On the Time Course of Perceptual Choice: A Model Based on Principles of Neural Computation

Marius Usher
James L. McClelland

Technical Report PDP.CNS.95.5
December 1995

Parallel Distributed Processing and Cognitive Neuroscience

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA

Western Psychiatric Institute and Clinic
University of Pittsburgh
Pittsburgh, PA

Neural and Behavioral Sciences
University of Southern California
Los Angeles, CA

MRC Applied Psychology Unit
Cambridge, England

On the Time Course of Perceptual Choice: A Model Based on Principles of Neural Computation

Marius Usher
James L. McClelland

Abstract

The time course of information processing is discussed in a model based on leaky, stochastic, non-linear accumulation of activation in mutually inhibitory processing units. The model addresses data from choice tasks using both time-controlled (e.g., deadline or response signal) and standard reaction time paradigms, and accounts simultaneously for aspects of data from both paradigms. In special cases, the model becomes equivalent to a classical diffusion process, but in general a more complex type of diffusion occurs. Mutual inhibition counteracts the effects of information leakage, allows flexible choice behavior regardless of the number of alternatives, and contributes to accounts of additional data from choice with conflict stimuli and word identification tasks.

Marius Usher and James L. McClelland, Department of Psychology, Carnegie Mellon University, and the Center for the Neural Basis of Cognition.

The research reported here was supported by Grants MH00385 and MH47566. We thank Mark Chappell and David Plaut for comments on a draft of this article.

Correspondence concerning this manuscript may be sent to the first author at Department of Psychology, University of Kent at Canterbury, CT2 7LZ, Canterbury, UK, or to the second author at Department of Psychology, Carnegie Mellon University, Pittsburgh PA 15213. Email may be sent via the Internet to marius@crab.psy.cmu.edu or mcclelland+@cmu.edu.

The process by which stimulus information accumulates in the cognitive system and gives rise to decisions and actions is a central topic in experimental cognitive psychology, but the nature of this process remains incompletely understood. We consider this process from the point of view of principles of neural information processing, and we show how these principles can be used in a productive way to guide and specify models that account for the time course of information processing. The principles that we will explore appear to characterize information processing in real neural systems as studied since the earliest days of single-unit recording studies (e.g., Creutzfeld et al., 1966; Ratliff, 1965): leaky, non-linear summation of graded signals; intrinsic variability; and lateral inhibitory interactions among processing units.

We will focus on two types of experimental conditions that are used in studies designed to characterize the time course of information processing: the time-controlled conditions and standard choice reaction time (RT) paradigm. In our usage, the time-controlled choice paradigm encompasses a variety of methods, in which processing time is experimentally controlled to generate a measure of performance (the signal detection variable d' or the probability for a correct response, P) as a function of the elapsed time since the stimulus onset. In some studies, a response signal is used to indicate when subjects must produce an immediate response (Corbett & Wickelgren, 1978; Dosher, 1976; Reed, 1973, 1976; McElree & Dosher, 1989; Wickelgren, 1977; Wickelgren & Corbett, 1977), or a short response deadline is used to force responses to occur at an early time after stimulus onset (Pachella & Pew, 1968; Pachella & Fisher, 1969).[1] In other studies, processing time is limited by controlling the duration of the display using a post-stimulus mask (Loftus, Busey & Sender, 1993).[2]

Interestingly, despite the complexity of the perceptual and cognitive systems, time-accuracy curves measured in such time-controlled experiments reveal a very simple regularity: In almost all studies, performance improves and then levels off following a shifted exponential approach to asymptote (Busey & Loftus, 1994; Corbett & Wickelgren, 1978; Loftus, et al., 1993; McElree & Dosher, 1989; Wickelgren & Corbett, 1977).

It is interesting to notice, however, that in experiments where perfect performance is obtained in unlimited time, the variable that behaves exponentially is the probability of correct response P (Busey & Loftus, 1994; Loftus, et al., 1993; Swenson, 1972), but in experiments where discriminability is reduced so that perfect performance cannot be obtained, the variable that follows the shifted exponential is d' (Corbett & Wickelgren, 1978; Wickelgren & Corbett, 1977; McElree & Dosher, 1989). This is intriguing, since the two performance variables, P and d', are not linearly related, and thus cannot both behave according to the same functional form, presenting a challenge for models dealing with the time course of information processing.

There are several existing theoretical models that have been applied to time-accuracy data, but none of them address this challenge completely, and there are often other difficulties as well. The cascade model of McClelland (1979), in which information accumulation is leaky and deterministic, provides a good account of the shifted-exponential shape of time-accuracy curves when d' is the dependent measure, but does not account for shifted-exponential functions when P is the dependent

---

[1] It should be noted that the use of a deadline does not necessarily force time controlled processing on all trials, particularly if the deadline is set long so that subjects may sometimes reach a high level of certainty about their responses before the deadline is reached; but in some cases (e.g., Pachella & Fisher, 1969), where quite short deadlines are used, RT distributions are quite compact in each deadline condition and the assumption that processing is time controlled seems to be at least a useful first approximation.

[2] In tasks where processing time is controlled by manipulating pay-offs associated with speed and accuracy (Swenson, 1972) it is not entirely clear whether processing is time controlled, or whether subjects adjust activation thresholds to produce different mean reaction times in different conditions. We will not, therefore, focus attention on such studies here.

variable. This model also has several other difficulties that appear to be due to the use of a deterministic information accumulation function (Ashby, 1982; Luce, 1986; Roberts & Sternberg, 1993; McClelland, 1993).

The idea that information accumulation is essentially a stochastic process is central to the approach we take here, and this is the key assumption of several stochastic choice RT models, such as the random-walk and diffusion models (Laming, 1968; Link & Heath, 1975; Ratcliff, 1978; Stone, 1960) or the accumulator models (Audeley & Pike, 1965; Laberge, 1962; Vickers, 1970; Vickers, Caudrey & Wilson, 1971). In these models, information accumulates as a stochastic process, but without loss or "leakage". The stochastic nature of the process helps account for many aspects of the relevant data, but these models are not completely successful in addressing all aspects of time-accuracy curves. Because of the lossless character of the information accumulation process, the simplest forms of the random walk and diffusion models predict that asymptotic accuracy will always be perfect, and thus cannot address shifted-exponential time-accuracy curves in d' at all (the d' measure always goes to infinity). When additional assumptions are introduced to produce finite asymptotic d' as in Ratcliff's (1978) diffusion model, d' does gradually approach an asymptote, but it does not behave exponentially as a function of time. A comparison of the shifted exponential and of the diffusion model to experimental data has revealed that the shifted exponential provides a better fit (McElree & Dosher, 1989).

Another approach that has recently been proposed to account for shifted-exponential time-accuracy curves is the two-stage information acquisition model of Loftus and Busey (Busey & Loftus, 1994; Loftus et al., 1993). In the first stage of this model, stimulus related activation is accumulated by convolving the input signal with a characteristic response function. In the second stage, a second system extracts information from the display, by randomly and independently sampling features with a rate proportional to the activation accumulated in the first system. Since features are sampled randomly, the probability of sampling new features is also proportional to the fraction of features not sampled yet, leading to an exponential convergence of P (which is related to the fraction of extracted features) to perfect performance. As remarked by Busey and Loftus (1994), as a "final puzzle", the model's result relies strongly on the fact that the rate with which information is extracted is exactly proportional to the accumulated activation in the first system (and not just some monotonic function of it). One of their concluding suggestions is that the neurophysiological mechanisms involve some other principles which imply an approximation to this behavior. We will show that such an approximation of the neurophysiological mechanisms can indeed be obtained, and that using signal detection theory, the simple law of exponential increase in performance can be derived both for d' and for P, under corresponding conditions.

The second type of experiments to be examined are standard choice reaction time experiments, where subjects are instructed to respond as fast as they can while maintaining a high degree of accuracy. In this case, our analysis focuses on the characteristics of RT distributions for correct and incorrect responses, and their dependency on stimulus variables that influence the probability of choosing each response alternative. As we shall explain below, efforts to understand such data in terms of classical stochastic models have not yet yielded completely satisfactory results. We will show that principles of neural information processing can be very valuable in the development of a more complete account. Of special interest are tasks where subjects have to deal with conflicting information, such as the Stroop or the Eriksen response competition tasks. As we shall show, characteristics of the data from these tasks indicate the essential role played by the key assumptions of our framework, particularly our assumption that there is mutual inhibition among units representing competing response alternatives. We will also show that a choice system based on these principles can also perform in a flexible way to select among a large number of alternative targets, as in perceptual identification tasks involving open-ended sets of response alternatives such as words.

We will begin by presenting a short description of the principles of our framework for information processing, and we will then show how the framework can be used to model time-accuracy curves in visual discrimination and memory retrieval tasks. Following that, we will examine the model's predictions for time-accuracy functions. We next examine standard choice reaction time tasks and follow this with a discussion of the role of lateral inhibition when there are multiple alternatives and/or multiple potentially conflicting sources of information.

From Neural Principles to Cognitive Models

Several principles of neural computation have emerged from the effort to model a variety of perceptual and cognitive tasks using ensembles of simple processing units (Anderson, Silverstein, Ritz & Jones, 1977; Cohen, Dunbar, & McClelland, 1990; Grossberg, 1978; McClelland, 1979, 1991, 1993; McClelland & Rumelhart, 1981). This framework, often called the Parallel-Distributed Processing (PDP) framework (Rumelhart, McClelland, and the PDP research group, 1986) assumes that perceptual and cognitive computations result from a parallel and cooperative process that involves large numbers of simple processing units. Five main principles for characterizing processing within such networks were discussed by McClelland (1993): these are the principles of graded, random, adaptive, interactive, and nonlinear processing. Although the principle that processing is adaptive is central to this approach, the details of the adaptation process are not yet fully clear, and will not be considered further in this article. In this section the other principles are described and their relation to principles of neurophysiology are considered.

Graded Propagation of Activation and Nonlinearity

The assumption of gradual propagation of activation has been motivated by contextual influences on processing (McClelland & Rumelhart, 1981). In respect to the dynamics of processing, the gradual activation framework (McClelland, 1979) suggested that the classical scheme of discrete processing stages should be replaced by a continuous (cascade) version. Support for this principle of continuous processing has recently been obtained through evoked potential studies of brain activation in information processing tasks (Gratton et al., 1988; Coles et al., 1985).

From a neurophysiological perspective, one can map the graded activation of a PDP unit to the mean firing rate of a population of neurons (Amit, 1989; Hertz, Krogh & Palmer, 1990). In neurophysiology, each neuron is characterized by two local variables: the local synaptic current and the unit's instantaneous firing rate, equivalent to the probability that the neuron will fire within a small time interval. One should also notice that both of these local variables are available at the cell level: the synaptic current is a characteristic of the input and can be measured at the cell body, while the response rate characterizes the cell's output. An important characteristic of the synaptic current is its passive exponential decay (Abbott, 1992; Amit & Tsodyks, 1991). For an isolated cell, the time constant for decay is determined by physiological factors to be very short (5-10 msec), unless modulated by recurrent excitation (as will be discussed below). This decay results in the "leakage" or dissipation of information over time, preventing the synaptic current from providing a lossless summation of incoming signals. This results in an exponential approach to an asymptotic synaptic current in response to a fixed input.

The synaptic current and firing rate are related by a nonlinear sigmoidal function which converts the synaptic currents into a firing rate. The exact form of this function, which depends on the physiological characteristics of the nerve cell, is less important. Often, logistic functions are used for their mathematical properties. Another function which captures most of the characteristics of neural response and its effect in networks, is the threshold-linear function. This function has the advantage of mathematical tractability, while preserving most of the essential nonlinear characteristics of the computation (Treves, 1990).

Our models are inspired by these characteristics of neurons. Each model unit can be thought of as corresponding to a population of neurons that collectively serve as its neural instantiation.[3] The unit has two variables: an input variable x, corresponding to the average synaptic current of the neural population (and to what is often called the net input to a unit in PDP models), and an output variable f, corresponding to the mean instantaneous firing rate of the neural population (and to what is often called the activation of the unit in PDP models). As in the cascade model of McClelland (1979) and in accord with the behavior of synaptic currents in neurons, we assume that x changes over time in response to a weighted sum of inputs from other units coupled with passive leakage or decay. To relate outputs to inputs we use the threshold linear function

$$\begin{array}{rcl} F(x) = & (x-\Theta), & x > \Theta \\ F(x) = & 0, & x \leq \Theta \end{array} \tag{1}$$

where Θ is a neural response threshold; in the following we will use $\Theta = 0$ for mathematical simplicity.[4]

Intrinsic Variability

Both human performance and single neurons show highly variable (probabilistic) behavior even under identical external stimulation. In the neuroscience literature, the issue of variability in the discharge patterns of nerve cells has recently been the focus of intensive research, both at the empirical (Softky & Koch, 1992, 1993) and the theoretical level (Shadlen & Newsome 1994; Usher et al., 1994, 1995). As emphasized by Shadlen and Newsome (1994), the variability in single cell discharge suggests that the code used by nerve cells is a graded rate code. At the cell population level, the rate can be defined as the number of spikes emitted by all members of the population within a short time period. Accordingly, one can think of the fluctuations in neural discharge as a source of noise or variability in this population rate variable. Interestingly, computational theories have demonstrated potential advantages of noisy computations over deterministic ones. For example noise can help to avoid local minima in the search for a global solution to a problem (Hinton & Sejnowski, 1983; Hopfield & Tank, 1985), and to destabilize spurious mixture states in attractor neural networks (see Amit, 1989). In addition, as shown by McClelland (1991) and by Movellan and McClelland (1995), noise is essential in order to explain the independent effects of context and stimulus information on perceptual identification (Morton, 1969; Massaro & Cohen, 1983) within an interactive activation model of word identification.

---

[3] For present purposes it is sufficient to think of each processing unit as corresponding to a non-overlapping set of neurons, each having identical input and output projections. This is a useful idealization of the real case, in which neurons are only partially redundant, and in which the populations that respond to different inputs are partially overlapping.

[4] The results developed in this paper do not depend on this choice, except for causing a delay in time-accuracy curves (Busey & Loftus, 1994), as we will discuss in the following sections.

Interactivity: Recurrent Excitation and Mutual Inhibition

The principle of interactivity in PDP models is motivated by the observation that in reading and other perceptual processes, decisions at all levels influence each other (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982; Rumelhart, 1977). From the neuroscience perspective, the abundance of interactivity is evident. For example, in the visual cortex, 80% of the input of cortical cells is provided by lateral connections from other cells within the same layer, or from connections from other layers further up in the processing stream, while only about 10% of the input to a cell originates within the earlier sensory layers (Douglas & Martin, 1990; Douglas, Koch, et al., 1995).

In this work we focus on two special cases of interactivity: recurrent self excitation and lateral inhibition. Recurrent excitation is thought to be due to mutually excitatory connections among neurons in the same population. As we shall see below, this recurrent self-excitation can play the role of modulating (increasing) the effective decay time constant of model units.

The neurophysiological data (e.g., Miller, Erikson & Desimone, 1995) is consistent with the idea that the strength of the recurrent excitation (and thus the effective time constant) increases along the ascending stream of processing. This provides a constraint for the rates of the units which in the cascade model (McClelland, 1979) were arbitrarily chosen. From the computational point of view the recurrent excitation can achieve two functions: the amplification of the incoming sensory input (Douglas & Martin, 1990; Douglas et al., 1995), and (when the recurrent effect is very strong) the maintenance of the neural activity during delay periods without stimulus (Amit, Brunel, & Tsodyks, 1994), obtaining thus a form of short term memory (See also Anderson, 1973; Anderson et al, 1977). In this work we will limit the discussion to the weaker effect of recurrent excitation, which, as we shall see, increases the effective time constant of passive decay even as it amplifies sensory responses.

The existence and role of mutual lateral inhibition is also well established. As for recurrent excitation, the physiological evidence for mutual recurrent inhibition is well documented (Chelazzi, Miller, Duncan & Desimone, 1993; Creutzfeld et al., 1966; Hess et al. 1975; Ratliff, 1965). Computationally, lateral competition allows contrast enhancement and the suppression of distractor activation (Cohen, Servan-Schreiber & McClelland, 1992; Feldman & Ballard 1982; Grossberg, 1976, 1978). In another context, Keele and Neill (1978) have proposed that lateral inhibition is a necessary component of selective attention mechanisms, being required to prevent uncontrolled spread of activation within the semantic memory space. Indeed, in attractor neural networks, lateral inhibition is essential in order to obtain sustained patterns of activity which do not spread to overlapping patterns (Hopfield, 1982; Amit, 1989).

Lateral but not Bottom-up Inhibition

In our model, between layer connections are purely excitatory; inhibition is used only within a layer, to implement competition among alternatives. Our decision to restrict the model in this way is motivated by computational considerations and inspired by aspects of the physiology of the neocortex. Computationally, there are distinct disadvantages to bottom-up inhibition. The problem arises whenever multiple alternatives are partially active at one layer of processing, and the information contained in these activations must be used at the next layer to select a response. It would be desirable to exploit the constraints imposed by these partial activations, but it can become difficult to do so if each partially activated alternative inhibits all of the alternatives it is inconsistent with at the next level.

Consider, for example, the interactive activation model of visual word recognition (McClelland

& Rumelhart, 1981), in which units at the word level receive inputs from position-specific letter-level units. If, in the first letter position of a word, the input is ambiguous, so that several letters are partially activated, it would be desirable for words consistent with one of these alternatives to receive some activation. In this way, words consistent with partially activated alternatives in all positions would receive support, and if only one word is consistent with one of the possibilities in each position, it will receive the strongest bottom-up excitation. However, under these conditions the presence of bottom-up inhibitory connections from letters to words can prevent this outcome, because the inhibitory influences the letters exert on words they are not consistent with can outweigh the excitatory effects that each letter exerts on the words that it is consistent with (McClelland, 1986). Suppose the stimulus WORD is viewed under degraded conditions, so that W, X, and V are all partially activated in the first position at the letter level, G, Q, and O are partially activated in the second position, R, P, and K are activated in the third position, and B, O and D are active in the fourth position.

In our example, the only word consistent with one of the possibilities in each position is WORD, and it would be useful if the mechanisms of perception were able to determine this. However, if each active alternative were to inhibit each word that it is not consistent with as much as it excites each word that it is consistent with, all of the word-level alternatives, including WORD, will receive net inhibitory input. For there to be a net excitatory effect in this case, it is necessary for the bottom-up inhibition to be less than half as strong as the bottom-up excitation. In general if N possible alternatives are active in a given position in one layer, the bottom-up inhibition exerted by each must be less than $1/(N-1)$ times the bottom-up excitation. The problem can be avoided entirely by eliminating bottom-up inhibition, and using lateral inhibition to suppress all but the most likely alternative. Many alternatives will receive some bottom-up excitation, but those that are the most consistent with the input will receive the most excitation. Lateral inhibition then allows this alternative to suppress all of the others, eventually dominating the computation.

In summary, lateral inhibition appears to be a powerful mechanism for selection, and we are not, by any means, the first to emphasize this (see, e.g., Grossberg, 1976, 1978). The problem is a very general one, and applies whenever the occasion arises to select the alternative that is most consistent with one or more sources of input. Such occasions arise continually, we believe, whenever one is processing any kind of natural stimuli (such as spoken or handwritten input) in the presence of and form of degradation (line noise in telephone circuits, faintness of signals due to distance, etc), and whenever one is attempting to exploit linguistic context, which generally supports multiple possibilities. In a later section we will specifically consider the case of using contextual and visual information in word recognition.

This differential computational utility of lateral as opposed to bottom-up inhibition may explain in part why the cortex of the brain appears to adhere to the same general constraint that we impose on our models. It is very well known that axonal projections from one brain area to another have purely excitatory effects, which within each local brain region there are both excitatory and inhibitory interactions. While it would be simplistic to identify conceptual levels of processing in models such as the interactive activation model with separate regions of the brain, it may be reasonable to construe each brain region, at least in the visual system, as deriving a particular type of representation of the input, based on inputs from many other regions. If so, the use of excitatory connections between areas and the use of inhibition to mediate competitive interactions within areas provides each region with the chance to select the overall representation that is most consistent with the possibilities suggested by information represented in every other area.

It may be worth noting that in the cortex, inhibition is mediated by interneurons. The so-called 'principal cells' (Shepherd, 1988) exert only excitatory influences on other neurons. They inhibit each other via their effects on the inhibitory interneurons, which can sum influences from many excitatory neurons, and then exert inhibitory influences back on the same excitatory neurons in a way that closely approximates the direct lateral inhibitory interactions that are assumed in our models. These interneurons can, of course, implement feed-forward as well as lateral inhibition, since they can receive excitatory input directly from other areas and 'invert the sign'. In our arguments are correct, this mechanism should not be used to implement a content-specific inhibition of one alternative by input that is inconsistent with it. Feedforward inhibition could, however, be used to implement a general input normalization operation (c.f., McNaughton & Morris, 1987) that would not be inconsistent with our proposals.

