Title
A Bayesian Sequential Sampling Model of Choice Reaction Time Incorporating Stimulus Onset/Duration Uncertainty

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 36(36)

ISSN
1069-7977

Authors
Meyer, Jordan
Zhang, Jun

Publication Date
2014

Peer reviewed

A Bayesian Sequential Sampling Model of Choice Reaction Time Incorporating Stimulus Onset/Duration Uncertainty

Jordan Meyer (jlmeyer@umich.edu) and **Jun Zhang (junz@umich.edu)**
University of Michigan, Department of Psychology
Ann Arbor, MI 48109 USA

Abstract

We propose a Bayesian sequential sampling model of choice reaction time (RT) which incorporates uncertainties about stimulus identity, onset, and duration. The model is the now-standard random-walk/drift-diffusion model, with a threshold-based response mechanism. The "substance" of the drift, however, is the posterior probability (belief) that a participant updates on a moment-to-moment basis during a trial — the update is done by combining the likelihood function on the evidence (modeling trial-dependent perception) with prior probability about stimulus identity, onset time, and duration (modeling trial-independent task knowledge). Response threshold, which equals the probability of correct response in choosing each alternative conditioned on prior knowledge and accumulated evidence, modulates speed-accuracy tradeoff. While sequential Bayesian updating without temporal uncertainty (regarding stimulus onset/offset) is trivial, we overcome the hurdle of incorporating the temporal prior into the dynamics of belief updating to derive an analytic expression for Bayesian belief. The advantage of the Bayesian formulation is to allow full control of where and how many free parameters appear: in likelihood functions, priors, or response threshold. Comparison of computer simulation of our model with human performance data (Smith, 1995) will be reported.

Keywords: Choice Reaction Times; Bayesian Modeling; Sequential Sampling; Stimulus Onset/Duration Uncertainty

Introduction

The study of the underlying mechanisms that mediate trial-by-trial variation in reaction time (RT) in simple choice tasks has been prolifically pursued in the field of mathematical psychology for over a century (cf. Luce, 1986, for a classic review). Here, we propose as a novel model for choice-RT experiments an exact Bayesian belief update procedure that dynamically combines sequential accumulation of evidence with prior knowledge about stimulus identity as well as temporal uncertainty in stimulus presentation.

The general leitmotif of most RT models is the assumption that information regarding the nature and onset of a stimulus is accumulated stochastically in discrete time steps from the beginning of a trial until some stopping point is reached, at which point a response is issued and an RT measured. Wald (1947) offered a seminal analysis of the statistical properties of such sequential sampling processes. In his SPRT ("Sequential Probability Ratio Test") model, a stimulus is presented at time zero (starting point), an evidence e_t is stochastically generated by the environment at time t, and the stimulus may be either a or b, with associated likelihood functions $l_a(e)$ and $l_b(e)$. Their probabilities prior to ($\Pr[\cdot]$) and after receiving e_t ($\Pr[\cdot|e_t]$) are related by:

$$\log\left(\frac{\Pr[a|e_t]}{\Pr[b|e_t]}\right) = \log\left(\frac{\Pr[a]}{\Pr[b]}\right) + X(e_t) \ , \qquad (1)$$

where the log likelihood ratio $X(e_t) = \log(l_a(e_t)/l_b(e_t))$ is a random number determined by the stochastic evidence e_t. Thus, the variable $\log(\Pr[a|e_t]/\Pr[b|e_t])$ can be seen as a random walk process with "drift rate" equal the expectation of $X(e_t)$, namely, $E_{\Pr}\{\log(\Pr[a|e_t]/\Pr[b|e_t])\}$ which, assuming a to be the true stimulus, equals

$$\sum_e \Pr[a|e_t] \log\left(\frac{\Pr[a|e_t]}{\Pr[b|e_t]}\right) = \mathrm{KL}(\Pr[a|e_t] \parallel \Pr[b|e_t]) \qquad (2)$$

where $\mathrm{KL}(\cdot \parallel \cdot)$ denotes Kullback-Leibler divergence.

Stone (1960) adapted the SPRT model to study choice reaction times. Further development (Thomas, 1975; Swensson & Green, 1977) derived the surprising prediction that RT distributions for correct and error trials should be identical. This has long been known to be empirically violated (e.g., Laming, 1968; Link, 1975). Ashby (1983) proposed an amelioration in the form of a constant bias term $\log(k)$ added to the drift rate $X(e_t)$ in (1). Later, the SPRT model was generalized to the case of an arbitrary number of hypotheses (Baum & Veeravalli, 1994; Draglia, Tartakovsky, & Veeravalli, 1999).

