Stochastic optimal control and the human oculomotor system *

Liam Paninski and Michael J. Hawken
Center for Neural Science, New York University,
4 Washington Place, New York, NY 10003, U.S.A.
(liam, mjh)@cns.nyu.edu
Suggestions and comments are welcome and should be directed to LP.
This work was presented in poster format at the Computational Neuroscience 2000 (CNS'00) meeting, held in Brugge, Belgium, July, 2000.

September 14, 2000

Abstract

Nervous systems seem to efficiently perform a large class of computationally difficult statistical tasks. We rigorously quantified the efficiency of the oculomotor system's solution to a prototypical example of this type of problem, with the hope of deriving a deeper understanding of how our brains solve such problems in general. Specifically, we observed the eye movements of human subjects attempting to track a visual target that moved stochastically across a computer screen. The subjects were able to perform surprisingly near the optimum under the conditions examined - an important step in verifying the introspective conviction that nervous systems are somehow designed to compute physically relevant probabilistic objects.

1 Introduction

Many of the major problems in perception and behavior can be viewed in terms of statistical decision and estimation theory. Pattern recognition, for example, is the underlying problem facing all sensory systems, and is most logically cast in terms of the familiar language of prior and conditional probabilities, cost functions, and so on. Behavioral systems, including those that do not directly control muscle activity but are instead associated with so-called "active perception," deal with a very interesting dynamical subset of these statistical problems. One can view the evolution of these systems in terms of a balance between two competing processes: on one hand, we continually gain information about the world - the state of the world is recursively conditioned (in the probabilistic sense) on the data afforded by our sensory systems; on the other hand, this information is constantly slipping away as the world changes with the passage of time [1] [2]. Moreover, the problems faced by these systems are sequential: what we do now influences what we'll be able to do (or even, in many cases, what we'll be able to perceive) in the future.

Now, much of our fascination with the nervous system is based on the feeling that brains seem to solve these problems (of which all but a trivial minority would otherwise seem to

*A shorter summary of this work is to appear in the CNS'00 meeting proceedings volume of the journal Neurocomputing, published by Elsevier.

be computationally intractable) with the greatest of ease: we hit baseballs, navigate through busy streets, and carry on conversations, all with a lack of conscious effort that belies the inherent difficulty of these tasks. The main question addressed here is this: how good, really, are we? How well can we compute the solutions to these dynamic statistical problems in real time? The hope is that, by examining the system at its quantitative limits, in a stripped-down laboratory setting, we can derive qualitative, fundamental understanding of the general mechanisms underlying the system's behavior under natural conditions [3]. As an analogy, we have kept in mind the gains in our understanding of retinal function following the psychophysical observation that humans seemed to be able to "count photons" - i.e., to reliably detect flashes of light weak enough to induce on the order of single photon captures at the photoreceptor surface. The insights gained in this important test system, of course, have generalized, and proven useful in studying a wide variety of other sensory modalities, all of which face similarly difficult information-processing problems [4].

In the same spirit, we analyzed the performance of an experimentally approachable, well-understood motor system - the system responsible for controlling the direction of our gaze - on one of the problems this system is designed to solve. This problem - tracking a visual stimulus which moves in a probabilistic fashion - has obvious ethological relevance and permits highly quantitative analysis in the laboratory setting. It is well known that the oculomotor system behaves both predictively and adaptively when confronted with this kind of task; these qualities make this system ideal for the study of the class of dynamic statistical problems discussed above. While the oculomotor system has been modeled from an optimal control point of view before (see [5] 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

2.1 Stimuli

Two subjects, heads fixed by a bite bar, viewed a monitor at a distance of 1.2 meters in a dark room. The visual stimulus consisted of a small black dot on a uniform gray field. The horizontal position of the dot as a function of time was given by a Markov chain on position and velocity (the vertical position was constant). Thus, given the position and the velocity of the stimulus dot at some arbitrary time 0, and knowledge of the underlying Markov chain, an ideal observer should be able to calculate the probability of seeing the dot travel any fixed arbitrary path in the future, and should therefore be able to predict the future behavior of the dot optimally given only this information.

2.2 Eyetracking

The horizontal position of one eye was recorded using a Purkinje-image infrared eyetracker [6]. The eyetracker linearly translated the instantaneous eye position into an analog DC voltage, which was then digitized and recorded in parallel on two separate machines. One copy of the signal was recorded on a CED1401 system (a signal-conditioning box attached to a PC) and saved for offline analysis (see [7] for details), while the other copy was monitored in real time by the SGI O2 generating the visual stimulus. This copy of the eye position signal was compared to the simultaneous position of the visual stimulus, and the difference between the two was used to generate an online error signal. (This signal was fed back to the subjects to allow them to purposefully adjust their tracking strategy, as described below.) The system was recalibrated by least squares every twenty trials in order to ensure stationarity of the recording apparatus.

2.3 Behavioral task

One of the major roles of the oculomotor system is to track visually attended objects - to stabilize, as much as possible, the foveated images of these objects as they (or we) move through the world. This makes the behavioral task we presented to this system (track the stimulus dot as it moves across the screen) a fairly natural choice. Of course, researchers have been employing variants of this tracking task for decades. What, then, is novel about our approach? We have already discussed the random nature of our stimuli: the Markovian parametrization we employ allows us great control over the aspects of the test signal that can and can not be predicted. This is an important difference between our stimuli and those which have appeared previously, such as narrowband Gaussian noise, or the sum of a small number of sinusoids, or purely deterministic signals such as square or triangle waves. The examples we present below should make the advantages of our stimulus description more clear.

The most important aspect of our approach, however, is that we are as explicit as possible about what we want our subjects to do. The fundamental idea behind this work is that the human oculomotor system (and, by extension, other motor systems) might perform certain tasks optimally. To examine this question experimentally, we need, first, to define our task, and second, to make every effort to ensure that the subject is performing the task we want him to perform. We cannot simply ask the subject to track the stimulus dot optimally; this task is not well-defined. Should we look at the subject's worst performance over all trials? His best? How should we define his performance even on a single trial? In keeping with standard statistical definitions of optimality, we chose to rank subject performance by their average error, across all trials, where the error per trial was basically a weighted time average of the instantaneous error (the difference between the location of the stimulus dot and the location of gaze) during the trial. Subjects were explicitly instructed to minimize this average error.

One could object that this task is in some sense "unnatural," and that we should instead determine how humans track objects without this excessive instruction. We agree that this question - how do we decide, given our knowledge of the world and the information available at the periphery of our vision, where to move our eyes? What could be the natural "error function" we might be trying to minimize as we watch a bird flying through the sky? - is very interesting. However, it is not the question we are addressing here; we have sacrificed some of the "naturalness" of our task in the interest of asking a question we could answer quantitatively given the methods currently available.

