

Theory of Rapid Behavioral Inferences Under the Pressure of Time

Ann M. Hermundstad^{1,*} and Wiktor F. Młynarski^{2,3,*}

¹*Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, Virginia 20147, USA*

²*Ludwig Maximilian University of Munich, Munich, Germany*

³*Bernstein Center for Computational Neuroscience, Munich, Germany*



(Received 29 May 2025; accepted 11 February 2026; published 5 March 2026)

To survive, animals must be able to quickly infer the state of their surroundings. For example, to successfully escape an approaching predator, prey must quickly estimate the direction of approach from incoming sensory stimuli and guide their behavior accordingly. Such rapid inferences are particularly challenging because the animal has only a brief window of time to gather sensory stimuli, and yet the accuracy of inference is critical for survival. Due to evolutionary pressures, nervous systems have likely evolved effective computational strategies that enable accurate inferences under strong time limitations. Traditionally, the relationship between the speed and accuracy of inference has been described by the “speed-accuracy tradeoff” (SAT), which quantifies how the average performance of an ideal observer improves as the observer has more time to collect incoming stimuli. While this trial-averaged description can reasonably account for individual inferences made over long timescales, it does not capture individual inferences on short timescales, when trial-to-trial variability gives rise to diverse patterns of error dynamics. We show that an ideal observer can exploit this single-trial structure by adaptively tracking the dynamics of its belief about the state of the environment, which enables it to speed its own inferences and more reliably track its own error, but also causes it to violate the SAT. We show that these features can be used to improve overall performance during rapid escape. The resulting behavior qualitatively reproduces features of escape behavior in the fruit fly *Drosophila melanogaster*, whose escapes have presumably been highly optimized by natural selection.

DOI: [10.1103/8m59-vf2m](https://doi.org/10.1103/8m59-vf2m)

I. INTRODUCTION

An animal’s survival depends on its ability to successfully interact with its surroundings. This requires continually inferring properties of the environment from incoming sensory stimuli to guide appropriate actions. Inference is thus a fundamental computation that must be performed by the nervous system, often in challenging circumstances that are limited by noise [1] and metabolic [2] or computational [3,4] constraints.

Of the many constraints that make inference challenging, strong time limitations are perhaps the most directly related to animal’s survival. For example, prey must be able to estimate the direction of an approaching predator and plan an effective escape strategy, often in a fraction of a second [5,6]. In such scenarios, accurate inference is challenging because the brain has limited time to gather sensory data that can itself be noisy and inaccurate. Because of the direct impact of rapid inference on life and death decisions, brains and other biological systems have likely evolved sophisticated

strategies to maximize inference accuracy when faced with strong time limitations [7,8].

Canonical theoretical approaches use the speed-accuracy-tradeoff (SAT) to describe the behavior of an ideal observer that must balance the time spent collecting sensory stimuli with the accuracy of the resulting inference [9,10]. The SAT has an intuitive interpretation: the longer the observer waits, the more sensory stimuli it can gather, and the more accurate the inference it can make. To determine when to stop collecting sensory stimuli and commit to an action, the observer must thus trade off the speed of its decision against the accuracy of its inference. This tradeoff has been shown to capture behavior across multiple scenarios and species, including in humans [11–14], rats [15], mice [16], bees [17], and nonneural organisms such as the slime mold *Physarum polycephallum* [18].

The SAT establishes a fundamental relationship between the speed and accuracy of inference, and suggests that animals should use all available time to achieve the highest possible accuracy. However, the SAT characterizes the performance of an ideal observer *on average*, and it ignores any structure observed in individual scenarios or trials. In particular, short sequences of stimuli can result in either much better or much worse performance than is captured by the average. In principle, this trial-to-trial structure could be exploited by detecting stimulus sequences that lead to a rapid reduction in inference error, thereby enabling the observer to reduce inference time and devote more time to coordinating actions. Alternatively, this structure could be exploited by detecting stimulus sequences that are likely to result in high inference

*These authors contributed equally to this work.

†Contact author: hermundstada@janelia.hhmi.org

‡Contact author: mlynarski@bio.lmu.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

errors that lead to inaccurate beliefs, thereby enabling the observer to take actions that maximize the chance of survival even when the underlying state of the environment is unknown (e.g., as observed in freezing behavior [6]). In both cases, the ability to exploit higher-order statistics of individual inferences could give animals a richer repertoire of possible actions and paths to survival, beyond waiting for more data or acting with inaccurate estimates, as suggested by the SAT.

