Decision-making through evidence integration at long timescales

Michael L. Waskom (mwaskom@nyu.edu) and Roozbeh Kiani (roozbeh@nyu.edu)
Center for Neural Science, New York University
New York, NY 10003 USA

Abstract

When multiple pieces of information bear on a decision, the best approach is to integrate the evidence provided by each one. Signatures of evidence integration have been identified in neuronal responses during decision formation, leading many investigations of how the brain produces complex behavior to focus on these computations. Yet, because evidence integration has most often been studied in simple tasks with short timescales of deliberation, it is unknown whether these models can provide a general account of decision-making. Here, we introduce a novel psychophysical paradigm with a long and discontinuous timescale of evidence availability. We show that choice behavior in this task reflects an evidence integration process that can extend over tens of seconds without loss of information due to memory leak or noise. Our results reveal that temporally-extended decisions approximate the normative computations used for rapid sensory discrimination, validating the generality of the evidence integration framework for modeling human cognition.

Keywords: decision-making; working memory; sequential sampling; integration time constant; integration noise; probabilistic inference; psychophysics

Background

Evidence integration models represent a powerful synthesis of cognitive neuroscience and computational theory (Shadlen & Kiani, 2013). They make strong quantitative predictions about the speed and accuracy of simple perceptual discriminations, and they provide a computational framework for exploring higher-level phenomena such as confidence, strategy, and flexibility. Evidence integration models also provide a basis for modeling and interpreting the neural mechanisms that underlie decision-making behaviors (Gold & Shadlen, 2007) and for translating insights across measurement modalities and model systems (Hanks & Summerfield, 2017). They are therefore regarded as a promising approach for studying the computational and neural basis of cognition.

Despite this promise, a major limitation of the evidence integration framework is that its relevance to decision-making behaviors beyond rapid perceptual discrimination is not clear. Perceptual discrimination tasks afford tight experimental control and embody many important aspects of decision-making, but they rarely demand integration at timescales longer than hundreds of milliseconds. In contrast, humans often deliberate for many seconds (or even much longer) while making decisions in daily life, and in doing so frequently consider multiple discrete pieces of information that bear on their choice. The disconnect between experimental tasks and naturalistic behavior is relevant because the long and discontinuous timescale of evidence availability in the natural world might prove challenging for the neural mechanisms thought to underlie integration of evidence in rapid discrimination tasks.

Experimental paradigm

To address this challenge, we developed a novel psychophysical paradigm that permits evaluation of evidence integration models in the context of long-timescale decisions. Our paradigm builds on the success of established perceptual discrimination tasks by providing evidence to the subject in the form of simple visual stimuli that can be parametrically controlled with high precision and whose basic encoding and representation in sensory cortex is well understood. Rather than expose these stimuli to subjects in a continuous stream, however, we ask subjects to make decisions on the basis of multiple brief exposures with variable strength that are separated in time by unpredictable gaps. Precise knowledge of the timing and strength of each sample permits detailed computational modeling of the mechanisms through which that evidence is used to determine the subject's choice.

A specific implementation of the task is shown in Fig. 1. During each trial, the subject sees brief (200 ms) presentations of a contrast pattern (a "sample") while maintaining central fixation (Fig. 1B). These samples are drawn, on any given trial, from one of two overlapping Gaussian distributions in log contrast space (Fig. 1B). The subject's task is to infer, on the basis of the samples seen in that trial, whether they had been generated from the "low" distribution or from the "high" distribution. We control the amount of information available for this inference by showing between one and five samples per trial (Fig. 1C). We additionally control when the evidence is available by separating each sample with an unpredictable gap of several seconds in duration (Fig. 1D). The data reported here were collected in sessions where the gaps were either "shorter" (1–4 s) or "longer" (2–8 s); these gaps extended the timescale of deliberation on many trials to tens of seconds (mean trial duration, 10.1s $\pm$ 5.6 s; range, 2.2–34 s).

Five human subjects were trained on this task until they reached an accuracy criterion ($>$76% correct; achieved over 2–4 sessions) and then performed, in total, 14,869 trials. This rich dataset allowed us to characterize the computations underlying decision-making behavior in our task.