Stimulus Uncertainty and Prior Knowledge

Other variants of sequential sampling models were proposed (e.g., Ratcliff and Smith, 2004; Busemeyer and Townsend, 1993; Usher and McClelland, 2001; Bogacz, Brown, Moehlis, Holmes, and Cohen, 2006). However, one curiosity is that most models assume the initiation of evidence accumulation is concurrent with the onset of stimulus presentation. But for such a process to be triggered, the system would already have known that a target has been presented. This seems to imply the unsavory conclusion that choice tasks consist of a stimulus detection stage in serial connection with a stimulus identification stage. Alternatively, one can allow for "premature" sampling, in which pure noise information is sampled before stimulus onset (e.g., Rouder, 1996). However, there has not yet been a principled way of incorporating prior knowledge about stimulus onset (and duration) uncertainty into RT models.

By uncertainty, we refer both to uncertainty regarding target identity, and uncertainty regarding its temporal presentation (onset time, offset time or duration). If a participant is trained extensively on a given task, it is usual for them to acquire a significant amount of task knowledge regarding the statistical structure of the stimulus. This knowledge can be embodied in a joint prior distribution over stimulus identity and stimulus onset-duration across trials. In the sequential sampling framework, a participant on a given trial receives a sequence of evidence that is (stochastically) related to the true stimulus state, and the participant uses this evidence to adjust their decision tendency over time. The participant will in general be *simultaneously* using their prior knowledge of the stimulus presentation to inform this adjustment. The dynamic interaction between information accumulation and prior knowledge about *temporal aspect* of the stimulus uncertainty during the decision process has not been well researched in the RT literature, and is the key theoretical motivation for the present work.

Previously, one Bayesian sequential sampling model was proposed that explicitly took account of temporal uncertainty in modeling detection RT. Stein and Rapoport (1978) modeled performance on a simple detection task (without a choice component) by combining prior probability of stimulus onset time with a Bayesian sequential sampling procedure. Unfortunately, the model did not extend to a choice paradigm (let alone with > 2 alternatives), nor was the stimulus onset distribution general enough (from their restricted case of Poisson distribution). Stimulus onset uncertainty in a choice task has been treated only with the above-mentioned *ad hoc* method of premature sampling. On the other hand, the study of the effect of the stimulus identity prior (i.e., stimulus presentation rate) dates back to the seminal studies of Hick (1952), Hyman (1953), and Crossman (1953), the latter two of whom proposed on information-theoretic grounds that the reaction time to the i-th stimulus with presentation probability p_i should obey

$$RT = c - d\log(p_i) \quad (3)$$

(with c, d constants) so that the mean RT is linearly related to the stimulus entropy $H = -\sum_i p_i \log(p_i)$. The prediction made by (3) has found some empirical support, for example in the study of an oculomotor task by Carpenter and Williams (1995).

Most interestingly, the question of how stimulus and onset/offset uncertainty interact and influence each other during the decision process is not well known empirically nor well understood theoretically. The main goal of the present work is to provide a Bayesian modeling framework for decision processes that systematically combines both types of prior knowledge with moment-to-moment stochastic evidence accumulation.

A Bayesian Sequential Sampling Model

Bayesian theory has proven to provide a powerful framework for decision making (DM) in uncertain environments . In general, a Bayesian DM agent may be interested in a particular phenomenon about which the agent maintains a complete set $\mathcal{H}$ of mutually exclusive hypotheses $h_i \in \mathcal{H}$. By making observations of its environment, the agent accumulates evidence regarding the relative truth or falsehood of each hypothesis in $\mathcal{H}$. Before such information accumulation process begins, the agent has a pre-established judgment regarding the probability of the truth of each h_i, denoted $\Pr[i]$. This is the prior distribution. Usually there is some information intrinsic to the phenomenon in question that allows for the establishment of nontrivial (i.e. non-uniform) priors. During the observation stage, the agent receives a particular evidence $e \in \mathcal{E}$. The probability of e having been observed given that h_i is true is denoted $\Pr[e|i]$ and referred to as the likelihood of the evidence; this is how the agent accounts for the evidence in light of its hypothesis h_i. At each time t, the agent uses the new evidence e_t to refine its estimation of the probability of its hypotheses from its previous estimate. At $t = 1$, the first datum e_1 is collected, and the posterior probability is calculated according to Bayes' theorem:

