

Using behavior to decode allocation of attention in context dependent decision making

Michael Shvartsman (ms44@princeton.edu)

Princeton Neuroscience Institute, Washington Rd., Princeton, NJ 08544 USA

Vaibhav Srivastava (vaibhavs@princeton.edu)

Department of Mechanical and Aerospace Engineering, 41 Olden St., Princeton, NJ 08544 USA

Narayanan Sundaram (narayanan.sundaram@intel.com)

Intel Parallel Computing Lab, 2200 Mission College Blvd., Santa Clara, CA 95054 USA

Jonathan D. Cohen (jdc@princeton.edu)

Princeton Neuroscience Institute, Washington Rd., Princeton, NJ 08544 USA

Abstract

We present a model of the dynamics of adaptive attention allocation in the AX Continuous Performance Test (AX-CPT), a simple context dependent decision making task of interest to the research communities concerned with cognitive control, schizophrenia, anxiety and aging (Braver et al., 2001; Cohen et al., 1999; Eysenck et al., 2007). We ground it in our recent theory of decision making under dynamic context, that assumes humans use sequential Bayesian inference to combine information from multiple sources in perception and (optionally) memory over time. The theory generalizes the well-known diffusion decision model of single-stimulus decision making (DDM; Ratcliff, 1978). Our first result is a new analysis that shows how memory encoding and retention can yield a variable initial condition for either a single- or multiple-stimulus decision, providing a theoretical grounding to the assumption of variability in initial condition previously shown to improve data fits for the DDM. Our second result is using this model to decode attention allocation from behavioral data in a novel quantitative payoff manipulation in the AX-CPT, showing that our model can capture the differences in how subjects encode and retrieve contextual information when the relative emphasis on task speed and accuracy is changed.

Keywords: decision making; context processing; sequential inference; Bayesian inference; wiener process; ornstein-uhlenbeck process;

Introduction

If there is one thing that defines human behavior, it is its pervasive adaptive nature. Such a viewpoint is perhaps best underpinned by the seminal work of Swets and colleagues on signal detection (e.g. Tanner and Swets, 1954), who showed that even the simple task of detecting a flash of light in a uniform background is amenable to strategic variation that trades off correct responses against false alarms in a way sensitive to reward and the statistics of the environment. The Signal Detection Theory model of this behavior assumed observers performed the equivalent of a fixed-sample likelihood ratio test when asked to respond, and was soon generalized to the case of variable sample size (Edwards, 1965; Laming, 1968; Stone, 1960), and to continuous time as the well-known Diffusion Decision Model (DDM; Ratcliff, 1978).

These models naturally capture a second tradeoff in decision making: that of speed against accuracy, and find support from a wide range of behavioral and neural data (e.g. Bogacz

et al., 2006; Bogacz and Gurney, 2007; Gold and Shadlen, 2007; Kira et al., 2015; Ratcliff et al., 2004; Turner et al., 2015; van Vugt et al., 2012).

We consider the above models to be *fixed context* models because they treat context, which we operationalize as additional task-relevant information, as fixed/known over the course of the decision. Our theory (Shvartsman et al., 2015) removes the assumption that the context is known, and considers the dynamics of processing the context simultaneously with an additional stimulus we term the decision *target*. The theory can be stated both as a sequential Bayesian inference model, and (in the continuum limit) as a nonlinear diffusion model, and is therefore a formal generalization of the models discussed above, as well as a previous Bayesian model of the Flanker task by Yu et al. (2009). In prior work, we used the same parameters from Yu and colleagues' paper to generate plausible behavior patterns in the AX Continuous Performance Task (AX-CPT), a task that differs from the Flanker task both in the time of stimulus presentation, and in response rules. The AX-CPT, which we describe in detail below, is arguably the simplest task where participants are required to combine information seen during different disjoint time intervals to make a decision, and is our focus in this paper.