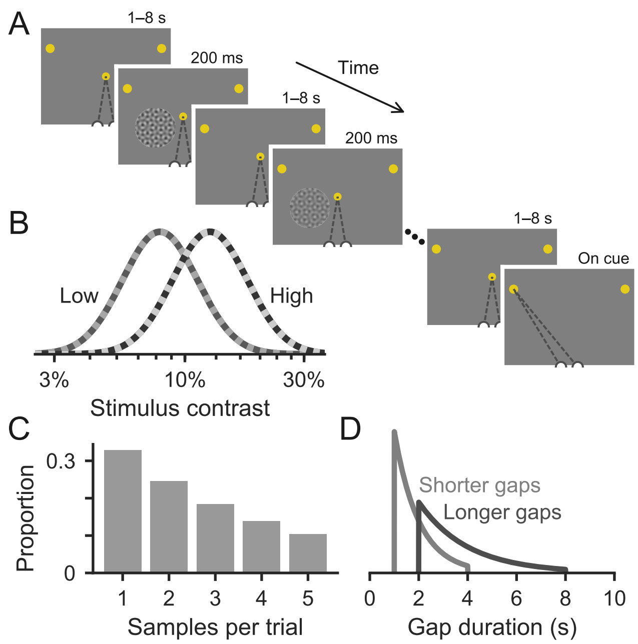

Figure 1: Experimental design. (A) Sequence of events in a trial. (B) Generating distributions for contrast samples. (C) Distribution of sample counts across trials. (D) Distributions of gap durations across trials in two timing conditions.

Computational models

We focus here on three questions. First, did subjects combine information from multiple samples before reaching a decision? Second, did information held in working memory decay or drift during the long gaps between samples? Third, were decisions based on the graded weight of evidence afforded by each sample? Together, these questions allow us to assess whether decisions were made through evidence integration at long timescales.

We answered these questions by implementing a set of computational models that formalize different approaches for using the evidence to make a choice. These models generate both quantitative predictions and qualitative signatures that are characteristic of different decision-making mechanisms. We focus here on four specific models: one implements the optimal policy of evidence integration and serves as a baseline for comparison, while the others provide insight into each of the three questions enumerated above. The models share a general structure of using each sample of evidence, denoted x, to update a decision variable, denoted V. The evidence is quantified by the log likelihood ratio (LLR) that the sample was generated from the high contrast distribution, and choice is determined by the sign of the decision variable at the end of the trial.

Linear Integration: Optimal performance in the task can be achieved by summing the continuous weight of evidence afforded by each sample. This computation can be formalized in a "Linear Integration" model defined by the following difference equation:

$$V_i = V_{i-1} + x_i + \xi_\eta \,. \tag{1}$$

Choice variability arises in the Linear Integration model because each sample is subject to Gaussian noise during sensory encoding, represented by ξ_η. The only free parameter in the model is σ_η, the standard deviation of ξ_η. Linear Integration is optimal in that it is limited only by variability in the stimulus generation and noisy perception of each sample; no other information is distorted or lost during deliberation (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006).

Extrema Detection: An alternate strategy, which can mimic integration of evidence in some settings, involves sequential sampling and commitment based on extreme values without using memory or integration. In this "Extrema Detection" model, each sample is compared against a decision threshold; if the threshold is exceeded, the process terminates in a commitment to the corresponding alternative, and further samples are ignored. Prior to commitment, intermediate samples of evidence are discarded and have no bearing on choice. If none of the samples produce a commitment by the end of the trial, then a response is generated randomly. The process leads to the following difference equation:

$$V_i = \begin{cases} +1 & V_{i-1} = 0 \quad \text{and} \quad x_i + \xi_\eta > +\theta_x \\ -1 & V_{i-1} = 0 \quad \text{and} \quad x_i + \xi_\eta < -\theta_x \\ V_{i-1} & V_{i-1} \neq 0 \quad \text{or} \quad -\theta_x < x_i + \xi_\eta < \theta_x \,. \end{cases} \tag{2}$$

As in the Linear Integration model, ξ_η represents Gaussian noise with standard deviation σ_η, and the other free parameter is θ_x, which represents the decision threshold.

Leaky Integration: The normative model in Eq. 1 does not require any information about the timing of the samples or the gaps between them. It is possible, however, that long deliberation timescales would cause deviations from normative integration due to limitations in working memory. To test for the presence of such limitations, we extended Eq. 1 into the temporal domain and explicitly modeled two ways that the integrator could lose information during the gaps: memory leak and memory noise. This "Leaky Integration" model can be defined by the following difference equation:

$$V(t) = V(t-\Delta t) - \Delta t \lambda V(t-\Delta t) + x(t) + I(t)\xi_\eta + \xi_d \,. \tag{3}$$

In this equation, λ represents the memory leak rate (the inverse of the integration time constant τ) and is scaled by time step Δt so that it is expressed in units of seconds. $x(t)$ represents the strength of evidence at time t and is 0 during the gaps. $I(t)$ is an indicator variable that specifies when a sample is visible and governs the influence of sensory noise,

$\xi_\eta \sim \mathcal{N}(0, \sigma_\eta)$. Finally, $\xi_d \sim \mathcal{N}(0, \sigma_d)$ represents memory noise, which accumulates throughout the gaps. When implementing the model, we used a time step equal to the sample duration (Δt = 200 ms).