$$\Pr[i|e_1] = \frac{\Pr[e_1|i]\Pr[i]}{\sum_j \Pr[e_1|j]\Pr[j]} . \quad (4)$$

Now, if at time $t = 2$ datum e_2 is generated independently of e_1, Bayes' theorem may be applied again to yield the new posterior probability

$$\Pr[i|e_1 e_2] = \frac{\Pr[e_2|i]\Pr[i|e_1]}{\sum_j \Pr[e_2|j]\Pr[j|e_1]} , \quad (5)$$

where the independence of e_1 and e_2 is assumed. The update of the form (5) may be applied iteratively at the reception of each new datum e_t, obviating the need for storage of data in memory.

Applied to the choice RT setting, evidence accumulation starts at the beginning $t = 0$ of each trial. As the Bayesian decision maker updates posterior beliefs over the set of hypotheses, he/she must evaluate the utility/cost of obtaining more evidence before terminating a trial by making a choice. Generally, a participant is instructed to perform a speeded choice task as accurately and quickly as possible, and rewarded accordingly. The determination of an optimal stopping rule for such a sequential sampling process that maximizes the expected reward is the familiar speed-accuracy tradeoff problem. In most sequential sampling sampling models, the threshold parameter and correct-choice rate are (generally monotonically) related but depend on other parameters. In contrast, a Bayesian sequential sampling model, in which posterior probability is the substance of drift-with-diffusion, has the advantage of equating the threshold parameter with the conditional probability of that choice being correct. This provides a special link between subjective threshold evaluation and external observed performance not present in other models, as shown below.

It must be realized that a special challenge arises in applying the vanilla Bayesian (5) updating to the data sampling process. It has been implicitly assumed that the environment remains *stationary*; that is, the denotation of environment status embodied in the hypotheses and likelihood functions must not change over the whole course of evidence accumulation. However, in a choice task in which the stimulus onset and offset vary probabilistically across trials, in each trial the environment undergoes a stochastic transition from a target-off state to a target-on state, then possibly another transition from target-on back to target-off. Because of this temporal uncertainty regarding the environmental state, the hypotheses of behavioral interest, i.e., whether the target (ii) has not yet appeared, (ii) is currently on, or (iii) was on in the past but has been turned off, cannot be directly ported into the Bayesian formulation. To work around this, we develop a formulation of an exact Bayesian model in terms of *mini-hypotheses*, from which we then can straightforwardly derive a behavior-level equivalent.

Basic Elements of the Bayesian Model

We consider the following general paradigm for a choice reaction time experiment, with arbitrary stimulus onset-offset. In each trial, a warning signal is given at time $t = 0$, beginning the trial. The k-th target (out of N possible ones) appears at some time $t = \tau_1 > 0$, and remains on until some time $t = \tau_2 > \tau_1$. The index k for the set of targets is general and refers to whatever stimulus attribute is relevant for the given task (e.g., spatial location, frequency, orientation, etc.). For simplicity, we will use the example of target location in a visual paradigm throughout, referring interchangeably to the k-th target and the target at the k-th location. Reaction time is defined as time lapse from τ_1 to the time step in which the participant makes a choice response (a judgment that "k-th target has already appeared"). The target location k, onset time τ_1, and offset time τ_2 vary probabilistically from trial to trial according to the design of the experiment. Importantly, it is assumed the participant has been extensively trained on the task, so that the prior probabilities for stimulus location, onset, and offset are effectively known.

The model consists of two basic elements: decision (accumulator) units and sensor (perceptual) units. The sensors reports evidence generated stochastically at each time step based on the true state of the environment. Each sensor is endowed with a particular tuning on the relevant dimension of the stimulus space, and its activation level depends (assumed to be deterministically) on the evidence stochastically generated by the environment. In this way, the sensors implement the likelihood functions in the Bayesian framework. The decision units *read-in* the likelihood evaluation from the sensors and combine it with prior knowledge about the stimulus uncertainty in the task (presumably store in memory representation) to update the posterior probabilities of each behavior-level hypothesis. This cycle is continued until the posterior probability for one of the $2N+1$ behavior-relevant hypotheses (to include onset-offset of N targets plus target absent) reaches a pre-determined threshold, at which time the corresponding response is made.