In this work, we propose a theory for how an ideal observer can exploit the structure of trial-to-trial variability to increase the speed of its inference and bound the expected magnitude of its error. Our key insight is that short sequences of stimuli generate diverse patterns of error on individual trials; the observer can exploit this structure by relying on the dynamics of its own uncertainty about the underlying state of the environment. As a result, the optimal adaptive observer violates the SAT. Using a simple observer model that mimics an escape task, we show that this adaptive inference strategy improves overall escape performance, and qualitatively captures features of escape behavior in *Drosophila melanogaster*. This suggests that these and other organisms might exploit statistical regularities to improve inferences under strong time constraints.

II. RESULTS

Our primary conceptual aim is to characterize trial-to-trial variability that arises when performing inference under strong time constraints, and identify statistical structure that could be exploited by an ideal observer to improve the speed and reliability of inference. We first demonstrate the existence of structured variability in a simple estimation task: inferring the mean of a Gaussian distribution from a limited number of samples. We then demonstrate how animals could exploit such variability to improve their performance in a time-constrained task that is crucial for survival: rapid escape.

A. Inference tasks exhibit structured variability in their trial-to-trial dynamics

In an inference task, an organism estimates a latent state of the environment θ from a sequence of incoming sensory stimuli $s_1, \dots, s_t$ [Fig. 1(a)]. The organism can be modeled as an ideal observer that performs optimal inference by computing the posterior distribution over the latent state, given the observed stimulus sequence (i.e., by computing $p(\theta|s_{1:t})$, where $s_{1:t}$ is the specific sequence of stimuli observed in a trial up to time point t). This posterior can then be used to compute a point estimate of the latent state $\hat{\theta}_t$, for example, by taking the posterior average [19].

As a simple illustrative example, we consider the problem of inferring the mean of a Gaussian distribution [Fig. 1(b)], where the observer has limited time to perform the inference, and can therefore only observe a maximum of $T_{\max} = 10$ stimulus samples. On average, the error in the observer's estimate of the mean decreases in time, obeying the SAT [Fig. 1(c), black dashed line]. However, a simple clustering analysis reveals that this average pattern does not capture the behavior observed on individual trials [Fig. 1(c), thin gray lines; Methods]. For example, a majority of trials (66%) converge

much faster and to a much lower error than the average [Fig. 1(c), blue], while the remaining trials (34%) converge more slowly and to higher values of error. This results in a counterintuitive violation of the SAT, where faster convergence leads to lower, rather than higher, error [Fig. 1(d)]. These patterns of error generalize across different inference scenarios with more or less environmental variability [parametrized here by the variance in the underlying latent state; Fig. 1(e)]. This variability is caused by the stochastic nature of short stimulus sequences. Purely by chance, some of these sequences are more representative of the underlying distribution, leading to a faster convergence of the observer's estimate and a lower inference error. We call such stimulus sequences “optimistic” [Fig. 1(f), blue lines]. In contrast, “pessimistic” sequences are not representative of the underlying stimulus distribution and can therefore be misleading about the value of underlying latent state [Fig. 1(f), red lines].

The possibility of performing rapid, low-error inferences depends on the relative frequency of optimistic stimulus sequences. In the mean estimation task considered here, such optimistic stimuli are located close to the mean of the distribution, within a bulk that is parametrized by a radius ϵ [Fig. 1(g), right]. The relative probability of observing an optimistic sequence—i.e., one in which all stimuli come from the bulk—decays rapidly with the length of the stimulus sequence [Fig. 1(g), left]. This rapid decay suggests that a potential inference strategy based on “fishing out” optimistic sequences is applicable only in very time-limited scenarios, where observations are limited to a small number of stimuli. While we relied here on a strict definition of a low-error sequence, this can in principle be relaxed, for example, by considering sequences in which a majority of stimuli (but not all) originate from the bulk.

It is worth emphasizing that this is one of the simplest scenarios that can be used to study the dynamics of inference. Despite its simplicity, individual trials exhibit structure that can be exploited to save time and bound error. Crucially, this structure is *intrinsic* to the statistics of short stimulus sequences, and would therefore be present in more complicated inference scenarios or under different model assumptions. In the remainder of the paper, we explore how this structured variability could be exploited by animals in an inference task directly related to their survival.