When deciders include additional contextual information in their decisions, a third tradeoff naturally arises, that of allocating processing (attentional or otherwise) between the available sources of information. In this paper we explicitly consider this tradeoff. In addition to being a formal generalization to the work on fixed-context decision making discussed above, it also bears some relation to work on attention allocation under multi-sensory integration (e.g. Sheppard et al., 2013), which it extends considering stimuli in memory. It is also complementary to work on the perception-memory tradeoff on longer timescales in the prospective memory paradigm (Einstein and McDaniel, 2005), and mechanistic work combining multiple sequential samplers in the ACT-R cognitive architecture (van Maanen et al., 2012).

Our main contributions are as follows: first, we provide additional insight into our theory of decision making under dynamic context presented in Shvartsman et al. (2015)

by providing analytical expressions for the posterior distribution over context after memory encoding and retention, and thereby provide a new motivation for the assumption of initial condition variability previously argued to be necessary to fit human data in the fixed-context DDM (e.g. Ratcliff and Rouder, 1998). Second, we apply the model to a dataset of humans performing the AX-CPT under a novel quantitative payoff manipulation, and use the model to develop a quantitative understanding of how participant strategy changes with task goals.

An overview of the theory and application to the AX-CPT

We assume that decision making can be understood as a sequential Bayesian inference process, specifically that the agent uses sequentially-drawn samples from external input and/or memory to compute the joint posterior probability over the identity of some true *context* and decision *target* over time. The agent maps from this joint posterior to a response probability using a known response rule, and uses a fixed threshold defined over the response probability to stop sampling and respond. We make a crucial distinction between our *theory* of decision making and individual *task models* that can be derived from the theory by applying it to specific tasks, as we do here. Our previous work has shown how our formalism can be used to derive different models when the additional context stimulus is either in perception or memory, under different response mappings (Shvartsman et al., 2015).

In this paper we focus on the AX Continuous Performance test (Servan-Schreiber et al., 1998), arguably the simplest decision making task that conditions responses jointly and uniquely on a perceptual and memory stimulus. The task is illustrated in Fig. 1: participants see one of two context stimuli (by convention labeled ‘A’ or ‘B’) followed by one of two targets (‘X’ or ‘Y’), and make one response (e.g. ‘left’) to AX and BY pairs, and the other (e.g. ‘right’) to AY and BX pairs.

Formally, we assume the agent conditions a decision based on her posterior belief over the identity of some unknown true *context* and some true *target*. We denote by C, G random variables representing the possible draws of context and target, and $r(\cdot)$ a function from the distribution $P(C, G)$ to a distribution over responses, which for the AX-CPT is the exclusive-or function:

$$r(P(C, G)) = \begin{cases} r_0, & \text{with } P(C = 'A', G = 'X') + \\ & P(C = 'B', G = 'Y') \\ r_1, & \text{with } P(C = 'A', G = 'Y') + \\ & P(C = 'B', G = 'X'). \end{cases} \quad (1)$$

The agent receives evidence samples s^C and s^G drawn i.i.d. from the environment, and updates her posterior distribution over the pair (C, G) using Bayes’ rule. We denote by t_c^{on} the time at which the context appears and t_c^{off} the time at which it disappears, and likewise $t_g^{\text{on}}, t_g^{\text{off}}$ for the target. This implies

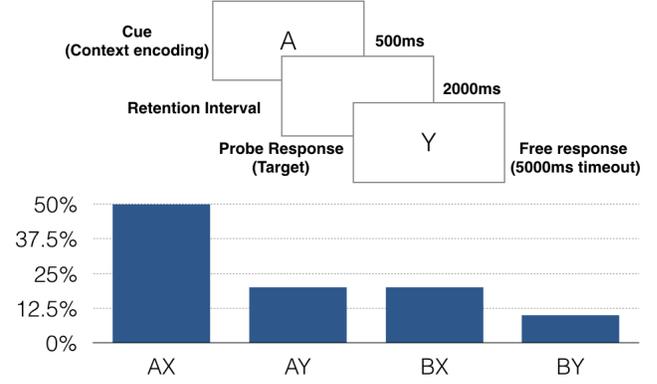


Figure 1: *Top*: sequence diagram for one trial of the AX Continuous Performance Test. *Bottom*: frequency of each trial type in the experiment. Actual letters seen were randomized for each subject.