The following notation is used to denote various aspects of the Bayesian model framework:

h_{k,τ_1,τ_2} – the hypothesis that "the k-th target appears at time τ_1 and disappears at time τ_2." These are the so-called mini-hypotheses mentioned earlier. The set of all such mini-hypotheses $\mathcal{H} = \{h_{k,\tau_1,\tau_2} : 0 < \tau_1 < \tau_2,\ 1 \le k \le N\}$ forms a complete, mutually exclusive partition of the space of hypotheses regarding target identity and temporality. In contrast, the behavior-level hypotheses maintained by the decision units are different; they are formulated as unions of mini-hypotheses as shown below.

$f_k(\tau_1, \tau_2)$ – the prior probability that h_{k,τ_1,τ_2} is true. This distribution is presumably acquired through task instruction or learned through previous task training. Of course, $\sum_k \sum_{\tau_1>0} \sum_{\tau_2>\tau_1} f_k(\tau_1,\tau_2) = 1$.

e_t – the particular datum drawn from $\mathcal{E}$ through stochastic generation by the sensor(s) during run-time $(t,\ t+1)$. The form or structure of such evidence e_t can be determined as appropriate for modeling purposes. For example, it may be uni- or multi-dimensional, with discrete or continuous values.

$E(t)$ – the entire sequence of evidences generated up to time t in a given trial. That is, $E(t) = \{e_0 e_1 \cdots e_{t-1}\}$.

$l_0(e_t)$ – the probability of evidence e_t being generated when no target has appeared. This is the sensory activation to the background noise.

$l_k(e_t)$ – the probability of evidence e_t being generated given that the k-th target is currently on. This is the sensory activation in the presence of the k-th target.

$l_{\bar{k}}(e_t)$ – the probability of evidence e_t being generated given that no target is currently on, but the k-th target was previously on. In the current implementation, this is taken to be equivalent to $l_0(e_t)$, but can in principle take alternative functional forms.

In a given trial, at each time step t, the probabilities of all mini-hypotheses h_{k,τ_1,τ_2} are evaluated according to the Bayesian recursive update formula (4) and (5), using the previous probabilities computed at time step $t-1$ and the likelihood of the evidence e_{t-1} generated between $t-1$ and t. The process continues until a response threshold is reached. Note that the mini-hypotheses are about target at a particular location with particular onset and offset times. Hence they are *stationary* during any trial – they are either true or false throughout the trial. As such, vanilla Bayes updating is adequate. However, the decision maker is not asked to act upon each mini-hypothesis; rather he/she is instructed to respond whenever a stimulus at a particular location has appeared. This leads to the behavioral model below.

Formulation of Behavioral Model

For modeling behavior, only the following hypotheses are task-relevant (where run-time t can be interpreted as "now")

$H_0(t)$ – No target has yet appeared as of time t.

$H_k(t)$ – The k-th target is currently on at time t.

$H_{\bar{k}}(t)$ – No target is currently on at time t, but the k-th target was previously on. That is, the k-target has been off as of time t.

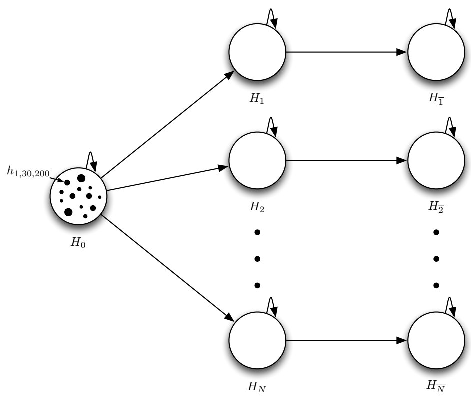

Figure 1: Markovian transition diagram in which mini-hypotheses (depicted as black dots) switch their memberships over the time-course of a trial into various behavior-level hypotheses (depicted as $2N+1$ nodes). Consider, for example, the mini-hypothesis marked $h_{1,30,200}$, stating that stimulus is presented at time $\tau_1 = 30$ at location 1 and then turned off at time $\tau_2 = 200$. Starting from its membership in H_0, which maintains that "target has not been turned on as of now", this mini-hypothesis $h_{1,30,200}$ will transition its membership, at run-time $t = 30$ during a trial, from H_0 to H_1, which states that "target has been turned on at location 1 as of now" (i.e., as of $t = 30$), and, at run-time $t = 200$, from H_1 to $H_{\bar{1}}$, which states that "target has been turned off at location 1 as of now" (i.e., as of $t = 200$). Note that these membership transitions occur on *each trial regardless* what and when a stimulus is presented on any particular trial. The various sizes of dots in H_0 schematically represent the associated posterior probabilities, which change during the progression of a trial according to sequential Bayesian updating. Change of posterior probabilities on a given trial depends on stimulus presentation of that trial.