B. An example escape task requires rapid inference under strong time limitations

Planning and coordinating an escape is a salient example of a task that animals must solve under strong time pressures. Here, we consider an important component of rapid escape: inferring the direction of an approaching predator from noisy sensory stimuli [6,7,20]. We study a simplified setting in which an animal must use a stream of incoming sensory stimuli s_t to infer a latent state θ that specifies the angular direction of an approaching predator [Figs. 2(a)–2(c)]. We assume that this latent state parameterizes the mean of von Mises distribution, and can take one of N discrete values that represent possible approach directions [Fig. 2(a)]. We model the animal as an ideal observer that maintains and updates a belief about the direction of approach, as summarized by the




FIG. 1. Individual inference trajectories exhibit different patterns of error. (a) In inference tasks, an organism must estimate a latent state of the environment θ from available sensory stimuli s_t . Here, we focus on scenarios where the state of the environment is unchanging in time. (b) To perform optimal inference, an ideal observer builds a posterior belief $p(\theta|s_{t \leq t})$ about the underlying latent state θ of the environment (e.g., the mean of the distribution of sensory stimuli) from incoming samples s_t . This belief can then be used to construct a point estimate $\hat{\theta}_t$ of the latent state that can be used to guide an appropriate action. Here, we assume that the observer begins with a Gaussian prior, and updated its posterior belief using i.i.d. stimulus samples drawn from a Gaussian distribution with the same mean (Methods). We take the observer's point estimate of the latent state to be the mean-squared error $\hat{\theta}_t = \langle (\theta - \hat{\theta}_t)^2 \rangle$. (c) The trial-averaged error in the observer's estimate (black dashed line) converges slowly. However, individual trials exhibit dynamics that can deviate strongly from the mean (thin gray lines). When clustered based on their error dynamics (blue and red lines), different clusters exhibit error that is higher (red) or lower (blue) than the mean. The pie chart depicts the fraction of all 10^6 trials within each of four clusters. (d) Error clusters violate the SAT. Clusters that converge earlier lead to lower error values, while clusters that converge later lead to higher error values. Individual error trajectories were first normalized based on the error at time $t = 1$, averaged across all trials; the convergence time of a cluster was then taken to be the time point where the normalized error—when averaged over trials within a cluster—was within 20% of its final value $t = T_{\max}$. The black dashed line denotes the final normalized error, averaged across all trials. (e) The structured variability of individual inference trials (upper panels, shown analogously to panel c) and corresponding violations in the SAT (lower panels, shown analogously to panel d) generalize across environments with more or less variability in the underlying latent state (parameterized here by the variance of the prior distribution). (f) Example stimulus sequences that are more (blue) or less (red) representative of the underlying stimulus distribution (“optimistic” and “pessimistic,” respectively), and that lead to lower or higher average error [colors correspond to the blue and darkest red clusters in panels (c) and (d)]. (g) Log-ratio of the probability of low-error stimulus sequences relative to the probability of all other sequences sampled from a Gaussian stimulus distribution. Low-error sequences are defined to lie within a radius ϵ of the distribution mean (right panel). The log-ratio is visualized as a function of the sequence length $T_{\max}$ for different values of ϵ measured relative to the standard deviation σ . The upper bound for ϵ is the point at which the bulk and the tails of the distribution have equal probability.

posterior distribution $p(\theta|s_{t \leq t})$. This belief can be used both to construct an estimate $\hat{\theta}_t$ of the direction of approach, and to decide when to stop performing inference and initiate an escape Fig. 2(b). Importantly, while the direction of approach is drawn from one of a discrete number of possible directions, the observer's estimate of direction $\hat{\theta}_t$ is continuous and equal to the posterior average Fig. 2(c). This estimation minimizes the squared inference error $(\theta - \hat{\theta}_t)^2$ [21].