the following update at timestep τ , assuming independence of the context and target evidence streams:

$$P_\tau(C, G | s^C, s^G) \propto P(s^C | C)P(s^G | G)P_{\tau-1}(C, G). \quad (2)$$

Before the target appears, the target likelihood is uniform over all the targets. Because there are only two responses (i.e. hypotheses), we can rewrite this update as a likelihood ratio update, and convert it into a nonlinear transformation of two diffusion processes. We refer the reader to the supplement of (Shvartsman et al., 2015) for the full derivation, and here only give the final expression:

$$\log Z = \log \frac{P_0(C = c_0, G = g_0)e^{z_c^t} e^{z_g^t} + P_0(C = c_1, G = g_1)}{P_0(C = c_0, G = g_1)e^{z_c^t} + P_0(C = c_1, G = g_0)e^{z_g^t}}. \quad (3)$$

Here, the target particle z_g^t is stationary until the target appears (t_g^{on}) and then evolves according to a Wiener process with drift:

$$dz_g^t = a_g dt + \sigma_g dW. \quad (4)$$

The context particle z_c^t evolves according to a Wiener process with drift from when it appears (t_c^{on}) until it disappears. Once the context disappears from perception, the memory system can provide continued samples \hat{s}^C after the stimulus goes away, but with some constant probability d at each time step it can start to provide uninformative noisy samples s^U . Assuming the agent has a good estimate of d , she also knows that the probability of receiving an informative sample exponentially decays with time, such that the time-varying likelihood distribution is the following mixture:

$$f(s^C) = (1 - \exp(-\lambda t))f(\hat{s}^C) + \exp(-\lambda t)f(s^U) \quad (5)$$

, where $f(\cdot)$ refers to the density of its argument. That is, the informative component of the likelihood decays – as does its

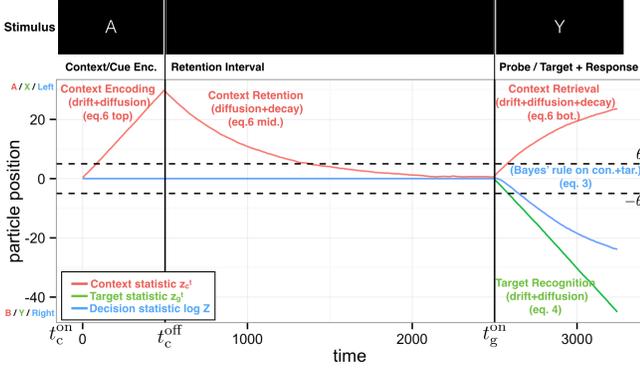


Figure 2: Average dynamics of the CDDM model for the AX-CPT. The ordinate does triple-duty in denoting the position of all three particles in their own spaces. In this figure, context almost (but not quite) decayed fully during the retention interval.

variance, if we make the standard assumption that evidence distributions are Gaussian. We model this retention and forgetting process in continuous time as a zero-mean Ornstein-Uhlenbeck (O-U) process, which has the same exponential decay property, though we leave the derivation of an explicit connection to future work. Finally, we model the retrieval process by switching the O-U process to a nonzero-mean process at t_g^{on} . Our O-U assumption could be alternatively motivated by assuming the memory system is implemented using a leaky competing accumulator model with large leak and mutual inhibition (Bogacz et al., 2006). The full expression for the context particle dynamics is as follows:

$$dz_c^t = \begin{cases} a_e dt + \sigma_c dW & \text{if } t_c^{\text{on}} \leq t \leq t_c^{\text{off}} \\ -\lambda z_c^t + \sigma_c dW & \text{if } t_c^{\text{off}} \leq t \leq t_g^{\text{on}} \\ -\lambda z_c^t + a_c dt + \sigma_c dW & \text{if } t_g^{\text{on}} \leq t \end{cases} \quad (6)$$

We assume by convention that $\sigma_g = \sigma_c = 1$ in order to make the model identifiable (but cf. Bitzer and Kiebel, 2015, for possible consequences of this choice). We therefore omit the Wiener noise coefficients in the notation that follows. Fig. 2 shows the average particle dynamics for the model. The full set of model parameters and their interpretation is given in Tab. 1 (we discuss parameters not related to the decision process itself later in the paper).

