



Alan Veliz-Cuba[†]
Zachary P. Kilpatrick[‡]
Krešimir Josić[§]

.O a a a c ca acc a d c a d c . C a c -
- a d ca d a c c c c c fi d
d ac a . H , a c a c a d . W a f f c d c
a c a a ca ca ab d c a ? T a d d
c c c a . O a a a d a b d c
d g a a a d b a , a d a c
d c acc a d b a a a a a
a c . A a b a a a b a c a
a , c a d , a a

evidence to make decisions. However, information about the state of the world is typically incomplete, and perception is noisy. Therefore, animals make choices based on uncertain evidence. The case of an observer deciding between two alternatives based on a series of noisy measurements has been studied extensively when the environment is static [8, 23, 32, 45]. In this case, humans [36] and other mammals [11, 24] can accumulate incoming evidence near optimally to reach a decision.

Stochastic accumulator models provide a plausible neural implementation of decision making between two or more alternatives [4, 43]. These models are analytically tractable [8], and can implement optimal decision strategies [10]. Remarkably, there is also a parallel between these models and experimentally observed neural activity. Recordings in animals during a decision task suggest that neural activity reflects the weight of evidence for one of the choices [24].

A key assumption in many models is that the correct choice is fixed in time; i.e., decisions are made in a static environment. This assumption may hold in the laboratory, but natural environments are seldom static [17, 34]. Recent experimental evidence suggests that human observers integrate noisy measurements near optimally even when the state of the environment changes. For instance, when observers need to decide between two options and the corresponding reward changes in a history-dependent manner, human behavior approximates that of a Bayes optimal observer [5]. An important feature of evidence accumulation in volatile environments is an increase in learning rate when recent observations do not support a current estimate [31]. Both behavioral and fMRI data show that human subjects employ this strategy when they must predict the position of a stochastically moving target [29]. Experimental work thus suggests that humans adjust evidence valuation to account for environmental variability.

However, the dynamics of decision making in changing environments has not been fully investigated. To address this question we extend optimal stochastic accumulator models to a changing environment. These extensions are amenable to analysis and reveal that an optimal observer discounts old information at a rate adapted to the frequency of environmental changes. As a result, the certainty that can be attained about any of the choices is limited. Our approach frames the decision-making process in terms of a first passage problem for a doubly stochastic nonlinear model that can be examined using techniques of nonlinear dynamics. Extending previous work, we also identify accurate piecewise linear approximations to the nonlinear model. This model also suggests a biophysical neural implementation for evidence integrators consisting of neural populations whose activity represents the evidence in favor of a particular choice. When the environment is not static, optimal evidence discounting can be performed exactly by populations coupled through excitation. We also show that the computation can be well approximated by appropriately tuned classical linear population models [10, 30, 41, 44].

2. We develop our model in a way that parallels the case of a static environment with two possible states. We therefore start with the derivation of the recursive equation for the log-likelihood ratio of the two states, and the approximating stochastic differential equation (SDE), when the underlying state is fixed in time.

To make a decision, an optimal observer integrates a stream of measurements to infer the present environmental state. In the static case, this can be done using sequential analysis [12, 45]: An observer makes a stream of independent, noisy measurements, $y_{1:n} = (y_1, y_2, \dots, y_n)$ at equally spaced times, $\mathbf{v}_{1:n} = (v_1, v_2, \dots, v_n)$. The

A


B

C

$$y = \ln(L_+/L_-)$$

time







Appendix D)

$$(4.1) \quad \pi_{n,i} = \frac{\Pr(\mathbf{1}_{1:n-1})}{\Pr(\mathbf{1}_{1:n})} \mathbf{1}_{i(n)} = \prod_{j \neq i} \frac{\xi_{ji}^{n-1,i} + \xi_{ij}^{n-1,j}}{\xi_{ji}^{n-1,i} + \xi_{ij}^{n-1,j}}$$