As in the previous section, we constrain the inference period to a maximum duration $T_{\max}$, which limits the number of stimulus samples that can be collected by the observer. Within this limit, the trial-averaged inference error again obeys the SAT, with longer inference times leading to lower average errors Fig. 2(d). This curve suggests that the observer should continue collecting stimulus samples up until the time limit $T_{\max}$, since only then will the average error be minimal. This

strategy appears to be increasingly more beneficial in more complex inference tasks (specified here by the number of latent states N), which require longer inference times to reach the same final error [Fig. 2(e), left]. Thus, the classic SAT would suggest that under stronger time limitations and in more complex scenarios, systems are doomed to suffer from high average error. However, this is precisely the regime in which there is more variability in the distribution of errors across individual trials [Fig. 2(e), right], and thus more potential structure to exploit. This broader distribution of errors suggests that the observer could benefit from estimating its own error to speed its inference. Moreover, if this error is likely to be low, any further reduction in inference time could be used to coordinate more precise actions. Below, we describe strategies for reducing inference time while bounding inference errors, even in seemingly challenging circumstances.




FIG. 2. An example inference task underlies rapid escape. (a) Example inference task: an animal infers the direction of an approaching predator to guide an effective escape. (b) To perform inference, an ideal observer builds a posterior belief $p(\theta|s_{\tau \leq t})$ about an underlying latent state θ of the environment from incoming sensory stimuli s_t . This belief can then be used to construct a point estimate $\hat{\theta}_t$ of the latent state that can be used to guide an appropriate action. Under time constraints, the observer must additionally decide *when* to stop performing inference and initiate the action. (c) In the example shown in panel (a), the latent state specifies the direction of approach and parametrizes the distribution of incoming sensory stimuli. These stimuli are used to infer the posterior probability distribution of different approach directions, which can in turn be used to estimate the approach direction and initiate an escape in the opposite direction. Here, we assume that the observer begins with a uniform prior belief over N possible approach directions. A fixed direction is drawn randomly from this distribution, and is used to parametrize the mean of a von-Mises distribution of sensory stimuli. The observer receives i.i.d. samples from this distribution, which are then used to update its posterior belief. We take the observer's point estimate $\hat{\theta}$ to be the circular mean of this posterior distribution. (d) On average, longer times lead to more accurate inferences. According to this well-known speed-accuracy tradeoff (SAT), an ideal observer should use all of its available time, up to a time limit $T_{\max}$, to make the most accurate inference. Error is averaged over 10^5 trials, and normalized by the average error on the first timestep. (e) Left: stronger time limitations (shorter $T_{\max}$) and more complex inferences (more latent states) lead to higher inference errors on average (“more difficult” region of heatmap). Color denotes the average error at $t = T_{\max}$, normalized by the initial error at $t = 1$. Right: under “more difficult” scenarios, there is a broader distribution of errors below the mean, and thus more structure to be exploited by an ideal observer. Color denotes the standard deviation of all errors below the mean, normalized by the standard deviation of the full error distribution.

C. The statistics of individual stimulus sequences differentiate patterns of inference error

Consistent with the simpler inference task considered in Fig. 1, we find that individual inference trajectories exhibit diverse patterns of error that decrease rapidly on some trials but increase on others [Fig. 3(a), lower panel]. Distinguishing between these scenarios is crucial for performing critical tasks such as planning an escape from a predator, where survival depends not on trial-averaged performance, but rather on maintaining performance above a certain threshold on individual trials. Clustering error trajectories across trials again revealed a diversity of error dynamics, with the largest cluster containing trials in which the inference error rapidly drops below the trial average [Fig. 3(a), upper panel]. The remaining clusters exhibit errors that significantly exceed the trial average; one cluster, for example, contains trials in which the inference error is high even at the time limit [Fig. 3(a), upper panel, yellow line].

As described earlier, this diversity of single trial dynamics arises from variability in stimulus sequences. Intuitively, stimulus sequences that are highly probable under the true latent state should rapidly yield an accurate inference. In contrast, stimulus sequences that are unlikely under the true latent state can be ambiguous and can lead to inaccurate inferences. We expressed this intuition earlier [see Fig. 1(g)] by quantifying the probability of stimulus sequences as a function of their distance from the distribution mean. Here, we generalize this measure and quantify the *excess surprise* S_e of a stimulus sequence $s_{\tau \leq t}$:

$$S_e(s_{\tau \leq t}; \theta = \theta_n) = -\frac{1}{t} \sum_{\tau=1}^t \log p(s_\tau | \theta = \theta_n) - H(s_t | \theta = \theta_n), \quad (1)$$

where $H(s_t | \theta = \theta_n) = -\int p(s_t | \theta = \theta_n) \log p(s_t | \theta = \theta_n) ds_t$ is the expected surprise, or entropy, of the stimulus