It is important to note that our definition of optimality depends strongly on the notion of error and on the statistical properties of the motion of the stimulus to be tracked. As we emphasized above, we had to make every effort to ensure that the subject's idea of what he was supposed to do - his definition of optimal tracking - was exactly identical to what we wanted him to do. The subject learned the statistical structure of each visual stimulus by observation: subjects completed a minimum of 200 trials on a given experimental condition before any data were analyzed, in order to give the subject ample opportunity to learn the probability distributions from which the stimuli were drawn. Similarly, the error metric we chose for each experiment could only be learned with practice, with feedback (it seems too much to ask to expect the subject to consciously minimize, say, mean-square vs. mean-absolute error, given only verbal instruction). We used the online error signal, generated at the SGI as described above, to deliver this necessary feedback in several different ways. First, we produced, in real time, an auditory feedback tone whose pitch was proportional to the running average of the error signal chosen for that experiment (e.g., mean-square distance between the direction of gaze and the dot position: $E(t) = \int_0^t (d(\tau) - e(\tau))^2 d\tau$). This auditory feedback gave subjects a very direct way to continually grade (and tune) their performance. Second, subjects received visual feedback on their performance after each trial; a number proportional to the total error over the course of the trial was displayed for two seconds following the end of each trial. Finally, the average of the error over all trials was displayed after each block of twenty trials (before recalibration); recall that this average is what the subjects were instructed to minimize. These steps gave us reasonable confidence that the subject was provided with all of the information necessary to perform the task as we intended.

The experiments were completely self-paced; the subject determined the start of each trial by the click of a button, and each trial lasted about seven seconds in total (including post-trial feedback). Our analysis included only the central 2.5 seconds of tracking, in order to exclude any possibly contaminating "end effects" at the beginning or the end of the trial.

3 Analysis

We would like to compare the subjects' performance on this tracking task to the best performance an oculomotor system could hope to achieve. Therefore, we need to discuss the constraints under which the oculomotor system operates. For our purposes, the eye moves in two ways: in "pursuit" and "saccade" mode [8]. Pursuit behavior is characterized by smooth motion, with speeds not exceeding about 10 deg/s and typically small accelerations. Saccades, on the other hand, are large, brief jumps of gaze direction. The important point is that the eye can not 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. These effective velocity constraints, and the fact that the eye does not respond to visual stimuli instantaneously but rather lags by some positive reaction time, lead to well-defined upper bounds on the performance of the human oculomotor system in this task. Much of the remainder of this section is devoted to describing these upper bounds in more detail.

At the same time, we have a natural null hypothesis to test against, corresponding to what we will call the "naive" strategy: the eye merely follows the stimulus, going wherever the dot was observed last. Between the upper bounds discussed above and the null point presented by this "naive" model, we have a natural scale by which we can judge the performance of our subjects. (See Figures 1 and 4.)

3.1 General problem; definitions

Let us state more precisely the problem we are asking the oculomotor system to solve. First, in colloquial language: we start with three ingredients: the (Markov) probability distribution from which we draw dot paths, the error metric we used for the given experiment, and our model of the biological constraints on the oculomotor system. We assume that the stimulus distribution and error metric are known exactly; that is, we ignore completely the problem of learning these objects from the finite data set the subjects are allowed to observe. Our job