Again after taking logarithms, $\ln \pi_{n,i} = \ln \frac{\Pr(\mathbf{1}_{1:n-1})}{\Pr(\mathbf{1}_{1:n})} \mathbf{1}_{i(n)}$, we can approximate the discrete stochastic process in (4.1) with an SDE:

$$(4.2) \quad dx = g(x)dt + \sigma(x)dW_t + \gamma(x)dx$$

where the drift has components $\gamma_i(x) = \frac{1}{\Delta t} E_{\xi} [\ln \pi_{n,i}]$

recorded neural activity and responses of monkeys performing two-alternative forced-choice decision tasks, where single trial stimuli have stationary statistics [24]. Even when reward rates are varied across trials, animals can adjust their behavior near optimally from trial to trial in ways that are well captured by mutually inhibitory models [20]. Interestingly, these networks also provide a plausible model of decision making in house-hunting honeybee swarms [33]. In previous studies, it has been shown that a single fixed point can be stabilized in linear population models, as long as the strength of mutual inhibition is weaker than the leak of individual populations [8, 10, 43]. As we will show, a complementary approach in linear population models is to consider a mutually excitatory network, with arbitrary leak in individual populations. As with the linear approximations discussed above, such models perform suboptimal inference in changing environments, but can approach the performance of the ideal nonlinear discounting process given by (3.3).

Optimal inference in dynamic environments with two states, $+$ and $-$, can be performed by mutually excitatory nonlinear neural populations with activities (firing rates) d_+ and d_- evolving according to

$$(6.1a) \quad \dot{d}_+ = -\lambda_+ d_+ + \alpha_+ (d_-) - \beta_+ (d_+) d_+$$

$$(6.1b) \quad \dot{d}_- = -\lambda_- d_- + \alpha_- (d_+) - \beta_- (d_-) d_-$$

where the transfer functions are $\beta_{\pm}(x) =$

$(x^+, x^-) = \frac{I_0}{\kappa^2 - \gamma^2} \cdot (\kappa \pm \gamma)$ and $\frac{I_0}{\kappa^2 - \gamma^2} \cdot (\gamma \pm \kappa)$ in either case. Stability of these fixed points is given by the nonzero eigenvalue $\lambda = -(\kappa + \gamma) < 0$, so these quasi-equilibria are always attractive. Note also that the reduced SDE for the difference $x = x^+ - x^-$ will take the form $dx = \left[-d(x) - (\kappa + \gamma)x \right] dt + \sigma dx$, where $-d(x) = -\frac{1}{2} \sigma^2 (x^+ - x^-)^2$ and $\sigma^2 dx = \sigma^2 (x^+ + x^-) dx$.

• • • • •

(2) alternatives and show that the log probability updates can be approximated by a nonlinear system of SDEs in the continuum limit. With the appropriate scaling of the probabilities, $\pi_i(\cdot) = \Pr(\cdot | \mathcal{I}_i)$, we can make precise the correspondence between the discrete and continuum models of posterior probability evolution. Lastly, we present a derivation for the stochastic integrodifferential equation that represents the log probability for a continuum of possible environmental states, $\theta \in \mathcal{E}$.

Note that throughout the appendices, we use notation involving a subscript n . This helps us define a family of stochastic processes indexed by the spacing between observations $\Delta t_n = \Delta t_{n-1}$. For instance, $\pi_{\Delta t, \pm}(\cdot)$ represents the probability of an observation, y_n , in environmental state \pm (or, in the language of statistics, when hypothesis \pm holds). This probability changes with the timestep Δt_n . This approach allows us to properly take the continuum limit $\Delta t_n \rightarrow 0$. However, for simplicity we refrain from using this notation in the main text.

where we have conditioned on the state of the environment, $(\psi^n) = \pm$ at time ψ^n .
Replacing the index

terms (analogous to the derivation of (A.1) and (A.2)), we can express each probability $\pi_{n,i}$ in terms the probability at the time of the previous observation, $\pi_{n-1,j}$:

$$\pi_{n,i} = \frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})} \Pr(\xi_n = i | \mathcal{F}_n) \prod_{j=1}^N \epsilon_{\Delta t, ij}^{x_{n-1, j} - x_{n-1, i}}$$