FIG. 3. Excess surprise and uncertainty differentiate patterns of inference error. (a) Upper: when clustered by error, different trials exhibit different patterns of average error that can violate the SAT. The largest cluster (purple, 68%) achieved consistently lower error compared to the SAT. The black dashed line denotes the average error across all 10^5 trials, normalized by the average error on the first timestep. The pie chart depicts the fraction of all trials within each of 6 clusters; individual cluster trajectories show normalized error averaged across the subset of trials that comprise each cluster. Lower: these deviations are apparent at the level of individual trials (thin solid lines) that exhibit a diversity of error dynamics over multiple timescales. (b) Upper: Excess surprise, an objective measure, differentiates “optimistic” and “pessimistic” sequences that are more or less surprising than expected on average, respectively. Over long timescales, excess surprise tends to zero for all error clusters (horizontal dashed line), but deviates from zero over short timescales. Cluster averages are performed as in the upper panel of (a). Lower: trials on which there were more optimistic sequences on average achieved lower average errors. The largest cluster of trials encountered the most optimistic sequences and achieved the lowest error. Each dot depicts the time-averaged value of normalized error and excess surprise, measured for each cluster trajectory shown in the upper panels of (a) and (b). (c) Same as panel (b), but shown for uncertainty rather than excess surprise. Upper: An ideal observer does not have access to the objective excess surprise of incoming stimuli, but can instead compute its own subjective uncertainty about the underlying latent state that generated those samples. Over long timescales, uncertainty tends to zero for all error clusters, but deviates from zero over short timescales. Lower: although the detailed dynamics of uncertainty differ from those of excess surprise, there average properties are similar: on trials in which the observer was more certain on average, the observer achieved lower average errors. (d) Over short timescales, the evolving dynamics of uncertainty and surprise deviate from zero (upper panel), but remain correlated with one another and with error (lower panels; markers indicate uncertainty and excess surprise at distinct time points; axis ranges match those in the upper panel). Thus, under time constraints, uncertainty can provide information about the excess surprise of incoming stimulus sequences, and, by consequence, the expected error.

distribution given the true value θ_n of the latent state θ . Excess surprise measures the difference between the average surprise of the observed stimulus sequence and the expected surprise of the full distribution. Because it is defined as a function of the full probability distribution, excess surprise does not depend on any additional parameters, such as the distance from the mean. A negative value indicates that the stimulus sequence was less surprising, and thus more probable, than expected on average; these are the “optimistic” stimulus sequences that unambiguously inform the observer about the true latent state that generated the sequence. In contrast, a positive value of excess surprise indicates that the stimulus sequence was more surprising, and thus less probable, than expected on average; these are the “pessimistic” sequences that can mislead the observer to form an incorrect belief about the true latent state. Excess surprise bears conceptual similarity to the notion of typicality in information theory [22]. Importantly, however, the typical set consists of very long sequences. We focus instead on the properties of

short sequences that can be observed during limited time intervals, which is reminiscent of ideas from large deviation theory [23].

Tracking the excess surprise over time and across trials confirms this intuition [Fig. 3(b), upper panel]. In the limit of long stimulus sequences, the excess surprise tends to zero (i.e., the surprise converges to the average), as expected. However, on short timescales, the excess surprise captures the observed variability in the error trajectories. Those trajectories that rapidly converge to very low error values (dark blue cluster) are generated by optimistic stimulus sequences with negative excess surprise. In contrast, when the error remains high across the duration of the trial, the underlying stimulus sequences are pessimistic, with positive excess surprise. Overall, the higher the average excess surprise of the observed stimulus sequences, the larger the average error [Fig. 3(b), lower panel].

These deviations from the average occur over short timescales and contain structure that can be exploited to

inform more accurate inferences under limited time. However, the relevant quantity—excess surprise—is an *objective* measure. This means that to evaluate the excess surprise of a given stimulus sequence, the observer would need to know the true distribution of stimuli $p(s_t|\theta_n)$, parametrized by the true value of the latent state θ_n . Because the purpose of inference is to determine the latent state θ_n , the observer cannot directly compute the excess surprise. Instead, the observer must rely on a *subjective* quantity to which it has direct access.