Representation: Local vs Distributed

For present purposes, we will be modeling cognitive processes using what are traditionally termed 'localist' assumptions in connectionist modeling. That is, we will assign a single dynamic variable to correspond to each response alternative, and will assume that both response choice and response timing depend on these variables. While we will continue in some contexts to use the term 'unit' to refer to such variables, we do not mean to imply that such variables represent individual neurons in the brain, and indeed we have already suggested that they may reflect the activity of populations of neurons in the brain. As such this suggestion introduces the idea that representation may be distributed over many neurons at the physiological level, even while it is treated as a single variable for the purpose of modeling the time-course of information processing in choice response situations (see Smolensky, 1986).

At a physiological level, it seems likely that individual neurons participate in the representation of particular inputs to varying degrees. Neurons generally appear to have fairly broad tuning curves, so that when a particular stimulus is shown different neurons will be activated to differing degrees (Tanaka and Saito, 1991; Miyashita, 1988). Furthermore, it seems likely that response alternatives may be represented by potentially overlapping populations. In general there are major advantages to using overlapping populations. Such overlapping patterns can encode the substructure of the items they represent and the similarities among different items, and affords an opportunity for generalization (See Hinton, McClelland, and Rumelhart, 1986, for discussion). Indeed, for structured mappings from stimuli to responses (e.g., pronunciation tasks, where the stimulus is a letter string and the response is a sequence of phonemes), the use of distributed representations allows generativity: A set of connections can be built up through experience on a body of examples (e.g., real words) that encode the regularities in the mapping and allow generation of responses both for familiar words and for novel but pronounceable nonwords. In such a system, alternative responses (e.g., "bark" and "gark") may involve many of the same active units, and only a few units that differentiate the two responses.

Our present analysis does not consider the effects of these overlapping patterns directly. We will consider cases in which an input partially activates two response alternatives, as it would if the optimal stimuli for each of these alternatives were overlapping input patterns. However, we will completely disregard the possibility of overlap of response representations, treating each response representation as completely distinct from the representation of each other response. We will return to the effects of relaxing this assumption in the General Discussion.

Another Important Simplification

The use of "localist" (i.e., single-valued) response variables in our simulations simplifies our analysis considerably. Further simplification is afforded by examining models in which all of the "action" occurs in a single layer of units, each representing a possible response alternative. We will assume a set of inputs to these units (and connection weights modulating the effects of these inputs on the unit corresponding to each alternative), but we will treat activations of these input units as occurring instantaneously, neglecting the dynamics of the input. This allows mathematical analysis and in particular detailed comparison with existing models in the classical reaction time literature.

Mathematical Formulation

Following these principles, we can quantitatively characterize the dynamic behavior for the synaptic input, x_i, of a unit i (for details see, Abbott, 1992, Amit & Tsodyks, 1991, Ermentrout, 1994). This is in fact a stochastic version of the cascade equation (McClelland, 1979):

$$\tau \frac{dx_i}{dt} = I_i - kx_i + \xi \tag{2}$$

This equation implies that the synaptic current of a unit, x_i, integrates the input received from other cells, I_i, with a characteristic rate constant k which reflects passive decay. The term ξ, which reflects intrinsic noise, is a stochastic Gaussian variable with zero mean and variance σ^2. Equation 2 reflects a leaky integrator. In the absence of noise this integrator converges exponentially (like a charging capacitor) to a steady state ($\frac{dx_i}{dt} = 0$) that is determined by $x_i^{asy} = \frac{I_i}{k}$; in the presence of noise the current x_i fluctuates around this "steady state". The input I_i can be decomposed into an internal (recurrent) component, I^{rec} (which originates from other cells in the same population) and is proportional to the unit's own response rate, $I^{rec} = \alpha f_i$, and an external source I^{ext}. Under a linear approximation ($f(x_i) = x_i$) Equation 2 leads to

$$\tau \frac{dx_i}{dt} = I_i^{ext} - (k - \alpha)x_i + \xi \ , \tag{3}$$

where I_i^{ext}, the input from outside the cell population, is the sum of the firing rates of other units weighted by their synaptic connections, ($\sum_j W_{ij} f_j$), and α is the weight of the self excitation.

As we observe in Equation 3 the self excitation has two effects on the dynamics: 1) it decreases the decay rate $k \to k - \alpha$, and 2) it amplifies the asymptotic response $x_i^{asy} = \frac{I^{ext}}{k-\alpha}$. In the very special case when $k = \alpha$, the systems becomes a perfect integrator.[5]

Equation 3 leads to a type of diffusion process, referred to as an 'Ornstein-Uhlenbeck' (OU) process (Ricciardi, 1977). The type of diffusion processes usually discussed in RT models (e.g., Ratcliff, 1978) are a special case of the OU process, which take place for a perfect integrator ($k = \alpha$). Recently, the general OU process has been introduced by Busemeyer and Townsend (1993) in decision making models.

Modeling Binary Choice Tasks

As we shall show later, our model can be applied to n-choice tasks. However, we first consider the case of two-alternatives, or binary choice, for ease of comparison with classical stochastic models

---

[5] In this case, as in the case where $\alpha > k$ a response-rate saturation non-linearity becomes essential to constrain the explosion of the response.

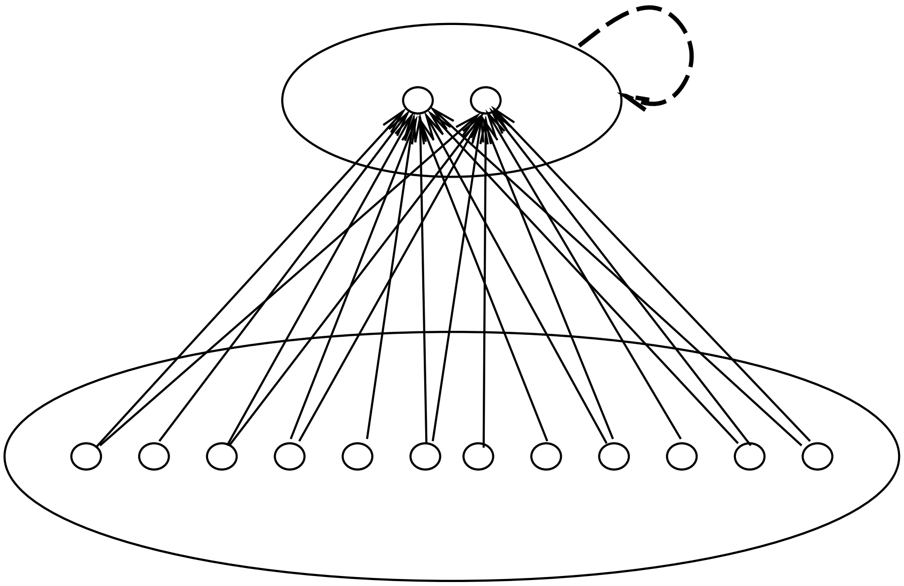

Figure 1: Network architecture for a binary choice task.

such as the classical random walk and diffusion models.[6] In what follows, we will use the terms *classical random walk* and *classical diffusion process* to refer to the pure forms of these models, in which their is neither "leakage" nor drift variance. The difference between the two models is simply that in the classical random walk one takes small discrete probabilistic steps of fixed size, either in a positive or negative direction, while in a classical diffusion process the process is thought to be fully continuous, subject to perturbation by Gaussian noise.

For our current purpose, we imagine a two-layer network, consisting of a set of input units indexed by j, over which the external input to the network is represented; and a set of output units indexed by i, one for each of the response alternatives, of which in this case there are two. The input units are thought of as having activations of 0 before the presentation of the stimulus, and then being set to values $f_j(\Phi)$ depending on the stimulus Φ.

We assume that the connection weights W_{ij} and the f_j are always ≥ 0 so that $I_i^{ext}(\Phi)$ will also be ≥ 0 for both output units. Denoting by S, the summed input, $S = \sum I_i^{ext}(\Phi)$, we consider the individual components of the input using the normalized components $\rho_i = \frac{I_i^{ext}(\Phi)}{S}$. With this normalization, the ρ_i sum to 1, so that in the case of binary choice, $\rho_2 = 1 - \rho_1$. To simplify the notation we choose $S = 1$; what this means is that the other parameters of the model (K,β, and α) are effectively scaled to the total summed input. Furthermore, we will assume that S remains constant across different stimulus conditions of the same experiment, so that the only variables that vary across conditions are ρ_1 and ρ_2.

When a stimulus is presented, the input units are clamped to the appropriate values, and the x_i of the response units are initialized at zero (we will also discuss later the effect of initialization to a random value with mean 0 and some variance). Activations of the response units are then allowed

---

[6] A random-walk process is defined as a trajectory which is initialized at the origin, and then at every time step moves to the right with probability p or to the left with probability $1 - p$. Random-walk models have been used to model choice RT (Laming, 1968; Stone, 1960) by assuming that a response is generated when the probabilistic trajectory crosses a response criterion, as further discussed in the following section.

to vary in accordance with Equation 3:

$$\begin{array}{rcl} \tau\frac{dx_1}{dt} & = & -kx_1 + \alpha f_1(x_1) - \beta f_2(x_2) \quad + \quad \rho_1 \quad + \xi_1 \\ \tau\frac{dx_2}{dt} & = & -kx_2 + \alpha f_2(x_2) - \beta f_1(x_1) \quad + \quad (1-\rho_1) \quad + \xi_2 \end{array} \tag{4}$$

We will use a threshold-linear type of non-linearity $f(x) = x$ for $x \geq 0$ and $f(x) = 0$ for $x < 0$. Interactivity comes into play in the model due to the recurrent self-excitatory weights (α) of the response units, and their mutual lateral inhibition (β).

When the activation of the two units is most of the time larger than zero,[7] Equation 4 can be approximated by

$$\begin{array}{rcl} \tau\frac{dx_1}{dt} & = & -(k-\alpha)x_1 - \beta x_2 \quad + \quad \rho_1 \quad + \xi_1 \\ \tau\frac{dx_2}{dt} & = & -(k-\alpha)x_2 - \beta x_1 \quad + \quad (1-\rho_1) \quad + \xi_2 \end{array} \tag{5}$$

In the following we will use the linear approximation 5 for mathematical analysis. However, in order to test the validity and the deviations of this approximation we will also perform numerical simulations for the non-linear version of these equations.

Parameters and Relating the Model to Data

Our interest in this paper mainly focuses on the general properties of the dynamics of choice, and so we do not provide detailed fits to the data from particular experiments. Rather, we seek to understand how various parameters affect the behavior of our models, and to understand whether particular parameter regimes provide explanations for the general trends that are found in the experimental data. As we shall see, different choices of parameters can at times produce very different results; and in some cases, there are clearly choices of parameters that lead to performance that is optimal (in ways we shall make explicit below) and others that lead to non-optimal performance that is actually more characteristic of the patterns seen in human data. In general, when we consider the model in the abstract, we will be concerned with understanding the effects of parameter variations. When we relate the model to experimental data, we choose parameters in ranges motivated by other aspects of our analysis that are also consistent with the patterns actually found in data. We will discuss these choices again below, where relevant.

Time-Accuracy Functions under Time-Controlled Conditions

Equations 5 and 4 fully characterize the system's dynamic behavior following the presentation of a stimulus. However, in order to also determine the moment when a response is generated, an additional type of assumption is needed. When the task is set up as a reaction time task, it is assumed that the subject does not respond until some preset criterion is required. Many types of response criteria have been proposed in the literature (Audeley & Pike, 1965; Laberge, 1962; Laming, 1968; Link & Heath, 1975; Ratcliff, 1978) and, as we discuss below, they lead to different results even when the same dynamics are assumed in the system. We begin our analysis by considering the predictions of the internal system dynamics independently of the separate and more problematic issue of the response criterion, by examining the data from experiments in which processing time is controlled, in order to factor out the influence of the response criterion.

First, we consider experiments using the response signal procedure (Corbett & Wickelgren, 1978; Dosher, 1976; Reed, 1973, 1976; McElree & Dosher, 1989; Wickelgren, 1977; Wickelgren &

---

[7] When one of the unit's activation reaches zero, due to lateral competition or to noise, it stops influencing the other positively activated unit (the β term disappears from that equation).

Corbett, 1977) in which subjects are required to respond immediately after a control signal. In cases where a binary choice between two alternatives is required, we assume that the response chosen will simply be the one for which the corresponding unit has accumulated the largest amount of activation at the moment the response signal triggers the response.[8] A similar rationale can be applied to experiments where the processing time is controlled by limiting the stimulus duration with a mask manipulation (Loftus et al., 1993). In this case, although the response time is not directly controlled, the optimal strategy for performing the task is to make the selection just as the mask degrades the stimulus representation, since from that point on the response units will be integrating noise rather than signal.

Many of the studies use a "yes-no" decision, where the task is to determine whether the stimulus previously occurred on a study list or "memory set". This task (and related tasks such as lexical decision) can be formulated in our framework in several different ways. For example, the two output units could be assigned to the "yes" and "no" responses, and the input units could correspond to detectors for particular items so that the activity of each one reflects the extent to which the input matches the particular item, based on inputs arising from previous layers. In this case the "yes" unit could simply sum the output of the units corresponding to items in the memory set, while the "no" unit might sum the output of units corresponding to items not on the memory set, or from a unit that simply reflects the presence of a stimulus regardless of its identity. The weight of this input can be chosen so that when the stimulus does not match any of the learned items the no-unit will tend to receive stronger input than the yes-unit. Within a distributed connectionist framework there are other possibilities that we will consider in the general discussion.

To obtain simple predictions for time-accuracy curves measured in the time-controlled experiments, one can subtract the two Equations 5, by defining a new variable $x = x_1 - x_2$ which reflects the excess activation of unit x_1 relative to that of unit x_2 (we use here the linear approximation, and we will explicitly test the non-linear effect in the following section):

$$\tau \frac{dx}{dt} = (2\rho_1 - 1) - (k - \alpha - \beta)x + \xi \; . \tag{6}$$

This makes the choice process equivalent to a classical 1D diffusion process, where $\nu = 2\rho_1 - 1$ is the diffusion drift (Ratcliff 1978; Ratcliff & Van Zandt, 1995) and σ is the noise variance. Unlike the classical diffusion process, however, the process described by Equations 6 is also characterized by a decay term with rate $K = k - \alpha - \beta$. Only in the special case when $K = 0$ does a classical diffusion process result.

In order to characterize the time-accuracy curves, a measure of d' has to be obtained. According to signal detection theory, d' is computed as the separation between the means of the signal and noise distributions, divided by their standard deviation (SD). For the binary and unbiased choice task discussed here, the signal distribution is the mean of the x variable (denoted as $\mu(t)$) when $\nu > 0$ ($\rho_1 > .5$), and the noise is the mean of $(-x)$ for the same diffusion drift ν (see Appendix), leading to

$$d'(t) = \frac{2\mu(t)}{SD(t)} \; . \tag{7}$$

The time evolution of the mean and the standard distribution of $x(t)$, as well as the time evolution of its distribution $P(x,t)$, can be obtained for the system 6, using the theory of stochastic

---

[8] Response biases can be obtained by introducing a bias parameter B; in this case response '1' could be chosen if $x_1 + B > x_2$. For simplicity we consider cases in which B is zero, noting that in most cases the dependent measure is d and changes in B should not have much effect.

processes (see Ricciardi, 1977, for the Fokker-Plank solution to Equation 6). If at time $t = 0$ the activation is initialized as $x(0) = 0$ then the distribution of x at a later time t is a Gaussian with distribution $P(x,t) = N[\mu(t), SD(t)]$, where the mean μ and the standard deviation σ are given by

$$\begin{aligned} \mu(t) &= \tfrac{\nu}{K}\tau(1 - \exp(-\tfrac{Kt}{\tau}))\,, \\ SD(t) &= \tfrac{\sigma}{\sqrt{K}}\sqrt{\tau(1 - \exp(-\tfrac{2Kt}{\tau}))}\,. \end{aligned} \tag{8}$$

From this (and using Equation 7) the time-accuracy function for the Ornstein-Uhlenbeck process is obtained:

$$d'_{OU}(t) = d_{asy}\frac{1 - \exp(-\frac{Kt}{\tau})}{\sqrt{1 - \exp(-\frac{2Kt}{\tau})}}\,, \tag{9}$$

where d_{asy}, the asymptotic value of d' obtained for unlimited time, is $d_{asy} = \frac{2\nu}{\sigma}\sqrt{\frac{\tau}{K}}$. It depends on the the diffusion rate, the standard deviation of the noise, and the decay rate constant. We notice that when K approaches zero, the integration time becomes infinite, as for a perfect integrator, and the asymptotic performance $d_{asy} \to \infty$, even for a highly degraded signal to noise ratio $\nu/\sigma << 1$. This is indeed expected since as the integration time increases the noise is progressively averaged out; an noiseless estimation of the stimulus can thus be obtained with infinite integration time.

The special case case $K = 0$ leads to a classical diffusion process. In this case the distribution $P(x,t)$ is also a Gaussian $N[\mu(t), SD(t)]$, whose mean increases linearly with time ($\mu(t) = \nu t$) and whose standard deviation increases as the square root of time ($SD(t) = \sigma\sqrt{t}$). In this case the accuracy increases with processing time to infinity, as $d' \propto \sqrt{t}$.

The bounded value of performance obtained in time-controlled experiments (Wickelgren, 1977; Reed, 1976; McElree & Dosher, 1989), can be explained by assuming that in addition to the noise that afflicts the activity within a trial, the drift parameter ν itself is also subject to variance from trial to trial. This model, introduced by Ratcliff (1978), is often simply called "the diffusion model" but we we call it the ***diffusion with drift variance*** (DDV) model to distinguish it from the classical diffusion model without drift variance. Taking this additional variance of χ into consideration, by averaging over a Gaussian drift distribution, one can show (Ratcliff, 1978) that the accuracy function of a DDV process becomes

$$d'_{DDV}(t) = \frac{d_{asy}}{\sqrt{1 + \frac{\sigma^2}{\chi^2 t}}}\,. \tag{10}$$

In Figure 2 we display the time-accuracy function for the OU process (Equation 9: dashed line) and for diffusion with drift variance (Equation 10: solid lines) for the same asymptotic value of performance. Although the two types of functions have some similar qualitative properties (a gradual increase to asymptotic value), one can clearly see that the OU curve (dashed line) is very close to a pure exponential (solid) line), while the curve of the DDV model (dotted line) approaches the asymptote faster at the beginning and more slowly toward the end (compared to the exponential) due to its power-law dependency on time.

Unlike the Wickelgren (1977) time accuracy curves, the graph in Figure 2 is not delayed; the exponential begins from $t = 0$. However, as mentioned before, $t = 0$ is taken as the moment when the stimulus instantaneously sets the activations of the input units. Presumably, there are additional processing stages between the actual sensory surface and the input layer of the model, as well as additional stages between output units and the muscle contractions effecting the overt

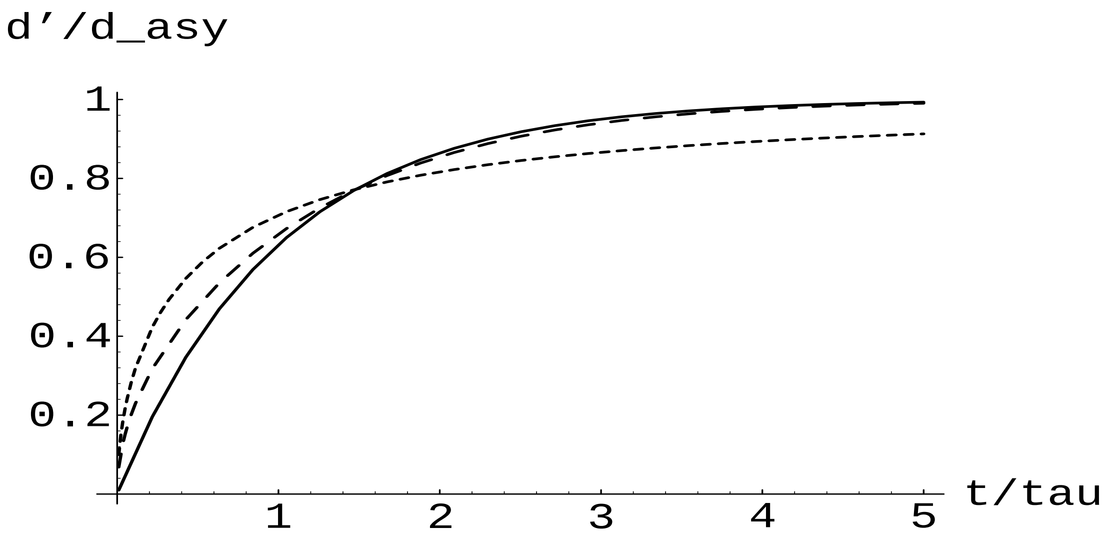

Figure 2: Time accuracy functions for the diffusion with drift variance (DDV) model ($\frac{\sigma}{\chi} = 1$; dotted line), the Ornstein-Uhlenbeck process ($K = .93$; dashed line), and Wickelgren-type exponential curve (solid). Observe that the DDV curve is farther away from the exponential curve than the OU curve, both at small and at large times. No value of $\frac{\sigma}{\chi}$ can bring the DDV curve closer to the exponential.

response. Such stages would of course introduce delays, leading to a shift in the rise of the time-accuracy curve, even if the propagation of signals through the layers of the system is cascaded rather than discrete (McClelland, 1979). The effect of cascading is to soften the initial rise of the time-accuracy curve. This would have the effect of bringing the output of a multi-layer system into very close correspondence with a shifted exponential, even if the rate-limiting process has the dynamics of the OU process characterized by Equation 5.