Transformed Evidence: The normative model in Eq. 1 makes full use of the graded weight of evidence afforded by each sample. Another potential deviation from normative integration in our task would therefore involve transformations applied to the evidence values before they enter into the integrator that cause them to deviate from the LLR. One extreme form of transformation would be to binarize the evidence from each sample as supporting either a "high" or "low" choice. This would cause the integration process to resemble counting with discrete values. More intermediate forms of transformation that curtail the weight of evidence in the tails of the distributions while amplifying evidence near the category boundary have also been proposed. This family of transformations can be modeled with a transfer function (Cheadle et al., 2014):

$$f(x, \sigma_\rho) = -1 + 2\frac{\int_{-1}^{x} \varphi(y \mid 0, \sigma_\rho)\, dy}{\int_{-1}^{1} \varphi(y \mid 0, \sigma_\rho)\, dy}, \tag{4}$$

where φ is a Gaussian PDF with standard deviation σ_ρ. This function transforms evidence with increasing strength for smaller values of σ_ρ and approaches a discrete step function as $\sigma_\rho \to 0$. Using the transformed evidence in the integration process alters Eq. 1 to

$$V_i = V_{i-1} + f(x_i + \xi_\eta, \sigma_\rho), \tag{5}$$

corresponding to a "Transformed Evidence" model.

Results

We fit the free parameters of each model by maximizing the likelihood of choice data given the sequence of samples on each trial. We then assessed model performance in several ways, including model comparison, consideration of qualitative model signatures, and inference on estimated parameter values. Here, we focus on comparing choice data and aggregate model predictions in terms of how well the models can explain improvements in accuracy with more samples (Fig. 2), a key feature for contrasting models of decision-making.

In the Linear Integration model, the improvement of accuracy with additional samples is jointly determined by the separation between the generating distributions and the sensory noise parameter σ_η. This function can be analytically derived from Eq. 1. Fig. 2A shows a qualitative match between model and data when fitting on all trials (solid line) and a nearly identical parameterization for trials with either longer or shorter gaps (dashed and dotted lines). Despite the qualitative agreement, it is apparently that the slope of the improvement is somewhat shallower for the data than for the model. This suggests that the Linear Integration model is slightly overestimating σ_η and that information is being lost in some other way.

Can this discrepancy be explained by sequential sampling without memory? It is also possible to derive how accuracy should improve with additional samples for the Extrema Detection model using Eq. 2. Fig. 2B shows the prediction of this model and indicates that the data cannot be explained by Extrema Detection, which systematically fails to account both for the overall level of accuracy and for the improvement in accuracy with additional samples. Indeed, in individual subject fits, the Linear Integration model provided significantly better match to the data than the Extrema Detection model for all subjects (bootstrap test; all $Ps < 0.05$). This result indicates that behavioral performance was supported by a process of combining information across samples.

We expected that, even if the task were solved through integration, the long gaps would cause some amount of information from early samples to be lost. Surprisingly, Fig. 2C shows that there was minimal influence of either leak or noise in the integrator memory. Simulating performance for the Leaky Integration model with the best fitting parameters produces a function that is indistinguishable from the Linear Integration prediction. In individual subject fits, only one subject had a leak rate that was significantly different from 0 (bootstrap test, $P <$ 0.05), and the integration time constant was estimated to be larger than 20 s in all subjects. Further, the memory noise parameter σ_d was close to 0 and approximately an order of magnitude smaller than the sensory noise parameter σ_η in all subjects. This panel also shows simulated performance for models with a moderate leak rate (dashed line, corresponding to τ = 2 s) or moderate memory noise (dotted line). These functions depart from the data in clear and distinct ways. Therefore, our experimental design can distinguish different parameterizations of the leaky integration model, but the data are most consistent with a process that approximates perfect integration and has minimal leak or accumulation of noise between samples.