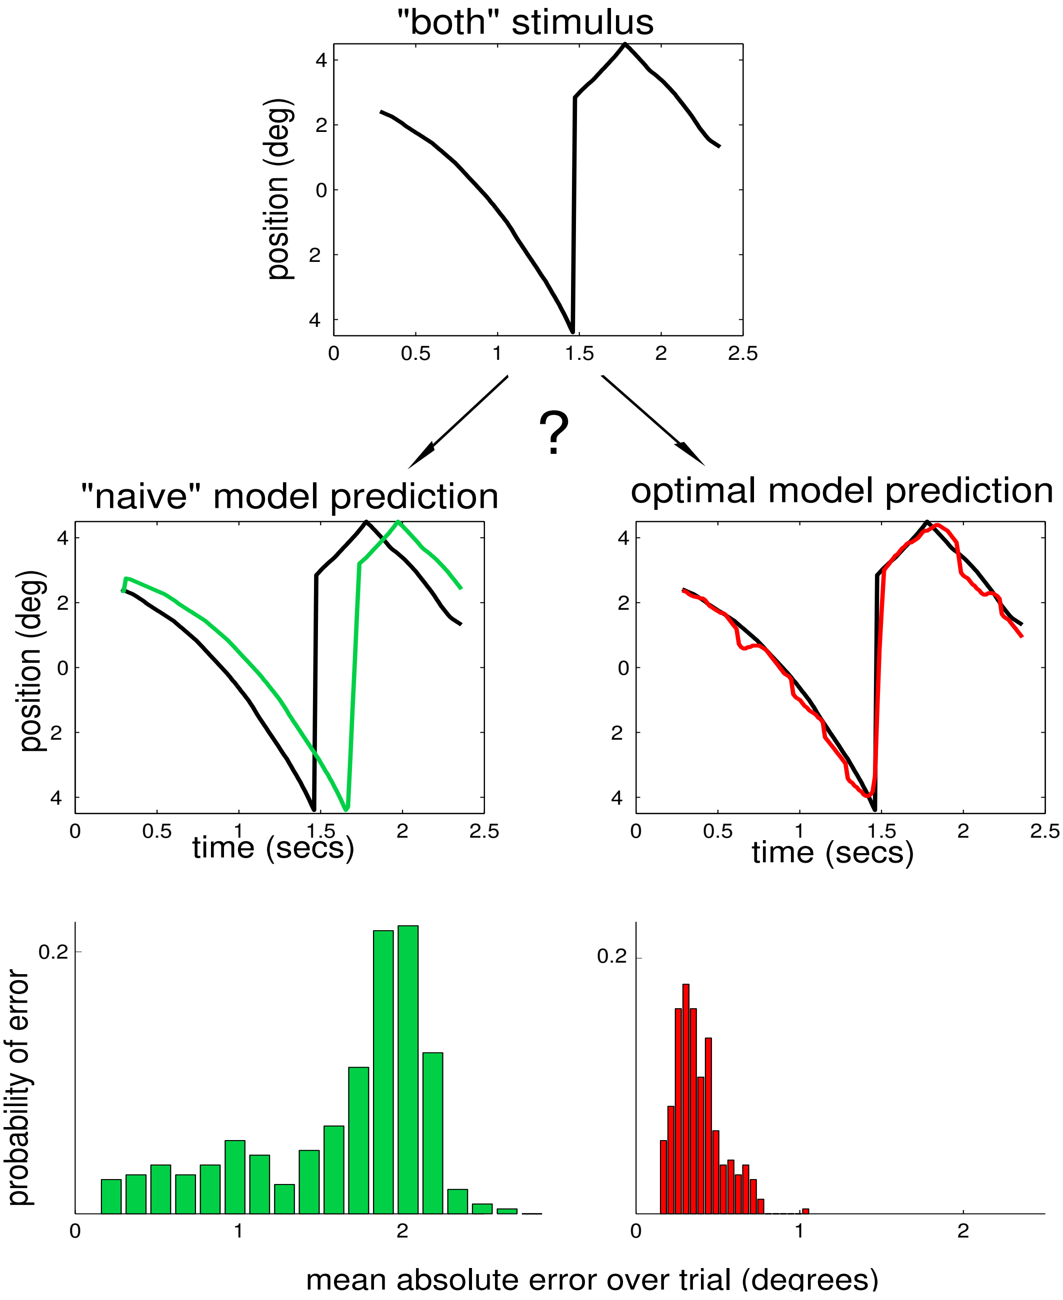

Figure 1: General experimental strategy; see Figure 4 for a comparison with the special case represented by the "saccade" condition. In each case, the subject tracks a simple visual stimulus whose position varies as a function of time, as represented in the top panel. We will compare the subjects' behavior to that of two models, as described in the text: the "naive" and the optimal. The comparison can be done qualitatively by examining sample paths (i.e., by looking at what the human subjects and the models do when presented with a given stimulus), or, more quantitatively, by examining the distributions of error incurred by the subjects and the models over several hundred trials.

is to find some 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 (unless stated otherwise, time is measured in steps of 14 ms, at the frame update rate of the monitor). We come up with eye paths by applying this rule recursively, updating our representation of the past behavior of the dot and the eye on each time step. The rule should be chosen to minimize the error function (the same error function the subject receives feedback about). This 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. The rest of this section lays out the mathematical details of how we want this rule to behave, how to find a good rule computationally, and how exactly we should compare the subject's performance to the optimum.

For simplicity, we only consider deterministic controls. By "optimal control," then, we mean a function

$$f^* : \mathcal{D}^- \times \mathcal{E}^- \mapsto \mathcal{E}_0, f^* \in \mathcal{F}, \tag{1}$$

where $\mathcal{D}^-$ and $\mathcal{E}^-$ denote the (product) spaces of possible past paths of the state of the world and the state of the control system, respectively ($\mathcal{D}$ stands for dot, and $\mathcal{E}$ for eye; $-$ and $+$ denote past and future behavior, respectively), $\mathcal{E}_0$ the space of possible present eye positions, and $\mathcal{F}$ the space of all admissible controls (i.e., those controls which the human oculomotor system can mimic, given the constraints under which the system operates). We discuss $\mathcal{F}$ in more detail below.

This f^* should ϵ-minimize an expected error term, something of the form

$$\int dp(d^+, e^+ | d^-, e^-) g(d^+, e^+), \tag{2}$$

where $p(d, e)$ is the probability measure induced in the natural way by f^* and $p(d)$, the distribution from which we draw the dot paths, and e^+, the future eye paths, are generated by recursively applying f^*: e^+ is given by $\{f^*(d_0, e_0), f^*(d_1, f^*(d_0, e_0)), \ldots, f^*(d_n, f^*(d_{n-1}, \ldots)), \ldots\})$. (Note that d_i and e_i refer to the state of the stimulus dot and eye at time i, and not just the dot and eye positions. In other words, e_0 tells us everything we need to know about the eye at the present time step: its position, its velocity, the time since the last saccade, and so on.) In this set of experiments, we used the L_p^p error, i.e., $g(d^+, e^+) \equiv \int_{\Re^+} (|d(t) - e(t)|^p)$, where p varied depending on the experiment. Since the motion of the dot is generated by a Markov chain, and is conditionally independent of the motion of the eye, $p(d^+, e^+ | d^-, e^-)$ may be replaced by something like $\prod p(d_n | d_{n-1})$.

3.2 Oculomotor model

We will be comparing the performance of an optimal f^* to that of our human subjects; therefore, we attempt to match the space of admissible controls, $\mathcal{F}$, to that of each individual subject. In other words, when we make our constraint space $\mathcal{F}$ explicit, we fix a model of the subject's oculomotor system. Our model is very simple: the system can either make saccades or pursue smoothly. We fix a minimum intersaccade interval, maximum saccade and pursuit velocities and accelerations, and a maximum saccade length, all fit to the behavior of the subject. Finally, we model the positive reaction time of the oculomotor system - the subject "receives" the state (i.e., location and velocity) of the dot some fixed reaction time RT after the dot was actually in that state, at time 0. Thus e^+, as explained below (2), should actually be $\{f^*(d_{-RT}, e_0), f^*(d_{-RT+1}, f^*(d_{-RT}, e_0)), \ldots, f^*(d_{-RT+n}, f^*(d_{-RT+n-1}, \ldots)), \ldots\}$. (In fact, we have two RT's - one for saccades and a shorter one for pursuit - but we'll ignore this to simplify the notation.)

3.3 Finite-horizon problem

Of course, we can not ask the system to look infinitely far into the future, and so we are led to treat a family of so-called finite horizon problems: fix N_geq1 and find f^* to ϵ-minimize

$$\begin{aligned}\int dp(d_{-RT},\ldots,d_{-RT+N-1},d_1,\ldots,d_N)\{|d_1-f(d_{-RT},e_0)|+&\\ |d_2-f(d_{-RT+1},f(d_{-RT},e_0))|+&\\ \ldots+&\\ |d_N-f(d_{-RT+N-1},f(d_{-RT+N-2},\ldots f(d_{-RT},e_0)))|\}.&\end{aligned} \tag{3}$$

N here represents the "horizon" - the number of time steps we'll allow ourselves to look ahead as we decide what to do at the present step. The greater the choice of N - the more distant the horizon - the more foresight we require of the system. We will see below that there is a direct correspondence between this depth variable and the computational complexity of the solution to the finite horizon problem at that depth. It should be clear, of course, that even when $N=1$, the optimal and "naive" models are fundamentally distinct.

