

On the effects of signal acuity in a multi-alternative model of decision making

Tyler McMillen¹ and Sam Behseta¹

¹Department of Mathematics, California State University at Fullerton, Fullerton, CA 92834

²Your second affiliation.

Keywords: Leaky accumulator, drift-diffusion model, neural network, signal acuity, multi-hypothesis sequential test, sequential ratio test.

Abstract

We consider the effects of signal ‘sharpness’ or acuity on the performance of neural models of decision making. In these models, a vector of signals is presented and the subject must decide which of the elements of the vector is the largest. McMillen and Holmes (2006) derived asymptotically optimal tests under the assumption that the elements of the signal vector were all equal except one. In this paper we consider the case of signals spread around a peak. The acuity is a measure of how strongly peaked the signal is. We find that the optimal test is one in which the detectors are passed through an output layer that encodes knowledge of the possible shapes of the incoming signals. The incorporation of such an output layer can lead to significant improvements in decision making tasks.

1 Introduction

Time-controlled decision processes often involve a trade-off between speed and accuracy. Generally, accuracy improves as more time is allotted to the decision, but fewer decisions can be made in a fixed duration block of trials. In a setting of successive trials in which a reward is given for each correct decision, the goal is generally to make the most correct decisions possible. For this reason we consider the optimal test to be the one that maximizes the *reward rate* (RR), the number of rewards per unit time. If the subject is free to make a decision at any time, the optimal test is the one that minimizes the *mean reaction time* (MRT)¹ as a function of the *error rate* (ER). If the subject is constrained to make a decision at a given time then the optimal test is the one that minimizes the ER. However, other considerations such as energy costs associated with the decision process may make a sub-optimal test preferable.

In this paper we consider tasks in which an observer must identify 1 out of N possible stimuli with an appropriate response. We consider these in the context of the *leaky competing accumulator model* (LCAM) of Usher and McClelland (2001) in which N accumulators accumulate evidence and inhibit the rest in proportion to their values.

Most studies of multiple choice tasks consider the alternatives as equivalent and equi-distant from each other. This approach goes back to Hick (1952), and is implicitly assumed in models such as those considered by Usher and McClelland (2001). However, this is contrary to experience in which many choices involve alternatives that are not all equi-distant from each other. In this paper we consider a notion of ‘distance’ between alternatives as it relates to the LCAM. This is a step toward understanding phenomena such as Hick’s Law (Hick, 1952; Teichner and Krebs, 1974; Usher et al., 2002; McMillen and Holmes, 2006) and the fall of in performance when the number of alternatives is increased past seven (Miller, 1956).

Our main result is that performance in choice tasks can be improved if the knowledge of the shape of the possible incoming signals is encoded. We take as our canonical example the angle discrimination task. Untrained monkeys and humans can discern between angles 15-20 degrees apart; monkeys can be trained to improve their discernment

¹In this paper we use RT to denote the time taken by a computational process or dynamical model to render a decision. This is more commonly called *decision time* (DT) in the psychological literature.

to 4-5 degrees (Ghose et al., 2002). Improvement is orientation specific: when monkeys learn to discern angles within 5 degrees of a 45 degree angle, they are not able to achieve this level of accuracy for angles around 20 degrees. We will see that there are different ways in which accuracy can be improved. Adjusting the tuning curves in the neurons sensitive to angles is one way to alter the acuity in the signal related to decision making centers in the brain. However, Ghose et al. (2002) note that changes in the tuning curves are not sufficient to explain the improvements in performance observed in monkeys in the angle discrimination task. Indeed, the modification of the tuning curves over time in response to training is much less than one might expect, based on the improved performance of the subjects.

We will show how improvement can be achieved by encoding the shape of the possible incoming signals. Moreover, we will also show that increasing the acuity of the signal curve can actually have an adverse effect on performance. This could go some way to explaining why tuning curves have the shapes they do (why they are not, for example, more sharply peaked). Finally, we show how such knowledge can be encoded in a simple neural network.

This paper is organized as follows. In the next subsections we give a definition of acuity and present a notion of distance between alternative hypotheses, which depends on the spread and overlaps of the various signals. In section 2 we describe the LCAM, which is a neural network model for testing between n alternatives. In the two sections that follow this we examine how various possible tests perform in simulations. In section 3 we examine such tests in the free response protocol. In this type of experiment the subject is free to respond at any time, and the goal is to make as many correct responses as possible in a sequence of trials. We describe the (asymptotically) optimal tests for multiple alternative testing, and how it is applied in the LCAM. It is here where we show how the possible shapes of the incoming signals can be encoded in a matrix defined in equation (11). A different task protocol, the interrogation protocol, in which the subject must respond immediately after the onset of stimulus with fixed duration, is described in section 4. Following the description of the results in these two protocols, we explain in section 5 how the tests we have described can be implemented in a simple neural network. We conclude with comments and suggestions for further research in section 6. The bulk of the mathematical details are provided in appendices.

1.1 Acuity as a measure of signal sharpness

The ultimate goal of statistical testing of any multi-alternative decision making process is to determine, on the basis of information present in the detectors, which of the signals is the largest. Suppose that n signals S_1, \dots, S_n are presented to a subject. The problem is then to determine which of these is the largest. McMillen and Holmes (2006) have explored the case in which all signals are equal save one, which is a quantity a greater than the rest, i.e. for some j (the ‘correct’ signal), $S_j = S_0 + a$, and $S_i = S_0$ for $i \neq j$. Then the problem is to determine which of the signals S_i is $S_i = S_j$ from the data present in the detectors. Here we consider the more general case in which the signals take on the values of a function evaluated at certain points. During a trial a vector of signals $S = (S_1, S_2, \dots, S_n)$ is presented, and the goal, as before, is to determine which element of S is the largest. Generally there will be one of the S_i ’s, say S_j that is greater than the rest. But the signals might have different values apart from just S_j . For example, those signals close to j might be larger than those farther away. Suppose, for example, that the signals are discrete points of a Gaussian shape as in Fig. 1:

$$S_i = S_i^{(j)} = S_0 + a \exp \left[-\frac{(i-j)^2}{2\phi^2} \right], \quad i = 1, \dots, n. \quad (1)$$

Then the largest signal is S_j . We call ϕ^{-1} the *acuity*² in the signal. This gives a precise meaning to the notion that signals with a larger acuity are more strongly peaked. In the limit as the acuity approaches infinity, this reduces to the previous case, $S_i = S_0 + a \delta_{ij}$ (Kronecker delta), but when $\phi^{-1} < \infty$, the signals are “spread” around the maximum.

Let

$$S^{(j)} = \left(S_1^{(j)}, S_2^{(j)}, \dots, S_n^{(j)} \right)$$

be the vector of signals as defined above, so that $S^{(j)}$ is the vector of signals with maximum in the j th signal. Throughout we will use superscripts for vectors and subscripts for components of vectors. For example, in Fig. 1 the dots are the signals $S_j^{(8)}$. If all possible signals have this form, then the task is to determine which of the $S^{(j)}$ s is equal to S . If it is assumed *a priori* that the signals have the shape given above then this knowledge can be used in the decision process. That is, if we can reduce the space of signals to a subset of all possible signals, this knowledge can be employed in some of the tests we will consider below.




Figure 1: Signals with a Gaussian form. The acuity is $\phi^{-1} = 1$

The above type of signal might represent the signals processed by the brain upon viewing a bar on a screen. It is known that, for example in cats and monkeys, the firing rates of certain pools of neurons are affected by the angle of a bar on a screen (Henry et al., 1974; Hubel and Wiesel, 1968). The tuning curve, which traces the firing rate as a function of the angle, has a Gaussian shape centered at a given angle. In the angle discrimination task the subject is asked to determine the angle the bar makes relative to a certain axis. In the simplest task, the subject must determine if the bar is horizontal or vertical. The problem is made more difficult by increasing the number of possible angles, or by decreasing the angle between possible bars. It would be much easier to discern if a bar is horizontal or vertical bar than discerning between bars that might lie at angles 75 or 90 degrees. The task is made more difficult still if the possible angles are $0, \pi/8, \pi/4, \dots, \pi$. Then the excitation of centers registering an angle of $\pi/8$ will also excite those corresponding to 0 and $\pi/4$, but very little excitation will occur in those registering $\pi/2$. Clearly, 'spillover' of excitation into neighboring angles is directly related to the acuity. Fig. 2 shows the angles for bars in increments of $\pi/8$, and the respective signals S corresponding to these with a finite acuity, assuming that the signals have a Gaussian shape.

The difficulty of the task thus depends on the number of possible alternatives, and the 'distance' between the alternatives, a function of signal acuity.

We will use the angle discrimination task as our canonical example. However, the results are more general, and can be applied to many different decision making experimental tasks. Indeed, tuning curves have been used as a first order description for virtu-

²Also known as *precision* in statistical literature.

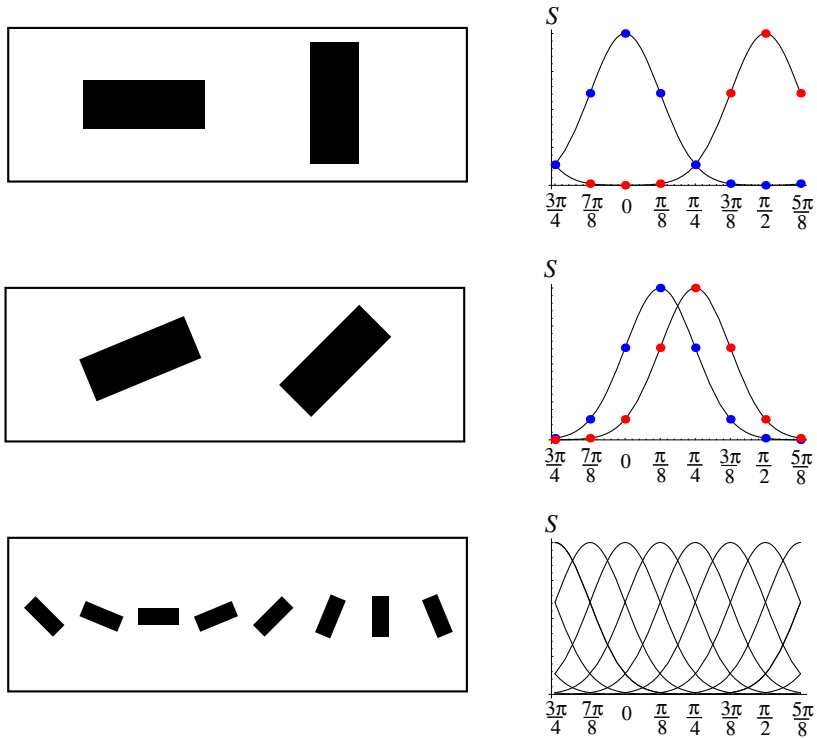


Figure 2: Bar angles (left panels) and corresponding signals (right panels) with acuity $\phi^{-1} = 1$.

ally all sensory systems. For example, tuning curves are used to determine the direction of wind in the cricket cercal system (Dayan and Abbott, 2001), as well as population-level analysis of movement direction in monkeys (Behseta et al., 2007; Georgopolous et al., 1986). As we will argue in section 6, the results we provide in this paper will extend beyond Gaussian functions for signals. One may recast the problem by allowing the signal peaks to be random variables, hence allowing for updating acuity through its posterior distribution.

1.2 Defining a notion of distance between alternatives

When presented with a single bar, it is more difficult to discern between the two bars in the middle panel of Fig. 2 than between those in the upper panel. The two in the middle panel differ by only $\pi/8$, whereas those in the upper panel differ by $\pi/2$. If the signal when presented is strongly peaked, i.e. has very large acuity, then the difficulty in discerning between the angles in the middle panel would be only slightly greater




Figure 3: Left panel: KL-divergence between hypothesis H_1 and H_j for inverse acuity $\phi = 0, 0.5, 1, 2$. The signals are as in eqn (4). Right panel: KL-divergence between hypotheses H_1 and H_5 as a function of the inverse acuity ϕ .

than discerning between those in the upper panel of Fig 2. In the limit of perfect acuity, i.e. when $S_m^{(i)} = S_0 + a \delta_{mi}$, there is no difference between the two cases—the distance between all alternatives is the same. The ‘distance’ between alternatives is therefore dependent not only on the distance between the means, but also on the spread of the densities and their overlaps. Furthermore, the distance between non-overlapping densities should not merely depend on their means.