One such subjective quantity is the observer’s uncertainty about the latent state, which can be measured using the entropy of the posterior distribution, $H(\theta|s_{\tau \leq t})$ [Fig. 3(c)]. As with excess surprise, uncertainty tends to zero in the limit of long stimulus sequences. On short timescales, the dynamics of uncertainty are correlated with the dynamics of both excess surprise and error [Fig. 3(d)]. For example, optimistic stimulus sequences belonging to the largest cluster lead to the most rapid decrease of the observer’s uncertainty; as described above, this is because they are most probable given the true latent state, and therefore quickly lead to a correct inference. In contrast, pessimistic stimulus sequences cause the observer to maintain high uncertainty across the duration of the trial. Overall, higher average uncertainty is generated by stimulus sequences with higher average excess surprise, which in turn leads to higher average error. When taken together, these relationships suggest that the observer can exploit its own uncertainty to indirectly assess the excess surprise of the observed stimulus sequence and, as a consequence, its own inference error.

D. An adaptive stopping rule improves performance in difficult inference scenarios

To exploit the observed structure in single-trial inference dynamics [Fig. 3], we designed a set of adaptive stopping rules that operate on the observer’s changing uncertainty about the underlying latent state [Figs. 4(a)–4(c)]. These stopping rules are parameterized by a dynamic uncertainty threshold Fig. 4(a); if the observer’s uncertainty drops below this threshold on a given timestep, the observer stops performing inference and commits to an action (we will refer to these as “converged” inferences). At the time limit $T_{\max}$, the inference stops for one of two reasons; either the observer’s uncertainty fell below the threshold at time $T_{\max}$ (but not before), or the observer ran out of time. These two scenarios have different implications for the actions that the observer should take conditioned on this inference; we will first discuss the inference dynamics, and return later to the component of action selection.

Given a particular instantiation of the uncertainty threshold, the adaptive stopping rule leads to a *distribution* of stopping times, each of which leads to a distribution of final errors Fig. 4(b). We use the distribution of normalized stopping times to compute the average stopping time $t_{\text{avg}} \in [0, 1]$, and we use the distribution of normalized final errors (measured across all stopping times) to compute the average final error $e_{\text{avg}} \in [0, 1]$. These are then used to optimize the parameters of the uncertainty threshold by minimizing the

following cost function:

$$C = \alpha e_{\text{avg}} + (1 - \alpha)t_{\text{avg}}, \quad (2)$$

where $\alpha \in [0, 1]$ is a parameter that controls whether the observer prioritizes the speed (α near 0) or accuracy (α near 1) of inference. Once optimized, a given value of α can be used to generate an error trajectory that measures the average final error at each stopping time against the classic SAT Fig. 4(c).

When optimized in this way, the adaptive stopping rules generate lower average error than the SAT for the same average stopping time Fig. 4(d). This improvement arises because adaptive stopping allows for a faster convergence of some error trajectories, thereby altering the underlying distribution of trajectories that converge on later trials—something that we dissect in the following section. This faster convergence is visible in the underlying distributions of stopping times [Fig. 4(e)], which are qualitatively different than the SAT: for α near 0, the observer stops quickly to prioritize inference speed; for α near 1, the observer tends to wait longer to prioritize inference accuracy [Fig. 4(e)]. When averaged at each stopping time, the resulting error trajectories strongly violate the SAT, maintaining near-constant errors until the time limit, when the error abruptly increases [Fig. 4(f)].

This behavior is even more apparent when we decompose inference trajectories into those that fell below the uncertainty threshold before or at the time limit (“early stops”) versus those that were forced to stop at the time limit (“late stops”) [Figs. 4(g) and 4(h), shown for $\alpha = 0.8$]. Early stops generate low errors whose distribution is consistent across time [Fig. 4(g), orange points]; for this task complexity ($N = 5$ latent states), the distribution is bimodal, corresponding to scenarios in which the observer was sufficiently certain to stop the inference process but was either correct (low error) or incorrect (high error) in its estimation of the latent state [Fig. 4(h); orange distributions]. Late stops, in which the uncertainty did not drop below the threshold and the observer was forced by the time limit to stop the inference process, show a qualitatively different distribution of error [Fig. 4(h), brown distribution]. We observe qualitatively similar behavior for more complex tasks with a near continuum of latent states [Fig. 4(i)], approaching the continuous example in Fig. 1. In this regime, the underlying error distributions are no longer bimodal, but the error of early versus late stops nevertheless follow different distributions. Taken together, these results highlight that different types of stops implicitly carry information about the underlying error, and could thus be used to guide distinct types of actions.