3.4 Dynamic programming approximation to finite-horizon optimal rule

Unfortunately, as a simple local variational approach shows, the minimization of term (3) is generally not computationally feasible. Whenever the allowed range of f depends on the eye position e in a non-trivial way (e.g., when we put bounds on the speed at which the eye can travel), the problem becomes one of global functional optimization in a very high-dimensional space. The constraints are what make the problem at once mathematically interesting and computationally intractable. So we have to approximate our problem, by (roughly) reducing the dependence of f^* at the present point, (d_{-RT},e_0) on possible future points (d_{-RT+i},e_{i-1}): find instead $\{f_1^*,f_2^*,_l dots,f_N^*\}$ to minimize

$$\begin{aligned}\int dp(d_{-RT},\ldots,d_{-RT+N-1},d_1,\ldots,d_N)\{|d_1-f_1(d_{-RT},e_0)|+&\\ |d_2-f_2(d_{-RT+1},f_1(d_{-RT},e_0))|+&\\ \ldots+&\\ |d_N-f_N(d_{-RT+N-1},f_{N-1}(d_{-RT+N-2},\ldots f_1(d_{-RT},e_0)))|\}.&\end{aligned} \tag{4}$$

This $\{f_n^*\}$ can be exactly computed using dynamic programming [9] [10]: f_N^* is trivially computed, which makes it easy to compute f_{N-1}^*, and so on up the tree of (d,e) paths. It is difficult to assess the validity of this approximation, but we see that we should be safe for small N, at least (for $N=1$, of course, the solution of 4 is an exact solution of 3).

3.5 Comparison to data

Now, since none of the relevant spaces are convex, we have no guarantees of uniqueness of the optimal control. This is important: we can not, in general, directly compare what

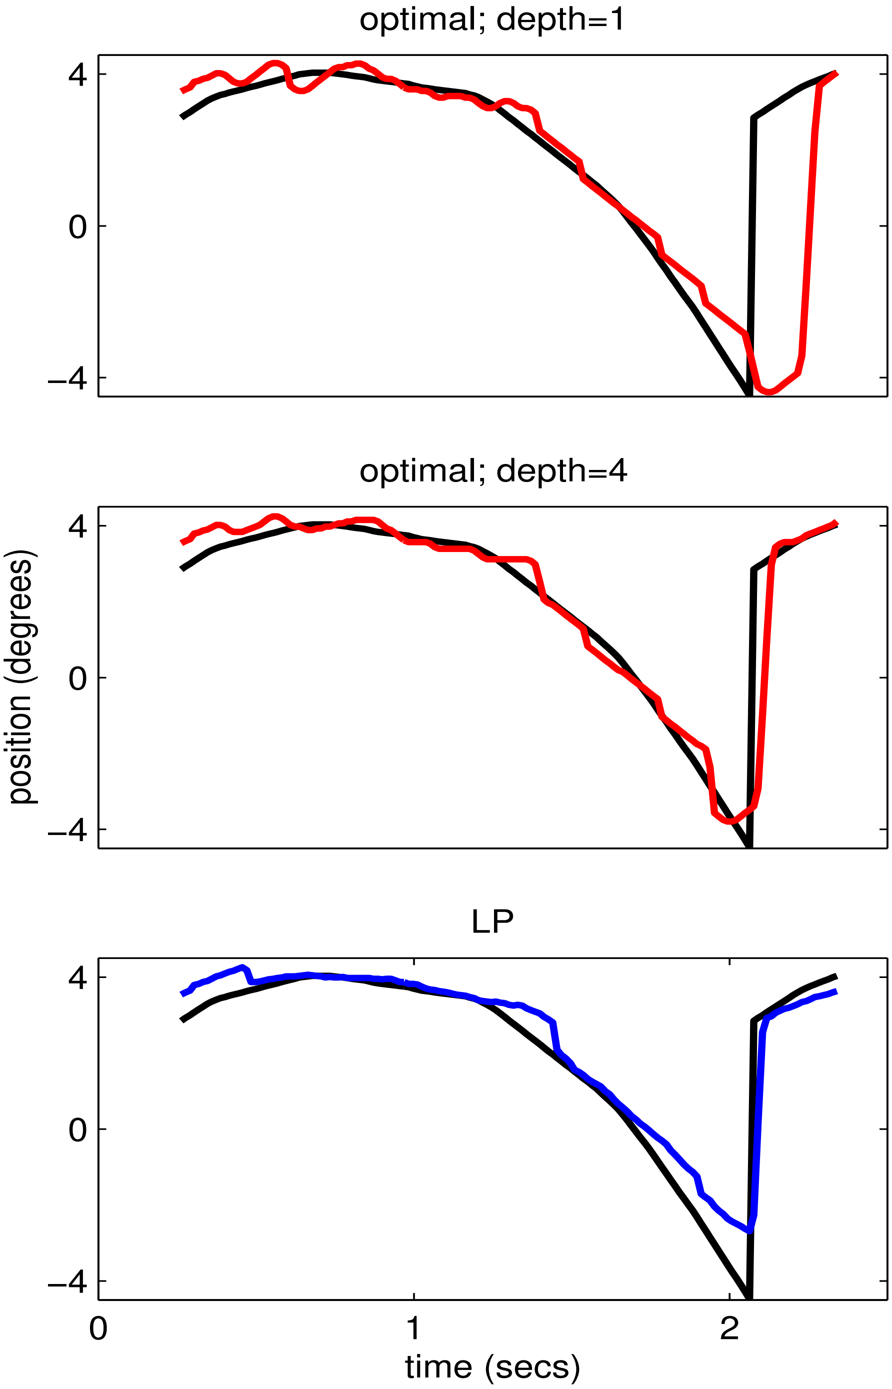

Figure 2: Example data trace (subject LP) and corresponding optimal model predictions. The top path was generated with the depth, or horizon, variable (see section 3.6 for details) set to 1, while the middle path had a depth of 4; subject LP's eye path appears in the bottom panel. Differences in the "foresight" of the two paths are obvious, especially at the site of the large rightward (upward in this graph) jump at the middle of the trial: the depth(1) path fails to anticipate this jump and makes a leftward (downward) saccade just before the jump, thus incurring a large error. The path of the subject's eye was typical in correctly anticipating the jump, slowing down well short of the far right of the screen in preparation for the well-timed, accurate saccade.