We rely on a number of properties of the AX-CPT that enable the analysis that follows (some of which are shared by other tasks as well). First, the context presentation and retention interval durations are both independent of the subject's actions and usually set a priori. Second, responses are only allowed when the target is on the screen. Given those two properties, the context particle distribution at the time of target appearance can be written in closed form, as follows.

First, define $\Delta t_e \triangleq t_c^{\text{off}} - t_c^{\text{on}}$ and $\Delta t_r \triangleq t_g^{\text{on}} - t_c^{\text{off}}$. Next, let z_c^e denote the position of the context particle at the end of encoding (t_c^{off}), and z_c^r denote the position of the particle when the target appears and retrieval begins (t_g^{on}). We can write $\mathbb{E}[z_c^r]$,

explicitly, recalling that it evolves as an O-U process, and that therefore its distribution is $\mathcal{N}\left(z_c^e e^{-\lambda \Delta t_r}, \frac{1}{-2\lambda}(e^{-2\lambda \Delta t_r} - 1)\right)^1$. We take care to use the law of total expectation, since z_c^e is itself a random variable with distribution $\mathcal{N}(a_e \Delta t_e, \Delta t_e)$.

$$\mathbb{E}[z_c^r] = \mathbb{E}[\mathbb{E}[z_c^r | z_c^e]] = \mathbb{E}\left[z_c^e e^{-\lambda \Delta t_r}\right] = a_e \Delta t_e e^{-\lambda \Delta t_r}. \quad (7)$$

We do the same for the variance, using the law of total variance:

$$\text{Var}[z_c^r] = \mathbb{E}_{z_c^e}[\text{Var}[z_c^r | z_c^e]] + \text{Var}_{z_c^e}[\mathbb{E}[z_c^r | z_c^e]] \quad (8)$$

$$= \frac{1}{-2\lambda}(e^{-2\lambda \Delta t_r} - 1) + \Delta t_e e^{-\lambda \Delta t_r}. \quad (9)$$

In the case where context is not retrieved from memory once the target appears, the resultant model is a DDM with initial condition variability (Ratcliff and Rouder, 1998), previously used to fit fast errors in decision making. Our analysis here formally provides a psychological interpretation for such fast errors: specifically, it suggests that data that is better fit using sampling or diffusion models with a variable initial condition may reflect subjects' memory of relevant contextual information. Furthermore, it provides an argument on theoretical grounds for using Gaussian rather than the uniform initial condition variability (e.g. Ratcliff and McKoon, 2008), especially if subjects may be using contextual information. Such a model may predict instantaneous decisions ('ultra-fast guesses') for some parameter settings that reflect strong memory encodings, which may be smoothed out to a multi-modal response distribution if nondecision time variability is included in the model.

A particularly elegant property of the model is that each component has both distinct psychological interpretation, and distinct effects on response time distributions. θ reflects the speed-accuracy tradeoff and therefore affects the shape of the RT distributions and the error proportion across all trial types. The three context parameters affect different portions of the RT distribution: a_e reflects encoding of the context, and contributes primarily to the early portion of the RT distributions due to how it affects the initial condition for the decision; λ reflects memory retention and contributes both to early and late portion of RT distributions because it persists both during retention and retrieval; and a_c reflects memory retrieval and primarily contributes to middle and late portions of the RT distributions because retrieval begins at target onset and saturates over time. All three context parameters will identically affect trial types that share a context (e.g. AX and AY). On the other hand, a_g reflects perceptual processing of the target stimulus and therefore affects the RT distribution throughout, but identically affects same-target pairs (e.g. AX and BX).