A pseudo-metric reflecting this intuition is the Kullback-Leibler divergence (Kullback and Leibler, 1951; Kullback, 1987), defined for probability measures P_i as follows:

$$d(P_i, P_j) = E_i \left[\log \left(\frac{dP_i}{dP_j} \right) \right], \quad (2)$$

where $E_i[\cdot]$ is the expectation under the measure P_i . If we consider the signals in noise, $s_i = S_i + c \eta_i$, where η_i is a standard normal random variable, and let H_i be the hypothesis that the signal is $S^{(i)}$, then

$$dP_i(s) = \prod_{m=1}^n \frac{1}{\sqrt{2\pi c^2}} \exp \left[-(s_m - S_m^{(i)})^2 / 2c^2 \right].$$

Therefore, $d(H_i, H_j) = \frac{1}{2c^2} |S^{(i)} - S^{(j)}|^2$. (3)

In the case of perfect acuity $S_m^{(i)} = S_0 + a \delta_{mi}$, $d(H_i, H_j) = (a/c)^2$ for all $i \neq j$: all alternatives are the same ‘distance’ apart.

In the angle discrimination task angle π is the same as angle 0, angle $\pi/8$ is the same as angle $9\pi/8$, etc. Thus the angles can be represented on a circle. In this case, when representing the signals, we must account for the ‘wrap around.’ Consider the case of a

given signal shape, and the task is to determine where the peak of the signal is. If the signal is on a circle, and has a Gaussian shape, then the signals take the following form:

$$\begin{aligned}\hat{S}_i &= S_0 + a \exp \left[-(i - \hat{n})^2 / 2\phi^2 \right], \\ S_i^{(j)} &= \hat{S}_{\{i-j\}},\end{aligned}\tag{4}$$

where \hat{n} is the smallest integer larger than $n/2$ and $\{i - j\} = (i - j + \hat{n}) \bmod n$. In other words, \hat{S} is the signal with its peak in the middle and the $S^{(j)}$'s are circular shifts of \hat{S} . These are the signals in the right panels of Fig. 2, for instance. In this paper we consider only signals of this type. We do not consider, therefore, important effects such as edge effects that can arise when the signals are confined to an interval (see Lacouture and Marley, 1995, 2004), but leave this for future research.

Figure 3 shows the distance between H_1 and other hypotheses using the signals in (4). As expected, the distance between a hypothesis and its next nearest neighbor decreases with the acuity. However, the distance between hypotheses that are separated by other hypotheses can *increase* as the acuity decreases. In the right panel of Fig. 3 we plot the distance between hypotheses H_1 and H_5 as a function of the inverse acuity ϕ . Since there are hypotheses between these two, the distance between H_1 and H_5 increases as the spread of the signals increases, up to a certain point, and then starts to decrease. There is a critical value of ϕ at which the KL-divergence is maximal. This critical value is important because it allows for the calibration of acuity for hypothesis discrimination. If the two signals are sufficiently separated, an increase in the spread of each signal contributes more information with which to base a decision on. Consider the top panels of Fig. 2. We see that when deciding if the bar is vertical or horizontal, we can use the fact that if the bar is horizontal the signal peak will be at angle 0, and also the fact that the angles $7\pi/8$ and $\pi/8$ will have a slightly smaller excitation. Thus the spread in the signal contributes to the decision and makes it easier, so long as the signals do not overlap too much. The task of deciding between angles $\pi/8$ and $\pi/4$, in the middle panels of Fig. 2, is more difficult because the signals corresponding to these two angles have a significant overlap.

It should be easier to discern between alternatives that are farther apart than those that are closer together. Thus, we expect that the spread of the signals can be an aid to decision making, but only if one can use the knowledge of the shape of the signals.

Notice that as the acuity decreases, the distance of all the alternatives eventually decreases. Depending on the task and the decision algorithm employed, there may be a finite optimal acuity value. In what follows we will see that performance in specific decision tasks is optimized at a positive value of the inverse acuity ϕ , if the knowledge of the spread of the signal can be utilized in rendering a decision. (See Figures 7 and 9.)

In this paper we are primarily interested in the behavior of decision processes as the number of alternatives varies. In particular we will be interested to see in what cases we can recover the classical Hick's Law (Hick, 1952), which says that if the accuracy is held at a high level, the MRT increases logarithmically with the number of alternatives N :

$$\text{MRT} = A + B \log N . \quad (5)$$

Teichner and Krebs (1974) is a review of the literature up until that time on decision processes in multi-alternative choice tasks. The authors compile data on thousands of experiments. They show that Hick's Law does indeed hold, to a first approximation, for a wide variety of tasks. However, it is not universal. In some experiments the MRT is, rather, fairly constant as the number of alternatives increases. Part of the explanation for the inapplicability of Hick's Law in all cases could be that it does not hold in general when the acuity of the signals for the various alternatives is too small. In those models where the acuity is assumed to be infinite, Hick's Law follows naturally (see Usher et al., 2002; McMillen and Holmes, 2006). But there is no reason to expect *a priori* that it will hold in cases, like those considered in the present paper, where the acuity is small.

One problem in testing decision processes as the number of alternatives varies is that there are several ways to order the alternatives. Clearly the manner in which the number of alternatives is increased will affect the outcome. Yet, this issue has received scant attention in the literature. Rather, the implicit assumption has usually been that all alternatives are an equal distance apart. Here we describe two possible orderings, which we call the *inside-out* and the *outside-in* orderings. If one performs an experiment with 2 alternatives, and increase the number to 3, how shall one do this? The top and middle panels of Fig. 2 show two cases. In the top panel the alternatives are a maximal distance apart. The two alternatives in the middle panel are much closer.

Starting with the alternatives in the top panel and then adding an alternative between them would correspond to the outside-in ordering; starting with angles $\pi/8$ and $\pi/4$, as in the middle panel, and then adding angle $3\pi/8$ as an alternative corresponds to the inside-out ordering. In the outside-in ordering, we start with the alternatives a maximal distance apart, and then add alternatives so that the distance between the new alternative and the existing alternatives is maximized. In the inside-out ordering, we start with the alternatives a minimal distance apart, and then add alternatives so that the distance between the new alternative and the existing alternatives is minimized. These two orderings are illustrated in Fig. 4. Here the maximal number of alternatives is 8. As we proceed from $N = 2$ to $N = 8$ alternatives we add alternatives that maximize (outside-in) or minimize (inside-out) the distance between the existing and new alternatives. In Fig. 4 only $N = 2, 3, 4$ are shown—the pattern continues until $N = 8$.




Figure 4: Two orderings of alternatives.

2 The leaky competing accumulator model

We consider a model of n leaky, competing units, each of which integrates an incoming signal S_i or stimulus with additive noise, decays at rate k , and is inhibited by all others (Hopfield, 1984; Grossberg, 1988; Usher and McClelland, 2001). The state variable x_i represents the ‘activations’ (e.g. mean soma potential) of the i -th group of neurons and $f(x_j)$ the firing rate of the j -th group of neurons. The number of accumulators n can be larger than the number of alternatives N . The function f is typically taken to be sigmoidal, often using the logistic form:

$$f(x) = \frac{1}{1 + \exp(-4g(x - b))} = \frac{1}{2} [1 + \tanh(2g(x - b))] , \quad (6)$$

where g denotes the maximum slope of f , which occurs at $x = b$, although sometimes piecewise linear functions are used (Usher and McClelland, 2001). The lower and upper

limits (0 and 1 for (6)) obtained for small and large input levels model the fact that neural activities in general, and firing rates in particular, are bounded above and below. For a single layer of competing units, the resulting nonlinear stochastic differential equations (SDEs) take the form:

$$dx_i = \left(-kx_i - w \sum_{j \neq i} f(x_j) + S_i \right) dt + c dW_i, \quad i = 1, \dots, n. \quad (7)$$

Here the W_i 's are independent, identically distributed Wiener processes, with unit variance, representing random fluctuations in the signal, intrinsic accumulator noise and unmodeled inputs, and we assume that the S_i 's are deterministic and constant during each trial. Also, k is the rate of decay, w is the weight of inhibition between the accumulators and c is the variance of the noise. The above is referred to as the *leaky competing accumulator model* (LCAM). Note that in the absence of decay and inhibition ($k = w = 0$), eqn (7) is simply a drift-diffusion equation where S_i is the drift and $c dW_i$ is the diffusion term. Such drift-diffusion models with $k = w = 0$ are often referred to as *race models*; they represent the simplest SDE models for decision processes.

Substantial support for models of this type, and the SDEs that emerge from them, come both from fits of human behavioral data (e.g. Ratcliff (1978); Ratcliff et al. (1999)) and direct neural recordings in behaving monkeys (Roitman and Shadlen, 2002; Ratcliff et al., 2003; Mazurek et al., 2003). For instance, in monkeys the firing rates of neurons in the lateral intraparietal sulcus (area LIP) have been related to competing accumulators that approximate the drift-diffusion process in oculomotor tasks (Gold and Shadlen, 2000, 2001; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002). Similar findings have been reported for frontal structures responsible for controlling eye movements (Hanes and Schall, 1996). For a recent review of theoretical work on such SDEs, see Smith (2000). In the case of two alternatives the theory, at least for the linearized model, is fairly complete. See Bogacz et al. (2006) for a complete description; a brief summary of results is presented in the appendix of McMillen and Holmes (2006). Here we will deal with an arbitrary number N of alternatives.

Linearization of f yields the following form, with various interpretations of the parameters (see Brown et al., 2005). Noting that $-kw_i - w \sum_{j \neq i} x_j = (w - k)x_i -$

$w \sum_{j=1}^n x_j$, we have

$$dx_i = \left(\lambda x_i - w \sum_{j=1}^n x_j + S_i \right) dt + c dW_i, \quad i = 1, \dots, n, \quad (8)$$

where $\lambda = w - k$ is the difference between inhibition and decay. Throughout the remainder of this paper we will deal with (8), where we have linearized the input or activation function (see Usher and McClelland, 2001). This linearized model is attractive for its mathematical tractability. Indeed, those results that can be obtained rigorously are mostly limited to the linearized model. The linearization is often justified by noting that the function f in (6) is nearly linear in the region near $x = b$, and that f can be approximated very closely by a piecewise linear function. The main qualitative difference between (7) and the linearized model (8) is the absence in (8) of upper and lower bounds on the inhibition term. Using a sigmoidal f as in (6) means that the $f(x_j)$ terms in (7) saturate at $f(x_j) = 1$ when x_j is large and are bounded below by $f(x_j) = 0$. This has the effect of attenuating the effects of inhibition on the accumulators. In the linearized model (8) the inhibition of the other accumulators on x_j is simply a linear function of the values of these other accumulators, and hence is unbounded. This provides us with a model that is a good approximation for those cases in which the accumulators remain in the linear range near $x = b$. It allows us to derive the results in the remainder of this paper. One line of future research is to analyze those situations in which results for the linear and nonlinear models diverge. This has been done to some extent in the case where the signals are assumed to have infinite acuity. For example see Usher and McClelland (2001); Brown et al. (2005); McMillen and Holmes (2006) for comments on the nonlinear system in which inhibitory inputs and stimuli are passed through a rate-limiting function f . We note that in most cases the linear and nonlinear models agree qualitatively in the ‘moderate’ ranges, i.e. where the number of alternatives is not too large and where the error rate is not too small. The linearized model (8) that we deal with in the remainder of this paper gives us a starting point from which we may derive rigorous conclusions that will hold, at least qualitatively and in most cases, in the nonlinear model.

The subject’s task is to determine which of the signals S_i is the largest. Depending on our knowledge of the possible ‘shapes’ of the incoming vector of signals, we can improve on performance over simply choosing the option corresponding to the largest

signal. Conventionally, each trial starts with initial data $x_i(0) = x_{i0}$ and ends either when the first accumulator passes its preset threshold b_i (for an absolute test under the free response protocol), or at a fixed time $t = T$ (in the interrogation protocol), in which case the accumulator with largest $x_i(T)$ wins. For unbiased trials with equal prior probabilities of H_i , $\pi_i = 1/n$, we conventionally take $x_i(0) = 0$ and set $b_i = \theta$, $\forall i$.

3 The free response protocol

Under the free response protocol, subjects may choose at any time after stimulus presentation and must therefore determine when sufficient knowledge has been gathered to make a decision. In the context of the LCAM (8) there are several different types of response criteria. The simplest is to set a threshold θ , and choose i if the i th accumulator crosses θ before any of the others. We refer to this as an *absolute* test, since it does not account for the magnitudes of other accumulator states. Tests which account for differences between states are referred to as *relative* tests. One such test is the *max-vs-next* test, which chooses i when the difference between x_i and the next greatest x_j reaches threshold. We will also see that the implementation of the asymptotically optimal test in the LCAM requires that we multiply the vector of accumulators by the matrix of possible signal vectors before comparing them.