These are the hypotheses maintained by the decision units during run-time t of a given trial. Our key observation is that these behavior-level hypotheses can be represented as sets of mini-hypotheses *whose memberships vary as a function of time t* (see Figure 1):

$$\begin{cases} H_0(t) = \{\bigcup h_{k,\tau_1,\tau_2},\ t < \tau_1 < \tau_2,\ 1 \le k \le N\} \\ H_k(t) = \{\bigcup h_{k,\tau_1,\tau_2},\ \tau_1 \le t < \tau_2\} \\ H_{\bar{k}}(t) = \{\bigcup h_{k,\tau_1,\tau_2},\ \tau_1 < \tau_2 \le t\} \end{cases} \tag{6}$$

Performing direct Bayesian updating on the set $\mathcal{H}$ of mini-hypotheses, the posterior probabilities associated with these behavior-level hypotheses can be formulated as sums of the mini-hypotheses that constitute them:

$$\begin{cases} P_0(t) = \Pr[H_0(t)|E(t)] = \sum\limits_{\substack{t<\tau_1<\tau_2 \\ 1\le k\le N}} \Pr[h_{k,\tau_1,\tau_2}|E(t)] \\ P_k(t) = \Pr[H_k(t)|E(t)] = \sum\limits_{\tau_1 \le t < \tau_2} \Pr[h_{k,\tau_1,\tau_2}|E(t)] \\ P_{\bar{k}}(t) = \Pr[H_{\bar{k}}(t)|E(t)] = \sum\limits_{\tau_1 < \tau_2 \le t} \Pr[h_{k,\tau_1,\tau_2}|E(t)] \end{cases} \tag{7}$$

Update of Posterior Beliefs As mentioned above, it is not possible to simply apply the vanilla Bayesian update procedure for the posterior probabilities of any of the $2N+1$ behavior-level hypotheses $P_i(t)$, since their environmental denotation is non-stationary; the mini-hypotheses of which they are made up may switch membership dynamically from one to another as t evolves, as demonstrated in Figure 1. Therefore, at each t, the update of the behavior-level probabilities must take into account both (a) the standard Bayesian update of all mini-hypotheses conditioned on new evidence e_t (equivalently, on $E(t+1)$) and (b) the transfer of probability due to certain mini-hypotheses switching membership. We will denote by $P_i^*(t)$ the *partially* updated probability of $H_i(t)$ conditioned on new evidence e_t; that is, $P_i^*(t)$ has updated with respect to (a), but not (b). Applying the standard Bayes' formula, we get:

$$P_i^*(t) = \frac{P_i(t) l_i(e_t)}{P_0(t) l_0(e_t) + \sum\limits_{k=1}^{N} P_k(t) l_k(e_t) + \sum\limits_{k=1}^{N} P_{\bar{k}}(t) l_{\bar{k}}(e_t)} \tag{8}$$

To complete the update, we must take into account the probability transfer due to all mini-hypotheses that switched membership after time t. It is straightforward to check, simply by definition, which mini-hypotheses are to switch at a given time step (again refer to Figure 1). The complete update of the behavior-level hypotheses can thus be written:

$$\begin{cases} P_0(t+1) = P_0^*(t) - \sum\limits_{\substack{\tau_2 > t+1 \\ 1 \le k \le N}} \Pr[h_{k,t+1,\tau_2}|E(t+1)] \\ P_k(t+1) = P_k^*(t) + \sum\limits_{\tau_2 > t+1} \Pr[h_{k,t+1,\tau_2}|E(t+1)] \\ \qquad\qquad - \sum\limits_{\tau_1 < t+1} \Pr[h_{k,\tau_1,t+1}|E(t+1)] \\ P_{\bar{k}}(t+1) = P_{\bar{k}}^*(t) + \sum\limits_{\tau_1 < t+1} \Pr[h_{k,\tau_1,t+1}|E(t+1)] \end{cases} \tag{9}$$