Thus, while differences between different trial types help uncover differences between context and target processing writ large, differences between different portions of the RT

¹Note that we use the $\mathcal{N}(\mu, \sigma^2)$ parameterization.

distribution may index differences in subcomponents of context processing. For example, high a_e will increase the number of very fast correct responses when the context predicts the correct target (e.g. AX) but also contribute fast errors when it does not (e.g. AY). In this way, the model may provide a more precise characterization of context processing differences than the behavioral indices in primary use today (e.g. Braver, 2012).

Parameter	Interpretation	Prior range
a_e	Context encoding drift	0 to 5
a_c	Context retrieval drift	0 to 5
a_g	Target recognition drift	0 to 5
λ	Memory decay	-0.05 to 0.2
θ	Decision threshold	0 to 20
μ_0	Mean of non-decision time	0 to 1000ms
σ_0	Standard deviation of non-decision time	0 to 200ms
p_c	Probability of contaminant RT	0 to 1
$P_0(\cdot)$	Prior distribution over trial types	Set to true trial distribution

Table 1: Model parameters, interpretation, and prior ranges.

A payoff adaptation experiment

To test the ability of our model to precisely estimate attention allocation from behavioral data, we gathered a dataset of humans performing the AX-CPT task on Amazon’s Mechanical Turk. Mechanical Turk is a web-based marketplace for “human intelligence tasks”, short web-based tasks that pay small sums of money. It is emerging as a standard method for high-throughput collection of data for psychological experiments (Crump et al., 2013). Our experiment was designed using custom front-end code, using psiTurk (McDonnell et al., 2015) for back-end and interfacing with Amazon. Response times were collected with the JavaScript high resolution timer API, which theoretically promises at least millisecond-level accuracy. Experiment code is available at <https://github.com/mshvartsman/axcpt-psiturk-coffeescript>.

We searched the parameter space of a preliminary version of the model to devise four different payoff schemes intended to differently prioritize speed and accuracy, illustrated in Tab. 2. The first two were designed to encourage participants to be particularly fast or accurate relative to each other. The other two were designed to encourage intermediate behavior, with either high or low overall cost.

Each subject started the experiment by reading the instructions, and then practicing the two A and B contexts separately. To test that they have learned the task, they next had to complete 10 correct trials in a row. If they failed to do so after 50 trials have passed, they were ejected from the experiment. If they succeeded, they completed 240 trials of the AX-CPT divided into 24 blocks. They were told their speed, accuracy,

Cond.	Error cost	Time cost (/s)	Intent
(A)	10	2	Prefer accuracy
(B)	2	10	Prefer speed
(C)	10	10	Balanced (high cost)
(D)	2	2	Balanced (low cost)

Table 2: Quantitative payoffs given to the participants. Correct responses were always worth 20 points. Both correct and error responses are penalized for time.

and points gained after every trial, and a running total every block. They earned \$1 for participating, plus \$1 for each 1000 points they earned in the experiment. We collected 20 subjects per condition, and most earned between \$3.50 and \$4.50.

Each participant’s actual letters were randomized, but in the remainder of the paper we label them as A, B, X, and Y according to their frequencies: each set of 240 trials contained 50% (120) AX trials, 20% (48) AY and BX trials, and 10% (12) BY trials, randomly distributed throughout the experiment. This is conventional in AX-CPT designs, and is what allows the AX-CPT to behaviorally index the allocation of attention. The reasoning is as follows: AY and BX trials have the same joint probability, but different probabilities conditioned on either knowing the true context (A or B) or target (X or Y). Therefore, any asymmetry between these two trials is argued to reflect differences in processing the two stimuli: specifically, better AY performance suggests more attention to the context, and better BX performance suggests more attention to the target (but cf. Lositsky et al. 2015 for evidence that this index may be overly simplistic).

Fig. 3 shows basic behavior summaries in the speed and accuracy conditions: the manipulation was successful in encouraging subjects to adjust their speed and accuracy, as evidenced by the mean behavior. Of more interest is that the effect was not uniform across the four trial types. When pressured to respond more quickly while sacrificing accuracy, subjects sacrifice accuracy in B trials more than in A trials, with the speedup primarily seen in BY, the rarest trial type. The two balanced conditions (not shown in the figure) for the most part showed intermediate behavior, with the high-cost condition patterning closer to the speed-encouraging condition, and the low-cost balanced condition patterning with the accuracy-encouraging condition. We focus the remainder of the analysis on the two extreme conditions, where the patterns are clearest.