There are advantages and drawbacks to the different types of tests. Generally, relative tests perform better in that they deliver lower MRTs for a given ER, but they are more complex to implement. The (asymptotically) optimal tests also require a matrix multiplication, as well as the computation of the matrix itself. In some situations energy costs or other considerations might dictate a sub-optimal test. In the following sections we examine the asymptotically optimal tests and compare these with absolute and relative tests.

3.1 Asymptotically optimal tests

Unlike in the case of two alternatives, for a test between multiple alternatives there is no optimal test. However, as explained below, the *multi-sequential probability ratio test*

(MSPRT) is an asymptotically optimal test, in the sense of minimizing mean reaction times in the limit as the error rate approaches zero.

Here we give an overview of the MSPRT and explain how it is applied to the LCAM. We confine the technical details to Appendix A. Sequential tests are tests on the histories of random variables. So suppose X_t is a vector-valued random variable, and that one of N hypotheses H_i must be chosen. Let π_i denote the prior probability of H_i , and let $X^t = \{X_s \mid 0 \leq s \leq t\}$ be the set of all values of X_s observed up to time t .

There are two differently-structured tests which give the same asymptotically optimal results (Dragalin et al., 1999): one determines when the posterior probability of a correct response crosses a threshold; the other test considers the ratios of posterior probabilities. These are called the δ_a and δ_b tests, respectively. In the case of two hypotheses both tests reduce to the optimal SPRT; in the case of $N > 2$ hypotheses, both tests achieve the same asymptotic optimality. These tests are defined in terms of the posterior probability of hypothesis H_i , and the generalized likelihood ratio between H_i and the remaining hypotheses. These are, respectively,

$$\Pi_i(t) = P(H_i \text{ is true} \mid X^t), \quad \text{and} \quad L_i(t) = \frac{P(H_i \text{ is true} \mid X^t)}{\max_{j \neq i} P(H_j \text{ is true} \mid X^t)}, \quad (9)$$

In other words, $\Pi_i(t)$ is the probability that hypothesis H_i is true given the entire history of the observed random variables from the beginning of the trial until the present time t .

In these sequential tests one sets constant thresholds a_i and b_i , $i = 1, \dots, N$, for each alternative. Then the δ_a test stops once a $\Pi_j(t)$ crosses a_j , at which time hypothesis H_j is accepted, and the δ_b test stops once an $L_j(t)$ crosses b_j , at which time hypotheses H_j is accepted. Thus the δ_a test accepts hypothesis H_i when the posteriori probability of hypothesis H_i crosses a threshold, and the δ_b test accepts hypothesis H_i when the ratio of the posteriori probability of hypothesis H_i to the posteriori probability of the next most likely hypothesis crosses a threshold. The thresholds a_i, b_i can be adjusted to achieve a given error rate: increasing the thresholds a_i and b_i will result in a lower probability of selecting hypothesis H_i incorrectly. The thresholds will generally depend on the desired accuracy of the test and the prior probabilities π_i of the different alternatives. In the case where all hypotheses are equally likely, $\pi_i = 1/N$, the thresholds will be the same, i.e. $a_i = a$, $b_i = b$.

The δ_a and δ_b tests are *asymptotically optimal* in that they give an expected value of the stopping time that is asymptotically the infimum of the stopping times of all sequential tests satisfying the same error bound (see Dragalin et al., 1999, Theorem 4.2). In other words, for a vanishingly small error rate these tests return a decision in the shortest time, on average.

Now consider the LCAM (8) in this context, with incoming signals taking the form of n possible vectors $S^{(i)}$. We make the following assumption:

$$\sum_{m=1}^n S_m^{(i)2} \text{ and } \sum_{m=1}^n S_m^{(i)} \text{ do not depend on } i. \quad (10)$$

The above assumption is true if, for example, the signal shape is the same regardless of which alternative is true and is a reasonable assumption in many different situations. Throughout the remainder of this paper we will assume that the assumption (10) holds. The calculations when this assumption does not hold are presented in Appendix B. We note that when the assumption (10) does not hold, the decision is biased toward the hypothesis corresponding to the largest sum of squares of the elements of the signal vector.

Turning now to the asymptotically optimal tests in the LCAM, the task is to determine, from the values of the accumulators $x_i(t)$, which of the signals is largest, or, equivalently, which of the vectors $S^{(i)}$ is the actual signal. We let H_i be the hypothesis that the signal is $S^{(i)}$. Then $\Pi_i(t)$ is the probability that the signal vector is $S^{(i)}$, given the history of the accumulators. The derivation of Π_i and L_i in the case when assumption (10) holds is in Appendix A. Let A be the $N \times n$ matrix of alternative signals ($A_{im} = S_m^{(i)}$):

$$A = \begin{pmatrix} S^{(1)} \\ S^{(2)} \\ \dots \\ S^{(N)} \end{pmatrix}, \quad (11)$$

and let

$$\mathbf{y} = \frac{1}{c^2} A \left(\mathbf{x} - \lambda \int_0^t \mathbf{x}(s) ds \right) + \log \boldsymbol{\pi}, \quad (12)$$

where $\log \boldsymbol{\pi}$ is the vector whose i th entry is $\log \pi_i$. Then

$$\log \Pi_i = y_i - \log \sum_{j=1}^n e^{y_j}, \quad (13)$$

$$\log L_i = y_i - \max_{j \neq i} y_j. \quad (14)$$

The asymptotically optimal tests, therefore, require a matrix multiplication of the accumulators. As we will show in section 5, this can be achieved by a second neural layer.

Note also that the asymptotically optimal tests depend only on the *difference* λ between inhibition and decay, and not on their absolute values. If $\lambda = 0$ and $S_m^{(i)} = a \delta_{mi}$, as in the case of infinite acuity (the signal is in one channel only), then $y_i = a/c^2 x_i$, and the tests are tests on the accumulators themselves. In that case, the δ_b test is a ‘max-vs-next’ test, which chooses hypothesis H_i when the difference between the i th accumulator and the next largest accumulator crosses a threshold. This case is studied in McMillen and Holmes (2006). Generally, though, the acuity is not infinite, and hence the matrix multiplication must be done to achieve asymptotic optimality.

We pause here to make an observation about the quantities (13-14) in the MSPRT, which are calculated from \mathbf{y} in (12). If $\lambda = 0$ and the prior probabilities are all equal ($\pi_i = 1/N$), then

$$\log L_i = \frac{1}{c^2} \left((A\mathbf{x})_i - \max_{j \neq i} (A\mathbf{x})_j \right). \quad (15)$$

Thus, in this case, the optimal test requires only a comparison of the values of the accumulators multiplied by the matrix A . However, if $\lambda \neq 0$, the calculation of Π_i and L_i requires an integration of the entire history of the accumulators. As we will show in section 5, the matrix multiplication can be easily done by a feed-forward neural layer. But it is not as easy to build a neural layer that keeps a running value of the integral $\int_0^t \mathbf{x}(s) ds$. Thus, it is simpler to build optimal neural decision networks for an LCAM where $\lambda = 0$. The case where $\lambda = 0$ is called the ‘balanced’ case, as it occurs when inhibition and decay are exactly balanced. Moreover, it is known that the absolute test on the accumulators achieves the best results when $\lambda = 0$ (Bogacz et al., 2006; McMillen and Holmes, 2006). Thus, the balanced case is preferable for several reasons.

Turning back to the MSPRT, we note that the matrix multiplication does come with some costs, and requires learning the matrix for a given task. This will be an advantage

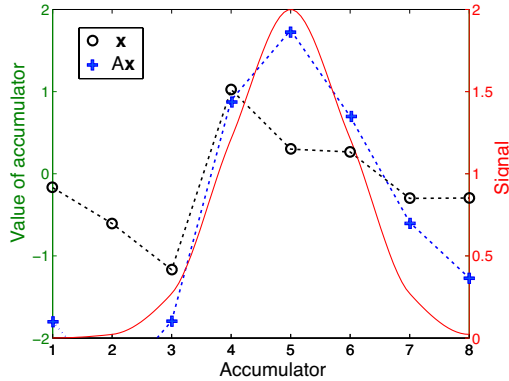


Figure 5: Values of accumulators (circles) and their transformed values (+’s), along with the underlying signal (right axis), for a single simulation. In this example, x_4 crosses the threshold before x_5 , but $(Ax)_5$ crosses before any of the other elements of Ax .

only if more correct decisions can be made by considering Ax instead of x when making a decision. Thus there is an advantage to multiplying the accumulators by the matrix A if and only if, on average, $(Ax)_i$ is the first to cross the given threshold more often than is x_i when H_i is the correct hypothesis. In Fig. 5 we show how this can occur. In this figure is shown the states of 8 accumulators a time t after a trial has begun in which H_5 is the correct hypothesis. We see that the fourth accumulator x_4 is larger than the rest. Thus, in the absolute test on the accumulators, the fourth hypothesis would be chosen, which would be incorrect. However, if we form A as in (11), then the fifth component of Ax is larger than the other components, and hence a test on the values of the transformed accumulators Ax chooses correctly.

3.2 Comparison of tests

In this section we compare different tests in order to explore the effects of pre-multiplying the signal by the matrix of possible signals, and also to see what advantage is gained performing the optimal test. In this section we consider the balanced case $\lambda = 0$, in which case y_i in (12) is a multiple of $(Ax)_i$ plus a constant. Thus, throughout, $w = k = 1$, unless otherwise noted. Results for different values of the parameters a, c, w, k are qualitatively similar for a wide range. Here we present results for the balanced case for

simplicity, but more importantly, because, as we mentioned, the absolute tests achieve the best results in the balanced case. Later, when we consider the interrogation protocol, we present results for a range of λ (see Fig. 8). In this section we present only the balanced case, while noting that in this, as in the interrogation protocol, the difference between the optimal tests and the absolute test on the accumulators is even more pronounced in the unbalanced case when $k \neq w$.

The four tests we consider, then, are as follows. The tests on the balanced model ($\lambda = 0$) choose hypothesis H_i when the following quantities cross a threshold:

$$\begin{aligned}
(A\mathbf{x})_i - \max_{j \neq i} (A\mathbf{x})_j & \quad (\delta_b \text{ test}) \\
x_i - \max_{j \neq i} x_j & \quad (\text{max-vs-next}) \\
(A\mathbf{x})_i & \quad (\text{absolute in transformed coordinates}) \\
x_i & \quad (\text{absolute})
\end{aligned} \tag{16}$$

The thresholds for the different tests will be different for a given accuracy level. The performance of tests is determined by how long it takes, on average, to make a decision at a given accuracy level. Therefore, we are interested in the MRT for a given ER. To calculate this, we adjust the thresholds for the individual tests until the ER is as desired, and then compute the MRT. All results are obtained through Monte Carlo simulations. (See Appendix D for a brief explanation of how to simulate the kinds of SDEs encountered in the present paper using Matlab. For a more general introduction, see Higham (2001) and Iacus (2008) for simulating SDEs in Matlab and R, respectively.) We will present results for only the four tests, since results for the δ_a and δ_b tests are indistinguishable. There is also a slight abuse of notation here relating to the selection of hypotheses. The hypothesis numbering does not necessarily correspond to the accumulator numbering. For example, suppose we number accumulators from 1 to 8 corresponding to the angles in the bottom panel of Fig. 2. That is, accumulator 1 is sensitive to angle $3\pi/4$, accumulator 2 to angle $7\pi/8$, etc. In a test on two hypotheses and eight accumulators, for example if we compare angle 0 to angle $\pi/2$ as in Fig. 2, hypothesis H_1 corresponds to angle 3 and H_2 to angle 7. In the third test in the list (16), the matrix A contains only the two rows corresponding to the two alternatives, and hence if $(A\mathbf{x})_1$ crosses the threshold before $(A\mathbf{x})_2$, we choose H_1 (angle 3), and vice versa. However, in the absolute test, we choose H_1 (angle 3) if x_3 crosses before x_7 ;

the other accumulators evolve but their values are not considered in making a decision. Comparisons of these tests are seen in Fig. 6.