It is now possible to eliminate the reference to the mini-hypotheses in the update equations of (9) by introducing an equivalent formulation in terms of "hazard" functions (cf. Luce, 1986). Let $g_k(t+1)$ denote the prior probability that the k-th target will appear at time $t+1$, given that no target has yet appeared. Similarly, let $g_{\bar{k}}(t+1)$ denote the prior probability that the k-th target will disappear at time $t+1$, given that the k-th target is currently on. It is fairly straightforward to see then that (9) can be reformulated as follows:

$$\begin{cases} P_0(t+1) = P_0^*(t) - \sum_{k=1}^{N} g_k(t+1) P_0^*(t) \\ P_k(t+1) = P_k^*(t) + g_k(t+1) P_0^*(t) - g_{\bar{k}}(t+1) P_k^*(t) \\ P_{\bar{k}}(t+1) = P_{\bar{k}}^*(t) + g_{\bar{k}}(t+1) P_k^*(t) \end{cases} \tag{10}$$

Equating (9) and (10), solving for $g_k(t+1)$ and $g_{\bar{k}}(t+1)$, and simplifying, we arrive at the following solutions for the hazard functions:

$$g_k(t+1) = \frac{\sum\limits_{\tau_2 > t+1} f_k(t+1, \tau_2)}{\sum\limits_{k=1}^{N} \sum\limits_{\tau_1 \ge t+1} \sum\limits_{\tau_2 > \tau_1} f_k(\tau_1, \tau_2)} \tag{11}$$

$$g_{\bar{k}}(t+1) = \frac{\sum\limits_{\tau_1 < t+1} f_k(\tau_1, t+1)}{\sum\limits_{\tau_1 < t+1} \sum\limits_{\tau_2 \ge t+1} f_k(\tau_1, \tau_2)} \tag{12}$$

where the final simplification yielding (12) is the result of an enforcement of a restriction on the priors:

$$\frac{f_k(1,t+1)}{\sum_{\tau_2 \geq t+1} f_k(1,\tau_2)} = \cdots = \frac{f_k(t,t+1)}{\sum_{\tau_2 \geq t+1} f_k(t,\tau_2)}$$

This enforcement is not necessary generally, and may be dispatched with as needed. Experimentally, it corresponds to the assumption of independence of onset time and duration for each target. Lifting the restriction significantly increases the complexity of the analytic form of (12).

Putting everything together, using (8) to eliminate $P_i^*(t)$ from (10), the iteration has the "canonical" form:

$$\begin{cases} P_0(t+1) = \dfrac{P_0(t)l_0(e_t)\left(1-\sum_{k=1}^{N} g_k(t+1)\right)}{P_0(t)l_0(e_t)+\sum_{k=1}^{N} P_k(t)l_k(e_t)+\sum_{k=1}^{N} P_{\bar{k}}(t)l_{\bar{k}}(e_t)} \\ P_k(t+1) = \dfrac{g_k(t+1)P_0(t)l_0(e_t)+P_k(t)l_k(e_t)\left(1-g_{\bar{k}}(t+1)\right)}{P_0(t)l_0(e_t)+\sum_{k=1}^{N} P_k(t)l_k(e_t)+\sum_{k=1}^{N} P_{\bar{k}}(t)l_{\bar{k}}(e_t)} \\ P_{\bar{k}}(t+1) = \dfrac{P_{\bar{k}}(t)l_{\bar{k}}(e_t)+g_{\bar{k}}(t+1)P_k(t)l_k(e_t)}{P_0(t)l_0(e_t)+\sum_{k=1}^{N} P_k(t)l_k(e_t)+\sum_{k=1}^{N} P_{\bar{k}}(t)l_{\bar{k}}(e_t)} \end{cases} \tag{13}$$

These equations derived here are generalizations of the standard Bayes' formulation to explicitly include the prior knowledge embodied in $g_k(t)$ and $g_{\bar{k}}(t)$. It is seen that the effects of evidence and prior knowledge on posterior probability is interactive; it is not possible to simply redefine the likelihood function to reduce (13) to the usual Bayes' formula. Note the prior probability does not simply enter as a bias term, contra Ashby (1983).