Model fitting details

While first passage times for diffusion and O-U processes are known, the nonlinear transformation that forms our model’s decision variable makes these analytics inapplicable. Therefore, we derive best fits by simulation, using the EP-ABC algorithm (Barthelmé and Chopin, 2014, as implemented at <https://github.com/sbitzer/pyEPABC>). It bears men-

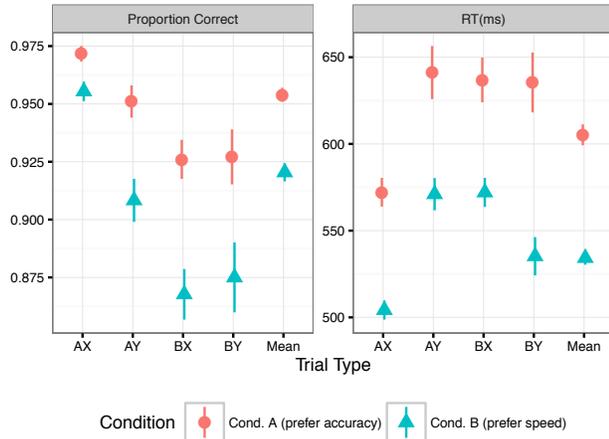


Figure 3: Mean accuracies and response times for the subjects across the four trial types and our two conditions of focus, showing slower speed and higher accuracy overall in the accuracy-focused condition, but also that these changes do not affect all trial types equally. Refer to Tab. 2 for payoff conditions. Error bars are standard errors.

tioning that our use of Bayesian methods for fitting is entirely orthogonal to our theoretical assertion that decision making in humans proceeds by Bayesian inference.

The original EP algorithm (Minka, 2001) approximates the joint posterior density of model parameters by a multivariate Gaussian, converting the Bayesian inference problem to an iterative optimization problem. Approximate Bayesian Computation (ABC) methods contend with performing Bayesian inference when simulation is possible but no likelihood is available, usually by assuming that parameters under which the simulator generates data ‘similar enough’ to the real data have high likelihood. EP-ABC as we apply it here assumes individual RTs are drawn i.i.d., which allows us to define ‘similar enough’ to mean ‘generates each individual data point to within 1ms’. We used EP-ABC with 5 passes over the dataset and a minimum of 3000 accepted samples per human data point, with EP hyper-parameter α set to 0.3.

The model as described above uses 5 parameters to describe the decision process (three drifts, one decay, one threshold). To describe actual response times, we add two additional assumptions. First, that there is a normally distributed offset to the decision time – a ‘non-decision time’ reflecting early perceptual processing, motor planning, and the motor response itself. This adds two parameters (mean and standard deviation of the non-decision time). Second, we assume that with some probability RTs are generated not from our model, but from a $\mathcal{U}(0, 5s)$ distribution of ‘contaminant’ RTs. We estimate the proportion of contaminant RTs from the data, adding another parameter. These assumptions are standard in parameter fitting for the fixed-context DDM (e.g. Ratcliff and Rouder, 1998) and provide a probability floor that makes it easier for EP-ABC to handle some unusually fast or slow RTs in our dataset. We used proper uniform priors over

plausible parameter ranges for all the parameters. Tab. 1 lists all of the parameters and their prior ranges.

To understand condition-level differences, we generated 10000 samples from the multivariate posterior density for each subject, and produced average parameter estimates for each condition based on these samples. This heuristic weights subjects’ posterior means in proportion to their covariances better than simple averaging of the means as point estimates, and without the ideal fully-hierarchical treatment that is extremely challenging in our setting. Because there is stochasticity in both the simulator and the fitting method, we repeated the model-based analysis twice. While there was some variability in the parameters estimated for each subject, the signs of the difference between parameters for the speed and accuracy condition was the same across both runs for all parameters, so we focus our interpretation on those differences. Model code is available at <https://github.com/mshvartsman/cddm>.