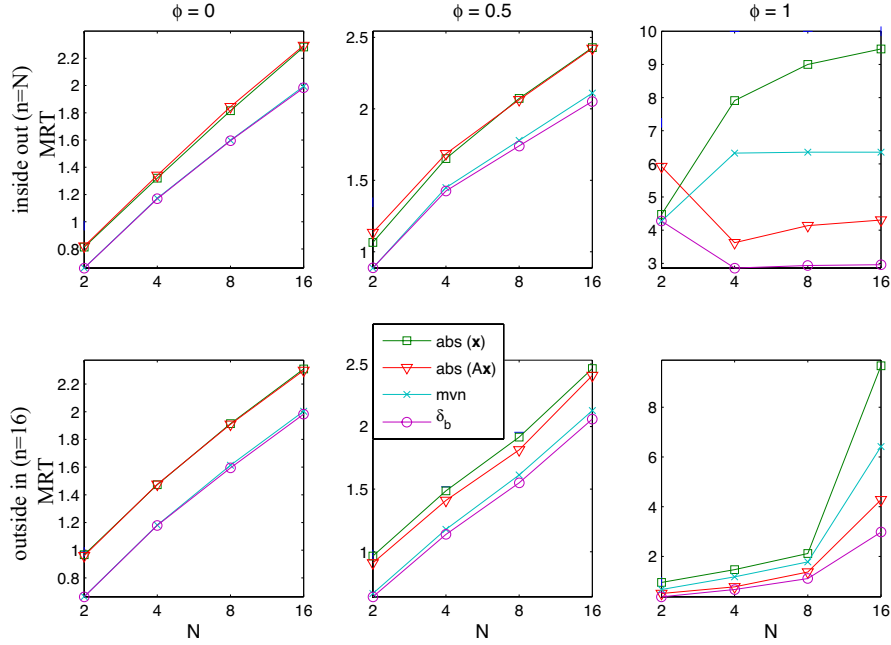


Figure 6: Comparison of MRT, as a function of the number of alternatives N for different values of the acuity. The ER is fixed at 0.05, and $a = 2$, $c = 1$, $w = k = 1$. The top panels show results for the inside-out ordering where the number of accumulators is the same as the number of alternatives ($n = N$); the bottom panels use the outside-in ordering and the number of accumulators is fixed at $n = 16$. When $N < 16$ all of the accumulators evolve but only those corresponding to the alternatives are considered when making a decision.

The results in Fig. 6 show that the effects of the type of ordering are only seen—as expected—at smaller values of the acuity, but that these differences are very marked when the acuity is small (ϕ is large). The behavior of the MRT as a function of the number of alternatives is qualitatively different depending on the ordering of alternatives when $\phi = 1$. We see that Hick’s Law holds, at least approximately, for both orderings when the acuity is very large, or when the signal is strongly peaked. For small acuity, Hick’s Law does *not* hold when the inside-out ordering is used, and it only holds for small numbers of alternatives when the outside-in ordering is used.

Perhaps the most striking fact about these results is seen in the upper right panel of Fig. 6, where the MRT declines in the inside-out ordering as the number of alternatives increases, when the acuity is small (ϕ is large), and when the tests that use knowledge of the shape of the signals are used. This is in fact, intuitive, since if the spread of the signals is large the difference between the signals is small. For tests that don't use the shape of the signals, this simply makes the test more difficult. However, if we know the possible shapes, then additional signals gives us more information about the shape, and we can use this to match the signal, in the manner of using a template. This is accomplished through the multiplication of the accumulators by the matrix A .

The max-vs-next test is the best one can do without knowledge of the signal matrix A . The differences in performance between the tests grow with the spread of the signal, or as the acuity declines. For an infinite acuity there is not much difference between the tests, but when the acuity is small the best tests return results much more quickly than the worst, for the same error rate. We see, for example, in the upper right panel of Fig. 6, that when there are 8 accumulators and the acuity is $\phi^{-1} = 1$, the time taken for the best test to achieve an ER of 0.05 is 1/3 of the time taken by the worst test to achieve the same ER. There is a significant advantage, then, to multiplying the accumulators by A , when the acuity is small. We see, in the outside-in ordering, for small ϕ , the max-vs-next test outperforms the absolute test on the transformed coordinates (i.e. choose i when $(Ax)_i$ crosses a threshold). But, as the acuity decreases (ϕ increases), the test on Ax performs better than the max-vs-next test. If the signals have a significant spread, it is much more beneficial to perform the pre-multiplication of signals by A than to implement a relative test.

Figure 7 shows the effects of the acuity on different tests. In the left panel we hold the acuity fixed at $\phi^{-1} = 1$ and vary the number of accumulators and alternatives. In these simulations the optimal δ_b test is used, but the results are qualitatively similar for the absolute test on Ax , that is, for tests in which knowledge of the shape of the signals is used. We see that performance increases with the number of accumulators, since more accumulators means that we have more information. In order to illustrate the effects of acuity on performance, we consider a test on two alternatives in the right panel of Fig. 7. Here the test is either between alternatives 1 and 2 or between alternatives 1 and 5. Eight accumulators are involved in the decision process. The MRT is plotted




Figure 7: MRT at $ER=0.1$. Left panel: acuity is fixed and the number of alternatives varies. Right panel: only two alternatives are considered (either 1 and 2 or 1 and 5), and the acuity is allowed to vary. The optimal δ_b test is used with $a = 2$, $c = 1$, $w = k = 1$.

against ϕ for these two cases. When these alternatives are close together (1 vs 2), an increase in ϕ causes an increase in the MRT. However, when there are accumulators ‘between’ the alternatives, when we compare 1 vs 5, an increase in ϕ actually leads to a decline in the MRT. In fact, there exists a finite optimal acuity value at which the MRT is almost half the MRT at $\phi = 0$. In other words, the spread of the signals can be an advantage if there are accumulators between those corresponding to the alternatives we are considering. (See also Fig. 3.) It is important to note that this can only be an advantage when we employ tests that utilize knowledge of the shape of the signals, as encoded in the matrix A .

4 The interrogation protocol

We now turn to the implementation of the interrogation protocol in which the subject is given a period of time T after stimulus presentation to arrive at a decision, which must be rendered at T . The problem, then, is to minimize the error rate for a process that ends at time T . This is a fixed sample size hypothesis test. We ask, what is the best strategy to follow? There are several possible strategies one might follow, depending on energy costs, and availability of computational ability. The posterior probability that

H_i is true is given by $\Pi_i(T)$ (eqn 13). Thus, the optimal test is to choose the largest of the Π_i s. However, this requires that we compute not only the value of the accumulators, but y , requiring a multiplication by the matrix A , and the calculation of $\lambda \int_0^T \mathbf{x} dt$. We will explore strategies apart from this optimal one, that do not involve all of these calculations. For instance, we may not know A , or we may not know λ . There are, then, four possible strategies we might follow, whose decision rules d are given by the following, noting that since the decision time is given no thresholds are involved:

$$(i) \quad d = \left\{ i : x_i(T) > \max_{j \neq i} x_j(T) \right\}, \quad (17)$$

$$(ii) \quad d = \left\{ i : x_i(T) - \lambda \int_0^T x_i(s) ds > \max_{j \neq i} \left[x_j(T) - \lambda \int_0^T x_j(s) ds \right] \right\}, \quad (18)$$

$$(iii) \quad d = \left\{ i : (A\mathbf{x}(T))_i > \max_{j \neq i} (A\mathbf{x}(T))_j \right\}, \quad (19)$$

$$(iv) \quad d = \left\{ i : y_i(T) > \max_{j \neq i} y_j(T) \right\}. \quad (20)$$

In (iv), y is as defined in (12). In other words, in strategy (i) one chooses the alternative corresponding to the accumulator with the largest value; in strategy (ii) one chooses the alternative corresponding to the accumulator with the largest value of $x_i - \lambda \int x_i(s) ds$; in strategy (iii) one chooses the alternative corresponding to the largest value of $(A\mathbf{x})_i$; and in strategy (iv) one chooses the alternative corresponding to the largest value of y_i . The optimal strategy is strategy (iv), because this is merely choosing the alternative with the largest posterior probability; since the $\log \sum \exp(y_j)$ term in the Π_i 's (eqn 13) does not depend on i , $\Pi_i - \Pi_j = y_i - y_j$.

We consider the case in which the prior probabilities are equal, $\pi_i = 1/n$, in which case the appropriate initial conditions are $x_i(0) = 0$ for $i = 1, \dots, n$. Formulas for the probability of a correct response in the interrogation protocol in the four cases are derived in Appendix C. The results are as follows: if signal i is the largest signal, the probability of choosing i at time T is, in the four cases,

$$(i) \ Pr(\text{correct}) = \frac{1}{2^{n-1}\sqrt{\pi}} \int_{-\infty}^{\infty} \prod_{j \neq i} \left\{ 1 + \operatorname{erf} \left(y + \frac{S_i - S_j}{c} \sqrt{\frac{\tanh(\lambda T/2)}{\lambda}} \right) \right\} e^{-y^2} dy, \quad (21)$$

$$(ii) \ Pr(\text{correct}) = \frac{1}{2^{n-1}\sqrt{\pi}} \int_{-\infty}^{\infty} \prod_{j \neq i} \left\{ 1 + \operatorname{erf} \left(y + \frac{S_i - S_j}{c} \sqrt{\frac{T}{2}} \right) \right\} e^{-y^2} dy, \quad (22)$$

$$(iii) \ Pr(\text{correct}) = \frac{1}{\sqrt{\pi^n |\hat{A} \hat{A}^T|}} \int_{-\infty}^{\infty} \left[\int_{\mathbf{z} \leq y + g(S, c, T, \lambda)} \exp \left(-\mathbf{z}^T (\hat{A} \hat{A}^T)^{-1} \mathbf{z} \right) d\mathbf{z} \right] e^{-y^2} dy, \quad (23)$$

$$(iv) \ Pr(\text{correct}) = \frac{1}{\sqrt{\pi^n |\hat{A} \hat{A}^T|}} \int_{-\infty}^{\infty} \left[\int_{\mathbf{z} \leq y + g(S, c, T, 0)} \exp \left(-\mathbf{z}^T (\hat{A} \hat{A}^T)^{-1} \mathbf{z} \right) d\mathbf{z} \right] e^{-y^2} dy, \quad (24)$$

where erf is the *error function* defined as $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-z^2} dz$. The matrix \hat{A} in formulas (23) and (24) is defined as follows. Let N be the number of possible signal vectors and n the number of accumulators. Let $[j] = j$ if $j < i$ and $[j] = j + 1$ if $j \geq i$.

Then

$$\hat{A}_{jk} = \frac{S_{[k]}^{([j])} - S_{[k]}^{(i)}}{S_i^{(i)} - S_i^{([j])}}, \quad 1 \leq j \leq N - 1, \quad 1 \leq k \leq n - 1, \quad (25)$$

and the function g in the limits of integration is

$$g(S, c, T, \lambda) = \frac{(S^{(i)}, S^{(i)} - S^{([j])})}{S_i^{(i)} - S_i^{([j])}} \frac{1}{c} \sqrt{\frac{\tanh(\lambda T/2)}{\lambda}}, \quad (26)$$

where (\cdot, \cdot) above is the standard inner product.

These formulas tell us a lot about the effectiveness of the four strategies. First, note that the probabilities of correct responses in cases (ii) and (iv) are obtained by letting $\lambda \rightarrow 0$ in the formulas for cases (i) and (iii), respectively. As $\operatorname{erf}(x)$ is a strictly increasing function and the correct probability in (23) is a strictly increasing function of g , the probability of a correct response in cases (i) and (iii) is maximized as a function of λ at the maximum of $\tanh(\lambda T/2)/\lambda$. As this is an even function with maximum at $\lambda = 0$, i.e. when decay is balanced with inhibition, the minimal ER in the interrogation protocol occurs when decay is balanced with inhibition, and the ER, and hence the RR, is an even function of λ . In strategies (ii) and (iv) the probability of a correct response thus depends only on the differences of the signal strengths, and not on the inhibition and decay.

In the case of perfect acuity the matrix $\hat{A} = I$ is the identity, and the above formulas (23) and (24) simplify to (21) and (22). But when \hat{A} is not the identity, the formulas

are not easy to evaluate. For quantitative comparisons with strategies (i) and (iii), we will find it simpler to evaluate the SDEs with Monte Carlo simulations, than through the numerical evaluation.




Figure 8: Comparison of strategies in the interrogation protocol. In the left panel the acuity is fixed at $\phi^{-1} = 4/3$, while λ is allowed to vary; in the right panel λ is fixed at $\lambda = 0.5$ and the acuity is allowed to vary. In both cases the probabilities of correct responses are computed for $n = 8$ alternatives, at $T = 2$. The correct signal is H_4 , the drift is $a = 2$ and the noise is $c = 1$.