Since we are only interested in comparing the magnitude of the probabilities, we can drop the common prefactor $\frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})}$ and use the fact that $\prod_{j=1}^N \epsilon_{\Delta t, ji} = 1$ (since $\epsilon_{\Delta t, ij}$ is a left stochastic matrix) to write $\epsilon_{\Delta t, ii} = \prod_{j \neq i} \epsilon_{\Delta t, ji}$ and obtain

$$(D.1) \quad \pi_{n,i} = \pi_{\Delta t, i}(n) \prod_{j \neq i} \epsilon_{\Delta t, ji}^{x_{n-1, j} - x_{n-1, i}} + \prod_{j \neq i} \epsilon_{\Delta t, ij}^{x_{n-1, j} - x_{n-1, i}}$$

where $\pi_{\Delta t, i}(n) = \Pr(\xi_n = i | \mathcal{F}_n)$. From (D.1), it follows that the log of the rescaled probabilities, $\ln \pi_{n,i}$, satisfies the recursive relation

$$\ln \pi_{n,i} - \ln \pi_{n-1, i} = \ln \pi_{\Delta t, i}(n) + \ln \prod_{j \neq i} \epsilon_{\Delta t, ji}^{x_{n-1, j} - x_{n-1, i}} + \prod_{j \neq i} \epsilon_{\Delta t, ij}^{x_{n-1, j} - x_{n-1, i}}$$

To derive an approximating SDE, we denote by $\Delta \ln \pi_{n,i} = \ln \pi_{n,i} - \ln \pi_{n-1, i}$ the change in the log probability due to an observation at time t_n . As before, we assume $\epsilon_{\Delta t, ij} := \sqrt{\Delta t} \epsilon_{ij} + o(\sqrt{\Delta t})$ for $t_n = i$ and drop the higher order terms, giving

$$\Delta \ln \pi_{n,i} = \ln \pi_{\Delta t, i}(n) + \ln \prod_{j \neq i} \sqrt{\Delta t} \epsilon_{ji}^{x_{n-1, j} - x_{n-1, i}} + \prod_{j \neq i} \sqrt{\Delta t} \epsilon_{ij}^{x_{n-1, j} - x_{n-1, i}}$$

Assuming $\Delta t \ll 1$, we again use the approximation $\ln(1 + x) \approx x$ for $|x| \ll 1$. We also assume that the change in the log probability, $|\Delta \ln \pi_{n,i}| \ll 1$, is small over the time interval, Δt so that

$$\Delta \ln \pi_{n,i}$$

The correlation of η_i 's is given by

$$\text{Corr}_\xi[\eta_i, \eta_j] := \text{Corr}_\xi[\mathbf{1}_{\mathbb{R}^+}(\eta_{\Delta t, i} - \bar{\eta}_{\Delta t, i}), \mathbf{1}_{\mathbb{R}^+}(\eta_{\Delta t, j} - \bar{\eta}_{\Delta t, j})]$$

Note that (D.3) is the multiple-alternative version of (B.3). Equivalently, we can write (D.3) as

$$\eta_{\Delta t, i} = \bar{\eta}_{\Delta t, i} + \sum_{j \neq i} \epsilon_{ij} e^{x_{t,j} - x_{t,i}} \epsilon_{ji}$$

where $\bar{\eta}_{\Delta t} := (\bar{\eta}_{\Delta t, 1}, \dots, \bar{\eta}_{\Delta t, N})$ follows a multivariate Gaussian distribution with mean zero and covariance matrix $\Delta_{\Delta t}$ given by

$$\Delta_{\Delta t, ij} =$$

We can also derive a continuum limit for the log-likelihood ratio for any two choices θ_1, θ_2 .

where $\mathbb{B}_{\Delta t} := (\mathbb{B}_{\Delta t, \theta})_{\theta \in \mathcal{A}, \mathcal{B}}$. For $\theta \in \mathcal{A}, \mathcal{B}$, $\mathbb{B}_{\Delta t, \theta}$ is a Gaussian process in the sense that any finite subset of points $\{\theta_1, \dots, \theta_n\} \in \mathcal{A}, \mathcal{B}$ have a multivariate Gaussian distribution with mean zero and covariance, $\mathbb{C}_{\Delta t, \theta\theta'}$ given by

$$\mathbb{C}_{\Delta t, \theta\theta'} = \frac{1}{\Delta t} \text{Cov}_{\xi} [\text{lr}_{\Delta t, \theta}(\cdot), \text{lr}_{\Delta t, \theta'}(\cdot)]$$