Fig. 2D further shows that the integration process used a moderately transformed weight from each sample, integrating continuous evidence but reducing the influence of samples that strongly favored either choice. Simulated performance for the Transformed Evidence model with the best fitting parameters shows a close correspondence to the data, and the estimated transformation parameter σ_ρ was larger than 0 in all subjects. The panel also shows predicted performance for a model with σ_ρ = 0 (dashed line), which implements a process of counting binarized representations of each sample. The prediction of the counting model fails to explain the data, because when σ_ρ = 0, binarization forces random guessing on many trials with even numbers of samples, and aggregate performance is identical to that on trials with the next smallest odd number of samples. This effect does not exist in the data (all $Ps <$ 0.05).

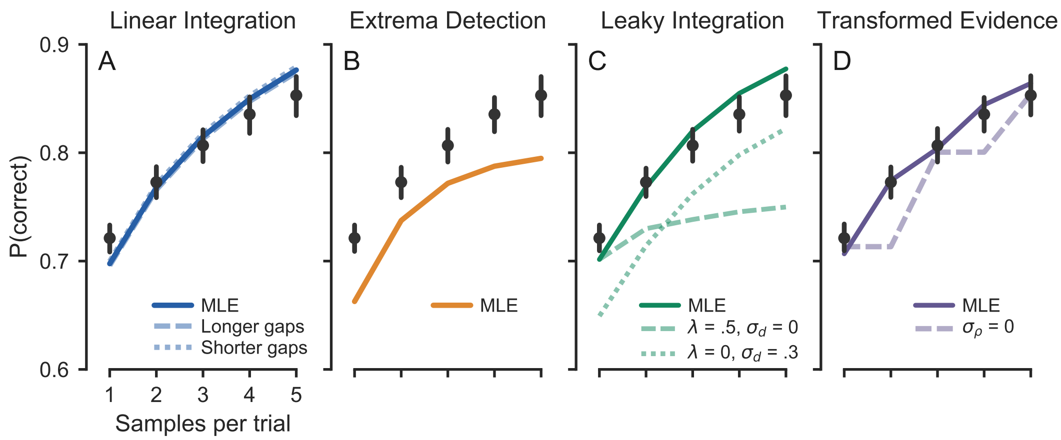

Figure 2: Model fitting results. Black points show means and 95% C.I.s. Solid lines show model predictions using maximum likelihood estimates (MLE) for free parameters, and dashed lines show predictions using alternate parameterizations.

Conclusions

Our novel experimental paradigm allows us to determine whether the evidence integration framework can account for decisions about discrete samples of sensory evidence that are separated by long and unpredictable gaps. These gaps extended the timescale of the deliberation process to an order of magnitude beyond standard perceptual discrimination tasks. We found that choice behavior was nevertheless consistent with a normative integration process that combined graded information from multiple samples with minimal memory leak or noise. A small departure from optimal performance could be explained by integration of transformed values that underweighted strong evidence. Therefore, our results support the generalization of many insights about decision-making that have been developed through the study of simple perceptual tasks to a broader class of more naturalistic decisions, addressing, for the first time, a long-lasting concern in the field.

While our results show that similar computational principles underlie decision-making at different timescales (Brunton, Botvinick, & Brody, 2013), they also prompt new questions about the neural implementation of the evidence integration process. In the standard model, evidence integration is implemented by persistent firing in networks of spiking neurons supported by local recurrent excitation (Gold & Shadlen, 2007). This mechanism is often modeled using attractor dynamics (Wang, 2008). Yet many standard parameterizations of these models would fail to maintain evidence across the gaps in our task, particularly at the longest durations. For example, networks that implement bi-stable point attractor dynamics would be subject to catastrophic leak during each gap and would be driven towards a choice based on a single strong sample, whereas networks that implement line attractor dynamics would be dominated by memory noise. We found instead that performance was largely invariant to a wide range of gap durations. One alternative explanation could be that evidence integration for long-timescale decisions is supported by distributed interactions between multiple large-scale systems including those supporting long-term memory. As this question comes into focus, we note that extending the timescale of the deliberation process also makes it accessible to study with functional MRI, adding a powerful new tool to the study of this fundamental cognitive ability.

Acknowledgments

This work was supported by R01-MH109180, the Simons Foundation, and a Pew Scholarship in Biomedical Sciences.

References

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.

Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, *340*(6128), 95–98.

Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Castañón, S. H., & Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron*, *81*(6), 1429–1441.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.

Hanks, T. D., & Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron*, *93*(1), 15–31.

Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, *80*(3), 791–806.

Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. *Neuron*, *60*(2), 215–234.