In a time-controlled version of the Sternberg task, McElree and Dosher (1989) fitted the experimental data with both the delayed-exponential curve proposed by Wickelgren (1977) and with the DDV curves (Equation 10); consistently within all experimental manipulations the exponential curves gave a better fit to the data. The DDV curves seem to approach asymptote more slowly than the data, leading to systematic distortions in the fit.

In Figure 2 we have shown that for $Kt > 1$, $d'(t)$ for the OU process is very close to an exponential. A more convenient way to test graphically the fit to the exponential is to use a log-linear graph. If some function $g(t)$ approximates an exponential approach to asymptote of the form $1 - \exp(-Kt)$, then $-\log[1 - g(t)]$ should be linear with time. Thus the adequacy of the exponential approximation can be assessed by inspecting the linearity of $-\log[1 - g(t)]$. In Figure 3.a we display $-\log(1 - \frac{d'(t)}{d_{asy}})$ for the OU diffusion (solid line) and for the DDV model. The graph shows a linear trend for the OU process. This can be understood by observing that for $t > 1/K$, d' in Equation 9 is dominated by the numerator leading to an exponential approach to asymptote. Clearly, the approach to asymptote is not exponential in the DDV case.

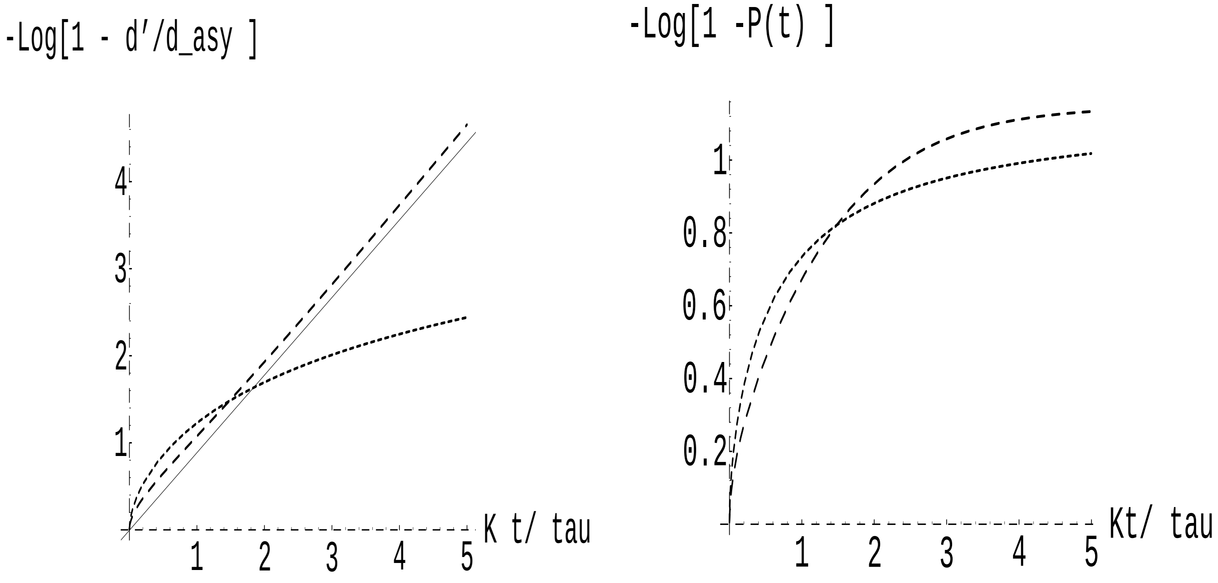

Figure 3: Time-accuracy curves for the Ornstein-Uhlenbeck (OU) model (dashed line) and the diffusion with drift variance model (DDV) (dotted line). a) $-\log(1-\frac{d'(t)}{d_{asy}})$, b) $\log(1-P_c(t))$. The stimulus discriminability was $d_{asy}=2$.

High Accuracy Data

In the situations just considered, asymptotic accuracy was assumed to be moderate, with d'_{asy} equal to a moderate value (say between 1 an 4). We now consider experiments in which the asymptotic accuracy is very high, corresponding to a very high d'-value. This is often the case in visual discrimination experiments where performance is controlled by a mask presented at a stimulus onset asynchrony (SOA), t, following the stimulus (Loftus et al., 1993), or when the duration of the stimulus is controlled (Busey & Loftus, 1994). Treating the mask as forcing a response at time t, we observe in Figure 2, when $t < T = 1/K$ the OU time-accuracy curve shows some deviation from a pure exponential (full and dashed lines). Those are however the time intervals where most of the data is collected in a time-accuracy experiment where the asymptotic accuracy is very high (since at longer times, performance is too close to 1 for accurate tracking of d').

In these cases, it is more useful to examine time accuracy curves in terms of the probability correct, P, or in terms of the probability of correct responses corrected for guessing, which for non-biased binary choices is $P_c = 2P - 1$ (Busey & Loftus, 1994; Loftus et al., 1993)). Those quantities are related to d' in accordance with (see Appendix):

$$P = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} \exp(-\frac{(d'/2)^2}{2}) dy \tag{11}$$

$$P_c = \frac{2}{\sqrt{2\pi}} \int_{0}^{x} \exp(-\frac{(d'/2)^2}{2}) dy \tag{12}$$

One can see that both P and P_c are non-linearly related to d'. As a consequence, either d' or the probability correct (but not both of them) can approach exponentially towards an asymptotic level, *within the same time range.* In light of this, it is not surprising that when d' is approximately

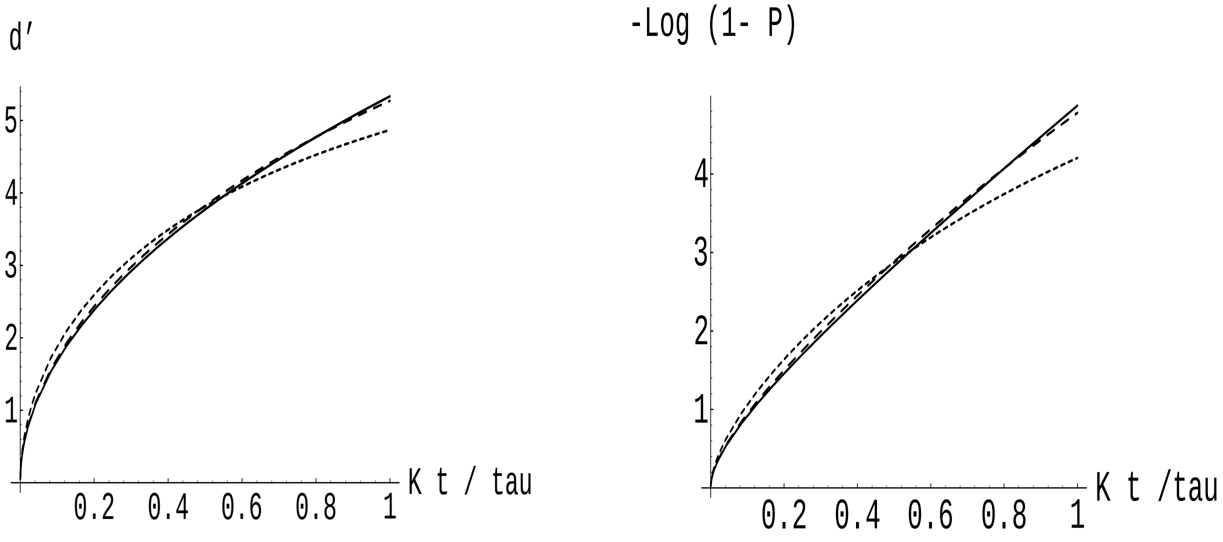

Figure 4: Time-accuracy curves for 3 processes with high discriminability: classical diffusion (without drift variance, solid line) OU process (dashed line, overlapping with the classical diffusion model); and DDV model (diffusion with drift variance, dotted line). a) $d'(t)$, b)$-\log(1-P_c(t))$. For $t<1$ the OU process is almost identical to the classical diffusion model resulting in the almost identical and linear curves for $-\log(1-P_c)$. As opposed to this, the DDV process does not approximate an exponential dependence on $P_c(t)$.

exponential for the OU process in the range of time $Kt<5$, as shown in Figure 3.a, P is not, as shown in Figure 3.b.

Surprisingly, at the relatively short time intervals $t<1/K$ that are used in high-accuracy experiments, an approximation for P (and for P_c) can be developed showing an exponential dependency as reported in experiments by Loftus et al. (1993). To see this, we note that for $t<1/K$ Equation 9 can be approximated (up to higher order corrections) by:

$$d'_{OU}=d_{asy}\frac{1-\exp(-Kt)}{\sqrt{1-\exp(-2Kt)}}\approx d_{asy}\sqrt{\frac{1-\exp(-Kt)}{1+\exp(-Kt)}}\approx d_{asy}\sqrt{\frac{K}{2}t}\ . \tag{13}$$

Thus the d' curve for the OU process increases at the beginning as the square root of time, as does a classical diffusion process without drift variance. The classical diffusion with drift variance, Equation 10, can also be approximated at short intervals by a $\sqrt{t}$ function. However in this case the validity range of the approximation is smaller (Figure 4.a.)

Using Equations 11,12,13 and the approximation (valid for $x>.5$)

$$\frac{2}{\sqrt{\pi}}\int_0^x \exp(-y^2)\approx 1-\frac{\exp(-x^2)}{\sqrt{\pi}}$$

an exponential increase in P_c and in P is obtained. As shown in Figure 4b, the OU process and the classical diffusion process lead to an exponential curve for $P_c(t)$ (within a time range covering values of P_c between 0 and .99) while the DDV process leads to a non-exponential curve. Typical experimental data (Loftus et al., 1993; Busey & Loftus, 1994) shows that $-\log(1-P_c)$ increases linearly with time; however the increase in performance is delayed ($P_c = 0$ for $t < t_0$). Notice that assuming the information is delayed while traveling through previous layers before reaching the choice representation is not enough to explain the delay in the accuracy as a function of SOA, if both the stimulus and the mask are delayed by the same amount. Two factors could contribute to the delay effect seen in actual data. First, theories on the nature of visual masking indicate that a noise visual mask can 'catch-up' with a previous stimulus (up to some SOA limit), due to an inhibitory interaction between mask-generated fast transient activity and stimulus-generated slower sustained activity (Breitmeyer & Ganz, 1976). This mechanism would lead to a shifted exponential curve for P_c as a function of SOA. However, a much smaller delay effect (of about 10-20 msec) is still found even in experiments where the stimulus duration is varied and where no mask at all is used (Busey & Loftus, 1994). This additional effect can be explained by assuming a threshold, $\Theta > 0$, in the response function of units, $F(x)$. If the response remains linear above the threshold, this would not change any of the other dynamical characteristics of the system. Accordingly, response units would only begin to respond when the synaptic current, x, exceeds the threshold Θ. The time required for the synaptic current on the response decision units to reach this threshold would then account for the additional delay in the time accuracy-function (weak and short sub-threshold stimuli might not reach this threshold at all), and would also account for priming effects of stimuli too brief to elicit above-chance responding themselves (Loftus et al., 1993).

Effects of Non-Linearity

In Equation 6 we assumed that the activation of both response units is positive. This allowed us to obtain an effective one dimensional diffusion process, in which lateral inhibition had the role of scaling down the rate constant and increasing asymptotic separation of curves associated with different stimuli. In order to test the validity of this result and to better understand the role played by the nonlinearity and lateral inhibition, we ran simulations of the diffusion process described by the two dimensional system in Equation 4. In each simulation trial the units were initialized at the origin and then the equations were numerically integrated, with the addition of Gaussian noise of variance σ.[9] We find that the result obtained in Equation 9 is correct despite the approximation. In Figure 5 we show $P(t)$ for four levels of competition ($\beta = 0, .2, .4, .6$) and the same net leakage, $(k-\alpha) = .2$, and stimulus discriminability.

The upper dashed-line corresponds to a competition level that exactly balances the leakage ($(k-\alpha) = \beta = .2$). An identical curve is obtained for $(k-\alpha) = 0, \beta = 0$ (a classical diffusion process) showing that the time-accuracy curves depend only on $K = k - \alpha - \beta$. As discussed above, balanced diffusion leads to an optimal decision process by averaging the noise, and thus perfect performance is reached (P tends to 1). Similarly, simulations with $(k-\alpha) = .2, \beta = 0$ (middle continuous line) produce identical speed-accuracy curves as simulations with $(k-\alpha) = .2, \beta = .4$ (middle, diamond-symbols) and with $(k-\alpha) = -.2, \beta = 0$ (not shown). This is somehow more surprising, implying a symmetry of the OU speed-accuracy curves to the inversion of K (from .2 to $-.2$). A negative value of K in Equation 6 has the implication that the system is dominated

---

[9] Numeric integrations were performed using Euler's method with steps of $\frac{dt}{\tau} = .1$. When integrating Equations 4, sometimes the synaptic current, x, of a unit becomes negative. In that case this negative value is tracked by the simulations, but its output is zero (due to the threshold non-linearity) and thus it does not affect any unit in the network).

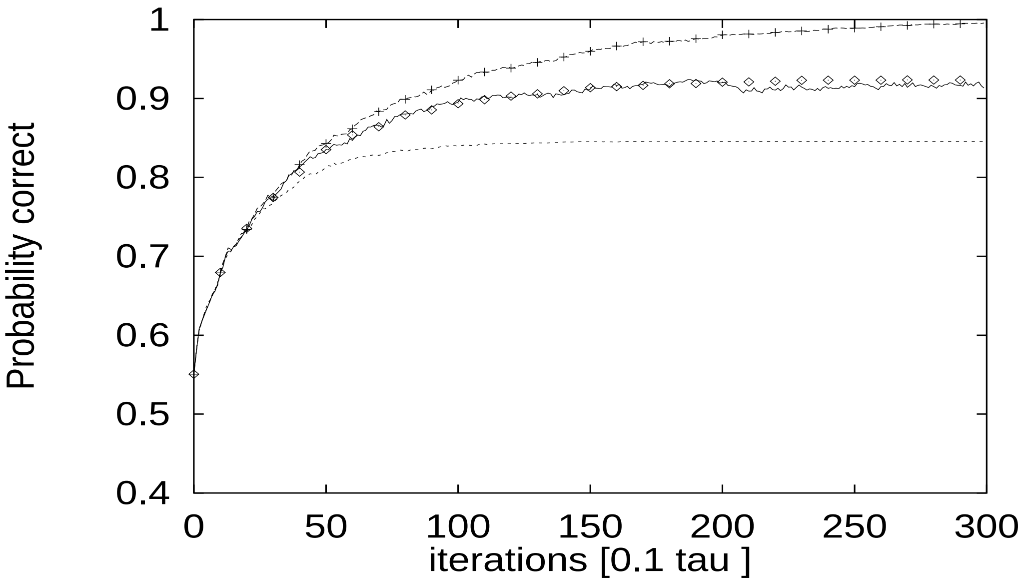

Figure 5: Dependency of the time-accuracy curves on the level of competition. The probability correct $P(t) = prob[x_1(t) > x_2(t)]$ is displayed as a function of time. Upper dashed line: balanced inhibition and leakage ($\beta = (k-\alpha) = .2$); overlapping '+'symbols: $\beta = 0, (k-\alpha) = 0$. Middle continuous line: $\beta = 0, (k-\alpha) = .2$; overlapping diamond-symbols: $\beta = .4, (k-\alpha) = .2$. Lower, dotted line: $\beta = .6, (k-\alpha) = .2$. Notice, conditions which have the same value of $|K|$, result in equivalent time-accuracy curves. Stimulus discriminability corresponds to $\nu = 2\rho_1 - 1 = .1$ and $\sigma = .5$, thus the signal/noise ratio is 0.2.

by divergence instead of decay (this can happen with strong recurrent excitation or strong lateral inhibition). This symmetry can be understood by checking that the OU d' time-accuracy functions Equation 9 are invariant under $[1-\exp(-\frac{Kt}{\tau})] \rightarrow [\exp(\frac{Kt}{\tau})-1]$:

$$\frac{1-\exp(-\frac{Kt}{\tau})}{\sqrt{1-\exp(-\frac{2Kt}{\tau})}} = \frac{\exp(\frac{Kt}{\tau})-1}{\sqrt{\exp(\frac{2Kt}{\tau})-1}}\,. \tag{14}$$

Recently, the OU diffusion with both negative and positive K value has been used by Busemeyer and Townsend (1993) in decision making, where the sign of the K coefficient relates to approach-avoidance characteristics of the decision. When a fixed response boundary is used (in terms of $x = x_1 - x_2$), it can be shown (Busemeyer & Townsend, 1993) that the sign of K achieves a speed-accuracy tradeoff: with positive K decisions have a higher probability of finding the correct stimulus but are slower, while for negative K values decisions are faster but less precise. From the neurophysiological perspective, positive K values (dissipation) may be characteristic of earlier perceptual brain areas, while negative K values might characterize brain areas that specialize in the sustaining activity in the absence of stimulation, such as the prefrontal cortex (Fuster, 1991; Funahashi, Chafee & Goldman-Rakic, 1993; Goldman-Rakic, 1992).

The symmetry of the OU speed-accuracy curves to K inversion implies therefore that the two strategies are equivalent in terms of what we might call their 'sensitivity-dynamics' (where sensitivity refers to the separation of distributions of activation generated by inputs from the distributions associated with the different response classes). Despite this symmetry however, the actual dynamics of the underlying activations is very different in these different situations (e.g., $k = .2, \beta = .4$ or $k = .2, \beta = 0$, or $k = -.2, \beta = 0$), as shown in Figure 6. When the system is dominated by decay (positive K) the value of x is bounded, and later inputs have more impact on the current trajectory. When K is negative earlier inputs are dominating the response and the trajectories diverge (in the absence of controlling non-linearities).

In order to illustrate those internal characteristics of the system, we display in Figure 6 the density distributions of x at three time intervals from the stimulus presentation, for the system described by Equation 4, with the parameters mentioned above. We observe that for the purely dissipative case, $k = .2, \beta = 0$, the distributions remain bounded, while for $k = -.2, \beta = 0$ they diverge. While in the second case the mean of the distribution moves faster towards the correct response side, the standard deviation increases in a corresponding manner, so that their ratio (which determines d') is the same in the two cases.

For $k = .2, \beta = 0.4$ we observe a new phenomenon. Despite the fact that d' and P are the same in the previous two cases, the distribution becomes bimodal. This is due to the fact that with inhibition, eventually one of the units is suppressed to zero. When this happens, due to the non-linearity, the effect of inhibition on the active unit disappears, and thus the active unit remains subject only to dissipation. Unlike a linear system, a system based on dissipation, competition and non-linearity is able to actually generate pattern segregation. Despite the fact that the time-accuracy curve is unaffected (the area at the right of origin remains the same) the non-linearity might make an important difference to actual RTs. The assumption under which the time accuracy curves discussed above were obtained was that the subject simply chooses the unit that is most active at the moment of decision. However, this assumption is an idealization; in cases where the difference is small, the decision itself might be subject to error, and/or it might itself take time. Because the segregating effect of competition tends to eliminate cases in which the difference in activation between the two alternatives is small, this process might well actually lead to a facilitation of actual responding.

Standard Choice RT Tasks

In this section we analyze our model's behavior in standard choice RT tasks. Modeling these tasks differs from modeling time-controlled experiments, in that to model standard choice RT it is necessary to make some assumptions about the conditions that must be satisfied in order to trigger a response. In many models, the concept of a criterion for responding is introduced.

It may seem that models that employ response criteria require some external decision making process that looks at the states of units and asks whether the criterion has been reached. However, there is an alternative possibility, which we propose is the one that is actually used by subjects in choice reaction time experiments. According to this approach, there is no separate external decision making process. Instead, we simply view responses as "downstream" events, outside of our model, that are triggered when activation within the model reach certain states. In this case, the critical value that must be exceeded to trigger a response is still a parameter that may be subject to strategic modulation, but it is not an abstract criterion that is applied by an external mechanism that then 'decides' to initiate a response.