Analysis and Discussion of the Model

Interpretation of Threshold for Stopping Rule At the beginning of each trial ($t = 0$), the initial values of the behavior-level hypotheses are set so that $P_0(t=0) = 1$, while the remaining $P_i(t=0) = 0$. The model then begins the cycle of evidence accumulation and probability update as described above. If we consider the vector $\mathbf{P} = (P_0, P_1 + P_{\bar{1}}, \ldots, P_N + P_{\bar{N}})$, we can view this process as a trajectory $\mathbf{P}_t$ within an (N+1)-dimensional probability simplex, with starting point $(1,0,\ldots,0)$. In the present model, whenever one of the target-positive terms of $\mathbf{P}_t$ of the form $P_k + P_{\bar{k}}$ exceeds a predetermined threshold θ_k, the corresponding response is made. Differences among the elements of the threshold vector θ reflect response biases determined by the Bayesian decision-maker. Such could naturally arise, for instance, in the case of asymmetric payoffs among responses.

It is worth noting that, in determining the stopping time (and thus the response), the posterior probabilities corresponding to the pairs of behavior-level hypotheses of the form $\{H_k, H_{\bar{k}}\}$ are pooled, instead of competing. This corresponds to the task demand of responding to the target location once the target is on and choosing the k-th alternative whether or not the target is subsequently turned off. With a small enough τ_2, this models a transient stimulus that is briefly turned on and then off; with large enough τ_2, this models a steady stimulus that is turned on and stays on. Our model incorporates both the *transient channels* (through H_k's) and *sustained channels* (through $H_{\bar{k}}$'s), while controlling the mixture of transient and sustained components through manipulating stimulus duration $\tau_2 - \tau_1$; this has been empirically and theoretically investigated in behavioral contexts (cf. Smith, 1995) and in neural context (Cleland, Dubin, & Levick, 1971).

As mentioned earlier, the threshold θ_k has a special interpretation. Because the sum $P_k(t) + P_{\bar{k}}(t)$ is exactly the posterior probability conditioned over all evidence in a trial that the target has already appeared at location k (but may have been turned off) by time t, this means that when θ_k is reached, it is precisely this same probability. Thus, if we count over many trials in which the k-th response was made, it can be shown that θ_k emerges as the percentage of correct trials (conditioned on the k-th response). Let a_k denote the probability with which the experimenter chooses to present the k-th target each trial, s_k denote the proportion of correct trials among all trials in which target k is presented, and r_k denote the proportion of trials in which the participant responds with the k-th alternative. Then it follows that:

$$a_k s_k = r_k \theta_k\ . \tag{14}$$

Summing over k in (14), we get:

$$\%\text{correct} = \sum_k a_k s_k = \sum_k r_k \theta_k\ . \tag{15}$$

Evidence Representation The model framework is intentionally general and non-committal regarding the mode of evidence representation and the sensor unit response. Ideally, any specific representational commitment should be made in light of the nature of the particular task, agent, and stimuli one wishes to model. However, for illustration purposes, we offer an example of minimal implementation. Consider the restricted case $N = 2$, a two-alternative forced choice task. The evidence generation can be accomplished by a single sensor unit, which may be either firing (F) or resting (R) at any time t. Thus, the evidence space is $\mathcal{E} = \{\text{F}, \text{R}\}$, while the environment may be in any one of three states: no target present, target 1 present, target 2 present. Suppose the sensor is slightly tuned to respond to target 1, slightly inhibited in response by target 2, and equally likely to be firing or resting when there is no target. Thus, its likelihood functions may take the following form:

$$\begin{aligned} l_0(e) = l_{\bar{1}}(e) = l_{\bar{2}}(e) &: \begin{cases} \Pr[\text{F}|0] = 0.5 \\ \Pr[\text{R}|0] = 0.5 \end{cases} \\ l_1(e) &: \begin{cases} \Pr[\text{F}|1] = 0.5 + \varepsilon \\ \Pr[\text{R}|1] = 0.5 - \varepsilon \end{cases} \\ l_2(e) &: \begin{cases} \Pr[\text{F}|2] = 0.5 - \varepsilon \\ \Pr[\text{R}|2] = 0.5 + \varepsilon \end{cases} \end{aligned} \tag{16}$$