4.1 Comparison of strategies

If the acuity is infinite ($S_m^{(i)} = S_0 + a \delta_{im}$), strategy (i) is equivalent to strategy (iii), and strategy (ii) is equivalent to strategy (iv). In this case, performance depends on λ in strategies (i) and (iii). If the acuity is large (concentrated signal) then there is little difference between the strategies. But, if the signal has a significant spread, there is a large advantage to using strategies (iii) or (iv). Figure 8 shows the probability of a correct response for 8 alternatives at time $T = 2$, following the four strategies outlined in the previous sections. For a small acuity, there is a significant advantage to multiplying the vector of accumulators by the matrix of possible signals. Figure 8 (right panel) shows the effect of acuity on the performance of the four strategies. In all four of the strategies, performance improves as acuity increases, if the time to make a

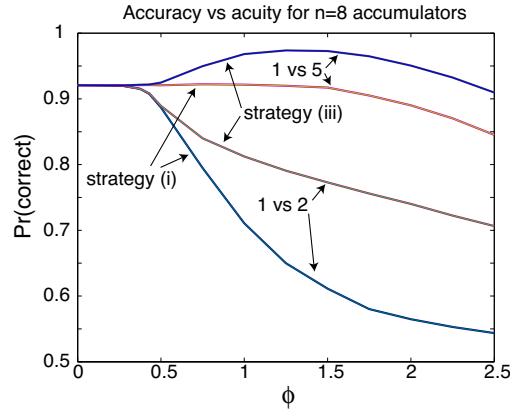


Figure 9: Probability of a two-alternative correct response as a function of acuity at $T = 1$. The number of accumulators is $n = 8$, but only those corresponding to the given alternatives are considered when making a decision.

decision is fixed. Furthermore, the performance converges for Strategies (i) and (iii), and Strategies (ii) and (iv), at moderate values of the acuity (around 2). In other words, for strongly peaked signals there is no advantage to multiplying the accumulators by A .

Figure 9 shows the effects of acuity on decisions between two alternatives when the number of accumulators is $n = 8$. When there are accumulators ‘between’ the alternatives there is a significant advantage to multiplying the accumulators by A , and the optimal value of the acuity is less than infinity. (See also Fig. 3.) In this case, again, the spread in the signal is an advantage.

5 Implementation in a neural network

As we have seen, there is a significant advantage to be gained by multiplying the accumulators by the matrix A of alternative signals. This can be accomplished in a simple network in which each accumulator x_i , $i = 1, \dots, n$, excites units z_j , $j = 1, \dots, N$, in an output layer with weight matrix W . In this model the accumulators x_i represent the sensory layer and the z_i ’s are the decision layer: a decision is made once the first z_i crosses a threshold. Figure 10 shows such a network with four accumulators and four output units. In this network the weight of the connection from x_i to z_j is W_{ji} . Then, if

we neglect noise in the output layer, the vector of output units \mathbf{z} is governed by

$$\tau_z \frac{dz_j}{dt} = -z_j + f_j \left(\sum_{i=1}^n W_{ji} x_i \right), \quad (27)$$

therefore, at the equilibrium state,

$$z_j = f_j \left(\sum_{i=1}^n W_{ji} x_i \right), \quad (28)$$

so $\mathbf{z} = \mathbf{f}(W\mathbf{x})$. The matrix multiplication $A\mathbf{x}$ is therefore accomplished by tuning the weight matrix so that $W = A$. In other words, the weight matrix W should be the matrix A of alternative signals to perform the required matrix multiplication.




Figure 10: Neural network implementation of matrix multiplication. The weight matrix W denotes the weights of the connections between the x_i 's and z_j 's. The weight of the connection between x_i and z_j is W_{ji} .

6 Comments and future directions

6.1 The place of the present work in the context of recent research

Given a decision task, such as deciding what angle a bar lies at on a screen, subjects improve their performance over time. The accumulator model, in which accumulators are preferentially sensitive to various alternatives, suggests several ways in which improvement can occur. For alternatives that are close together, in the sense of there being no or few accumulators 'between' them, the way to improve is to increase the acuity in

the signal. To discern between angles that are close together, for example, this would mean tightening the tuning curves of neurons. However, if the task is to decide quickly between angles that are a moderate distance apart, this can be facilitated by adjusting the acuity in the signal. There is a finite optimal acuity value where the spread of the signal can be an aid to decision making.

Significant improvements in performance can be achieved if the incoming signals are pre-multiplied by the matrix A of possible signals. This performance improvement is such that it becomes beneficial to increase the spread of the signals in some cases. This may perhaps go some way to explaining why tuning curves are not always at their maximal sharpness. Why is it that they can be adjusted at all? We showed that optimal performance occurs at positive values of ϕ only if there are accumulators ‘between’ the alternatives under consideration and only if the signal is multiplied by A . This suggests an explanation for the improvement in angle discrimination tasks in monkeys. Presumably, when the test is first taken, the subject uses a criteria without any (or, with little) knowledge of the spread of the signals. After a few trials, the subject learns the shape of the incoming signals, and can use this to her benefit. One possible mechanism for learning, then, is a shift from a test that does not rely on a knowledge of the shape of the signals, to one that does. For example, the subject could initially begin with a max-vs-next test, and then shift to a δ_b test, once the matrix A has been learned.

Additionally, the subject can learn by adjusting the spread of the signals ϕ , by altering the tuning curves of neurons. This will be an advantage if the alternatives are separated by accumulators. The effects of these adjustments on improvement are limited by the number of available accumulators. For the angle discrimination task, for example, if there are no accumulators sensitive to angles between 22 and 26 degrees, then the subject will gain an advantage in tests on these angles by increasing the acuity of signals as much as possible. But in tests between 15 and 30 degrees, it will be an advantage to decrease the acuity. If this model is an accurate description of decision making, then this explains why the acuity can be adjusted. For the most common discrimination tasks it is an advantage to have a moderately small acuity. Training a subject in a particular task can result in a sharpening of the tuning curves. But, the improvements in performance are limited by the number of accumulators.

The work of Beck et al. (2008) provides a ground-breaking foundation both for the

neurophysiological, as well as the theoretical aspects of multiple alternative decision making. The dominating proposal here is the idea of sensory evidence accumulation without loss of information and its single most important consequence that the most likely action is chosen through dynamic attractors. This theory is showcased for multiple or a continuum of alternative decisions through single neuronal recordings from the lateral intraparietal cortex during the performance of a random-dot motion discrimination task (Roitman and Shadlen, 2002; Churchland et al., 2008). The optimal evidence accumulation then is obtained via linear integration of neural activity and eventually is modelled with a seemingly straightforward Bayesian formulation across trials and over time. The authors demonstrate the command of their Bayesian model when the object of statistical inference is set around predicting the firing patterns of the so-called Poisson-like cortical neurons (Shadlen and Newsome, 1998), which are involved in the decision making process. The Bayesian model validates the frequently reported trade-off between accuracy and speed in decision making. Additionally, the information is shown to increase proportionally with time and stimulus coherence.

Furman and Wang (2008) show that neural behavior, when choosing alternatives, can be affected by the number of choices and their similarities. The authors report that the overlap between similar alternatives contributes to a burst in neuronal activity, and hence decreases the response time. The similarity has an adverse effect of significantly increasing the error rate. Niwa and Ditterich (2008) use a variation of the random-dot task on humans to study the competition between multiple integrators to the decision threshold level. They propose a simple computational model to account for the accumulation of sensory evidence in the integrators. Their model asserts that faster responses are observed for higher coherence in the stimuli, and the complexity of alternatives has a proportional effect on response time. In contrast to the leaky, competing accumulator model of Usher and McClelland (2001), integrators in Niwa and Ditterich are perfect.

As discussed in section 3.1, and as shown in Fig. 6, we explore the relationship between the number of alternatives and the MRT in conjunction with signal acuity. Similar to the results given in different panels of Fig. 6, Furman and Wang report that response time increases in direct proportion to the number of alternatives. The authors also track neuronal responses to the stimuli's motion coherence. Similarity then is partly viewed through coherence. Because of the latter, it would be nontrivial to make a direct

connection between our work and theirs. The challenge is to design a simulation study in which the uncertainty due to coherence is properly explained.

Similar to the model in Niwa and Ditterich, our proposed neural network in section 5 would account for the accumulation of evidence in the integrators. As we argue below, due to its Bayesian nature, this feature of our modeling approach is more compatible with the work of Beck et al. (2008).

We take time in articulating the results in Beck et al. (2008), mainly due the fact that we find our proposed neural network of section 5 as analogous to their methodology. To be precise, Beck et al. (2008), formulate the evidence accumulation of a population of M neurons as

$$Pr(s|r(t_N)) = \prod_{n=1}^N \frac{r(t_n|s, c)}{r(t_n|c)}, \quad (29)$$

in which $r(t) = \{r_1(t), \dots, r_M(t)\}$ is the vector of neuronal activity so that $r_i(t)$ is the spike count of neuron i at time t , c refers to all nuisance parameters in the statistical model, and s depicts the direction of the stimulus. The choice of a flat prior on s is a reflection of the noninformative nature of the implemented Bayesian thought process. The stimulus s is then estimated by $\hat{s} = \text{argmax}_s Pr(s|r(t_1 : t_{stop}))$, where $\hat{s} = \text{argmax}_s$ is the maximum likelihood estimator for s , and t_{stop} is a stopping time. Because of the uniformity of the prior on s , the posterior point estimator coincides with the maximum likelihood estimator, and therefore, it emphasizes on the role of the Bayesian approach in this setting, merely as a catalyst, having an identical frequentist interpretation. The existence of t_{stop} in the maximum likelihood estimator calculation captures the essence of evidence accumulation over time, and it also facilitates the proper accounting for the fact that an optimal decision is being made at a stopping time at which the evidence exceeds a threshold boundary. Finally, an optimal network having the structure $Pr(s|r(t_n)|s, c) = \Phi(r(t_n)c(t_n)) \exp(h(s).r(t_n))$ is followed, where Φ and h are real functions. This is the familiar exponential family (see chapter 9 in Wasserman (2004)). The Poisson-like nature of spike activities is subsequently encoded in the link function h , where h may be taken as the logarithm of the acuity of participating neurons in the network.

In the graphical model presented in section 5, we essentially build an algorithmic counterpart to the Bayesian dynamic attractor model of Beck et al. First, we note that as

long as the inherent model for the network follows the exponential form, the Gaussian model in (1) being an example, the Bayesian paradigm of Beck et al. is implemented through the dynamic feed-forwarding nature of our network. Second, the role of time related priors is played by the weight matrix A . Casting the problem this way, we take into account the accumulation of evidence as a function of time, and with no loss of information. To make an analogy to Beck et al., we propose a Bayesian network whose stopping values (outputs) are defined when the process of evidence accumulation is exhausted at the point that the decision threshold level is crossed. The choice of components in \hat{A} (see (70) in the appendix section) makes it possible to preserve the noninformative priors of Beck et al.

To elaborate, we let:

$$z_j \rightarrow f_j \left(\sum_{i=1}^n W_{ji} \xi(x_i) \right), \quad (30)$$

where f is the activation function and ξ is a basis function that depends on model parameters and will be determined during network's training. We let $y_n = f_j(\sum_{i=1}^n W_{ji} \xi(x_i))$. The network is then trained by minimizing the expected error of the weight parameters (here the elements in the matrix A). That is, the statistical estimation process revolves around finding W_{ji} such that $E(w) = \frac{1}{2} \|y_n - t\|^2$ is minimized, where t is some target function. To conform to the Gaussian function in equation (1), we let the target function having a Gaussian form with $Var(t) = \phi$, or the precision ϕ^{-1} . Stated this way, the maximum likelihood estimator for W s are obtained by minimizing the sum-of-squares function of the above. This will lead us to ϕ_{ML} , the maximum likelihood estimator for ϕ through:

$$\frac{1}{\phi_{ML}} = \frac{1}{n} \sum_1^n \{\hat{y}_n - t\}^2, \quad (31)$$

where \hat{y}_n is the maximum likelihood solution for y_n . Finally, we note that the sum-of-squares function is minimized by solving $\nabla E(w) = 0$, where ∇ represents the gradient of the error function in our proposed feed-forward network. The maximum likelihood solution for w can be obtained using the backpropagation method of Rumelhart et al. (1988).