In this paper, we will not actually simulate the response triggering process *per se*. Instead, we

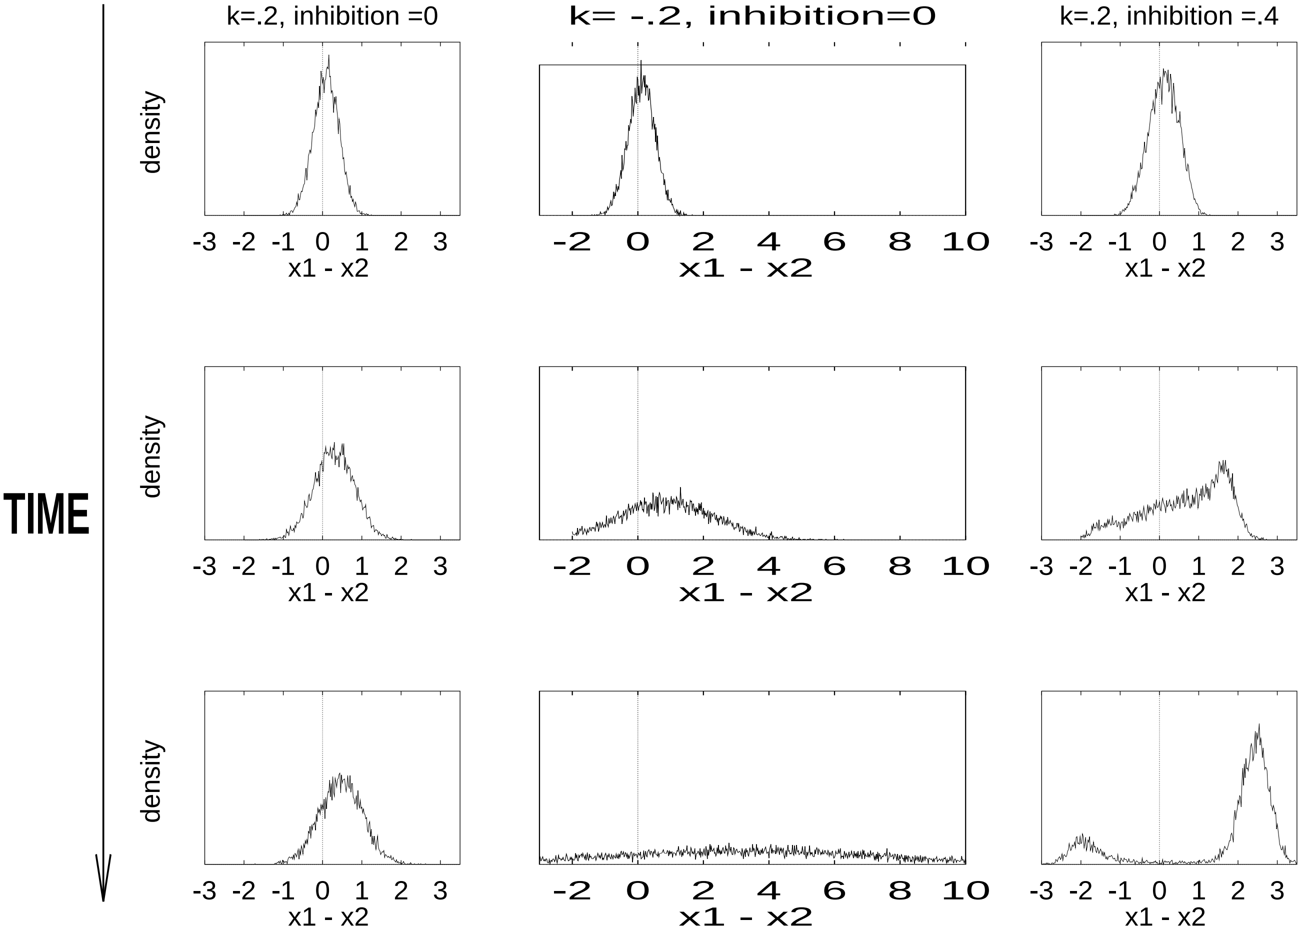

Figure 6: Density distribution for $x_1 - x_2$. Left column: Dissipative case $(k - \alpha = .2)$. Center column: self-expanding, $(k - \alpha) = -.2, \beta = 0$. Right column, dissipation and inhibition $(k - \alpha) = .2, \beta = 0.4$. Top panels: $t = \tau$, Middle panels: $t = 5\tau$, Bottom panels: $t = 19.5\tau$ for the left and right panels (asymptotic distributions) and $t = 10\tau$ for the middle-low panel.

will adopt the classical assumption that for each response unit there is a critical value, such that when the activation of the unit reaches that value, a response is triggered. With this assumption in mind, we can now ask how our leaky, competing accumulator model compares relative to other possible alternative models.

The first point to note is that a special case of our model, in which there is neither leakage nor inhibition (that is, in which $k = \beta = 0$), is equivalent to a continuous time version of what we will call the classical accumulator model (also known as the recruitment model; Audeley & Pike, 1965; Laberge, 1962; Vickers, 1970; Vickers et al, 1971; Wilding, 1974). Historically, several related models have been labeled by the name of 'accumulator' model. For example, Audely and Pike (1965) used this label to refer to a counter model, where at discrete moments in time a unit is added, either to the counter representing the first choice alternative (with probability p) or to the counter corresponding to the second choice alternative (with probability $1 - p$), and a decision is reached when one of counters reaches an absolute criterion. Later, Vickers (1970) used the same label for a model where, in addition to the assumptions described above, the amount of activation that is added to the counters is not a constant but rather a stochastic Gaussian distributed variable (of variance, say δ). A discrete-time accumulator model with probability p and noise variance δ, is equivalent (in the limit of small time steps) to a 2-dimensional diffusion process (as in Equation 5), with drifts $\rho_1 = p$, $\rho_2 = 1 - p$ and a noise-standard deviation of $\delta + p(1 - p)$. In any case, according to this class of models, choice reaction time responses depend on a race between two independent accumulators. As soon as one accumulator reaches a critical value, called an *absolute response criterion*, the response is triggered, regardless of the state of the other accumulator(s) at that time. This is equivalent to the assumption of our model that a response is triggered when the state of one of the response units reaches a critical value.

The classical accumulator model has been the subject of intense analysis and comparison to the classical random walk (where a trajectory initialized at the origin moves at every time step either to the right with probability p, or to the left with probability $1 - p$) or its continuous counterpart, the classical diffusion model (Laming, 1968; Link, 1975; Link & Heath, 1975; Ratcliff, 1978; Stone, 1960). A key result is that the classical random walk (and similarly the classical diffusion model) is an optimal decision making model, in the sense that if the criterion is chosen to achieve some specific desired error probability, it reaches a decision faster on average than any other system (Luce, 1986; Wald & Wolfowitz, 1948). The classical accumulator model is at a distinct disadvantage in this regard, since in this case the decision is based on less information. The reason for the difference can be seen most clearly by comparing the classical random walk with the accumulator model of Audely and Pike, since both involve accumulation of signals supporting one alternative (with probability p) or the other (with probability $1 - p$). In the classical random walk, the decision is based on the cumulated difference between the number of signals supporting each of the two alternatives, while in the classical accumulator the decision is based on the absolute amount of support for each alternative. An equivalent level of accuracy can be achieved by the classical accumulator model, but only by adopting a criterion that takes longer to reach.

One way to increase the efficiency of our model would be to replace the absolute criterion in that model with a difference criterion. In fact, the special case of our model using two independent, non-leaky units with $k = \beta = 0$ becomes equivalent to the classical diffusion model rather than the classical accumulator model if response selection is based on the difference in activation between the two response units. According to this criterion, response-'1' is selected when the difference in the activation of the response units, $x = x_1 - x_2$, exceeds a critical value. With the addition of some decay ($k > 0$), a one dimensional OU-diffusion model (Busemeyer & Townsend, 1993) is obtained. With variation from trial to trial in the input variable ρ_1 (and corresponding variation in $\rho_2 = 1 - \rho_1$), the model becomes equivalent to the diffusion model with drift variance (Ratcliff, 1978).

Based on these points, it seems at first glance that triggering a response in our model when the state of one of the response units reaches a critical value places the system at a distinct disadvantage relative to the classical random walk and diffusion models. However, we will show in what follows that this disadvantage is almost completely eliminated when lateral inhibitory interactions between alternative response units are introduced. The use of lateral inhibition essentially allows each unit's state to reflect the state of all of the others, in a way that generalizes flexibly, independently of the number of alternatives; an absolute criterion can then be applied independently to the resulting activations of all response units. As we shall see, the use of lateral inhibition makes it possible to achieve a time-efficiency far superior to that of independent accumulators with absolute criteria, and nearly equivalent to what would be possible using the optimal, difference-based criterion.

In order to demonstrate this point, we performed simulations, using the mutually inhibitory or competing accumulator model (described by Equation 15) and the independent accumulator model (Equation 16). We begin by considering the case of two alternatives. For non-leaky independent units, we consider both the absolute criterion and the difference response criterion; as previously described, the latter is more efficient. For the competing units, we consider only the absolute criterion. Initially, we consider both models under conditions that led to optimal time-accuracy curves, which, as we saw before (Figure 5, occurs when $K = 0$. For independent units this amounts to assuming neither leakage nor inhibition, and for mutually inhibitory units it amounts to assuming that the inhibition parameter β exactly balances the leakage k. The goal is to see whether the incorporation of lateral inhibition in the competing model allows it to approach the optimally efficient choice strategy, obtained with the use of a difference criterion with independent accumulators. To address this issue, speed-accuracy tradeoff functions were produced by varying the response criteria; requiring a larger absolute activation (or a larger difference between the two units activation) leads to greater accuracy, but at a cost of greater time. One can then plot the resulting speed-accuracy tradeoff functions, to see how they compare in terms of the time each requires to reach a particular level of accuracy.

The left panel of Figure 7 shows the resulting time-accuracy curves. As expected, for non-leaky independent response units, the difference criteria (upper line with x's) produce better results than absolute criteria (dashed line with square symbols). The next question is, how does the performance of the competing choice system with absolute criteria compare to the performance of these other systems? The competing leaky system (solid line with diamonds) is far better than the independent (leaky or non-leaky) system with absolute criteria (lower line with '+' symbols and dotted line with square symbols, respectively) and almost as good as the independent (non-leaky) system with difference criteria (upper line). It appears that balanced, leaky-competing units can approach the optimal behavior of non-leaky units with a difference criterion.[10]

Another benefit of the use of lateral inhibition should also be noted. As we saw earlier in our discussion of time-accuracy curves from time-controlled experiments, lateral inhibition can counteract the effects of leakage, helping to preserve information against loss. The use of a criterion based on the difference in activation of two leaky units cannot do this. We examined this issue by considering the effects of leakage in the independent accumulator models (for a leakage, $k = .5$),

---

[10] The adequacy of the approximation to optimality is affected by the magnitude of the lateral inhibition. Using $\beta = .7$ (and a corresponding value for k) led to the best results in these simulations. For a given value of k, performance is closest to optimal when β equals k.

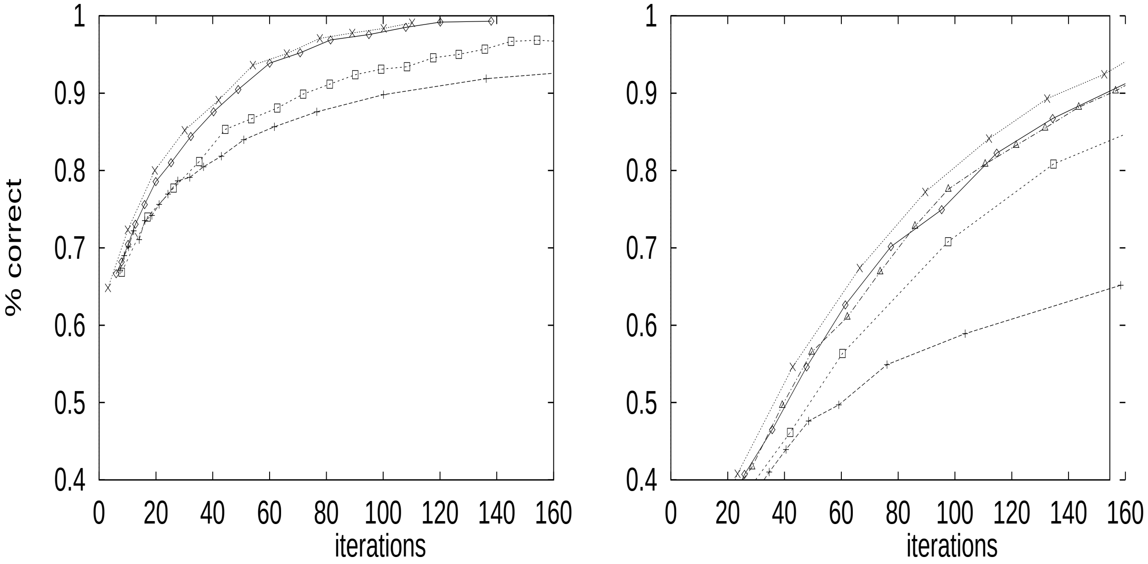

Figure 7: Left panel: Time accuracy curves for choice reaction time responses with two alternatives. The solid line with diamonds shows the case of leaky, competing accumulators ($k = \beta = .5$), with an absolute activation criterion for responding. Note that this produces results almost as good as the optimal performance that is achieved by independent, non-leaky accumulators ($k = \beta = 0$) and the use of a difference criterion, as shown in the upper line with 'x's. Two other models that do less well are also shown: Independent non-leaky accumulators with an absolute criterion ($k = 0, \beta = 0$, dotted line with squares); and independent, leaky accumulators with an absolute criterion ($k = .5, \beta = 0$; lower dashed line with '+'). Right panel: Time accuracy curves for a nine-alternative choice. Leaky, competing accumulators ($k = \beta = .5$; solid line with diamonds), are compared to leaky non-competing accumulators ($k = .0, \beta = 0$) using three different criteria: the absolute criterion (dotted line with squares); the max-vs-next criterion (solid line with x's) and the max-vs-ave criterion (solid line with triangles). See text for further explanation of the criteria. Also shown is the case of independent, leaky accumulators ($k = 0.5; \beta = 0$; dashed line with '+') with the absolute criterion. The 'correct' unit received input $I + \Delta = .55$ and all others received input $I = .45$. The noise χ had standard deviation $\sigma = .5$.

using both absolute and difference criteria. For clarity, only the case of absolute criteria is shown (lower dashed line with '+'). Although the difference-based criterion produced better results than the absolute criterion, it did not perform as well as the model using leaky, competing units.

We next considered the use of leaky, competing accumulators in a multi-choice situation. Here we performed simulations for Equations 15 with $n = 9$ choice alternatives (one unit receives input $I + \Delta$ and all of the other eight receive an input of I). We compared leaky, competing accumulators with absolute response criteria to non-leaky, independent accumulators under three different types of criteria: an absolute criterion, in which a response is made when the activation of any unit reaches threshold, and two possible generalizations of the difference criterion that is optimal for the two alternative case, specifically a *most-vs-next* criterion, in which a response is made when the activation of the most active detector exceeds the activation of the next most active other detector by a criterial amount, and a *most-vs-ave* criterion, in which the response is made when the activation of a detector exceeds the average of the activations of the other detectors by a criterial amount. The results are shown in the right panel of Figure 7.

Considering first the non-leaky, independent accumulators, the max-vs-next criterion (solid line with 'x') is the most efficient. We are not aware of analytic results on whether this represents the maximum possible efficiency; however, among the cases considered here it is certainly the best. The max-vs-ave criterion (solid like with triangles) is not quite as good, but both are much more efficient than the *absolute* criterion (dotted line with squares). Interestingly, the leaky, competing accumulator model (solid line with diamonds) produces performance that is indistinguishable from that of the non-leaky independent accumulators with the max-vs-ave criterion, and far better than the absolute criterion. Once again the use of leakage without inhibition produces poorer performance; as before we show only that case of the absolute criterion (dashed line with '+').

It is important to note that the efficiency of the model based on leaky competing units is relatively independent of parameters. The model converges to the (inefficient) classical accumulator model when k and β are close to 0, but efficiency increases rapidly as k and β increase, and very similar results are obtained for a range of values of k and β from 0.2 to 0.7, as long as $k = \beta$. Also, it is not necessary to rescale k and β as the number of alternatives increases; the same values can be used for $N = 2$ and $N = 9$ (as shown in Figure 7.

Based on these results, it appears that leaky, competing integrators can produce performance that is fairly close to optimal in choice reaction time tasks in a very general and flexible way. It is true that this approach is not perfectly optimal, but it is far better than what can be achieved with an absolute criterion, if inhibitory interactions are not used.

The non-interacting systems based on the use of a difference criterion can produce optimal performance when there are only two alternatives, and the most-vs-next and most-vs-ave generalizations of the difference criterion perform well on multi-choice decisions. However, it seems likely that the slight advantages of independent, non-leaky accumulators using the difference criterion and its most-vs-next generalization could only be obtained in the brain as some cost in complexity. While a simple difference between two alternatives would be relatively easy to compute, its extension to a generalized max-vs-next computation seems relatively complex. Even the computation of max-vs-ave values requires normalization of the influence of each alternative for the total number of alternatives, so that an average can be computed. While the use of lateral inhibitory interactions among leaky accumulators is not strictly identical to either of these two approaches, it achieves much the same effect. The information about the state of other accumulators that is necessary for efficiency is propagated to each alternative by the lateral inhibition, with the advantage that information is then localized in a each response unit, whose individual state then determines the triggering of the response.

In any case, an optimality analysis is quite a separate matter from an account for the details of human performance. We saw previously that human performance time-controlled studies does not in fact conform to optimality: in that case, optimal performance is obtained with a non-leaky integrator, or, as we saw, with leaky competing accumulators with balanced decay and inhibition, leading to $K = 0$. The fact that human time-accuracy curves asymptote at finite d' suggests that in fact there is information loss in the human information processing system, and the form of time-accuracy curves is consistent with the assumption that K is not equal to 0. In what follows, we will adopt the assumption that in fact inhibition, β is larger than decay, k, since as we showed earlier this leads the choice variables to become segregated while they remain bounded.

Accounting for Latency Probability (LP) Functions

Up to now we have compared our leaky, competing accumulator model with absolute response criteria to other models on the basis of optimality and flexibility, but we have yet to consider how it might fare in relation to experimental data. In this section we consider data derived from a paradigm that was popular some years ago in efforts to compare and contrast the predictions of the classical accumulator and random-walk models. We first consider the contrasting predictions of the classical models concerning the way in which reactions times should vary with stimulus characteristics, such as the degree to which the stimulus favors one response alternative over the other. We note that neither classical model was completely successful, and then consider two models that fare better: Ratcliff's diffusion with drift variance model (DDV) and our competing, leaky integrator model.

To make the discussion concrete, we consider a task that has sometimes been used in this literature: deciding whether a visual stimulus is presented to the right or to the left of the display center (e.g., Wilding, 1974). By continuously varying the position of the stimulus from the most left to the most right position, one can manipulate the probability that the subject will respond "left" or "right", monotonically; it is possible in this case to order the stimuli on a one-dimensional continuum (the x-axis) from stimuli that almost never lead to a 'right'-response (left-most stimuli; $P = 0$) to stimuli that almost always lead to a 'right'- response (right-most stimuli; $P \approx 1$.[11] One can then measure the way in which latency depends on response probability by plotting latency probability (LP) functions, which are defined as the mean RT for a given response to a particular stimulus, plotted against the probability of that response. Sometimes, LP functions are displayed in terms of a parameter that is monotonically related to response probability. For the data, one such value (as in the example presented here) is the linear distance from the center of the screen (Wilding 1974). For the models, a value often used is the drift parameter of the diffusion process, or, for the discrete random walk, the probability p of a step toward the given response (Vickers et al. 1971).

The classical random walk model and the classical accumulator model differ in their predictions for the shapes of LP functions. The classical random-walk/diffusion models predict that the LP functions are symmetrical, so that the reaction time will be slowest for the most ambiguous stimuli, and will get faster toward the extremes, with both the high and the low probability responses occurring more quickly (see Figure 8). This result seems highly counter-intuitive at first, but can be understood in the following way. Consider a random walk initialized at the origin; at every

---

[11] Often, when the task is highly symmetric, response time data can be collapsed across the two responses for mirror-symmetric stimuli. For example, in the Wilding task, the common practice is to combine left and right responses from mirror-symmetric points along the stimulus continuum.

