motor systems in general, which tend to face similar problems). We have provided some hints that humans perform difficult, real-time, probabilistic tasks with surprising efficiency, and this is in nice qualitative agreement with the large body of work that exists on the sensory side of this problem. Our results also provide upper bounds on the processing noise in the oculomotor system, in the sense that any reasonable neural model of the oculomotor system must be capable of this kind of near-optimal behavior.

References

[1] W. Bialek, Physical limits to sensation and perception, Ann. Rev. Biophys. Chem. 16 (1997) 455–478.
[2] M. Pavel, Predictive control of eye movement, in: E. Kowler (Ed.), Eye Movements and their Role in Visual and Cognitive Processes, Elsevier, Amsterdam, 1990, pp. 71–114.
[3] H.D. Crane, C.M. Steele, Generation-V dual-Purkinje-image eyetracker, App. Opt. 24 (1985) 527–537.
[4] R. Carpenter, Eye Movements, CRC Press, Boca Raton, 1991.