6.2 Lines of future research

The linearized LCAM (8) along with the neural network implementation in section 5 of the matrix multiplication represent in some sense the simplest model of decision making that incorporates mutual inhibition, decay, and noise, as well as knowledge of the possible shapes of the incoming signals. There are many ways that this model can be modified and extended to incorporate different effects.

Perhaps the most obvious extension would be an examination of the nonlinear LCAM (7). Unfortunately there is little that can be said analytically about such nonlinear systems of equations, particularly when the number of equations is greater than two. A systematic study of the various effects we have considered would involve substantial numerical simulations over a broad range of parameters. As we have noted, these are likely to agree qualitatively with the linear model, but it would be interesting to see where the linear approximation breaks down.

Another logical extension of this model is to consider a continuum of alternatives. This would involve allowing the number of alternatives to approach infinity. Of course, there are not infinitely many accumulators in any one brain, or any finite collection of brains for that matter. And, although there are infinitely many possible angles a bar may make on a screen, as we have noted, the ability of humans and monkeys to discern between two angles is limited to angles that differ by no less than about 4 degrees, even after significant training. There are thus obvious limits to how many alternatives can be considered by any one brain. It will be interesting to see if such limits must be incorporated into models like the LCAM, or if they are a consequence of these models.

We have seen that the matrix A encodes the knowledge of possible signals and that this can be encoded in a simple neural network. As we have mentioned, encoding the network with the matrix A is one way to improve at various tasks. Experiments that look for such neural training, whether at the neural or behavioral level, will be of great interest. In future work we intend to explore the modulation of the weights of such a network as a mechanism for improving at tasks such as the angle discrimination tasks. Simen et al. (2006) have shown how simple neural mechanisms can adapt thresholds in the two alternative forced choice task in order to maximize the reward rate. Their results and those of the current paper suggest how such learning might take place in a multiple choice setting. Note also that the network in section 5 is the simplest possible—a

feed-forward system without feedback or mutual inhibition. In connection with learning algorithms for training such a network, it would be interesting to see if inhibition and/or feedback in this layer may be an aid to learning the set of possible signals or implementing the matrix multiplication.

Additional lines of research include experiments that test for performance on different types of ordering of alternatives, such as the inside-out and outside-in orderings we have suggested. We also intend to explore the effects of signals drawn from various distributions.

Finally, in a work in progress, we attempt to extend the decision-making strategies explained in this paper to the case when the signals are randomly distributed. Following the notation in section 1.2, when there is noise in the signals, $\frac{s_i - S_i}{c}$ will follow a Normal(0, 1) distribution. Subsequently, by letting S_i follow a Gaussian distribution, similar to equation 1, we will have a natural Bayesian formulation with a conjugate prior on the means. This will allow us to further investigate the effect of acuity through sampling from its posterior distribution. A plausible approach then would be to consider a Bayesian hierarchical model (Gelman et al., 2003) on s_i 's, through which Φ is being updated. This alternative approach however, will impose a series of non-trivial challenges, chief among which would be the modification of the SDEs (7) and (8).

Acknowledgements We would like to thank two anonymous reviewers for their careful reading of the text and very constructive suggestions.

Appendix: Mathematical details

A Derivation of posteriori probabilities under assumption (10)

Let X_t be a vector-valued random variable allowed to evolve in time under one of N hypotheses H_i . Let π_i denote the prior probability of H_i , and let (Ω, \mathcal{F}) be a measure space on which $\{X_t, t \in \mathbb{R}\}$ is observed, and \mathcal{F}_t the sub- σ -algebra of \mathcal{F} generated by $X^t = \{X_u, 0 \leq u \leq t\}$ observed up to time t . Let P_i be the probability measure under

hypothesis H_i , and $P_i^t(\cdot)$ the restriction of the measure P_i to the σ -algebra \mathcal{F}_t . The log-likelihood ratios are defined as follows:

$$Z_i(t) = \log \frac{dP_i^t}{dQ^t}(X^t), \quad i = 1, \dots, N, \quad (32)$$

where $dP_i^t/dP_j^t(X^t)$ is the Radon-Nikodym derivative of the measure P_i with respect to a dominating measure Q^t , restricted to \mathcal{F}_t . If $Q^t = P_j^t$, the corresponding process will be denoted $Z_{ij}(t)$, and is the log-likelihood ratio of posterior probabilities given the information in \mathcal{F}_t . We denote a sequential test by $\delta = (\tau, d)$, where τ is a stopping time and $d \in \{1, 2, \dots, N\}$ is the decision.

Two differently-structured tests that give the same asymptotically optimal results are called the δ_a and δ_b tests (Dragalin et al., 1999). These are defined in terms of the posterior probability Π_i and generalized likelihood ratio L_i that are given by

$$\begin{aligned} \Pi_i(t) &= \frac{\pi_i \exp[Z_i(t)]}{\sum_{j=1}^N \pi_j \exp[Z_j(t)]} = P(H_i \text{ is true} | X^t), \\ \text{and } L_i(t) &= \frac{\pi_i \exp[Z_i(t)]}{\max_{j \neq i} \pi_j \exp[Z_j(t)]} = \frac{P(H_i \text{ is true} | X^t)}{\max_{j \neq i} P(H_j \text{ is true} | X^t)}, \end{aligned} \quad (33)$$

be the posterior probability of hypothesis H_i , and the generalized likelihood ratio between H_i and the remaining hypotheses, respectively. In the case of two hypotheses both tests reduce to the optimal SPRT; in the case of $N > 2$ hypotheses, both tests achieve the same asymptotic optimality.

For constant thresholds a_i and b_i , the stopping times for the two tests are as follows:

$$\begin{aligned} \delta_a: \quad \tau_i &= \inf \{t : \Pi_i(t) \geq \exp(a_i)\} \\ &= \inf \left\{ t : Z_i(t) - \log \left(\sum_{j \neq i} \frac{\pi_j}{\pi_i} \exp[Z_j(t)] \right) \geq a_i \right\}, \\ \delta_b: \quad \nu_i &= \inf \{t : L_i(t) \geq \exp(b_i)\} \\ &= \inf \left\{ t : \min_{j \neq i} \left[Z_{ij}(t) - \log \left(\frac{\pi_j}{\pi_i} \right) \right] \geq b_i \right\}, \end{aligned} \quad (34)$$

which are the Markov ‘accepting’ times for the hypothesis H_i . $\delta_a = (\tau_a, d_a)$ and $\delta_b = (\nu_b, d_b)$ are then the tests with stopping times and decisions given by

$$\begin{aligned} \tau_a &= \min_{1 \leq i \leq n} \tau_i, & d_a &= j \text{ if } \tau_a = \tau_j, \\ \nu_b &= \min_{1 \leq i \leq n} \nu_i, & d_b &= j \text{ if } \nu_b = \nu_j. \end{aligned} \quad (35)$$

The δ_a and δ_b tests are *asymptotically optimal* in the following sense. Let R_i be the probability of accepting hypothesis H_i incorrectly. Then, for the δ_a and δ_b tests there exist bounds \bar{R}_i such that among all tests that satisfy similar error bounds, the $\delta_{a,b}$ tests minimize the expectation of reaction (see Dragalin et al., 1999, Theorem 4.2):

$$R_i(\delta_{a,b}) \leq \bar{R}_i, \quad (36)$$

$$E_i[\tau_a, \nu_b] \sim \inf_{\{\delta | R_i \leq \bar{R}_i\}} E_i[\tau], \text{ as } \max_j \bar{R}_j \rightarrow 0, \quad (37)$$

where the notation $x_\gamma \sim y_\gamma$ as $\gamma \rightarrow \gamma_0$ means that $\lim_{\gamma \rightarrow \gamma_0} (x_\gamma/y_\gamma) = 1$. Equation (37) implies that the tests $\delta_{a,b}$ give an expected value of the stopping time that is asymptotically the infimum of the stopping times of all sequential tests satisfying the error bound in (36). This is the multi-dimensional analog of the optimality result for the SPRT, but unlike that case, the result for $n > 2$ only holds in this asymptotic setting. For $n = 2$ we may replace ‘ \sim ’ in equation (37) with ‘ $=$ ’, and conclude that equality holds for *all* values of $\bar{R}_i > 0$.

We now apply these results to the linearized LCAM (8). The log-likelihood ratios $Z_{ij}(t)$ are determined by the infinitesimal increments:

$$dZ_i(t) = \log p_i(d\mathbf{x}), \quad (38)$$

where p_i is the probability density function of the increments $d\mathbf{x}$ which are Gaussian distributed,

$$p_i(d\mathbf{x}) = \prod_{m=1}^n \frac{1}{\sqrt{2\pi\sigma_m^{(i)2}}} \exp \left[-\frac{(dx_m - \mu_m^{(i)})^2}{2\sigma_m^{(i)2}} \right], \quad (39)$$

where $\mu_m^{(i)}$ and $\sigma_m^{(i)2}$ are the mean and variance of the m th component of the increment of \mathbf{x} , given that hypothesis H_i is true. Hence

$$\begin{aligned} \mu_m^{(i)} &= \left(\lambda x_m - w \sum_{j=1}^n x_j + S_m^{(i)} \right) dt, \\ \text{and } \sigma_m^{(i)2} &= c^2 dt, \end{aligned} \quad (40)$$

where $\lambda = w - k$. Letting $\gamma_m = \lambda x_m - w \sum_{j=1}^n x_j$, and substituting (40) into (38),

we obtain:

$$\begin{aligned}
dZ_i &= \log \left(\frac{1}{\sqrt{2\pi c^2 dt}} \right) \\
&\quad - \frac{1}{2c^2 dt} \sum_{m=1}^n \left[dx_m^2 - 2\gamma_m dx_m dt + \left(\gamma_m^2 - 2w \left(\sum_{j=1}^n x_j \right) S_m^{(i)} + S_m^{(i)2} \right) dt^2 \right] \\
&\quad + \frac{1}{c^2} \sum_{m=1}^n [S_m^{(i)} dx_m - \lambda S_m^{(i)} x_m dt] .
\end{aligned} \tag{41}$$

We use the assumption (10), and the fact that all terms that do not depend on i will cancel when we subtract terms. The only terms in (41) that do depend on i are those in the third row of (41). Therefore, after integrating (41), we have

$$Z_i(t) = C + \frac{1}{c^2} \sum_{m=1}^n S_m^{(i)} \left(x_m - \lambda \int_0^t x_m(s) ds \right) , \tag{42}$$

where C is a function of time but does not depend on i . We form the $N \times n$ matrix of alternative signals:

$$A = \begin{pmatrix} S^{(1)} \\ S^{(2)} \\ \dots \\ S^{(N)} \end{pmatrix} . \tag{43}$$

Then the logs of the posterior probability and generalized likelihood ratio are

$$\begin{aligned}
\log \Pi_i &= \frac{1}{c^2} \left((A\mathbf{x})_i - \lambda \int_0^t (A\mathbf{x})_i dt \right) - \log \sum_{j=1}^n \frac{\pi_j}{\pi_i} \exp \left[\frac{1}{c^2} \left((A\mathbf{x})_j - \lambda \int_0^t (A\mathbf{x})_j dt \right) \right] \tag{44} \\
\log L_i &= \frac{1}{c^2} \left((A\mathbf{x})_i - \lambda \int_0^t (A\mathbf{x})_i dt \right) - \max_{j \neq i} \left\{ \frac{1}{c^2} \left((A\mathbf{x})_j - \lambda \int_0^t (A\mathbf{x})_j dt \right) + \log \frac{\pi_j}{\pi_i} \right\} \tag{45}
\end{aligned}$$

B Relaxation of assumption (10)

We write eqn (41) as

$$\begin{aligned}
dZ_i &= \log \left(\frac{1}{\sqrt{2\pi c^2 dt}} \right) - \frac{1}{2c^2 dt} \sum_{m=1}^n [dx_m^2 - 2\gamma_m dx_m dt + \gamma_m^2 dt^2] \\
&\quad + \frac{1}{c^2} \sum_{m=1}^n [S_m^{(i)} dx_m - \lambda S_m^{(i)} x_m dt] + \frac{dt}{2c^2} \sum_{m=1}^n \left[S_m^{(i)} \left(S_m^{(i)} - 2w \sum_{j=1}^n x_j \right) \right] \tag{46}
\end{aligned}$$

Therefore,

$$Z_i(t) = C(t) + \frac{1}{c^2} \sum_{m=1}^n S_m^{(i)} \left(x_m - \lambda \int_0^t x_m(s) ds \right) - \frac{t}{2c^2} \left(\sum_{m=1}^n S_m^{(i)2} \right) + \frac{w}{c^2} \left(\sum_{m=1}^n S_m^{(i)} \right) \int_0^t \left(\sum_{j=1}^n x_j(s) \right) ds, \quad (47)$$

where $C(t)$ does not depend on i . Let

$$\mathbf{y} = \frac{1}{c^2} A \left(\mathbf{x} - \lambda \int_0^t \mathbf{x}(s) ds \right) + \log \boldsymbol{\pi}, \quad (48)$$

$$\text{and } u_i(t) = -\frac{t}{2c^2} \left(\sum_{m=1}^n S_m^{(i)2} \right) + \frac{w}{c^2} \left(\sum_{m=1}^n S_m^{(i)} \right) \left(\sum_{j=1}^n \int_0^t x_j(s) ds \right). \quad (49)$$

Then

$$\log \Pi_i = y_i + u_i - \log \sum_{j=1}^n e^{y_j + u_j}, \quad (50)$$

$$\log L_i = y_i + u_i - \max_{j \neq i} (y_j + u_j). \quad (51)$$

We see that under assumption (10), the u_i 's drop out of the above formulas. When this assumption does not hold, the decision is biased toward the hypothesis corresponding to the largest sum of squares of the elements of the signal vector. Interestingly, the u_i terms depend on the inhibition w , but not on the decay.