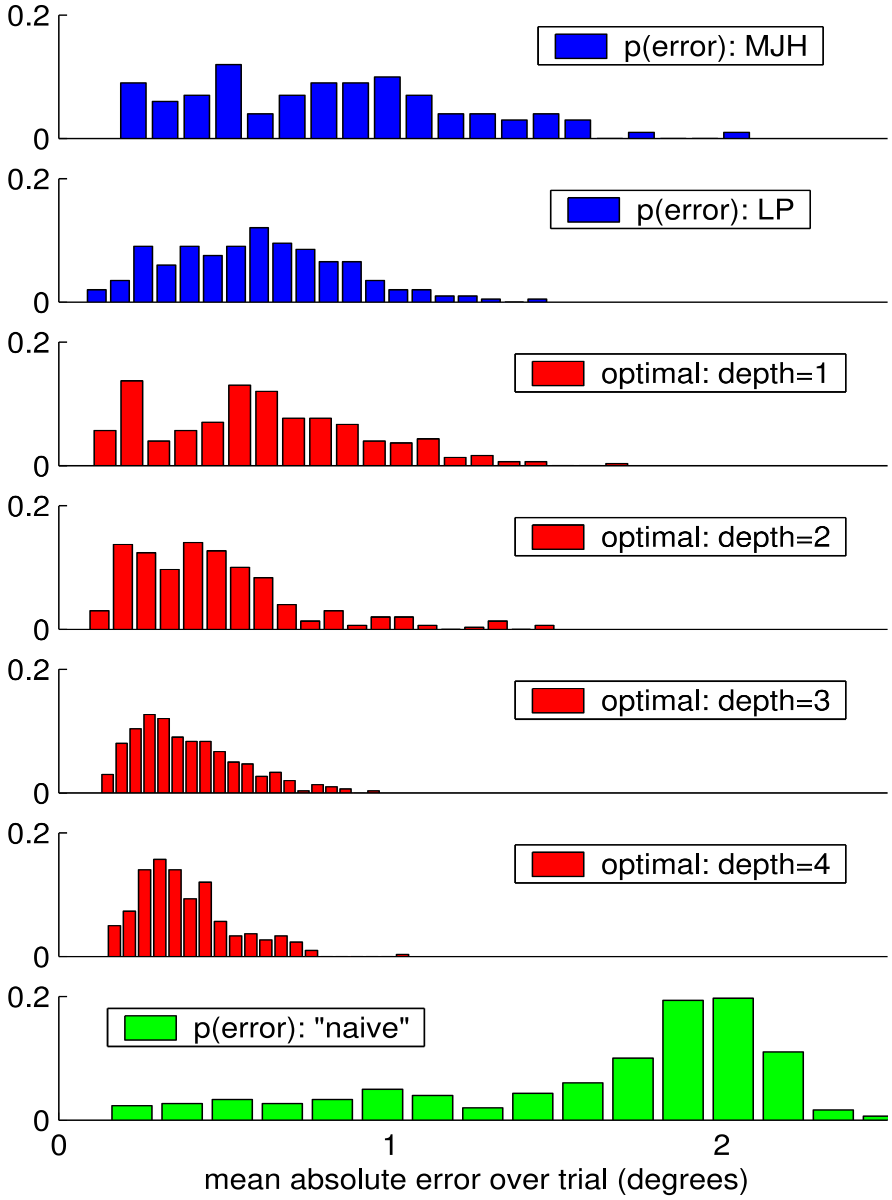

Figure 3: Error distributions for the experiment shown in Figure 2 condition. Here we can perform the comparison of the subjects' performance to that predicted by the optimal and "naive" models (as detailed in section 3.5). It is clearly easy to reject the "naive" null hypothesis; the data seem to agree best with the optimal (depth=1) model.

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 optimal "in the mean." In other words, we'll compare the distribution of errors for an optimal control and subjects over a large ensemble of trials. A similar distribution-based procedure can be used to test against the null hypothesis corresponding to the "naive" strategy.

3.6 "Saccade" stimulus condition

A great deal of the complexity of the above discussion stems from the sequential nature of the task the eye faces: the limited dynamic range of the oculomotor system forces us to act with foresight, and the computational cost of our solution to the problem increases exponentially with this foresight. It turns out to be possible to avoid these problems, by choosing a set stimuli which do not engage the constraints of the oculomotor system; as we noted above, without these constraints, the computational problem is trivial.

Thus, in the "saccade" stimulus condition, the movement of the dot was discrete - the dot jumped from point to point on the screen with no smooth motion intervening. Each pause was of the same duration, so that the temporal structure of each stimulus was completely predictable; the natural timestep, then, is given by the length of time between each jump, about 250 ms. This time interval is long enough to allow the eye to observe the new position of the dot and complete a saccade large enough to reach any part of the screen in this time. Thus we have effectively removed all dynamical constraints from the problem; to choose an optimal path, we only need to decide what's best right now, given the last observed dot position, without worrying about the impact our present choices might have on our distant future.

This condition presents several advantages. First, finding this well-defined best strategy is a computationally trivial task (we used the mean-square error metric in this set of experiments; therefore, we needed only calculate the conditional mean dot position at the next time step, given the current position, $E(D1|D0)$). The uniqueness of the optimal strategy allows us to compare the behavior of the subject to the optimum on a pathwise, trial-by-trial, basis. This permits a much more detailed test of the optimality hypothesis. Finally, with the mean-square error metric, deviations from optimality can be conveniently decomposed into a "bias" and "variance" term, corresponding to a systematic deviation, on average, from the optimal strategy and the variation around this mean, respectively. This sheds further light on exactly where and how the system might not be performing optimally.

In the "saccade" experiments discussed here, the visual stimuli had an especially simple statistical structure. The stimulus dot always started at the far right edge of the screen and jumped to the left on each time step. The size of each jump was determined by choosing a random number from $[0,1]$ according to a probability distribution, $p_\theta(x)$ and multiplying this number x by a scale factor, $D0$, corresponding to the distance from the current dot position to the far left of the screen. Thus the dot hops monotonically leftward, making smaller and smaller jumps as the dot approaches the left edge of the screen. The major point here is that $D1/D0$ - the scaled position of the dot on the next time step - was drawn from a probability distribution on $[0,1]$, $p_\theta(x)$, which was constant throughout the whole experiment. (Note that θ was drawn randomly at the beginning of the experiment, so that subjects had to learn p_θ, and not simply apply some precomputed rule to track optimally. In other words, θ acts as a "parameter," in the classical statistical sense, which must be learned from data.) This allows