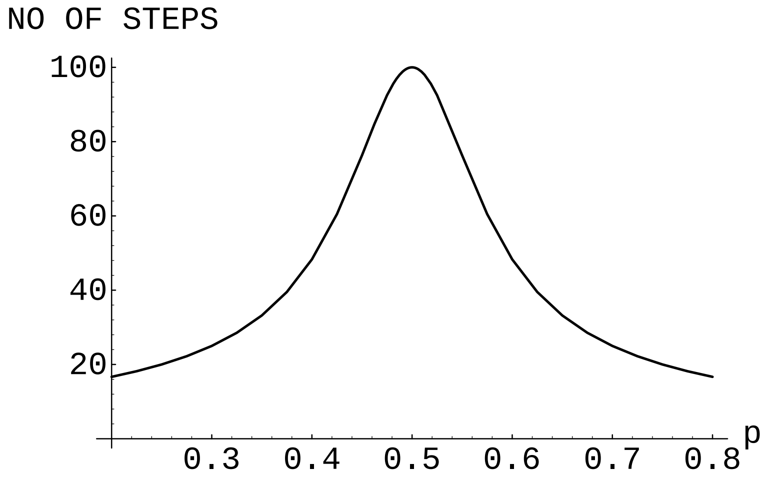

Figure 8: Mean number of steps to reach the 'right' boundary at $N = 10$, as a function of the random-walk probability, p. The mean number of steps, calculated following Cox and Miller (1965) is $\frac{N(p^N-q^N)}{(p-q)(p^N+q^N)}$, where $q = 1 - p$.

time step, the process moves either to the right (with probability p) or to the left with probability ($q = 1 - p$) and the process is terminated when it has progressed N steps, either to the right (A) or to the left (B). One can check (Audeley & Pike 1965) that the ratio of the conditional probabilities of reaching either A or B given that the response time is n time steps, are independent of n: $(P(A|n)/P(B|n)) = (p/q)^N$. Thus the two RT distributions have the same temporal dependency and differ only by a factor that reflects the fraction of correct responses. On the other hand the accumulator models predict that the mean RT has a decreasing relationship with probability; or equivalently, responses become slower and slower as they become less and less probable (Audeley & Pike 1965, Vickers et al. 1971). In this case, at every time step a unit is added either to a left counter or to a right counter. Thus the ratio of the conditional probabilities given that the process has finished after n steps is $(P(A|n)/P(B|n)) = (p/q)^{N-n}$. Assuming $p > q$, the relative probability for a correct response decreases with time, and thus errors tend to be on average slower than correct responses.

The above analysis shows that the choice between a difference-based criterion and an absolute criterion can have empirically testable consequences. The contrasting predictions of these two models led to a a series of articles that examined LP functions in binary choice RT experiments (Audeley & Mercer, 1968; Jamieson & Petrusik, 1977; Pike 1968; Vickers et al, 1971; Wilding, 1974), and recently Ratcliff and Van Zandt (1995) have reopened this line of investigation. In general these studies report asymmetric LP functions which increase as the response probability decreases down to a fairly low value and then turn downward, as shown in the results for the mean RT, in the middle-left panel of Figure 9. An additional observation from this data (Vickers et al., 1971) is that the LP curves have a more asymmetric characteristic for the slower subjects, tending to increase as probability decreases until turning down again at the very lowest probabilities. The degree of asymmetry decreases with practice (Jamieson & Petrusik, 1977). Participant 4 of Ratcliff and Van Zandt appears to represent the extreme case of this generalization, in that this subject was by far the fastest overall in their experiment and produced a highly symmetric LP curve.

The actual shape of empirically obtained LP functions seems to rule out both the classical models. The classical random walk/diffusion model predicts that LP functions will always be

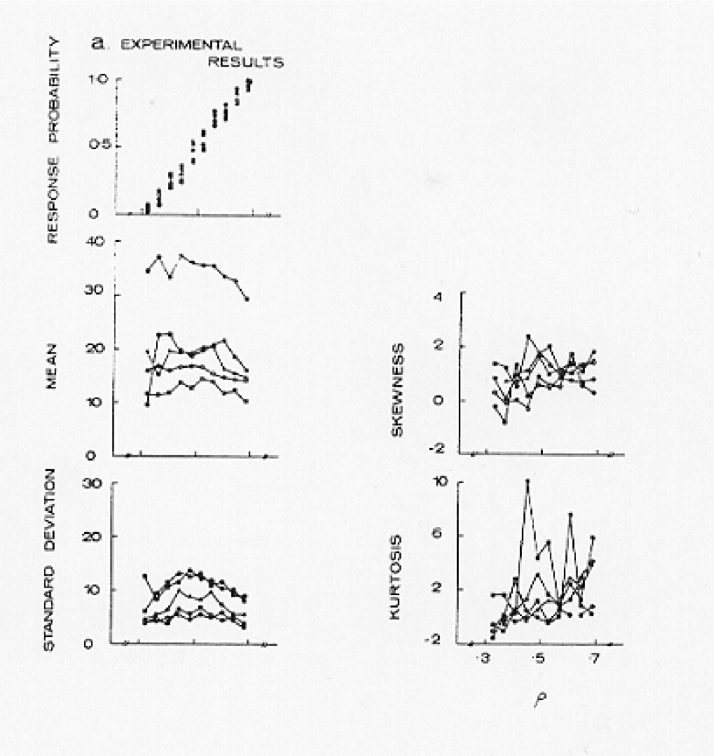

Figure 9: Latency-probability functions from a typical experiment. Subjects watched a stream of red and green flickering lights and had to decide as fast as possible which of the lights is predominant. The relative frequency, f, of the red lights was controlled (Vickers et al. assumed that this frequency controls the probability p of the accumulator model). The response probability, the mean, standard deviation, skewness and kurtosis of the RT are plotted as a function of the red-light frequency. From Figure 1 of "Discriminating between the frequency of occurrence of two alternative events", by Vickers, Caudrey & Wilson, 1971, *Acta Psychologica, 35*, p. 157. Copyright 1971 by Elsevier Science B. V. Permission Pending.

symmetric, while the classical accumulator model predicts LP functions which increase towards the low response probabilities. The typically asymmetric, but still U-shaped curves that are generally obtained seem inconsistent with both approaches. Also, the classical accumulator model makes incorrect predictions about the skewness of RT distributions (Wilding, 1974).

Modifications can be introduced into these classical models to produce asymmetric LP functions. The addition of drift variance into the classical diffusion model, producing Ratcliff's (1978) diffusion with drift variance (DDV) model, can significantly increase the error RTs leading to LP functions corresponding to the experimental data.[12] As Ratcliff and Van Zandt have shown, drift variance increases error RT (relative to correct RT) for the following reason. For every stimulus-condition, a mixture of different drifts is generated across trials. For a classical diffusion process with fixed drift, the probability of making an error (e.g., responding '1' when the drift is negative) decreases as the drift becomes more negative, while the mean RT for such errors decreases (see Figure 8). When drift variance is introduced, each condition produces a mixture of response times. Since the less-negative drifts result in a higher proportion of errors, the averaging operation that is introduced with the incorporation of drift variance is biasing the result toward the less negative drifts and thus it increases their RT relative to the classical model with the same drift. For correct responses, on the other hand, since the response probability of a classical diffusion process increases with the positive drift (while the mean RT of the process decreases, see Figure 8) the averaging arising from drift variance leads to a faster RT compared to the classical process with the same drift. As a consequence, drift variance leads to a slow-down of errors relative to correct RTs.

In sum, the diffusion model with drift variance can account for asymmetric but still U-shaped LP functions and may be adequate to capture other aspects of the data from LP experiments. However, the fact that this model does not provide a fully adequate account for time-accuracy curves, as discussed previously, leads us to consider how our leaky, competing accumulator models fare in accounting for latency-probability data. It might have been supposed that such a model would be most likely to succeed if it adopted a difference criterion, as in the DDV model. However, in accordance with optimality results of the previous section, we will show that with leaky competing accumulators, we can give a good account of the general form of LP functions using an absolute activation criterion instead of a difference criterion.

To assess the adequacy of the leaky, competing integrator model represented by Equations 4 with an absolute response criterion we have performed numerical simulations. For these simulations, we assumed that the summed input to the two response units is the same for all of the stimuli, so that what varies among the stimuli is simply the distribution of this sum between the two response units. As before we fix the constant summed input at 1, so that we can think of the stimuli as falling along a continuum in which ρ_1 varies from 0 to 1 (reflecting the manipulation of red-light frequency in the experiment of Vickers et al., 1971), while $\rho_2 = 1 - \rho_1$ varies from 1 to 0. Using this approach, we can simulate the probability, latency, and other characteristics of the "1" response as ρ_1 varies from 0 to 1.

Results of two simulation runs are shown in Figure 10. The mean RTs show the typical form reported by Vickers et al. (1971) and by Wilding (1974) (consistent also with the reports of Audeley & Mercer, 1968, Pike, 1968, and Ratcliff & Van Zandt, 1995). The SDs also have an inverted U-

---

[12] Two other types of modifications can be considered: i) addition of variance in the starting point of the drift process (called "initial condition variance"), and ii) addition of decay to the diffusion process (OU diffusion). Neither of these modifications alone can account for the obtained LP curves. Initial condition variance leads to a asymmetry, but now it is the low-probability responses that are generally faster than the high-probability responses, just the reverse of what is found. Introducing decay has only a very mild effect of increasing error RTs, and is not by itself sufficient to produce the markedly skewed LP functions that are typically found.

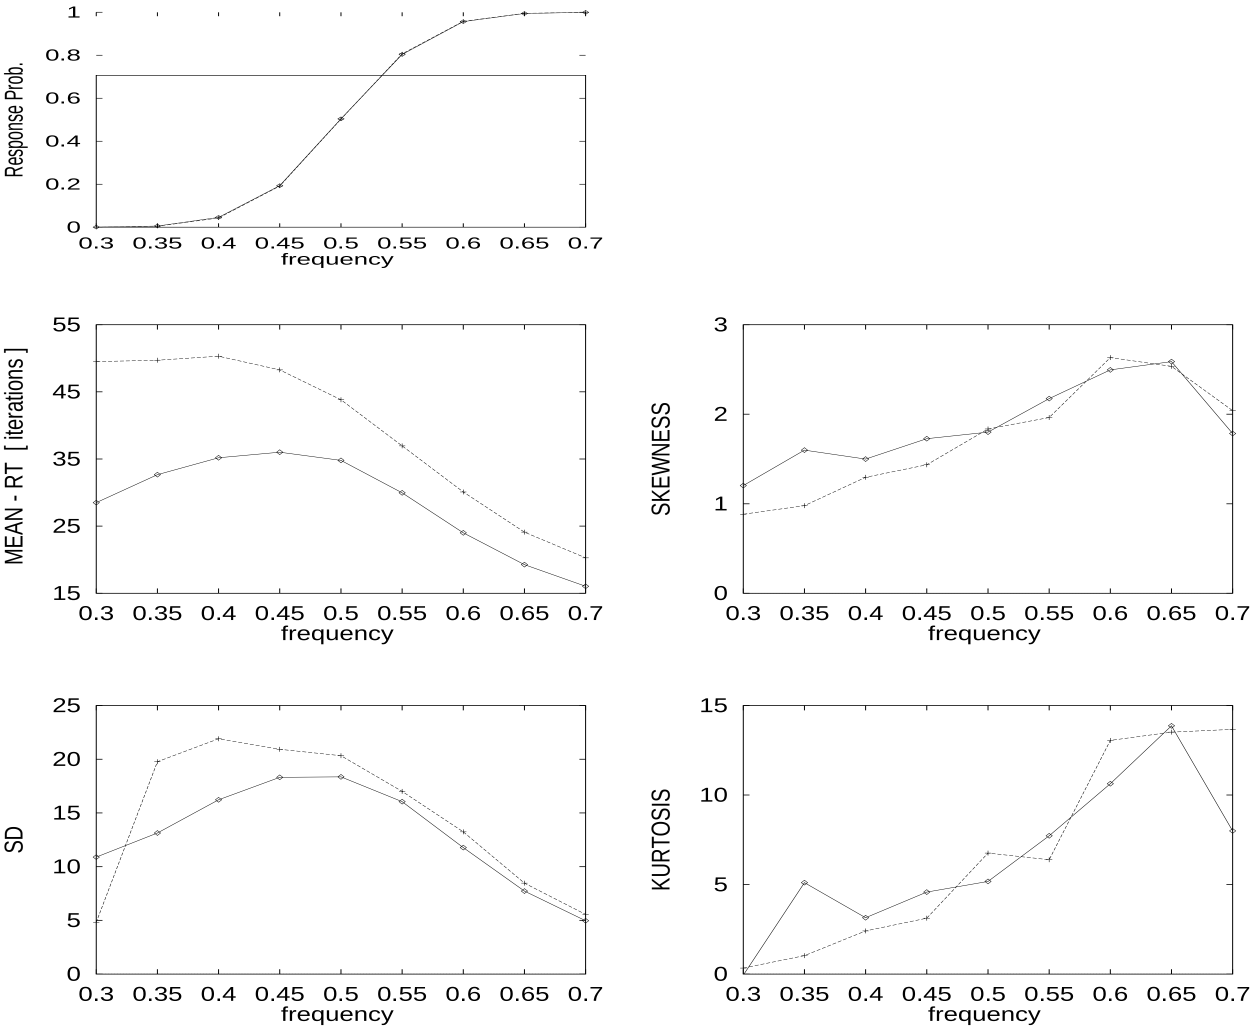

Figure 10: Simulated LP curves (response probability, mean number of $(.1\tau)$ iterations, SD, skewness and kurtosis, as a function of ρ_1 – the 'red light frequency', in Vickers et al., 1971), based on Equations 4, using absolute response thresholds for each response, using $(k = .2, \beta = .75, \sigma = .5)$, for a higher $(\theta = 1.057$; dashed lines with diamonds) and a lower $(\theta = .85$; solid lines with crosses) response criterion. Compare with Figure 9.

shape but with some decreasing tendency, and the skewness and kurtosis are increasing with ρ_1, as found in the data. The lateral competition also corrects the discrepancies found by Wilding (1974) with the independent accumulator model, where the skewness was significantly smaller than in the data. Also, note that we can capture the observation by Vickers (1971) that slower subjects show more nearly monotonic LP curves than faster subjects, simply by assuming that the slower subjects are adopting a more conservative criterion. This may also reflect the tendency of the LP curves to shift from near monotonic to near-symmetrical with practice, reported by Jamieson and Petrusik (1977).[13]

The leaky competing accumulator model's characteristics can perhaps be best understood by contrasting it with the classical accumulator model, since this is essentially a special case of our model in which both $(k-\alpha)$ and β equal 0. We display the RT distributions and hazard functions[14] for these models in Figures 11 and 12, for responses made in two different stimulus conditions, corresponding to different probabilities of making the emitted response. The left columns correspond to "incorrect" responses (e.g., green response when the the probability of a green light is 0.11) and the right columns correspond to "correct" responses (green response when the probability of a green light is 0.9).

We observe that the RT distributions generated by the leaky-competing accumulators are more skewed than the ones of the pure accumulators (Figures 11, 12). The leaky competing accumulator distributions are more similar to actual experimental data (Luce, 1986; Ratcliff & Van Zandt, 1995; Wilding, 1974). The hazard functions show another major difference between the two models. For the leaky-competing accumulators, the hazard function is monotonically rising to an asymptote, in the case of incorrect (infrequent) responses, but is peaked for correct (frequent) responses (Figure 11). For ambiguous stimuli ($\rho_1 = .5, P = .5$, not shown), the hazard function is marginal in between those two.

The distinction between monotonically increasing and peaked hazard functions is well documented experimentally in simple detection tasks (Luce, 1986); weak sub-threshold stimuli lead to monotonically increasing hazard functions, while strong stimuli typically lead to peaked functions. This suggests a mapping between identification and simple detection tasks; correct responses in identification tasks (which are frequent) correspond to strong stimuli in simple detection, while incorrect responses in identification tasks (which are infrequent) to weak stimuli in simple detection. There is not much experimental data which directly reports hazard functions for identification tasks. However, recently Ratcliff and Van Zandt (1995) have reported hazard function for a binary choice task which are peaked for frequent response and monotonic for infrequent responses (as in Figure 11). The hazard functions for the non-interacting accumulators (Figure 12), on the other hand, are always monotonic, missing therefore an important characteristic of RT data.

In summary, it appears that the leaky competing accumulator model can account for aspects of the LP data missed by both the classical accumulator model and the classical random walk model (equivalent to a difference-of-accumulators model). In some sense the model appears to capture aspects of both. On the one hand, it uses an absolute response criterion, like the classical accumulator model. This tends to allow the model to account for a tendency for error responses to be slower than correct RTs. On the other hand, it incorporates lateral inhibition, which tends to

---

[13] While one would expect shifts in criteria to be accompanied by changes in error rates, the effects can be slight, since relatively symmetric curves can be obtained with strict criteria that lead to very low "error" rates (i.e., criteria that lead to probabilities of choosing the low probability alternative in the range of 1-2% or less, when the stimuli are at one or the other extreme of the continuum.

[14] The hazard function gives, for each reaction time, the probability of a response at that time given that no response has occurred before that time. See Luce (1986) for a full discussion.

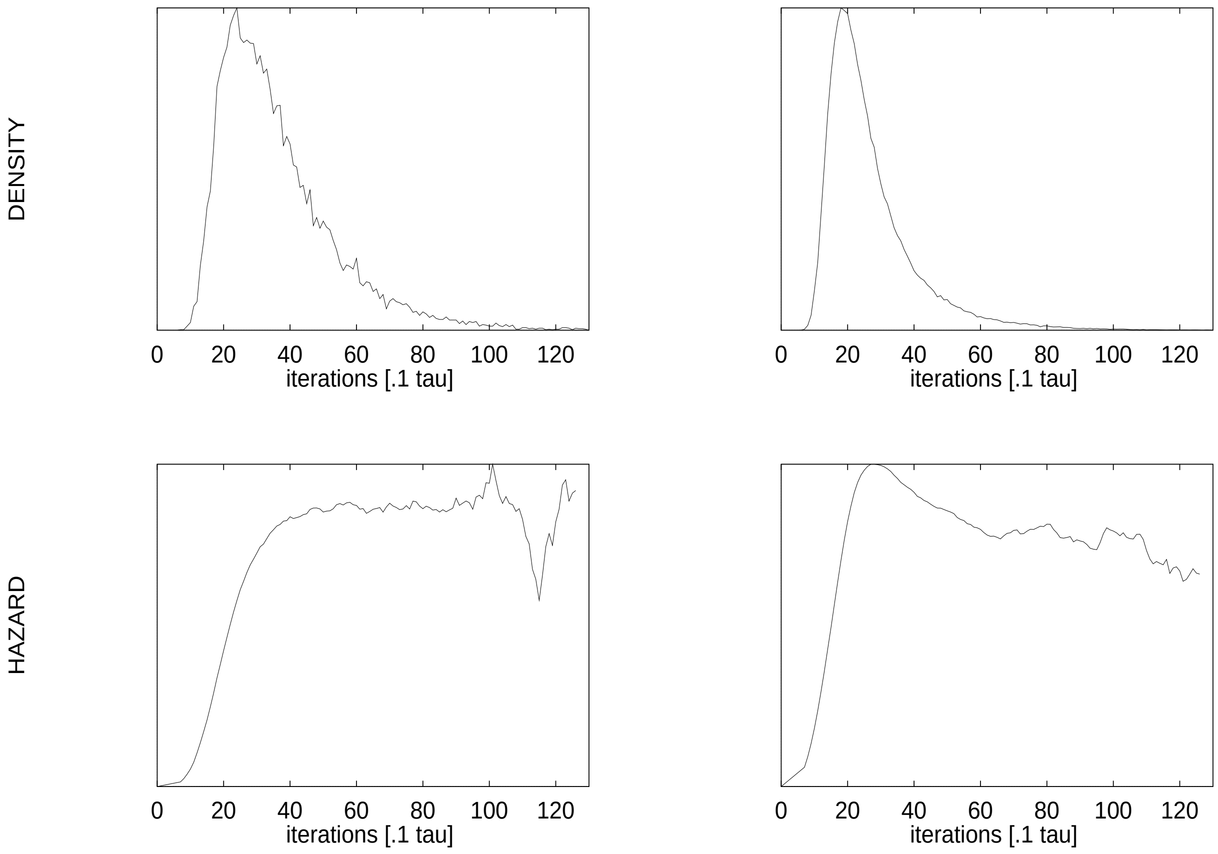

Figure 11: RT distributions (upper row) and hazard functions (bottom row), for correct responses (right column), incorrect responses (left column) for the leaky-competing accumulators. Notice the fact that RT distributions are skewed and the hazard functions are monotonic for errors but are peaked for correct responses. The fluctuations in the hazard functions at long response times reflect the low density of sampled data in the extreme tail of the RT distribution.

have the effect of making each activation variable (x_1 and x_2) reflect to some extent the activation of the other, just as the use of a difference criterion would. The presence of both of these tendencies at the same time allows the model to account for a family of curves, ranging from monotonic to symmetrically U-shaped.

Based on the above analysis, the leaky, competing accumulator model with absolute criteria and Ratcliff's DDV model both appear consistent with many aspects of the general patterns found in Choice RT data, particularly those observed in LP experiments. In the next section we consider evidence consistent with a role for lateral inhibition in at least some choice tasks.