We can make a further simplification by subtracting from both sides of (50-51) those parts of the sums of the signals that do not depend on i . Let

$$\sum_{m=1}^n S_m^{(i)2} = q_0 + q_i, \quad (52)$$

$$\sum_{m=1}^n S_m^{(i)} = r_0 + r_i, \quad (53)$$

where q_0 and r_0 are the minimums of the corresponding sums. Since the terms involving q_0 and r_0 do not depend on i , we can let

$$u_i(t) = -\frac{1}{c^2} \left(2t q_i - w r_i \sum_{j=1}^n \int_0^t x_j(s) ds \right). \quad (54)$$

Then (50-51) holds with this new definition of u_i . We see that the posterior probability depends on the variation in the values of the sums and sums of squares of the possible signals.

C Derivation of probability of a correct response in the interrogation protocol

C.1 Strategy (i)

We define the random variables $V_{ij}(t) \stackrel{\text{def}}{=} x_j(t) - x_i(t)$, and the constants $a_{ij} \stackrel{\text{def}}{=} S_i - S_j$ for $j = 2, \dots, n$ so that (21) becomes:

$$Pr(\text{correct}) := P_i(n, T) = P \left(\bigcap_{j \neq i} (V_{ij}(T) < 0) \right). \quad (55)$$

From (8) we see that the random variables V_{ij} satisfy the equations

$$dV_{ij} = (\lambda V_{ij} - a_{ij}) dt + c (dW_j - dW_i), \quad (56)$$

for independent Wiener processes dW_j , and we may integrate to obtain:

$$V_{ij}(t) = \frac{a_{ij}}{\lambda} (1 - e^{\lambda t}) + c \int_0^t e^{\lambda(t-s)} [dW_j(s) - dW_i(s)]. \quad (57)$$

We further define random variables $W_0(t)$ and $U_{ij}(t)$, for $j \neq i$ as follows:

$$W_0(t) \stackrel{\text{def}}{=} c \int_0^t e^{\lambda(t-s)} dW_i(s), \quad (58)$$

$$\begin{aligned} U_{ij}(t) &\stackrel{\text{def}}{=} V_{ij}(t) + W_0(t) \\ &= \frac{a_{ij}}{\lambda} (1 - e^{\lambda t}) + c \int_0^t e^{\lambda(t-s)} dW_j(s), \end{aligned} \quad (59)$$

so that W_0 and the U_{ij} 's are independent processes. The U_{ij} 's have means $\mu_{ij}(t)$ and variances $\sigma^2(t)$, and W_0 has mean 0 and variance $\sigma^2(t)$, where

$$\mu_{ij}(t) = \frac{a_{ij}}{\lambda} (1 - e^{\lambda t}) \quad \text{and} \quad \sigma^2(t) = \frac{c^2}{2\lambda} (e^{2\lambda t} - 1). \quad (60)$$

Since $V_{ij} = U_{ij} - W_0$, the probability of making a correct response can now be written as

$$P_i(n, T) = P \left(\bigcap_{j \neq i} [U_{ij}(T) < W_0(T)] \right). \quad (61)$$

The random variable W_0 has the density $f(w)$ of a normal random variable with mean 0 and variance σ^2 , and U_{ij} has the distribution $F_{U_{ij}}$ of a normal random variable with

mean μ_{ij} . Since the U_{ij} 's and W_0 are independent,

$$\begin{aligned} P_i(n, T) &= \int_{-\infty}^{\infty} \prod_{j \neq i} F_{U_{ij}}(w) f(w) dw \\ &= \int_{-\infty}^{\infty} \prod_{j \neq i} \left\{ \frac{1}{2} \left(1 + \operatorname{erf} \left[\frac{w - \mu_{ij}}{\sqrt{2\sigma^2}} \right] \right) \right\} \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{w^2}{2\sigma^2} \right] dw. \end{aligned} \quad (62)$$

Making the change of variables $y = w/\sqrt{2\sigma^2}$, we obtain:

$$P_i(n, T) = \frac{1}{2^{n-1}\sqrt{\pi}} \int_{-\infty}^{\infty} \prod_{j \neq i} \{1 + \operatorname{erf} [y + g(a_{ij}, c, \lambda, T)]\} e^{-y^2} dy, \quad (63)$$

where g in the argument of the error function takes the form

$$g(a_{ij}, c, \lambda, T) = -\frac{\mu_{ij}(T)}{\sqrt{2\sigma^2(T)}} = \frac{a_{ij}}{c} \sqrt{\frac{\tanh(\lambda T/2)}{\lambda}} \quad (64)$$

where we have used the definitions of $\mu(T)$ and $\sigma^2(T)$ in (60). Therefore,

$$P_i(n, T) = \frac{1}{2^{n-1}\sqrt{\pi}} \int_{-\infty}^{\infty} \prod_{j \neq i} \left\{ 1 + \operatorname{erf} \left[y + \frac{S_i - S_j}{c} \sqrt{\frac{\tanh(\lambda T/2)}{\lambda}} \right] \right\} e^{-y^2} dy, \quad (65)$$

C.2 Strategy (ii)

In this case, the calculations are similar to the previous section, except now we define $V_{ij}(t) \stackrel{\text{def}}{=} x_j(t) - x_i(t) - \lambda \int_0^T [x_j(s) - x_i(s)] ds$. Then, as before, $P_i(n, T) = P \left(\bigcap_{j \neq i} V_{ij}(T) < 0 \right)$, and eqn (56) becomes

$$dV_{ij} = -a_{ij} dt + c (dW_j - dW_i). \quad (66)$$

In other words, when we include the λ term in the decision, it drops out of the probability calculation. In this case, following the calculations of the previous section, but with $\lambda = 0$, we find that the probability of a correct response is

$$P_i(n, T) = \frac{1}{2^{n-1}\sqrt{\pi}} \int_{-\infty}^{\infty} \prod_{j \neq i} \left\{ 1 + \operatorname{erf} \left(y + \frac{S_i - S_j}{c} \sqrt{\frac{T}{2}} \right) \right\} e^{-y^2} dy. \quad (67)$$

C.3 Strategy (iii)

We define $V_{ij}(t) \stackrel{\text{def}}{=} (A\mathbf{x}(t))_j - (A\mathbf{x}(t))_i$. Then the probability of choosing correctly, given that signal $S^{(i)}$ is presented is

$$P_i(n, T) = P \left(\bigcap_{j \neq i} V_{ij}(T) < 0 \mid S = S^{(i)} \right). \quad (68)$$

If $S = S^{(i)}$,

$$dV_{ij} = [\lambda V_{ij} + (S^{(j)} - S^{(i)}, S^{(i)})] dt + c (S^{(j)} - S^{(i)}, d\mathbf{W}) , \quad (69)$$

where (\cdot, \cdot) is the dot product and \mathbf{W} is a vector of Wiener processes. Let \hat{A} be the $(N-1) \times (n-1)$ matrix formed by removing the i th column from the matrix of rows of $S^{(j)} - S^{(i)}$, and $-\tilde{S}^{(i)}$ the column that we remove. That is,

$$\begin{pmatrix} S^{(1)} - S^{(i)} \\ \dots \\ S^{(i-1)} - S^{(i)} \\ S^{(i+1)} - S^{(i)} \\ \dots \\ S^{(N)} - S^{(i)} \end{pmatrix} = \begin{pmatrix} \underbrace{\hat{A}_1}_{\text{1st column}} & \underbrace{\hat{A}_2}_{\text{2nd column}} & \dots & \hat{A}_{i-1} & \underbrace{-\tilde{S}^{(i)}}_{\text{i-th column}} & \hat{A}_i & \dots & \underbrace{\hat{A}_{n-1}}_{\text{n-th column}} \end{pmatrix}, \quad (70)$$

where \hat{A}_k is the k th column of \hat{A} . Likewise, let V be the vector of V_{ij} 's and \hat{S} the vector of $(S^{(j)} - S^{(i)}, S^{(i)})$, with the i th elements (of 0) removed. Then eqn (69) can be written in vector form as

$$dV = (\lambda V + \hat{S}) dt + c (\hat{A} d\hat{\mathbf{W}} - \tilde{S}^{(i)} dW_i), \quad (71)$$

where $\hat{\mathbf{W}} = (W_1, \dots, W_{i-1}, W_{i+1}, \dots, W_n)^T$. (Note that $\tilde{S}^{(i)}$ is a vector and W_i is a scalar Wiener process.) Making the further transformations $V_j \rightarrow V_j/\tilde{S}_j^{(i)}$, $\hat{S}_j \rightarrow \hat{S}_j/\tilde{S}_j^{(i)}$, $\hat{A}_{jk} \rightarrow \hat{A}_{jk}/\tilde{S}_j^{(i)}$, the equation becomes

$$dV = (\lambda V + \hat{S}) dt + c (\hat{A} d\hat{\mathbf{W}} - \mathbf{1} dW_i), \quad (72)$$

where $\mathbf{1}$ is the vector of 1's. The difference, then between the above equation (72), and eqn (56), is the matrix \hat{A} , which in (56) is simply the identity. To summarize, the

quantities appearing in eqn (72) are

$$\hat{S}_j = \frac{(S^{([j])} - S^{(i)}, S^{(i)})}{S_i^{(i)} - S_i^{([j])}}, \quad \hat{A}_{jk} = \frac{S_{[k]}^{([j])} - S_{[k]}^{(i)}}{S_i^{(i)} - S_i^{([j])}}, \quad 1 \leq j \leq N-1, \quad 1 \leq k \leq n-1, \quad (73)$$

where $[j] = j$ if $j < i$ and $[j] = j + 1$ if $j \geq i$.