Results

Tab. 3 shows the combined parameter estimates. We remark on a number of properties: first, the contaminant proportion is low, and similar across the fits. Second, the threshold is higher in the accuracy-focused than in the speed-focused condition, consistent with the idea that it governs the speed-accuracy tradeoff. Non-decision times are shorter and less variable in the accuracy-focused than in the speed-focused condition – perhaps an indicator of greater focus overall. This point is also supported by the fact that the accuracy condition shows a higher total drift (summed across the three drift variables) than the speed condition, and by the fact that it shows slightly less memory decay.

Parameter	Accuracy condition	Speed condition
a_e	3.1701	2.4623
a_c	1.0576	0.8088
a_g	0.7648	1.1681
λ	0.0756	0.0834
θ	10.9656	6.1589
μ_{t_0}	406.2522	435.1687
σ_{t_0}	76.7926	94.8328
p_c	0.0513	0.0530

Table 3: Fit parameter values in aggregate across the two subject groups. See text for fit details.

Most interesting, however, is how this drift is allocated. In the accuracy-focused group, both context-involving drifts (encoding a_e and retrieval a_c) are higher than the corresponding drifts in the speed-focused group, and the pattern is reversed in the target drift a_g . That is, when incentivized to be more accurate, participants rely more on their memory and contextual information and slightly less on the perceptual stimulus in front of them. Since the task is designed such that context and target information is equally useful in being

able to make the correct response, we suspect that this has to do with the effort involved in encoding, maintaining, and reinstating the memory of the contextual rule – something that is probably worth doing more of when accuracy is more important.

Discussion and Conclusions

In this work, we applied our theory of decision making under dynamic context to the AX Continuous Performance Test. First, we showed how the simple memory encoding and retention dynamics of the model map to variability in the initial condition of a diffusion decision process, providing both further theoretical grounding for the use of variable initial conditions in data fits, and a better understanding for the psychological origin of fast errors in decision making.

We next applied our model to estimate allocation of attention between perception and memory in a novel quantitative payoff manipulation in the task as measured on Amazon’s mechanical turk. This manipulation succeeded in not only changing participants’ speed-accuracy tradeoff, but also their attention allocation tradeoff between perception and memory. The ability to manipulate attention allocation continuously rather than using discrete task dimensions (e.g. deadlines, distractors) as used previously may pave the way to more quantitative mapping out of the attention allocation strategy space.

Finally, our method of measuring the relative allocation of attention between the cue and probe stimulus in AX-CPT is among the first model-based efforts to understand this tradeoff, which has previously been measured as a behavioral index (the difference between AY and BX performance; e.g. Braver, 2012). This behavioral index, like our model, indicates that subjects in our accuracy-preferring group focus more on context, as indexed by a reduction in BX errors. However, the behavioral index has no way of making the distinction between encoding-oriented and retrieval-oriented contextual effects, a distinction that our model captures. Such multidimensional, quantitative characterization of strategic variability in the processing of perceptual and memory stimuli may pave the way to a richer understanding of context processing both in normal and in diseased populations.

References

Barthelmé, S. and Chopin, N. (2014). Expectation Propagation for Likelihood-Free Inference. *Journal of the American Statistical Association*, 109(505):315–333.

Bitzer, S. and Kiebel, S. J. (2015). The Brain Uses Reliability of Stimulus Information when Making Perceptual Decisions. In *Advances in Neural Information Processing Systems 28*, pages 1–9.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and 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):700–65.

Bogacz, R. and Gurney, K. N. (2007). The basal ganglia and

cortex implement optimal decision making between alternative actions. *Neural Computation*, 19(2):442–77.

Braver, T. S. (2012). The variable nature of cognitive control: a dual mechanisms framework. *Trends in Cognitive Sciences*, 16(2):106–13.

Braver, T. S., Barch, D. M., Keys, B. a., Carter, C. S., Cohen, J. D., Kaye, J. A., Janowsky, J. S., Taylor, S. F., Yesavage, J. a., Mumenthaler, M. S., Jagust, W. J., and Reed, B. R. (2001). Context processing in older adults: Evidence for a theory relating cognitive control to neurobiology in healthy aging. *Journal of Experimental Psychology: General*, 130(4):746–763.