Physiological Evidence for Lateral Inhibition

There is some physiological evidence that is suggestive that the mechanism of lateral inhibition is at work in some choice reaction time tasks. Before we consider the evidence itself, we consider what effects we should expect to see in the activations of processing units in our model, under the assumption that bottom-up excitatory influences tend to activate multiple alternatives, and lateral

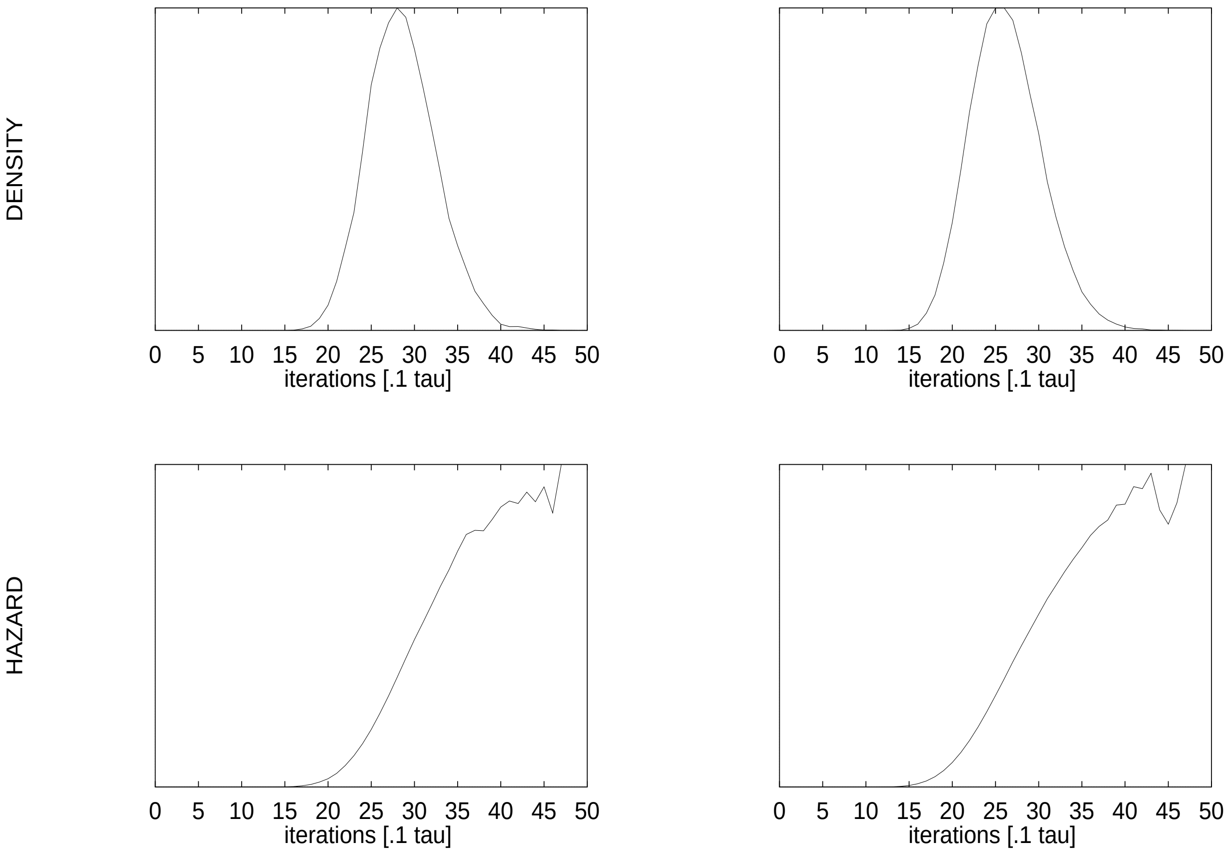

Figure 12: RT-distributions (upper row) and hazard functions (bottom row), for correct responses (right column), incorrect responses (left column) and guessing (very low discriminability stimuli; middle column), for pure accumulators. Notice the fact that RT distributions are less skewed and the hazard functions are monotonically increasing, in all the stimulus/response conditions.

inhibition then acts to increase the contrast among the activated alternatives. For this purpose we consider the deterministic (noiseless) system involving lateral inhibition:

$$\begin{aligned} \tau \frac{dx_1}{dt} &= -x_1 + I_1 - \beta x_2 \\ \tau \frac{dx_2}{dt} &= -x_2 + I_2 - \beta x_1 \end{aligned} \tag{15}$$

comparing it to a simpler system involving independent accumulators:

$$\begin{aligned} \tau \frac{dx_1}{dt} &= -k x_1 + I_1 \\ \tau \frac{dx_2}{dt} &= -k x_2 + I_2 \end{aligned} \tag{16}$$

Let us apply these two systems to the Eriksen flanker task, where subjects perform a two choice letter discrimination, for example, between the letters S and H. It is found that when incompatible letter stimuli are presented simultaneously with the target as flankers (e.g., HSH) RT is slowed, compared to the case where neutral flankers are shown (e.g., XSX). Furthermore, the RT slows in proportion with the feature overlap between the neutral flankers and the incompatible target-

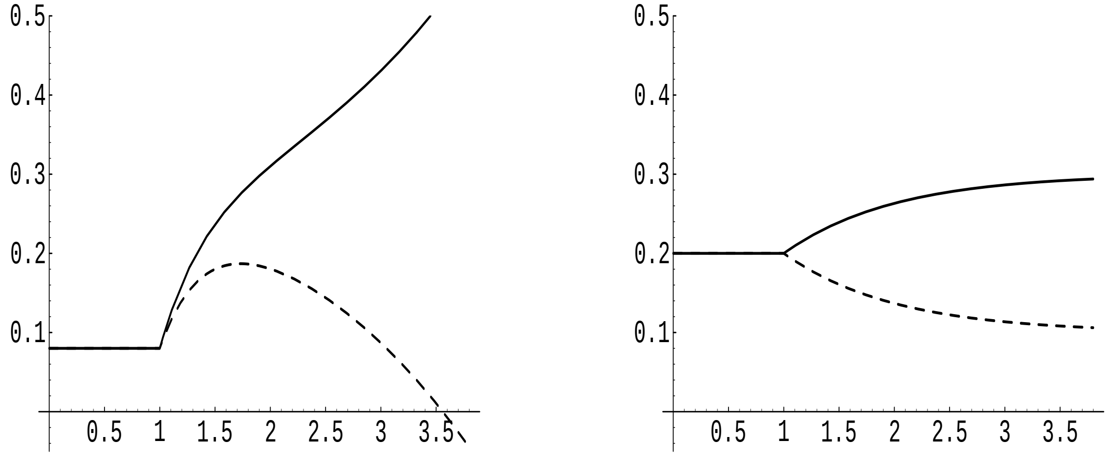

Figure 13: Dynamics of lateral and feedforward inhibition. Correct unit activity (solid lines) and incorrect unit activity (dashed lines). The stimulus is applied at $t = 1$. Before this both units receive a background input $I_0 = .2$ and after stimulation the inputs are increased by $I_c = .5$ and $I_2 = .4$. A) Lateral inhibition (Equation 15) with $I_1 = .7$, $I_2 = .6$. B) Feedforward inhibition: $I_1 = .3$, $I_2 = .1$.

letter (Eriksen & Eriksen 1974, 1979; Eriksen & Shultz, 1979). This happens despite the fact that subjects know in advance the location of the target and are instructed to ignore the flankers.

A system like that shown in Equations 15 can account for the effect of incompatible flankers in the Eriksen task by the use of lateral inhibition. In this case, the activation of the incorrect alternative should initially increase when an incompatible flanker is presented; suppression of this alternative should only occur once the unit for the correct alternative becomes partially activated.

A system like that shown in Equations 16 can account for the effect as well. This can be done by assuming that the stimulus for each alternative provides inhibitory input to the response unit for the other alternative. In this case the inhibitory influence is 'feed-forward' rather than lateral.

To mimic the inputs received by the response units in the Eriksen task, for the case of an incompatible stimulus, we assume that before $t = t_0$ both units receive a background input I_0, and that for $t > t_0$, in the lateral inhibition case, I_1 is increased to $I_0 + I_c$ and I_2 to $I_0 + I_i$ $(I_c > I_i)$; for the feedforward inhibition case we assume that I_1 is increased to $I_0 + (I_c - I_i)$ and that I_2 decreases to $I_0 + (I_i - I_c)$.

The response units in these two systems exhibit very different activation trajectories. Despite the fact that in both cases stimulation ultimately leads to an increase in activity of one unit and to suppression of the other, the trajectory of the incorrect unit is not monotonic when lateral inhibition is used, showing a transient activation followed by suppression (see Figure 13). Such transients are a characteristic property of systems with lateral inhibition, and are due to the fact that the amount of inhibition depends on the activity of the units. Initially activity is low and therefore the additional input dominates, but as the activity increases, a point when the inhibition from the stronger alternative outweighs the excitatory input to the second (weaker) unit is reached. One should also notice that the lateral inhibition scheme achieves a strong amplification of the difference in activation of the two units.

There is another obvious way to account for the effect of incompatible flankers in the Eriksen task: This would be to use a difference criterion to choose when to respond. The use of a difference criterion by itself, however, does not lead to non-monotonic activation curves like those seem with lateral inhibition. Indeed one might expect, if there were true "difference" detectors operating, that their trajectories would look rather like those produced by the effects of bottom up inhibition—indeed this difference criterion might be implemented using a form of feed-forward inhibition.

Important characteristics of the dynamics of activation in the Eriksen task have been revealed by an experiment by Gratton, Coles, Sirevag, Eriksen and Donchin (1988), who measured ERP signals at several sites over the scalp, while subjects performed in a target identification task with several flanker stimuli. By using a sophisticated vector filter procedure (Gratton, Coles & Donchin, 1989) they isolated single-sided measures of activity over the left and the right motor areas (corresponding to movement of the right and left hands, respectively). These traces, which when averaged over trials, are analogous to the activation trajectories in Figure 13, and show a clear transient in the incorrect response channel following the stimulus presentation.

Another, independent finding consistent with the lateral inhibition mechanism is provided by a recent experiment by Chelazzi, Miller, Duncan & Desimone (1993), who recorded spike trains of cells from the inferotemporal cortex of monkeys during the performance of a visual search task. In the experiment, monkeys were first shown a cue shape which they had to use for initiating a visual search. Following a delay, a visual display containing the old cue shape and a new one was presented, and the monkey had to respond by moving its eyes toward the cued shape. As can be seen in Figure 15, following the presentation of the choice stimulus, the activity in cells that respond to either of the two stimuli increased. However, after rising at about the same rate for about 100 msec, the activity in the cued cell continued to rise, while the activity of the distractor was suppressed (see also Usher & Niebur, in press, for an explicit model addressing this data).

Behavioral evidence of lateral inhibition can also be predicted in some cases. Such cases hinge on detecting the non-monotonic activation of the incorrect alternative that is predicted by the lateral inhibition account. Consider the version of the Eriksen task used by Gratton et al (1988), in which the target (H or S) is flanked by a total of four flankers, all of which are either compatible (HHHHH or SSSSS) or incompatible (SSHSS or HHSHH) with the target. These investigators obtained time accuracy curves for the incompatible condition that were non-monotonic: There was an initial decline in accuracy below chance, followed by recovery to almost perfect performance. Just such a pattern can easily arise in a model that relies on lateral inhibition (Cohen, Servan-Schreiber & McClelland, 1992). According to this model, each letter in the display initially excites a set of position-specific detectors, and these in turn send excitation to the appropriate response units. At first, there are four sets of position-specific detectors supporting the incorrect alternative, and only one set supporting the correct alternative, thus giving the response unit for the incorrect alternative greater initial activation. If, however, input from attention favors the set of detectors in the center location, and if the sets of detectors are all mutually inhibitory, then the set of detectors in the center position can eventually suppress all of the others; when this happens, the support for the incorrect alternative is eliminated, causing the correct response to dominate the incorrect response later on in processing (see Cohen et al, 1992, for details).

In summary there is both physiological and behavioral evidence consistent with the non-monotonic pattern of activation predicted by models that exploit lateral inhibition. However, it is important to note that it may be possible to account for such non-monotonicities in other ways. Our point is not to suggest that the findings rule out other explanations. Indeed, Gratton et al proposed that information about identity regardless of position travels faster than information about identity in

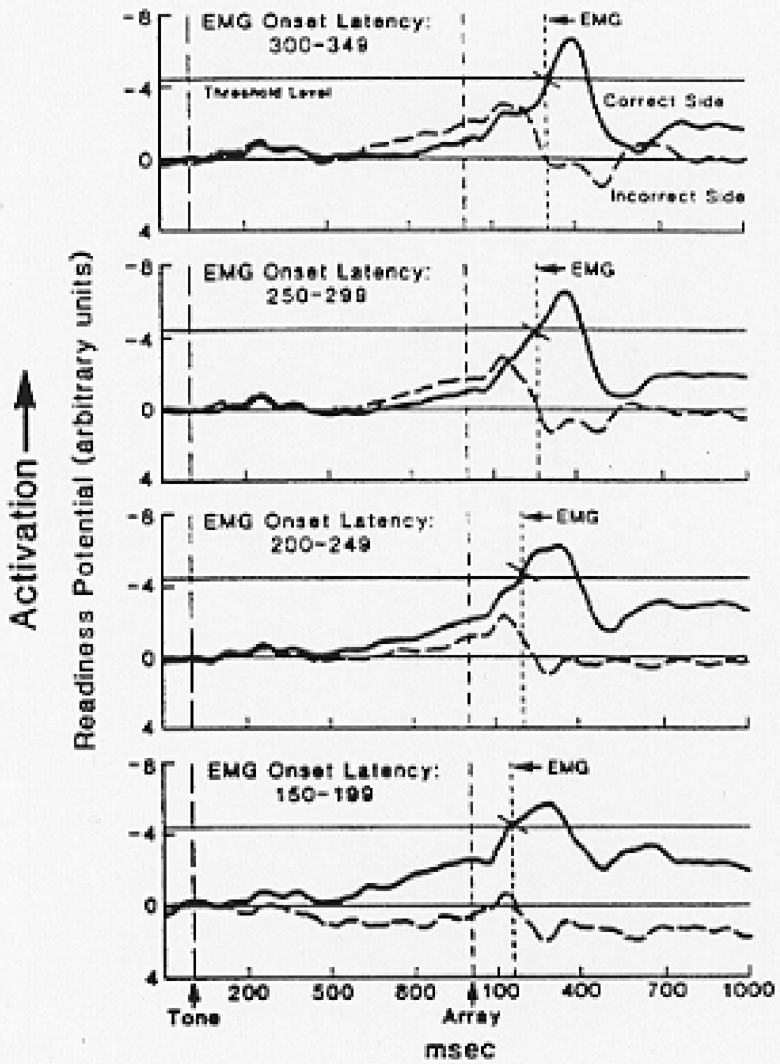

Figure 14: ERP activation in the Eriksen task. From Figure 6 of "Pre-and post-stimulus activation of response channels: A psychophysiological analysis". by Gratton, Coles, Sirevag, Eriksen & Donchin, 1988. *Journal of Experimental Psychology: Human Perception and Performance, 14*, p.340. Copyright 1988 by the American Psychological Association. Permission pending.

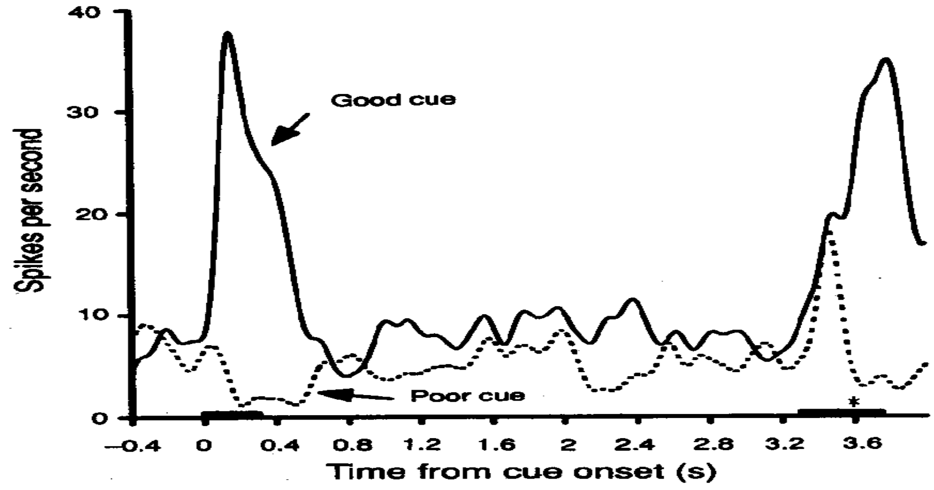

Figure 15: (a) Cell response in monkey IT cortex, as reported by Chelazzi, Miller, Duncan & Desimone (1993). Shown is the average firing rate of a cell when it responds to the cued shape (solid line) or to the distractor shape (dotted line). Black horizontal bars on the abscissa indicate presentation of cue and choice stimuli. Average time of saccade onset ($297ms$) is indicated by an asterisk. From Figure 2 of "A neural basis for visual search in inferior temporal cortex", by Chelazzi, Miller, Duncan & Desimone, 1993, *Nature, 363*, p. 345.Copyright 1993 by Macmillan Magazines Ltd. Permission Pending.

a specific position. Instead, the point is just to note that models that incorporate lateral inhibition can in some cases predict strongly U-shaped time-accuracy curves, and such curves can indeed be seen in physiological and behavioral data.

Visual Word Recognition: The Intersection Principle

Word recognition is a typical task that requires selection among a very large number of possible targets, and as such it is a good candidate for assessing the adequacy of lateral inhibition as a basis for response selection. Previously we considered the benefits of lateral rather than bottom-up inhibition in the interactive activation model of visual word recognition. Another case illustrating positive properties of lateral inhibition is brought out by the pattern of data reported by McClelland and O'Regan (1981).

In their experiment, McClelland and O'Regan tested the benefit in RT that subjects can obtain in a word identification task from two additional sources of prior information: a visual one (in the form of a very brief parafoveal preview) and semantic context. For the visual source, there was a *strong preview condition* in which the preview was the same word as the target; a *weak preview condition* in which the preview was a word visually similar to the target, and a *control preview condition*, in which the preview was a row of 'X's. Similarly, there was a *strong context condition*, which strongly suggested the actual target word; a *weak context condition*, in which the target word tended to be one of many that were consistent with the context; and a *neutral context condition*, which provided no constraint on the possible identity of the target.

McClelland and O'Regan found that neither the weak preview paired with a neutral context

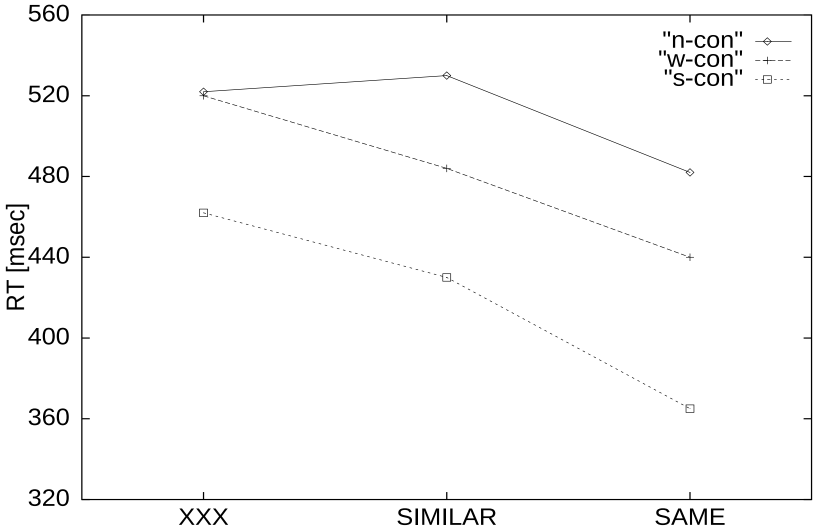

Figure 16: Joint effects of context and preview information on the time it takes to read a word target aloud. Solid line (diamonds): neutral semantic context, Dashed line (crosses): weak semantic context. Dotted line (squares): strong semantic contexts. Visual preview condition is shown on the x-axis and recognition RT on the y-axis. Adapted from Figure 2 of "Expectations increase the benefit derived from parafoveal visual information in reading words aloud", by J. L. McClelland and J. K. O'Regan, 1981, *Journal of Experimental Psychology: Human Perception and Performance, 7*, p. 639. Copyright 1981 by the American Psychological Association. Permission Pending.

nor the weak context paired with a neutral preview were sufficient to speed the recognition RT, compared to neutral-context/neutral-preview baseline. However, a combination of a weak preview together with a weak context did speed the identification of the target word, by similar amounts. This effect shows that some non-linear interaction between the two sources of partial information must be taking place. This would in fact be expected if the word recognition system exploits the intersection of information sources (Humphreys, Wiles & Dennis, 1994). According to this idea, each weak source activates several words in addition to the target. When two such sources intersect, the correct alternative will receive support from both sources. Given that context selects a set based on semantic relationships, and a preview selects a set based on visual relationships, and given that semantic and visual relationships are approximately independent (i.e., visual similarity is only very weakly correlated with semantic similarity), other words will rarely be consistent with both sources of information, so the correct alternative will receive two sources of support, and therefore can win out over all the other alternatives. As we show next, a lateral inhibition scheme can implement this intersection principle, and can explain why the intersection-target is boosted when two sources of support intersect, even though there is little or no benefit when only one source is available.