Following the strategy of Appendix C.1, we define $U = V + \mathbf{1}W_0$, where $W_0(t)$ is defined by eqn (58), and U satisfies

$$dU = (\lambda U + \hat{S}) dt + c \hat{A} d\hat{\mathbf{W}}. \quad (74)$$

The elements of U are not correlated with W_0 , but they *are* correlated with each other, and the correlation of elements of U is related to the acuity in the signal. The probability of a correct response is therefore

$$P_i(n, T) = \int_{-\infty}^{\infty} F_U(\mathbf{1}w) f(w) dw, \quad (75)$$

where F_U is the joint distribution of U , and $f(w)$ is the density of a normal random variable with mean 0 and variance σ^2 . The joint distribution F_U is calculated from the joint density:

$$f_U(\mathbf{u}) = \frac{1}{\sqrt{(2\pi)^{n-1} |\boldsymbol{\sigma}|}} \exp\left(-\frac{1}{2}(\mathbf{u} - \boldsymbol{\mu})^T \boldsymbol{\sigma}^{-1}(\mathbf{u} - \boldsymbol{\mu})\right), \quad (76)$$

$$\text{where } \boldsymbol{\mu} = \frac{\hat{S}}{\lambda} (e^{\lambda t} - 1), \quad (77)$$

$$\text{and } \boldsymbol{\sigma} = \frac{c^2}{2\lambda} (e^{2\lambda t} - 1) \hat{A} \hat{A}^T \quad (78)$$

is the covariance matrix. Note that $\boldsymbol{\sigma} = \sigma^2 \hat{A} \hat{A}^T$, where σ^2 is the same as in eqn (60). Hence, the joint distribution of U is

$$F_U(\mathbf{w}) = \int_{\mathbf{y} \leq \mathbf{w}} f_U(\mathbf{y}) d\mathbf{y}. \quad (79)$$

Making the change of variables $\mathbf{z} = (\mathbf{y} - \boldsymbol{\mu})/\sqrt{2\sigma^2}$, in the integral for F_U ,

$$F_U(\mathbf{w}) = \frac{1}{\sqrt{\pi^{n-1} |\hat{A} \hat{A}^T|}} \int_{\mathbf{z} \leq \frac{\mathbf{w} - \boldsymbol{\mu}}{\sqrt{2\sigma^2}}} \exp\left(-\mathbf{z}^T (\hat{A} \hat{A}^T)^{-1} \mathbf{z}\right) d\mathbf{z}. \quad (80)$$

Substituting this into the expression eqn (75), and making the further change of variables $y = w/\sqrt{2\sigma^2}$, we have

$$P_i(n, T) = \frac{1}{\sqrt{\pi^n |\hat{A} \hat{A}^T|}} \int_{-\infty}^{\infty} \left[\int_{\mathbf{z} \leq y - \frac{\boldsymbol{\mu}}{\sqrt{2\sigma^2}}} \exp\left(-\mathbf{z}^T (\hat{A} \hat{A}^T)^{-1} \mathbf{z}\right) d\mathbf{z} \right] e^{-y^2} dy. \quad (81)$$

C.4 Strategy (iv)

In 3.1 we computed the posterior probabilities $\Pi_i(t)$ of the hypotheses H_i , given $\mathbf{x}(s)$, $0 \leq s \leq t$, which, for equal prior probabilities, are given by

$$\mathbf{y} = \frac{1}{c^2} A \left(\mathbf{x} - \lambda \int_0^t \mathbf{x}(s) ds \right) \quad (82)$$

$$\log \Pi_i = y_i - \log \sum_{j=1}^n e^{y_j}, \quad (83)$$

Since $\log \sum_{j=1}^n e^{y_j}$ does not depend on i , choosing the most likely hypothesis is equivalent to choosing the largest y_i . Following this strategy, if we set $V_{ij}(t) \stackrel{\text{def}}{=} c^2 [y_j(t) - y_i(t)]$, then eqn (68) holds, with this new definition, and the equation for the V_{ij} 's is the same as in (69), but with $\lambda = 0$. The same calculation as in the previous section therefore follows, but with $\lambda = 0$. Therefore, strategy (iv) achieves the maximal probability of a correct response, as a function of λ , as achieved in strategy (iii).

D Fast Monte Carlo simulations using Matlab

The SDEs considered in this paper are Ornstein-Uhlenbeck (OU) processes (Gardiner, 2004):

$$dX = (BX + S) dt + c dW. \quad (84)$$

In the LCAM (8), $B = \lambda \mathbf{I} - w \mathbf{1}$, where \mathbf{I} is the identity matrix and $\mathbf{1}$ is the full matrix of 1's. Sample paths of the OU process (84) are approximated by

$$X_{i+1} = CX_i + \Delta S + c\sqrt{\Delta}N_{i+1}, \quad (85)$$

where $X_i \approx X(i\Delta)$, (notice that X_i is a *vector*) N_i is a vector of standard normal random variables, Δ is the time step, and $C = \mathbf{I} + \Delta B$. Naively, we could evaluate (85) in a 'for' or 'while' loop. In the free response protocol, given a threshold z , we can calculate the reaction time of a single simulation by the following (assuming 1 is the correct hypothesis):

```
% N - number of alternatives
% DT - time step
% B - linear part of SDE
```

```

% c - variance of noise

C = eye(N) + DT*B;
y = zeros(N,1);
step = 0;
while max(y) < z
    step = step + 1;
    y = C*y + S*DT + sqrt(DT)*c*randn(N,1);
end

step * DT      % Reaction time
y(1)==max(y)   % 1 if decision is correct, 0 if error

```

In order to determine the MRT for a given ER, we must first determine the threshold. This is done by calculating the error rates for a range of thresholds, determining which threshold gives the desired error rate, and then calculating the mean reaction time for this threshold.

Matlab's optimized array computations may be utilized to vastly improve the performance of Monte Carlo simulations. The key is to write all of the calculations in terms of arrays. We can compute all of the X_i 's for $i = 1, \dots, m$ in a single matrix operation. We directly sum (85) to obtain

$$X_i = C^i X_0 + \Delta \left(\sum_{j=0}^{i-1} C^j \right) S + c\sqrt{\Delta} \sum_{j=0}^{i-1} C^j N_{i-j}. \quad (86)$$

If Δ is sufficiently small, C can be decomposed into $C = UDU^T$, where D is diagonal and U is an orthogonal matrix (Horn and Johnson, 1985). We use the fact that since U is orthogonal, $U^T N_{i-j}$ is also a vector of independent standard normal random variables (Gardiner, 2004). Therefore,

$$X_i = UD^i U^T X_0 + \Delta U \left(\sum_{j=1}^i D^{j-1} \right) U^T S + c\sqrt{\Delta} U \sum_{j=1}^i D^{j-1} N_{i-j}. \quad (87)$$

We can compute an $n \times$ npts matrix where the i th column is the vector X_i by the following Matlab operations:

```

% N - number of alternatives
% npts - number of iterations of X_i to calculate
% DT - time step
% B - linear part of SDE
% y0 - initial condition
% c - variance of noise

C = eye(N) + DT * B;
[U, D, V] = svd(C);
dd = diag(D);
ds = dd(diag(1:N) * ones(N, npts)).^...
      (ones(N, npts) * diag( 0:(npts-1) ));
ds2 = dd(diag(1:N) * ones(N, npts)).^...
      (ones(N, npts) * diag( (npts-1):-1:0 ));
drift = DT * U * diag(U'*S) * cumsum(ds, 2);

y = U*diag(V*y0)*D*ds + drift + c*sqrt(DT)*U* ...
      (cumsum(ds2.* randn(N, npts),2)./ ds2);

```

y is then an $N \times npts$ matrix whose i th column is X_i , i.e. $X_i = y(:, i)$. The key point is that the computation of a sample path is contained in only the last line of the above. So to run many simulations one only need evaluate this line the desired number of times, for example in a ‘for’ loop. If the initial condition is the same for all sample paths, the first term can be absorbed into “drift.” If only the last element of the sample path is required, “cumsum” can be replaced with “sum” to slightly increase performance. To calculate $\int x_i(s)ds$, however, these terms will be needed. The above algorithm is generally about ten times faster than a simple ‘for’ loop. Also, note that if one is calculating a first passage time problem, these can be done in ‘blocks’, and a search through the block made, which is still about five times faster than a ‘while’ loop.

References

- Beck, J., Ma, W., Kiani, R., Hanks, T., Churchland, A., Roitman, J., Shadlen, M., Latham, P., and Pouget, A. (2008). Probabilistic population codes for bayesian decision making. *Neuron*, 60:1142–1152.
- Behseta, S., Kass, R., Moorman, D., and Olson, C. (2007). Testing equality of several functions: Analysis of single-unit firing-rate curves across multiple experimental conditions. *Statist. Med.*, 26(21):3958–3975.
- Bogacz, R., Brown, E., Moehlis, J., Hu, P., Holmes, P., and Cohen, J. (2006). The physics of optimal decision making: A formal analysis of performance in two-alternative forced choice tasks. *Psych. Rev.*, 113(4):700–765.
- Brown, E., Gao, J., Holmes, P., Bogacz, R., Gilzenrat, M., and Cohen, J. (2005). Simple neural networks that optimize decisions. *Int. J. Bifurcation and Chaos*, 15(3).
- Churchland, A., Kiani, R., and Shadlen, M. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6):693–702.
- Dayan, P. and Abbott, L. (2001). *Theoretical Neuroscience*. Cambridge, Massachusetts: MIT Press.
- Dragalin, V., Tartakovsky, A., and Veeravalli, V. (1999). Multihypothesis sequential probability ratio tests, part I: Asymptotic optimality. *IEEE Trans. Inform. Theory*, 45:2448–2461.
- Furman, M. and Wang, X. (2008). Similarity effect and optimal control of multiple-choice decision making. *Neuron*, 60:1153–1168.
- Gardiner, C. (2004). *Handbook of Stochastic Methods*. Berlin: Springer-Verlag, 3rd edition.
- Gelman, A., Carlin, J., Stern, H., and Rubin, D. (2003). *Bayesian Data Analysis*. Chapman and Hall, second edition.
- Georgopoulos, A., Schwartz, S., and Kettner, R. (1986). Neuronal population coding of movement direction. *Science*, 243:1416–1419.

- Ghose, G., Yang, T., and Maunsell, J. (2002). Physiological correlates of perceptual learning in monkey v1 and v2. *J. Neurophysiol.*, 87:1867–1888.
- Gold, J. and Shadlen, M. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404:390–394.
- Gold, J. and Shadlen, M. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5:10–16.
- Grossberg, S. (1988). Nonlinear neural networks: principles, mechanisms, and architectures. *Neural Networks*, 1:17–61.
- Hanes, D. and Schall, J. (1996). Neural control of voluntary movement initiation. *Science*, 274:427–430.
- Henry, G., Dreher, B., and Bishop, P. (1974). Orientation specificity of cells in cat striate cortex. *J. Neurophysiol.*, 36:1394–1409.
- Hick, W. (1952). On the rate of gain of information. *Quart. J. Exp. Psych.*, 4:11–26.
- Higham, D. (2001). An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Rev.*, 42(3):525–546.
- Hopfield, J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA*, 82:3088–3092.
- Horn, R. and Johnson, C. (1985). *Matrix analysis*. Cambridge Univ. Press.
- Hubel, D. and Wiesel, T. (1968). Receptive fields and functional architecture of the monkey striate cortex. *J. Physiol.*, 195:215–243.
- Iacus, S. (2008). *Simulation and Inference for Stochastic Differential Equations, with R examples*. Springer-Verlag, New York.
- Kullback, S. (1987). The kullback-leibler distance. *The American statistician*, 41:340–341.
- Kullback, S. and Leibler, A. (1951). On information and sufficiency. *Ann. Math. Statist.*, 22(1):79–86.

- Lacouture, Y. and Marley, A. (1995). A mapping model of bow effects in absolute identification. *J. Math. Psych.*, 39:383–395.
- Lacouture, Y. and Marley, A. (2004). Choice and response time processes in the identification and categorization of unidimensional stimuli. *Perception & Psychophysics*, 66(7):1206–1226.
- Mazurek, M., Roitman, J., Ditterich, J., and Shadlen, M. (2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex*, 13(11):1257–1269.
- McMillen, T. and Holmes, P. (2006). The dynamics of choice among multiple alternatives. *J. Math. Psych.*, 50(1):30–57.
- Miller, G. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psych. Rev.*, 63:81–97.
- Niwa, M. and Ditterich, J. (2008). Perceptual decisions between directions of visual motion. *J. Neurosci.*, 28(17):4435–4445.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psych. Rev.*, 85:59–108.
- Ratcliff, R., Segraves, M., and Cherian, A. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of simple two-choice decisions. *J. Neurophysiol.*, 90:1392–1407.
- Ratcliff, R., Van Zandt, T., and McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psych. Rev.*, 106:261–300.
- Roitman, J. and Shadlen, M. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.*, 22(1):9475–9489.
- Rumelhart, D., Hinton, G., and Williams, R. (1988). Learning internal representations by error propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, volume 1: Foundations, pages 318–362. MIT Press.
- Shadlen, M. and Newsome, W. (1998). The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J. Neurosci.*, 18:3870–3896.

- Shadlen, M. and Newsome, W. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.*, 86:1916–1936.
- Simen, P., Cohen, J., and Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural networks*, 19:1013–1026.
- Smith, P. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *J. Math. Psych.*, 44:408–463.
- Teichner, W. and Krebs, M. (1974). Laws of visual choice reaction time. *Psych. Rev.*, 81:75–98.
- Usher, M. and McClelland, J. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psych. Rev.*, 108:550–592.
- Usher, M., Olami, Z., and McClelland, J. (2002). Hick’s law in a stochastic race model with speed-accuracy tradeoff. *J. Math. Psych.*, 46:704–715.
- Wasserman, L. (2004). *All of statistics*. New York: Springer-Verlag.