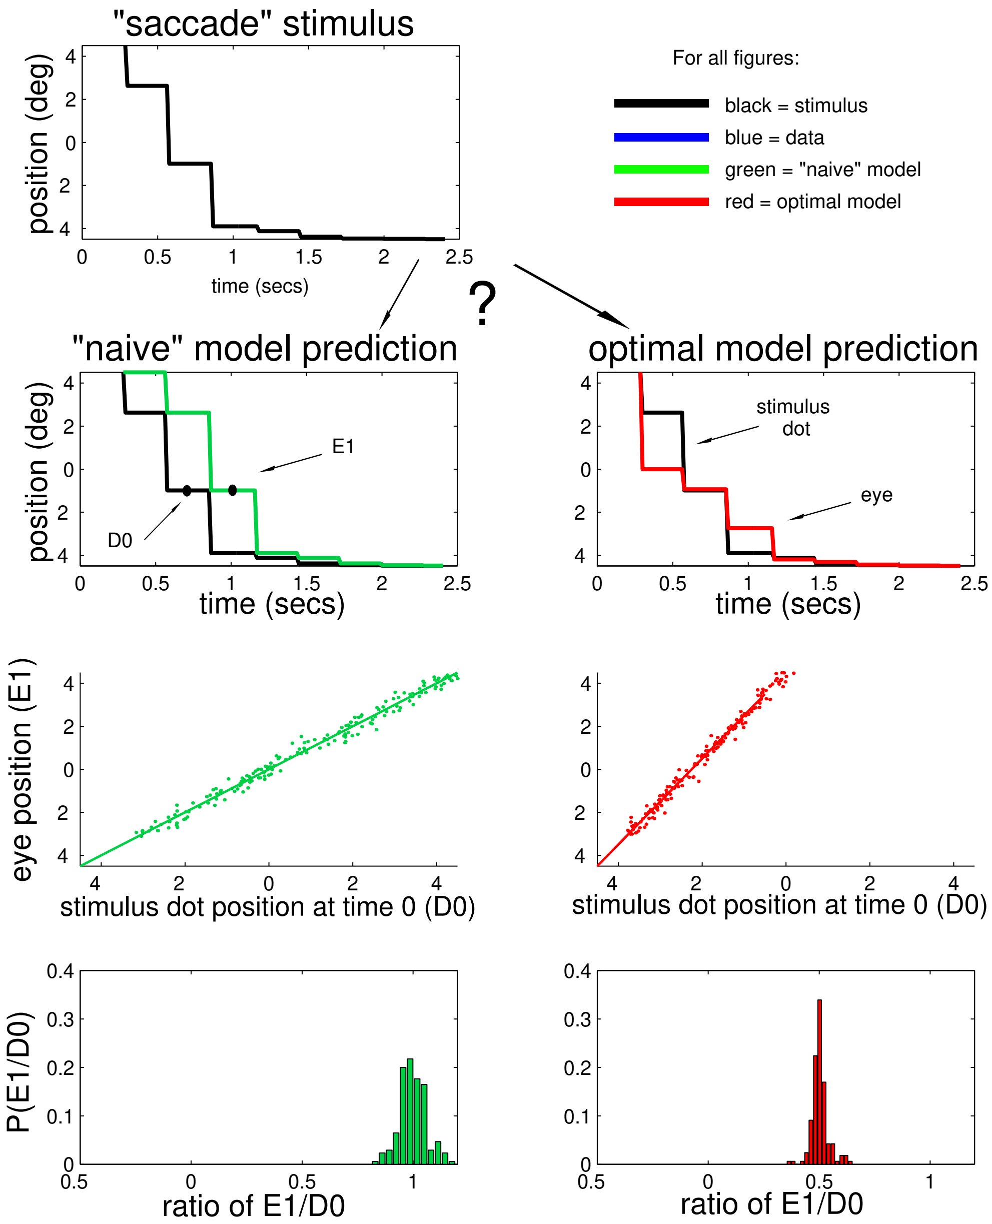

Figure 4: The experimental strategy for the "saccade" condition. As in the general case, there are two well-defined ways to perform this task: one can track the stimulus optimally, following the conditional mean of the stimulus given the information available at time 0, or one can follow the "naive" strategy of placing one's eye wherever the dot was last seen. To evaluate the predictions of these two tracking models, we examine the distribution of eye positions, $E1$ (observed at the midpoint of the static phase of the stimulus, each ~ 250 ms), given the position of the stimulus dot at the last time step, $D0$. More specifically, we collapse the data along the $D0$ axis by examining $p(E1/D0)$. The bottom two rows show the empirical distributions, $p(E1/D0)$, expected under each model. Note that the "naive" model predicts that the ratio $E1/D0$ should be concentrated at unity, where the optimal model predicts a clustering at $E_\theta(x)$, the mean of $E1/D0$ under θ.

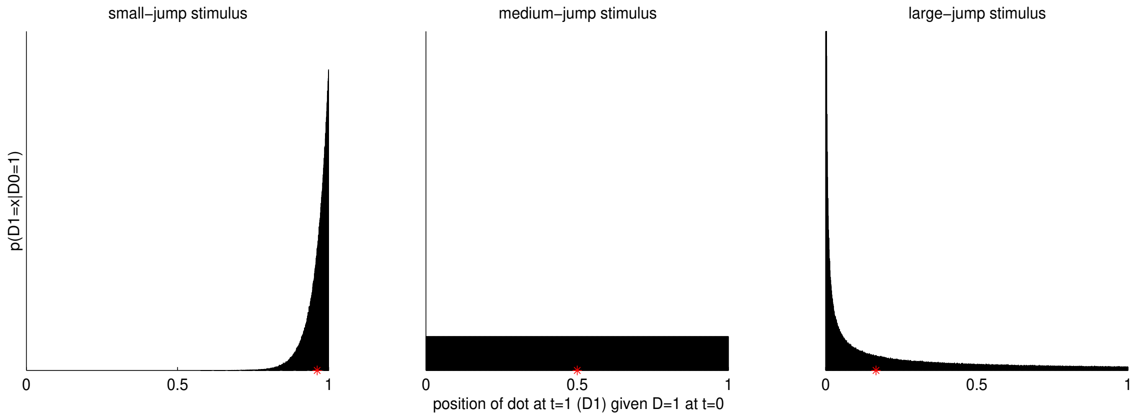

Figure 5: $p_\theta(x)$ for three values of theta examined in Results section. p_θ gives the probability distribution of the jump, properly scaled by the current position, $D0$. See Analysis (section 3.6) for details. The columns of Figures 6 and 7 are organized similarly, with the small-jump stimulus (smaller jumps are most probable, i.e., the mean of $p(D1/D0) > .5$) on the left and the large-jump stimulus on the left.

us to combine data over the whole trial to examine $p(E1/D0)$, instead of having to examine each conditional distribution $p(E1|D0)$ individually. In other words, we can treat our data $\{E1_i/D0_i\}$ as i.i.d. samples from a fixed probability distribution, thus greatly increasing the statistical power at our disposal. The hypothesis, now, is that if we change p_θ in such a way that the mean of x under p_θ shifts to the left or the right, we should see the mean of p($E1/D0$) shift accordingly. Under the null hypothesis, corresponding to the "naive" strategy, we should see no change, as $p(E1/D0)$ should be centered around 1 regardless of the statistical structure of the visual stimulus. This is all drawn out schematically in Figure 4.

4 Results

There is a clear trend in the data away from the null hypothesis represented by the "naive" strategy. In the "saccade" stimulus condition, for example, the conditional distribution of eye positions given the most recent dot position, $p(E1|D0)$, is clearly θ-dependent (i.e., the subject is clearly using the statistical structure of the stimuli to predictively track the dot), and thus distinct from that expected if the subject were following the "naive" strategy (Fig. 1). In some cases, the subject tracks the conditional mean (the optimal strategy) with negligible bias, although the variance of $E1/D0$ is fairly high in all cases (this variance increases with the expected jump size). We discuss possible sources of this variance below.

In the more general stimulus conditions, subjects performed about as well as the optimal (depth=1) model predicted. Eye position traces consistently showed predictive behavior of the type seen in Figure 2.