Consider an n-choice generalization of the 2-choice lateral inhibition system in Equation 15 (where n is the size of the vocabulary). We assume that each 'weak' source alone sends input to the correct word-alternative and to $k-1$ other alternatives. For simplicity we choose all these inputs to be equal to each other, and denote their strength by Δ. Second, we assume that the two sources add their input contributions to the target units, linearly. Thus when two sources of weak input exist, each one projects input of strength Δ to $k-1$ different word-units, while the correct

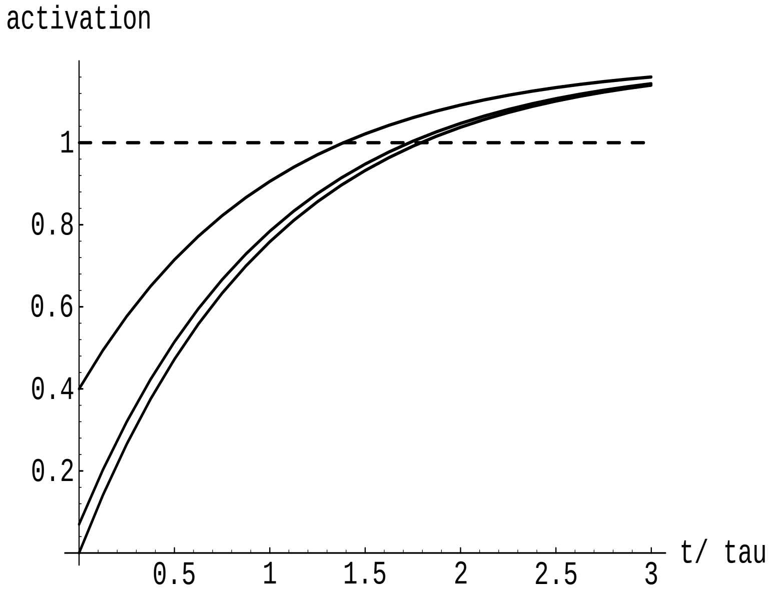

Figure 17: Qualitative illustration of priming effect in McClelland and O'Regan (1981) experiment. The word input has an asymptotic value of 1.2 a time constant of 1, and the response criterion is 1. The lower curve, which starts at zero, represents the control. The second curve represents a single weak information source, characterized by $x(k)$, $\Delta = .2$ and $k = 4$ (Equation 17). The third curve represents a combination of two weak sources leading to a starting point activation of 2Δ.

alternative (which is activated by both sources) receives an input of 2Δ.

In order to estimate the priming benefit under each condition, we assume that when the word is presented a strong signal projects to the target unit. The starting point of this unit activity, however, depends on its pre-stimulus activation which is a function of the preview and context information. We focus attention here only on the four conditions crucial to the interaction, namely those involving neutral and/or weakly constraining sources of information, since a single strong constraint appears to be sufficient to produce a facilitation in isolation. In order to obtain the simplest qualitative account for the effect, we neglect here the dynamics of the preview and context information, and assume that when the word target is presented the corresponding unit is in a steady state of activation that characterizes the pre-stimulus condition.

It is easy to calculate the steady state solutions (obtained by requiring that $\frac{dx_i}{dt} = 0$, for all $0 < i < N$) in the n-choice version of Equation 15. Consider first the case of a single weak source. All the units that do not receive input are inhibited, and are thus suppressed to zero (we assume threshold linear units, so that a negative synaptic current x_i results in a zero activation, or firing rate). All the units that receive input, will have the same steady state (due to the equal input values) of $x(k)$, which satisfies:

$$x(k) = \frac{\Delta}{1 + \beta(k-1)} \tag{17}$$

Without lateral inhibition, each unit would have a steady state activation of Δ that reflects the input. The effect of the lateral inhibition is a reduction in the magnitude of this activation in proportion with the number of activated units and with the strength of the inhibition.

Consider next the steady state of the units in the condition of two weak intersecting sources.

Again, the units that receive no input are suppressed to zero. All the non-intersecting units that receive input will have a steady state denoted as x_N and the intersecting unit has a steady state of x_I, which satisfy:

$$\begin{array}{rll} X_I = & 2\Delta & -\beta(K-1)x_N \\ X_N = & \Delta & -\beta x_I - \beta(K-2)x_N \end{array} \tag{18}$$

(where $K = 2k - 1$). The solution for the non-intersecting units, x_N is $X_N = x(K)\frac{1-2\beta}{1-\beta}$. Thus the presence of the intersecting unit leads to an additional suppression of the K non-intersecting units by a factor of $\frac{1-2\beta}{1-\beta}$ (the value of '2' in this expression reflects the fact that we use only two intersecting sources). Thus, for any value of $.5 < \beta < 1$, all the non-intersecting units are suppressed to zero (negative synaptic current leads to zero firing rate), and as a consequence the steady state activation of the intersecting unit is restored to its value that reflects the input alone, $x_I = 2\Delta$ (Equation 18). In Figure 17, we display the deterministic time course of activation for a leaky-integrator, under the impact of the same stimulus (reflecting a common asymptote and time constant) but from different starting positions which reflect the pre-stimulus condition. We observe that the first curve (which reflects a control condition; zero initial activation) and the second curve (reflecting a single weak information source) are very similar, and intersect the response criterion at almost the same time. For two weak sources, a significant speed-up results.

Discussion

We have shown that a model for perceptual choice guided by neurophysiological principles could be constructed. Unlike models which assumed a deterministic accumulation of activation and where response variability was due to the variability of the response-threshold or of the initial conditions (Anderson, 1991; Grice, 1972; McClelland, 1979), in the model we developed, intrinsic variability plays a major role in the dynamics. This model explains various characteristics of the time-controlled and standard choice RT experimental data, within a unified framework. In particular we have shown that the model can explain both exponential d' time-accuracy curves, which have been mainly reported in memory retrieval experiments where the asymptotic accuracy is less than perfect, and exponential P-time-accuracy curves, which have mainly been reported in perceptual discrimination tasks where asymptotic accuracy can reach 100%. We propose that the perceptual choice and the memory retrieval systems operate according to the same neural information principles and that the difference in the time-accuracy curves results from differences in task difficulty. Our model predicts, therefore, that time-accuracy curves for very difficult perceptual choice tasks will show the same type of exponential increase in d' towards a finite accuracy level. An experiment by Swenson (1972), who measured time-accuracy curves for very difficult visual discrimination tasks, indicates that this may be the case; in this experiment Log(1-P) is linear for easy discrimination but levels off to an asymptote at hard discrimination (Swenson did not fit the d' curves with exponential functions). It is also interesting to note that our stochastic leaky accumulator system approximates a random-independent feature extraction theory discussed by Busey and Loftus (1994).

Most of the theoretical studies on the topic of information accrual and choice RT were conducted within the framework of diffusion or random walk models. This motivated Ratcliff and Van Zandt (1995) to pursue a direct comparison of diffusion and connectionist models. We think that our study demonstrates that the question is not whether connectionist models are more (or less) correct than diffusion models. As we have shown, our neural model can, under specific conditions, conform exactly to one or another of the classical stochastic reaction-time models. For example if we assume a 'difference' response criterion and a perfectly balanced inhibition-excitation-decay, we obtain a classical diffusion or random walk (Laming, 1968; Link & Heath, 1975; Stone 1960). These classical models, and more recent variants such as Ratcliff's (1978) diffusion with drift variance model are very useful for mathematical analysis, and are an important domain for continued study. The domain is a broad one, however, and some source of constraint on the search among the range of possible model variants is required. Our use of principles from physiology can help guide the selection among various types of stochastic models. Within the last thirty years a variety of stochastic models, yielding contradictory predictions have been intensively studied. Indeed, after comparing all the existing diffusion models to the experimental data and finding that none of them could explain important features of it, Wilding (1974) has suggested that some new approach is required, since otherwise there are endless modifications that can be devised in an arbitrary way to the stochastic models. We believe that the neural approach might contribute to the search for such an approach. For example, it strongly motivates the use of the Ornstein-Uhlenbeck diffusion process, which was essential in order to explain the exponential curves in performance, within both high and low accuracy experiments. This OU diffusion processes has been used by Smith (1995) in order to explain RT distributions for simple detection tasks within a psychophysically principled model, and by Busemeyer and Townsend (1993) in decision making.

Adaptive Setting of Input Strengths

A very important issue for stochastic models that can be addressed in a neural instantiation is the physical reality of model parameters, such as the diffusion drift. In our framework, the drift parameter is not just a mathematical abstraction, but reflects the relative strength of synaptic projections between a stimulus representation (input layer) and a response representation (ρ_1 and ρ_2, in Figure 1). While modeling the discriminability effects on RT (the LP functions), we assumed that when the stimulus moves on a one-dimensional continuum, the synaptic projections ρ_1 and ρ_2 satisfy a normalization condition that they sum to a fixed constant value. This assumption has the mathematical benefit of simplifying the variation among the inputs in the different conditions from 2 dimensions to one. Interestingly, Ratcliff and Van Zandt (1995) found that the diffusion drift parameters, fitted from the data, were linearly related to the *a priori* probability with which stimuli are associated with response alternatives. While we have not focused here on issues of learning connection weights, it is interesting to note that connection-based learning rule can lead to connection weights for our model that capture this relationship, and at the same time justify our assumption that ρ_1 plus ρ_2 sum to a constant.

We will illustrate how this would be possible if there is a single input unit representing each value on the stimulus continuum used in a latency probability experiment. We use this localist representation assumption just for simplicity of illustration. There are generalizations (e.g., Movellan & McClelland, 1993), that would have the same effects if both the input and the output representations are distributed. As an example we will consider the continuum used by Ratcliff and Van Zandt (1995). In this experiment, subjects were shown a square 10x10 array in which some number of array locations contained asterisks and the others were blank. On a given trial the actual display was generated from one of two distributions, one that tended to produce a high density of asterisks and another that tended to produce a lower density, and the task is to decide which distribution actually generated the stimulus. In this case the continuum in question is the number of asterisks, which varied from 0 to 100, so for the present discussion we assign one input unit to each possible number of asterisks. We also assign an output unit to represent the 'high' response and another to represent the 'low' response. Given these assumptions, it is possible to construct a learning rule that will set the weights from each input unit to each output unit so it reflects the probability that the input (i.e., the number of asterisks) came from the distribution identified with the output. For example, the weight from the '26 asterisks' input unit to the 'high density' response unit would be set to the probability that 26 asterisks occurs as a sample from the high density distribution.

The rule assumes that learning takes place when the subject has already seen the input (so that the correct input unit is activated), and has also made the response and received feedback about which response is correct (i.e., which distribution generated the input); and that he uses this information to set the activations of the corresponding units to 1 and all other units to 0. Suppose that on a given trial 26 asterisks were shown, followed by feedback indicating that the input had been generated by the high density distribution. In this case, the '26 asterisk' input unit and the 'high density' output unit would be activated. The learning rule then simply adjusts the weight to each input unit i from each response unit j according to the following variant of the Hebbian learning rule:

$$\Delta w_i j = \epsilon(a_i a_j - w_i j) \ . \tag{19}$$

In this expression, a_i represents the activation of the response unit, a_j represents the activation of the input unit, $w_i j$ represents the weight, $\Delta w_i j$ represents the change in the weight, and ϵ represents the learning rate constant. If this rule is applied repeatedly, the expected values of the weights between the input and the response units converge to values that reflect the probabilies (p, $1-p$) that the stimuli are associated with each of the two responses (Rumelhart and Zipser, 1986).

In short, a simple learning rule provides a mechanism that implements the assumption that $\rho_1 + \rho_2 = S$, where S is a constant; in this case, the ρ_i for each number of asterisks are just the weights $w_i j$ for the input unit corresponding to that number. The learning rule also implements the linear dependency between the drift parameter ν of the diffusion model and the *a priori* stimulus-response probabilities in the study of Ratcliff and Van Zandt (1995), or the relative red/green flickering frequencies in the experiment by Vickers et al (1971). The diffusion drift parameter ν is linearly related to our parameter ρ_1 (specifically, $\nu = 2\rho_1 - 1$). The Hebbian learning rule sets the weights so that the inputs ρ_1 and ρ_2 will equal the corresponding p and $1-p$ that the input came from each of the two response distributions, thus accounting for the linear relationship between ν and p.[15] It should be noted, however, that the learning rule as we have described it only really works well if ϵ is small; if it is too large, the weights are continually buffeted about by the latest random sample, and indeed the variance of the weights around the real probabilities decreases as ϵ decreases. (If ϵ systematically decreases, approaching 0 over time, each $w_i j$ will gradually converge exactly to the appropriate probability.) Thus, it would be important to avoid making large changes to the connection weights after each trial of the experiment (See McClelland, McNaughton & O'Reilly, 1995, for a full consideration of the reasons for avoiding large, trial-to-trial changes in connection weights that are directly involved in information processing).[16]

---

[15] The actual value of the ν parameter of the model is subject to a further scaling factor and potentially to a bias factor that would affect the slope and intercept of this linear relationship.

[16] Ratcliff and Van Zandt (1995) found systematic sequential dependencies from trial to trial in their experiment. One might be tempted to assume that such dependencies reflected trial-to-trial changes in connection weights. However, the results indicated that they were not consistent with the learning rule given above, since the response on trial n reflected the response, but not the feedback, on trial $n-1$. This finding, together with the observation that weight changes from trial to trial should be kept quite small, suggests that the sequential dependencies may have been based on something other than changes in the connection weights. One possibility is that response units retained some of the activation they acquired on the previous trial (for some subjects) or were strategically suppressed (for others).

Relation of the Leaky, Competing Integrator Model to Distributed Attractor Networks

In light of the results of our analysis, we return to an issue that we touched on earlier in our analysis, namely the relationship between our leaky, competing integrator model and models that exploit distributed representations. As previously indicated, it is likely that the brain employs distributed representations, in which the population of neurons that represents a particular response or decision may overlap with the set that represents another response or decision. Our use of a single variable to represent each population dramatically simplifies the representational issues, and probably simplifies the dyamics as well. The simplification appears to be useful in that it leads to models that can account for many aspects of empirical results, and focuses attention of key relationships between aggregate variables that, we assume, represent aspects of the behavior of the entire population of neurons. But the simplification does raise questions about how aspects of our models might be implemented in the brain, whether the real neural implementation would introduce important complications in the dynamics, and whether there are substantive differences between the behavior of our model and the behavior we would expect in a truly distributed model. Several specific aspects of our model deserve mention in this regard: Recurrent self-excitation, organization of excitatory and inhibitory connections, and the concept of the response criterion.

Concerning self-excitation, we should emphasize that the self-excitation parameter α of our model does not reflect a connection from a neuron to itself, but rather the net effect of the ensemble of connections among the units participating in a particular pattern of activation. In this latter form, such connections would likely be learned connections, reflecting learned associative relationships among elements that belong together in a pattern. There are many distributed connectionist models exemplifying this approach, such as Anderson's BSB (Anderson, 1991), the Hopfield net (Hopfield, 1982), and several other models (Chappell & Humphreys, 1994; Herrmann et al., 1993; Ruppin & Yeshuron, 1991; McClelland & Rumelhart, 1986; Plaut & Shallice, 1993). Our model differs from many of these models, however, in that we have focused our discussion on cases involving relatively weak recurrent excitation, sufficient to cause an amplification of the input, but not to sustain it indefinitely after the input is removed. Weak recurrent activation (Douglas et al., 1995), is common in most cortical areas. In a few specific brain areas, such as the prefrontal cortex, there can be sustained states of activity (Funahashi, Chafee & Goldman-Rakic, 1993; Fuster, 1991; Goldman-Rakic, 1992; Miller, Erikson & Desimone, 1995) which may be due to the use of a stronger form of recurrency, closely related to attractor neural networks (Zipser et al., 1993). A few such models have qualitatively discussed RT within this framework (Anderson, 1991; Chappell & Humphreys, 1994; Ruppin & Yeshuron, 1991). For example, Anderson (1991) has proposed an account for RT distributions within the Brain State in the Box (BSB) model. Within this model, activation propagates deterministically within the system under the influence of external input and of recurrent connections and a response is generated when the activity pattern within the network reaches some typical distance from an attractor (the equivalent of the response criterion). Variability is induced in this model by the variance in the initial conditions of the units. The time-controlled paradigm or the impact of discriminability on RT (the LP functions), have not yet been modeled within this framework. Surely, one can also modify the BSB model by adding intrinsic noise to its dynamics. It might be interesting to notice that this will lead to a stochastic diffusion process (Equation 6) with negative K; due to the strong recurrence assumed in BSB and other attractor neural networks (e.g., Hopfield (1982)) a small activation of a pattern is amplified (this process is controlled by the saturation non-linearity). Whether in fact the perceptual choice system operates with decay or amplification, or within a perfectly balanced regime (and whether the higher cortical areas such as the prefrontal cortex are essential in explaining RT in perceptual and memory choice tasks), is an important topic which deserves further investigation. It is possible that different brain areas which subserve different choice processes operate under different diffusion parameters. It is also possible that the diffusion parameters can be varied depending on task demands, producing a range of different effects depending on circumstances.

Of particular importance in our model is the role played by lateral inhibition. We have shown that the lateral inhibition allows the model to segregate the decision variables and to implement a choice process that can be flexibly generalized to any number of alternatives. The model simulated by Equations 4 with absolute thresholds incorporates positive features of both the classical diffusion and the independent accumulator models. We have also noted that lateral inhibition can compensate for the loss of information in leaky accumulators: A leaky, competing information integration process that is allowed to run without response thresholds produces time-accuracy curves identical to those of perfect integrators. When response boundaries are present, we have shown that a mechanism of balanced lateral inhibition and leakage generates performance which is superior to classical accumulators with absolute thresholds, and which approaches optimal choice. A main characteristic of lateral inhibition systems, transient activation of partially-supported but dominated alternatives, is typically found when subjects perform a selection under conflicting information. Moreover, we have seen that such system can guide selections among a large number of alternatives by intersecting partial sources of information. Such a process might play an essential role in visual and auditory language processing, where the perceptual system must integrate various sources of partial information on-line.

It may be noted that many of the distributed models exhibit dynamics that are similar to those seen in our model, without making the restrictions we have adopted, in which between-layer connections are strictly excitatory and within-layer connections are inhibitory. For example, Kawamoto (1988) used the BSB model (Anderson et al, 1977) to simulate the settling of an attractor network into a state representing one of the two alternative intreptations of an ambiguous word such as WIND. In his simulations, the input (taken to represent the spelling of the word) tended initially to produce a pattern partially overlapping with the appropriate representation of both meanings of WIND. As the settling process continued, one interpretation was eventually reached, and as it began to dominate the representation the other was correspondingly suppressed. A curve very similar to the left-hand panel of Figure 13 resulted. Such a pattern represents an abstract competition among patterns, but is not necessarily a sign of lateral inhibition as construed literally, since such patterns arise even when, as in the BSB model, inhibition has no special meaning, in that activation values both above and below zero are propagated to other units, and positive and negative activations have equal status. Positive weights make units cause others to adopt the same sign, while negative weights make units cause others to take opposite signs. If the sender is negative, the effect of a negative weight therefore becomes excitatory. In fact, competition-like effects (that is, the pattern in the left-hand panel of Figure 13) can even be produced in networks where all of the connection weights are positive.

This observation raises the question as to whether the information processing principle that we have associated with the neural concept of lateral inhibition is really a more abstract principle of competition. The observation can be extended further by noting that in Kawamoto's model, the inputs governing the influence of the stimulus on the initial activations of units are not restricted to positive values. This suggests that the same information processing result — initial partial activation of multiple alternatives followed by suppression of all but one — can, at least in some cases, arise without adhering to the literal physiological constraint of between-level excitation and within-level inhibition. The question that remains is whether there are aspects of the behavior of the BSB model and other models that do not adhere to the physiological constraints that would be improved if they were made more consistent with the physiology. For example, such models could be modified to use sparse distributed representations, in which units tend to be quiescent most of the time and communicate with other neurons only when their activation increases above a baseline. There are now several models in the neural network literature that incorporate these properties (e.g.., Amit et al., 1994; Buhmann, Divko, & Schulten, 1989; Chappell & Humphreys, 1994; Tsodyks & Feigelman, 1988), but to our knowledge they have not yet been applied in much detail to the modeling of the time-course of information processing.