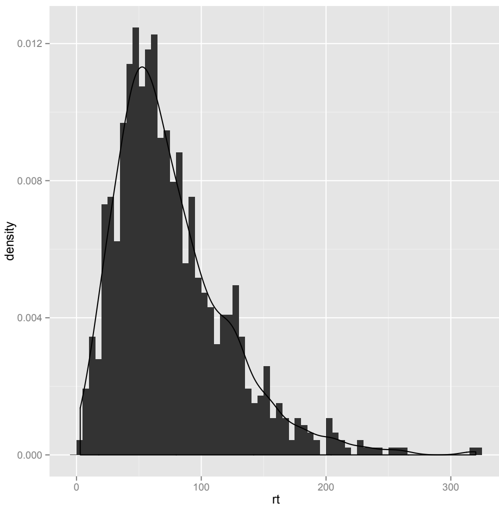

Figure 2: Characteristic shape of RT distributions for correct responses when stimulus presentation probabilities are equal on a uniform temporal prior. Note the heavy tail in the RT density function.

where ε is a small constant representing the sensor's tuning. It is easy to extend such a representation to multiple sensors, complex activation functions, interactions between units, etc.

Simulation Results We simulated a variant of the two-choice model presented above, and were able to reproduce many empirical data patterns of choice RT experiments, such as the prediction of (3) that median RT is dependent via a log-linear transformation upon prior stimulus probability (Carpenter & Williams, 1995), the speed-accuracy tradeoff, and the dependency of both correct and error RT distributions on stimulus probability (Link, 1975; Green, Smith, and von Gierke, 1983; see Figure 2). We also were able to numerically simulate and confirm the special Bayesian significance of the threshold θ_k parameter, as described above.

Acknowledgments

Supported by ARO grant W911NF-12-1-0163 and AFOSR grant FA9550-13-1-0025 (PI: Jun Zhang).

References

Ashby, F. G. (1983). A biased random walk model for two choice reaction times. *Journal of Mathematical Psychology*, *27*(3), 277–297.

Baum, C. W. & Veeravalli, V. V. (1994). A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory*, *40*(6), 1994–2007.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, *113*(4), 700–765.

Busemeyer, J. R. & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, *100*(3), 432–459.

Carpenter, R. H. S. & Williams, M. L. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, *377*(6544), 59–62.

Cleland, B. G., Dubin, M. W., & Levick, W. R. (1971). Sustained and transient neurones in the cat's retina and lateral geniculate nucleus. *The Journal of Physiology*, *217*(2), 473–496.

Crossman, E. R. F. W. (1953). Entropy and choice time: the effect of frequency unbalance on choice-responses. *Quarterly Journal of Experimental Psychology*, *5*, 41–52.

Draglia, V. P., Tartakovsky, A. G., & Veeravalli, V. V. (1999). Multihypothesis sequential probability ratio tests. i. asymptotic optimality. *IEEE Transactions on Information Theory*, *45*(7), 2448–2461.

Green, D. M., Smith, A. F., & von Gierke, S. M. (1983). Choice reaction time with a random foreperiod. *Perception & Psychophysics*, *34*(3), 195–208.

Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, *4*(1), 11–26.

Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of experimental psychology*, *45*(3), 188–196.

Laming, D. R. J. (1968). *Information theory of choice-reaction times*. New York: Wiley.

Link, S. W. (1975). The relative judgment theory of two choice response time. *Journal of Mathematical Psychology*, *12*(1), 114–135.

Luce, R. D. (1986). *Response times: their role in inferring elementary mental organization*. New York: Oxford University Press.

Ratcliff, R. & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological review*, *111*(2), 333–367.

Rouder, J. N. (1996). Premature sampling in random walks. *Journal of Mathematical Psychology*, *40*(4), 287–296.

Smith, P. L. (1995). Psychophysically principled models of visual simple reaction time. *Psychological Review*, *102*(3), 567–593.

Stein, W. E. & Rapoport, A. (1978). A discrete time model for detection of randomly presented stimuli. *Journal of Mathematical Psychology*, *17*(2), 110–137.

Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, *25*(3), 251–260.

Swensson, R. G. & Green, D. M. (1977). On the relations between random walk models for two-choice response times. *Journal of Mathematical Psychology*, *15*(3), 282–291.

Thomas, E. A. C. (1975). A note on the sequential probability ratio test. *Psychometrika*, *40*(1), 107–111.

Usher, M. & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological review*, *108*(3), 550–592.

Wald, A. (1947). *Sequential analysis*. New York: Wiley.