5 Discussion

We hope we have been successful in imparting the generality of the concepts here: nervous systems face difficult statistical problems, and before we investigate the mechanisms underly-

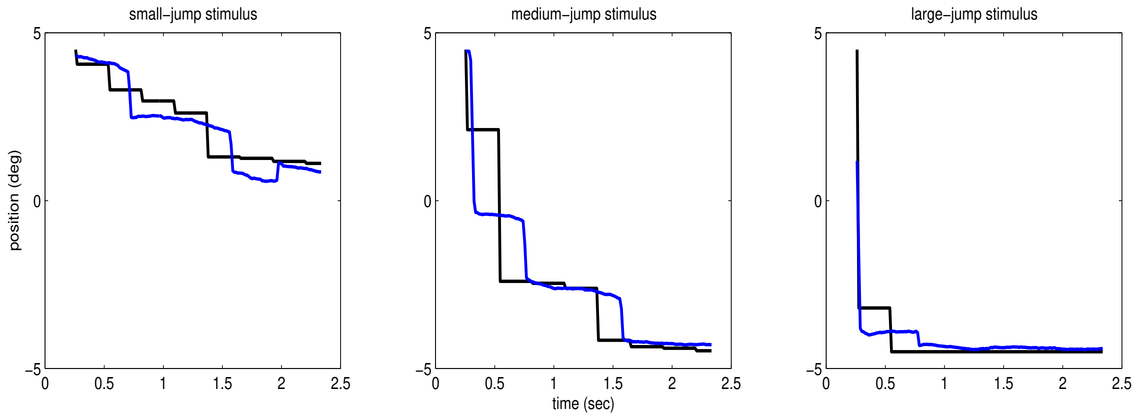

Figure 6: Some sample data traces from the "saccade" condition. Note the discrete nature of the eye movements, and that the frequency with which the eye is observed to the left (below) the stimulus dot. These traces are typical in that the medium- and large-jump stimuli reach the leftmost position - the large-jump stimulus does so within a couple jumps - but the short-jump stimulus usually only makes it about halfway across the screen by the end of the trial.

ing our performance in these types of situations, in can be extremely instructive to determine how good our performance really is, in a rigorous quantitative fashion. This approach forces us to think in a deep, mathematical way about the essence of the problem we're interested in, and leads experimentally to sharp constraints on the design of any model we might wish to build that might mimic the behavior of the real system. Thus, the "ideal observer" idea is useful in examining time-varying behavior; we believe that the optimal stochastic control approach presented here nicely unifies previous work examining predictive vs. reactive components of the oculomotor system in particular. We have provided some hints here that humans perform difficult, real-time, probabilistic tasks near optimally, and this is in nice qualitative agreement with the large body of work that exists on the sensory side of this general problem. Our results 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.

That said, it should also be clear just how small a step we have been able to take here. There is much work remaining to be done, both conceptually and empirically, before we have reached the lofty goals we set forth in the introduction, with our comparison to the arc of the work done in the retina. We'll structure our discussion of this future work by starting with the questions that are most easily addressed empirically, moving gradually towards a more theoretical level until we end on a purely conceptual note.

First, the near-optimal performance we observed here has to break down somewhere, as we make the task more and more difficult (e.g., by increasing the state-space of the Markov chain generating the stimulus). We would like to see exactly where this dropoff in performance occurs, as this would point to a strong bound on this system's computational resources. The space of Markov chains on position and velocity is quite large, of course (and even this large class does not exhaust the set of stimulus distributions we might be interested in); do we have any good way to explore this space systematically? One possible approach would be to exploit the natural ordering on the generative Markov chains given by the "foresight" parameter