Finally, taking the limit $\Delta t \rightarrow 0$, and assuming that the limits

$$(F.3) \quad \mathbb{g}_{\theta}(\mathbf{x}) := \lim_{\Delta t \rightarrow 0} \mathbb{B}_{\Delta t, \theta}(\mathbf{x}) \quad \text{and} \quad \mathbb{C}_{\theta\theta'}(\mathbf{x}) := \lim_{\Delta t \rightarrow 0} \mathbb{C}_{\Delta t, \theta\theta'}(\mathbf{x})$$

are well defined, we obtain the system of SDEs

$$(F.4) \quad d\mathbf{x} = \mathbb{g}_{\theta}(\mathbf{x})d\mathbf{w} + \mathbb{C}_{\theta\theta'}(\mathbf{x}) \int_a^b \epsilon_{\theta\theta'} e^{x_{\theta'} - x_{\theta}} \epsilon_{\theta'\theta} d\theta' d\mathbf{w}$$

or equivalently the system of SDEs

$$d\mathbf{x} = \mathbf{g}(\mathbf{x})d\mathbf{w} + \mathbf{C}(\mathbf{x})d\mathbf{W}_t + \mathbf{x}(\mathbf{x})d\mathbf{w}$$

where $\mathbf{g}(\mathbf{x}) = \mathbb{g}_{\theta}(\mathbf{x})_{\theta \in \mathcal{A}, \mathcal{B}}$ and $\mathbf{C}(\mathbf{x}) \mathbf{C}^T(\mathbf{x}) = \mathbf{C}(\mathbf{x})$ are defined using the limits in (F.3), $\mathbf{x}(\mathbf{x}) = \int_a^b (\epsilon_{\theta\theta'} e^{x_{\theta'} - x_{\theta}} \epsilon_{\theta'\theta}) d\theta'$, and the components of \mathbf{W}_t are independent Wiener processes.

While we have formally taken the limit of the discrete (F.2), it is important to note that establishing the well-posedness of stochastic integrodifferential equations is not straightforward. Conditions for the existence and uniqueness of solutions to certain nonlinear stochastic partial differential equations (SPDEs) are demonstrated in Chapter 7 of [15]. This approach considers the solutions to SPDEs to be random processes that take their values in a Hilbert space of functions. Recently, this concept has been extended to provide general conditions on the constituent functions of stochastic neural fields to ensure the existence of solutions [19, 25]. The form of stochastic neural fields is closely related to (F.4), since both types of equation possess a linear drift and a convolution defining a nonlocal coupling between their state variables. It may be possible to utilize these previous approaches to establish the existence and uniqueness of solutions to (F.4) in future studies.

We thank Eric Shea-Brown, Shawn Gu, and the anonymous reviewers for helpful comments.

REFERENCES

[1] R. P. Adams and D. J. MacKay, *Bayesian Online Changepoint Detection*, arXiv:0710.3742, 2007.

[2] P. Armitage, *Sequential analysis with more than two alternative hypotheses, and its relation to discriminant function analysis*, *J. R. Stat. Soc. B*, 12 (1950), pp. 137–144.

[3] G. Balázsi, A. van Oudenaarden, and J. J. Collins, *Cellular decision making and biological noise: From microbes to mammals*, *Curr. Biol.*, 21 (2011), pp. 910–925.

[4] J. M. Beck, W. J. Ma, R. Kiani, T. Hanks, A. K. Churchland, J. Roitman, M. N. Shadlen, P. E. Latham, and A. Pouget, *Probabilistic population codes for Bayesian decision making*, *Nat. Neurosci.*, 11 (2008), pp. 1142–1152.

[5] T. E. Behrens, M. W. Woolrich, M. E. Walton, and M. F. Rushworth, *Learning the value of information in an uncertain world*, *Nat. Neurosci.*, 10 (2007), pp. 1214–1221.

[6] P. Billingsley, *Probability and Measure*, Wiley, New York, 1995.

- [35] R. Ratcliff, *A theory of memory retrieval*, *Psychological Review*, 85 (1978), pp. 59–108.
- [36] R. Ratcliff and P. L. Smith, *A comparison of sequential sampling models for two-choice reaction time*, *Psychological Review*, 111 (2004), pp. 333–367.
- [37] Y. A. Rozanov, *Probability Theory: A Concise Course*, Dover, New York, 1977.
- [38] T. D. Seeley, S. Camazine, and J. Sneyd, *Collective decision-making in honey bees: How colonies choose among nectar sources*, *Behavioral Ecology & Sociobiology*, 28 (1991), pp. 277–290.
- [39] A. N. Shiryaev, *Optimal Stopping Rules*