Cohen, J. D., Barch, D. M., Carter, C., and Servan-Schreiber, D. (1999). Context-processing deficits in schizophrenia: Converging evidence from three theoretically motivated cognitive tasks. *Journal of Abnormal Psychology*, 108(1):120–133.

Crump, M. J. C., McDonnell, J. V., and Gureckis, T. M. (2013). Evaluating Amazon’s Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE*, 8(3):e57410.

Edwards, W. (1965). Optimal Strategies for Seeking information : Models for Statistics , Choice Reaction Times , and Human Information Processing. *Journal of Mathematical Psychology*, 329:312–329.

Einstein, G. O. and McDaniel, M. a. (2005). Prospective Memory. Multiple Retrieval Processes. *Current Directions in Psychological Science*, 14(6):286–290.

Eysenck, M. W., Derakshan, N., Santos, R., and Calvo, M. G. (2007). Anxiety and cognitive performance: Attentional control theory. *Emotion*, 7(2):336–353.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–74.

Kira, S., Yang, T., and Shadlen, M. N. (2015). A Neural Implementation of Wald’s Sequential Probability Ratio Test. *Neuron*, 85(4):861–873.

Laming, D. R. J. (1968). *Information theory of choice-reaction times*. Academic Press, London.

Lositsky, O., Wilson, R. C., Shvartsman, M., and Cohen, J. D. (2015). A Drift Diffusion Model of Proactive and Reactive Control in a Context-Dependent Two-Alternative Forced Choice Task. In *The Multidisciplinary Conference on Reinforcement Learning and Decision Making*, pages 103–107.

McDonnell, J. V., Martin, J. B., Markant, D. B., Coenen, A., Rich, A. S., and Gureckis, T. M. (2015). psiTurk.

Minka, T. P. (2001). Expectation Propagation for Approximate Bayesian Inference. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, UAI’01, pages 362–369, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.

Ratcliff, R., Gomez, P., and McKoon, G. (2004). A diffusion

- model account of the lexical decision task. *Psychological Review*, 111(1):159–82.
- Ratcliff, R. and McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, 20(4):873–922.
- Ratcliff, R. and Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, 9(5):347–356.
- Servan-Schreiber, D., Bruno, R. M., Carter, C. S., and Cohen, J. D. (1998). Dopamine and the mechanisms of cognition: Part I. A neural network model predicting dopamine effects on selective attention. *Biological Psychiatry*, 43(10):713–722.
- Sheppard, J. P., Raposo, D., and Churchland, A. K. (2013). Dynamic weighting of multisensory stimuli shapes decision-making in rats and humans. *Journal of Vision*, 13(6):1–19.
- Shvartsman, M., Srivastava, V., and Cohen, J. D. (2015). A Theory of Decision Making Under Dynamic Context. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R., editors, *Advances in Neural Information Processing Systems 28*, pages 2485–2493. Curran Associates, Inc.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3):251–260.
- Tanner, W. P. J. and Swets, J. A. (1954). A decision-making theory of visual detection. *Psychological Review*, 61(6):401–409.
- Turner, B. M., van Maanen, L., and Forstmann, B. U. (2015). Informing cognitive abstractions through neuroimaging: The neural drift diffusion model. *Psychological Review*, 122(2):312–336.
- van Maanen, L., van Rijn, H., and Taatgen, N. (2012). RACE/A: An Architectural Account of the Interactions Between Learning, Task Control, and Retrieval Dynamics. *Cognitive Science*, 36(1):62–101.
- van Vugt, M. K., Simen, P., Nystrom, L. E., Holmes, P., and Cohen, J. D. (2012). EEG oscillations reveal neural correlates of evidence accumulation. *Frontiers in Neuroscience*, 6(July):106.
- Yu, A. J., Dayan, P., and Cohen, J. D. (2009). Dynamics of attentional selection under conflict: toward a rational Bayesian account. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3):700–17.