Lastly we consider the concept of the criterion for responding. As we have noted, the assumption that there is a criterion for responding can suggest some external decision making process that looks at the states of units and asks whether any have reached a critical value. This method of triggering of responses is trivial if there is a single unit associated with each response, but it is more difficult to extend to the responses represented as distributed patterns. This would require the external decision-maker to determine whether the network had reached one of its coherent distributed attractor states, imposing a substantial computational burden on the decision maker. However, the problem goes away if, as in our approach, the exteral decision maker is dispensed with entirely, in favor of a more direct implementation of the actual process of response initiation, in which responses triggered when the activations of the population corresponding to the response reaches a critical level of activation. The triggering of an action potential when depolarization at a cell body reaches a critical value is an example of such a mechanism. Crucially, though, the activation that is triggered need not be the activation of a single neuron; recurrent connections among neurons in a population could lead to the triggering of an aggregate neural response that ultimately leads to the muscular contractions that subserve behavioral responses. The triggering input in such cases need not originate from a single neuron or even a dedicated population of neurons; it can instead be produced by a distributed pattern of activation that produces sufficient input to trigger a response by the use of weighted connections. In this case, there would be no special difficulty associated with using a distributed pattern of activation to represent the input to the trigger. Furthermore, the activation that is triggered need not be the activation of a single neuron; recurrent connections among neurons in a population could lead to the triggering of a population response.

It remains for further research to delineate more exactly how the principles we have explored using single variables to represent each response alternative carry over to fully distributed models. Based on the arguments we have presented in this section, we expect that more fully distributed models can exhibit most of the same properties. Whether new properties relevant to understanding the process of choosing an appropriate response in a choice reaction time paradigm emerge from these more distributed representations remains to be seen.

Conclusion

In this paper we have argued for the usefulness of formulating models of information processing in terms of principles derived from neurophysiology. We have found that sensitivity to these principles can suggest and/or motivate particular assumptions about the nature of human information processing, and that models based on these assumptions appear to capture many aspects of the dynamics of information processing. Given how useful these assumptions have been, it seems likely that the strategy of formulating models in terms of principles derived from neurophysiology will continue to be productive.

References

- Abbott, L. F. (1992). Firing rate models for neural populations. In O. Benhar, C. Bosio, P. Giudice, & E. Tabet (Eds.), *Neural networks: From biology to high energy physics.* Pisa: ETS Editrice.

- Amit, D. J. (1989). *Modeling brain function: The world of attractor neural networks.* Cambridge, UK: Cambridge University Press.

- Amit, D. J., Brunel, N., & Tsodyks, M. V. (1994). Correlations of cortical Hebbian reverberations, *Journal of Neuroscience, 14*, 6435-6445.

- Amit, D. J., & Tsodyks, M. (1991). Quantitative study of attractor neural network retrieving at low spike rates: I Substrate - spikes, rates and neuronal gain. *Network, 2*, 259-273.

- Anderson, J. A. (1973). A theory for the recognition of items from short memorized lists. *Psychological Review, 80*, 417-438.

- Anderson, J. A. (1991). Why having so many neurons, do we have so few thoughts? In W. E. Hockley, & S. Lewandowsky (Eds.) *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 477-507). Hilldale, NJ: Erlbaum.

- Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception and probability learning: Some applications of a neural model. *Psychological Review, 84*, 413-451.

- Ashby, F. G. (1982). Deriving exact predictions from the cascade model. *Psychological Review, 89*, 599-607.

- Audeley, R. J., & Mercer, A. (1968). The relation between decision time and the relative response frequency in a blue-green discrimination, *British Journal of Mathematical and Statistical Psychology, 21*, 183-192.

- Audley, R. J., & Pike, A. R. (1965). Some stochastic models of choice, *British Journal of Mathematical and Statistical Psychology, 18*, 207-225.

- Breitmeyer, B. G. & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing. *Psychological Review, 83*, 1-36.

- Buhmann J., Divko R., & Schulten K., (1989). Associative memory with high information content. *Physical Review, A39*, 2689-2692.

- Busey, T. A., & Loftus, G. R. (1994). Sensory and cognitive components of visual information acquisition, *Psychological Review, 101*, 446-469.

- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory. *Psychological Review, 100*, 432-459.

- Chappell, M., & Humphreys, M. S. (1994). An autoassociative neural network for sparse representations - analysis and application to models of recognition and cued recall. *Psychological Review, 101*, 103-128.

- Chelazzi, L., Miller, E., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature, 363*, 345-347.
- Cohen, J., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes. *Psychological Review, 97*, 332-361.
- Cohen, J., Servan-Schreiber, D., & McClleland, J. L. (1992). A parallel distributed processing approach to automaticity. *American Journal of Psychology, 105*, 239-269.
- Coles, M. G., Gratton, G., Bashore, T. R., Eriksen, C. W., & Donchin, E. (1985). A psychophysiological investigation of the continuous flow model of human information processing. *Journal of Experimental Psychology: Human Perception and Performance, 11*, 529-553.
- Corbett, A. T., & Wickelgren, W. A. (1978). Semantic memory retrieval: Analysis by speed and accuracy tradeoff functions. *Quarterly Journal of Experimental Psychology, 30*, 1-15.
- Cox D., R., and Miller H., D., (1965). The theory of stochastic processes. New York: Wiley.
- Creutzfeldt, O., Fuster, J. M., Hertz, A., & Straschill, M. (1966). Some problems of information transmission in the visual system. In J. C. Eccles (Ed.), *Brain and the conscious experience.* Berlin: Springer-Verlag.
- Dosher, B. A. (1976). The retrieval of sentences from memory: A speed-accuracy study. *Cognitive Science, 8*, 291-310.
- Douglas, R. J., & Martin, K. A. (1990). Neocortex. In G. M. Shepherd (Ed), *The synaptic organization of the brain*, pp. 389-438. New York: Oxford University Press.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A., & Suarez, H. (1995). Recurrent excitation in neocortical circuits. *Science, 269*, 981-985.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a non-search task. *Perception & Psychophysics, 16*, 143-149.
- Eriksen, C. W., & Eriksen, B. A. (1979). Target redundancy in visual search: Do repetitions of the target in visual search impair performance? *Perception & Psychophysics, 26*, 195-205.
- Eriksen, C. W., & Schultz, D. W. (1979). Information processing in visual search: A continuous flow conception. *Perception & Psychophysics, 25*, 249-263.
- Ermentrout, B. (1994). Reduction of conductance-based models with slow synapses to neural nets. *Neural Computation, 6*, 679-695.
- Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties, *Cognitive Science, 6*, 205-254.
- Funahashi, S., Chafee, M. V., & Goldman-Racik, P. S. (1993) Prefrontal neural activity in rhesus-monkeys performing a delayed-response task. *Nature, 365*, 753-756.
- Fuster, J. (1991). The prefrontal cortex and its relation to behavior. *Progress in Brain Research, 87*, 201-211.
- Goldman-Racik, P. S. (1992). Working memory and the mind. *Scientific American, 267*, 111-117.

- Gratton, G., Coles, M. G.,& Donchin, E. (1989). A procedure for using multi-electrode information in the analysis of components of the event related potential: Vector filter. *Psychophysiology, 26*, 222-232.
- Gratton, G., Coles, M. G., Sirevag, E. J., Eriksen, C. W., & Donchin, E. (1988). Pre-and poststimulus activation of response channels: A psychophysiological analysis. *Journal of Experimental Psychology: Human Perception and Performance, 14*, 331-344.
- Grice, G. R. (1972). Application of a variable criterion model to auditory reaction time as a function of the type of catch trial. *Perception & Psychophysics, 12*, 103-107.
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding, *Biological Cybernetics, 23*, 121-134.
- Grossberg, S. (1978). A theory of visual coding, memory and development. In E. L. J. Leeuwenberg, & H. F. Buffart (Eds.), *Formal theories of visual perception.* New York: Wiley.
- Herrmann M., Ruppin, E. and Usher, M. (1993) On the Dynamic Activation of Memory. *Biological Cybernetics, 68*, 455-463.
- Hertz, J., Krogh, A., & Palmer, R. G, (1990), Introduction to the theory of neural computation. Redwood, CA: Addison-Wesley.
- Hess, R., Negishi, K., & Creutzfeld, O. (1975). The horizontal spread of intracortical inhibition in the visual cortex. *Experimental Brain Research, 22*, 415-419.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In: D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*, Vol. 1, pp. 77-109. Cambridge, MA: MIT Press.
- Hinton, G. E., & Sejnowski, T. J. (1983). Optimal perceptual inference. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 448-453.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science, USA, 79*, 2554-2558.
- Hopfield, J. J., & Tank, D. W. (1985) Neural computation of decisions in optimization problems. *Biological Cybernetics, 52*, 141-152.
- Humphreys, M. S., Wiles, J., & Dennis, S. (1994). Toward a theory of human memory – Data-structure and access processes. *Behavioral and Brain Sciences, 17*, 655-667.
- Jamieson, D. G., & Petrusik, W. M. (1977). Preference and time to choose. *Organizational Behavior and Human Performance, 19*, 56-67.
- Kawamoto, A. (1988). Distributed representations of ambiguous words and their resolution in a connectionist network. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology, and Artificial Intelligence.* San Mateo, CA: Morgan Kaufmann.
- Keele S. W., & Neill W., T. (1978). Mechanisms of attention, In E.C. Cartrette & M. P., Friedman (Eds.), *Handbook of perception, Vol. 9*, 3-47. New-York: Academic Press.

- LaBerge, D. A. (1962). A recruitment theory of simple behavior. *Psychometrika, 27*, 375-396.
- Laming, D. R. (1968). *Information theory of choice-reaction times.* New York: Academic Press.
- Link, S. W. (1975). The relative judgement of two-choice response time. *Journal of Mathematical Psychology, 12*, 114-135.
- Link, S. W., & Heath, R. A. (1975). A sequential theory of psychological discrimination, *Psychometrica, 40*, 77-105.
- Loftus, G. F., Busey, T. A., & Senders, J. (1993). Providing a sensory basis for models of visual information acquisition, *Perception and Psychophysics, 54*, 535-554.
- Luce. D. R. (1986). *Response times.* New York: Oxford University Press.
- Massaro, D. W., & Cohen, M. M. (1983). Phonological constraints in speech perception. *Perception & Psychophysics, 34*, 338-348.
- McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review, 86*, 287-330.
- McClelland, J. L. (1986). The programmable blackboard of reading. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2.* Cambridge, MA: MIT Press.
- McClelland, J. L. (1991). Stochastic interactive activation and the effect of context on perception. *Cognitive Psychology*, 23, 1-44.
- McClelland, J. L. (1993). Towards a theory of information processes in graded, random, interactive networks. In D. E. Meyer, & S. Kornblum (Eds.), *Attention & Performance XIV: Synergies in experimental psychology, artificial intelligence and cognitive neuroscience.* Cambridge, MA: MIT Press.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why There are Complementary Learning Systems in the Hippocampus and Neocortex: Insights from the Successes and Failures of Connectionist Models of Learning and Memory. *Psychological Review, 102*, 419-457.
- McClelland, J. L., & O'Regan, J. K. (1981). Expectations increase benefit derived from parafoveal visual information in reading words aloud. *Journal of Experimental Psychology: Human Perception and Performance, 7*, 634-644.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of contexts effect in letter perception. *Cognitive Psychology, 23*, 1-44.
- McClelland, J. L., Rumelhart, D. E., & the PDP Research Group, (1986). *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 2.* Cambridge, MA: MIT Press.
- McElree, B., & Dosher, B. A. (1989). Serial position and set size in short-term memory: The time course of recognition. *Journal of Experimental Psychology, 118*, 346-373.

- McNaughton, B. L. & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences, 10*, 408-415.

- Miller, E. K., Erickson, C. A., & Desimone, R. (1995). Comparison of prefrontal and inferior temporal neurons during performance of a memory task. *Society of Neuroscience Conference Abstracts, 267.14*

- Miyashita, Y. (1988). Neuronal correlate of visual associative long term memory in primate visual cortex. *Nature, 335*, 817-820.

- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review, 76*, 165-178.

- Movellan, J. R. and McClelland, J. L. (1993). Learning Continuous Probability Distributions with Symmetric Diffusion Networks. *Cognitive Science, 17*, 463-496.

- Movellan, J. R., & McClelland, J. L. (1995). *Stochastic interactive activation, Morton's Law, and optimal pattern recognition.* Technical Report PDP.CNS.95.4. Pittsburgh, PA: Carnegie Mellon University, Department of Psychology.

- Pachella, R. C., & Fisher, D. (1969). Hick's law and the speed-accuracy trade-off in absolute judgement. *Journal of Experimental Psychology, 92*, 378-384.

- Pachella, R. C., & Pew, R. W. (1968). Speed-accuracy tradeoff in reaction time: Effect of discrete criterion times. *Journal of Experimental Psychology, 76*, 19-24.

- Pike, A. R. (1968). Latency and relative frequency of response in psychophysical discrimination, *British Journal of Mathematical Psychology, 21*, 161-182.

- Plaut, D. C. & Shallice, T. (1993). Deep Dyslexia: A Case Study of Connectionist Neuropsychology. *Cognitive Neuropsychology, 10*, 377-500.

- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*, 59-108.

- Ratcliff, R., & VanZandt T. (1995). *Comparing connectionists and diffusion models of reaction time.* Unpublished manuscript, Northwwestern University, Evanston, IL.

- Ratliff, F. (1965). *Mach bands: Quantitative studies on neural networks in the retina.* San Francisco, CA: Holden-Day.

- Reed, A. V. (1973). Speed-accuracy trade-off in recognition memory. *Science, 181*, 574-576.

- Reed, A. V. (1976). List length and the time course of recognition in immediate memory. *Memory & Cognition, 4*, 16-30.

- Ricciardi, L. (1977). *Diffusion processes and related topics in biology.* New York: Springer-Verlag.

- Roberts, S., & Sternberg, S. (1993). The meaning of additive reaction time effects: Tests of three alternatives. In D. E. Meyer, & S. Kornblum (Eds.),*Attention and performance XIV.* Cambridge, MA: MIT Press.

- Rumelhart, D. E. (1977). Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and Performance VI.* Hilsdale, NJ: Erlbaum.
- Rumelhart, D. E. & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60-94.
- Rumelhart, D. L., McClelland, J. L., & the PDP Research Group, (1986). *Parallel distributed processing: Explorations in the microstructure of cognition* Vol. 1. Cambridge, MA: MIT Press.
- Rumelhart, D. L., & Zipser, D. (1986), Feature discovery by competitive learning. In D. L. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition*, Vol. 1, pp.151-193. Cambridge, MA: MIT Press.
- Ruppin, E., & Yeshuron, H. (1991). Recall and recognition in an attractor neural network of memory retrieval, *Connection Science* 3, 381-399.
- Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical T organization. *Current Opinion in Neurobiology, 4*, 569-579.
- Shepherd, G. M. (1988). The basic circuit of cortical organization. In M. S. Gazzaniga (Ed.), *Perspectives in memory research.* Cambridge, MA: MIT Press.
- Softky, W. R., & Koch, C. (1992). Cortical cells should fire regularly, but do not. *Neural Computation, 4*, 643-645.
- Softky, W. R., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience, 13*, 334-350.
- Smith, P. L. (1995). Psychophysically principled models of visual simple reaction time. Psychological Review, 102, 567-593.
- Smolensky, P. (1986). Neural and conceptual interpretation of PDP models. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*, Vol. 2, pp. 390-431. Cambridge, MA: MIT Press.
- Stone, M. (1960). Models for choice reaction time. *Psychometrika, 25*, 251-260.
- Swensson, R. G. (1972). The elusive trade-off: Speed versus accuracy in visual discrimination tasks. *Perception & Psychophysics, 12*, 16-32.
- Tanaka, K., & Saito, Y. (1991). Coding visual images of objects in the inferotemporal cortex of macaque monkey. Journal of Neurophysiology, 66, 170-189.
- Treves A. (1990). Threshold linear formal neurons in autoassociative nets. *Journal of Physics, A, 23*, 2631-2650.
- Tsodyks M. V., & Feigelman M. V. (1988). The enhanced storage capacity in neural networks with low activity level. *Europhysics Letters, 40*, 101-105.

- Usher, M., & Niebur, E. (in press). Modeling the temporal dynamics of IT neurons in visual search: a mechanism for top-down selective attention.*Journal of Cognitive Neuroscience.*

- Usher, M., Stemmler, M., Koch, C., & Olami, Z. (1994). Network amplification of local fluctuations causes high rate variability, fractal firing patterns and oscillatory local field potentials. *Neural Computation, 6*, 795-836.

- Usher, M., Stemmler, M., & Olami, Z. (1995). Dynamic pattern formation leads to 1/f noise in neural populations. *Physical Review Letters, 74*, 325-329.

- Vickers, D. (1970). Evidence for an accumulator of psychophysical discrimination. *Ergonomics, 13*, 37-58.

- Vickers, D., Caudrey, D., & Wilson, R. (1971). Discrimination between frequency of occurrence of two alternative events, *Acta Psychologica, 35*, 151-172.

- Wald, A., & Wolfowitz, J. (1948). Optimum characteristic of sequential probability ratio test. *Annals of Mathematical Statistics, 19*, 326-329.

- Wickelgren, W. A. (1977). Speed accuracy tradeoff and information processing dynamics. *Acta Psychologica, 41*, 67-86.

- Wickelgren, W. A., & Corbett, A. T. (1977). Associative interference and retrieval dynamics in yes-no recall and recognition. *Journal of Experimental Psychology: Human Learning and Memory, 3*, 189-202.

- Wilding, J. M. (1974). Effects of stimulus discriminability on the latency distribution of identification responses. *Acta Psychologica, 38*, 483-500.

- Zipser, D., Kehoe, B., Littlewort, G., and Fuster, J. (1993). A Spiking Network Model of Short-Term Active Memory. *Journal of Neuroscience, 13* 3406-3420.

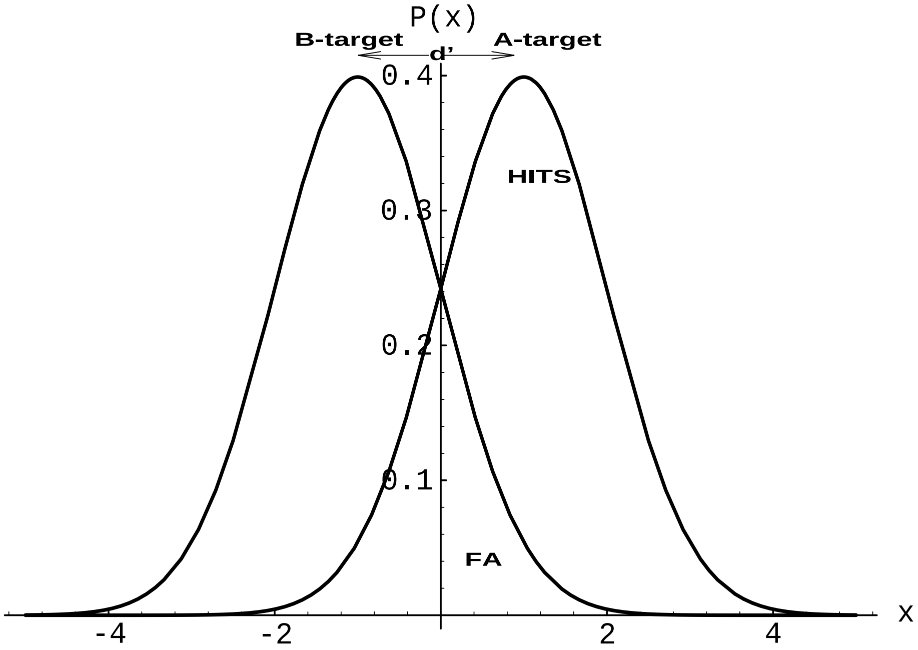

Figure 18: Density distribution for the difference in activation of the two response units, for the two choice stimuli.

Appendix: Derivation of d' and correct-response-probability, P.

Assume a binary choice with two target stimuli 'A' and 'B'. Denote by x the difference in activation accumulated by the two response units (positive x reflect A-choices, and negative x reflect B-choices). Due to noise, the activation for A and B stimuli, at time t following the input presentation is given by two Gaussians shown in Figure 18. The mean of the Gaussians $\mu(t)$ and their standard-deviations SD, are given in accordance to the solutions of the Fokker-Plank Equation (see Ricciardi, 1977) by Equations 8. Thus d'—the separation between the two distributions—is $d' = \frac{2\mu(t)}{SD(t)}$. The probability for a correct response is the integral from zero to infinity over the A-target Gaussian (if $x > 0$, A is chosen), which relates to d' as:

$$P(t) = \frac{1}{\sqrt{2\pi}} \int_0^\infty \exp[-\frac{(x-\mu(t))^2}{2SD^2(t)}] = \frac{1}{\sqrt{2\pi}} \int_{-d'/2}^\infty \exp(-\frac{x^2}{2SD^2(t)}) \, .$$