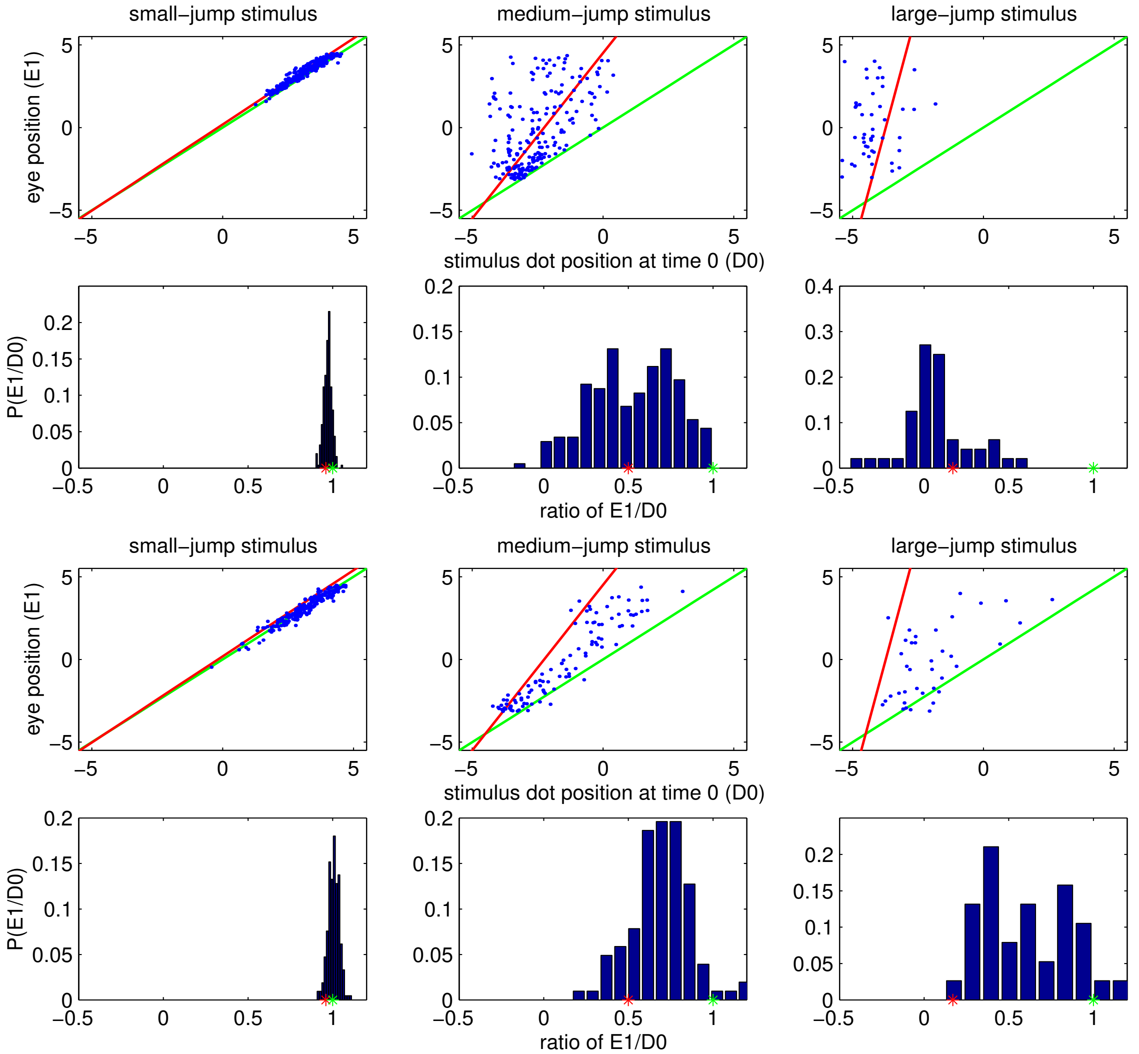

Figure 7: Summary "saccade" data, subjects LP (top two rows) and MJH (bottom two rows). The rows are laid out following the bottom two rows of Figure 4: the top row shows a scatter plot of eye positions versus the stimulus dot positions at the previous time step, while the bottom row marginalizes this data over $D0$ after dividing by the scale factor DO. Columns are as in Figure 5. The asterisks located on the independent axis of the histograms on the bottom row represent the mean $E1/D0$ ratio expected under the "naive" and optimal strategies, as in Figure 4. Note that the data appear to track the optimal mean quite well, although there is sizeable variance around this mean. (The apparent scarcity of data in the large-jump stimulus plot is a result of the fact that the stimulus dot jumps to the leftmost side of the screen much more quickly than in the other two stimulus conditions. In fact, more trials were run under the large-jump stimulus than under either of the other two conditions.)

defined in section 3.3: we can rank the "complexity" of different stimulus distributions by something like the slope of the mean error incurred by the optimal tracking strategy as a function of this foresight parameter (roughly, one Markov chain is more difficult to track than another if one requires more foresight to track it adequately). This ordering clearly depends on the constraints on the oculomotor system, as it should: it should be harder to track a dot that makes more unpredictable large jumps than one which moves relatively slowly and smoothly, for example. On a more trivial level, we can make the task more ethologically relevant by increasing the complexity of the display, by adding distractor stimuli, allowing the target to move in more than one dimension, making the motion of the target dependent on the behavior of the eye, and so on.

Before we make the task harder, perhaps we should address the deviations from optimal performance we have observed here. Two of the major noise sources in this context are quite well-known, and most likely are of simple physiological origin. The first can be called saccade endpoint noise, a variability in the final position of the eye following the saccade which is proportional to the length of the saccade (analyzed in a more general context in [11]). The second obvious source of noise is in the variability of the timing of saccades: small timing variability translates directly into large errors when the size of a (semi-)predictable jump is large. The timing of saccades near these large jumps had a standard deviation on the order of 50 ms. (However, this is most likely an overestimate of the timing variability of the system: note that, in the "non-saccade" conditions, the "correct" jump time is poorly defined - it is a random variable whose distribution depends on the last observed velocity of the stimulus, $\sim RT$ time units before the actual occurrence of the jump. In the "saccade" condition, on the other hand, the interjump interval is too short to rule out possible timing dependencies between jumps. An additional experiment, testing e.g. the ability of a human observer to synchronize the timing of his saccades to a predictable cue, is needed to address this question in more detail.) This leads us to the idea of refining the model of the oculomotor system that lies at the heart of the optimal analysis, in order to include noise sources similar to those seen in the data, just as we incorporated constraints on eye velocity, reaction time, etc. This would, of course, have the effect of shifting the optimal histograms in Figure 3 to the right, indicating that our human observers are operating with more foresight than is apparent.

The other major probable source of error is certainly as general, although probably less relevant to the more classical, non-probabilistic assays of oculomotor function. The attentive reader might have noticed a certain lack of discussion of any aspects of learning, which is clearly necessary for optimal performance. Clearly, even an optimal computer will perform suboptimally if it operates with an incorrect model of the world. We ignored the problem of learning the stimulus distribution in the analysis above: we assumed that the optimal system had explicit access to the Markov chain generating the stimulus. Thus, one might argue that any deviations from optimality that remain after the more trivial noise sources discussed above could be accounted for, not by a failure to compute optimally, but by a failure to learn optimally. We went to great lengths here to define optimal performance in a consistent manner: can we define "optimal learning" in a similarly rigorous sense? Unfortunately, this is much more difficult mathematically: we have to introduce some additional concepts from statistics (sufficiency, consistency, identifiability, priors) and learning theory (concept classes, learnability) to discuss this idea in a satisfactory way. It suffices to say that this is an interesting area for further theoretical research, but quantitative empirical applications might well prove too difficult for now.

As we mentioned above, we would like to end on a completely abstract level; our final thoughts will require much more theoretical development before any practical applications are possible, but we can at least hope that the reader finds the discussion conceptually useful. The question lies at the foundations of our work: what is the definition of a neural computation? Much of what we have done here could be considered as an empirical analysis of the computational capacity of the oculomotor system, but from a theoretical standpoint we have not even defined the objects of interest (computational resources, complexity, etc.). We hope that the utility of such a well-defined neural computational theory (along the lines of the Shannon-Kolmogorov or the Turing-Church-et al. theories) is obvious; perhaps this work will inspire future work in this fundamental area.

6 Acknowledgements

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 interesting discussions.

References

[1] Thomas Cover and Joy Thomas. *Elements of Information Theory.* John Wiley and Sons, 1991.

[2] A. Shiriaev, editor. *Collected Works of AN Kolmogorov: Information theory and the theory of algorithms.* Kluwer Academic, 1993.

[3] Fred Rieke, David Warland, Rob de Ruyter van Steveninck, and William Bialek. *Spikes: Exploring the Neural Code.* MIT Press, 1997.

[4] William Bialek. Physical limits to sensation and perception. *Ann. Rev. Biophys. Biophys. Chem.*, 20(S2), 1987.

[5] M. Pavel. Predictive control of eye movement. In Eileen Kowler, editor, *Eye Movements and their Role in Visual and Cognitive Processes*, pages 71–114. Elsevier, 1990.

[6] H. Crane and C. Steele. Generation-V dual-Purkinje-image eyetracker. *Applied Optics*, 24:527–537, 1985.

[7] Dario Ringach, Michael Hawken, and Robert Shapley. Binocular eye movements caused by the perception of three-dimensional structure from motion. *Vis. Res.*, 36:1479–1492, 1996.

[8] R. Carpenter. *Eye Movements.* CRC Press, 1991.

[9] Dimitri Bertsekas and Stephen Shreve. *Stochastic Optimal Control: The Discrete-Time Case.* Academic Press, 1978.

[10] W. Fleming and R. Rishel. *Deterministic and stochastic optimal control.* Springer-Verlag, 1975.

[11] Cristopher Harris and Daniel Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394:780–784, 1998.