E. Adaptive stopping rules exploit the statistics of individual inferences

To gain a more mechanistic understanding of this adaptive stopping rule, we examined how it altered the distributions of error and excess surprise over the course of inference [Fig. 5]. On a given timestep, different trials exhibit a wide range of errors and values of excess surprise. The adaptive stopping rule can be viewed as subselecting a fraction of trials from this full distribution [Fig. 5(a), bright orange distributions]; the remaining trials then constitute the distribution that can be subselected on the next timestep [Fig. 5(a), brown




FIG. 4. An adaptive stopping rule exploits predictable patterns of error. Gray box: Schematic of adaptive stopping criteria. (a) We design an adaptive stopping rule, captured by a time-dependent uncertainty threshold, that is used to terminate the inference process up until a maximum time limit of $T_{\max}$. If the observer's uncertainty drops below this threshold before $T_{\max}$, then the inference process terminates early. At $t = T_{\max}$, the inference process can be terminated either because the observer's uncertainty dropped below the threshold ("early" inferences), or because the observer reached the time limit ("late" inferences). (b) A given uncertainty threshold will generate a distribution of stopping times, and a distribution of errors at each stopping time. These are used to compute the final average normalized error e_{avg} and the average normalized stopping time t_{avg} (dashed lines). (c) We optimize uncertainty thresholds to minimize a cost $C(\alpha)$ that weighs the average inference error against the average inference time (low versus high values of α prioritize fast versus accurate inference, respectively). A given value of α will produce a distribution of errors and stopping times [as in panel (b)] that can be used to compute the average error at each stopping time (dotted black line). This error trajectory can deviate from the classic SAT (solid black line). (d)–(h) Performance using adaptive stopping rules optimized for $N = 5$ latent states and $T_{\max} = 10$ samples. (d) Higher values of α lead to lower average errors and higher average stopping times; for a given average stopping time, the average error achieved by the adaptive uncertainty threshold (colored circles) is lower than that achieved by the classic SAT (black line). Quantiles span 20%–80% of the distributions of errors (vertical quantiles) and stopping times (horizontal quantiles), measured across all 10^5 trials for a given value of α . The gray shaded region similarly spans 20%–80% of the distributions of errors for the classic SAT, computed at each possible stopping time and averaged across all 10^5 trials. The inset shows the optimal uncertainty thresholds for each value of α . (e) Different values of α generate different distributions of stopping times; higher values of α generate distributions that are shifted toward $T_{\max}$. (f) Adaptive stopping rules generate error trajectories that violate the SAT; average error trajectories are nearly flat for $t < T_{\max}$, and they increase abruptly at $t = T_{\max}$. Each point along a given curve is computed by averaging the final normalized error of all trials that stopped at the given stopping time. (g) Median inference errors for the adaptive strategy, separated into early stops (i.e., those trajectories that converged before or at $t = T_{\max}$) versus late stops (i.e., those that did not converge). Shown for $\alpha = 0.8$ (orange) and compared to the SAT (black). Quantiles span 20%–80% of the error distributions for each stopping time. The gray shaded region spans 20%–80% of the distributions of errors for the classic SAT. (h) Corresponding distributions of inference errors for $\alpha = 0.8$. The adaptive stopping rule generates a bimodal distribution of errors; a small set of high errors correspond to cases in which the observer is certain but of the wrong latent state; the




FIG. 5. The adaptive stopping rule exploits optimistic stimulus sequences that generate low uncertainty. (a), (b) Distributions of error (left column) and excess surprise (right column) for different stopping rules. (a) Under a fixed stopping rule based on the SAT, an ideal observer waits until a fixed time limit to stop inference. Across trials, the corresponding distributions of error (left column) and excess surprise (right column) evolve over time (black distributions) until this time limit. In contrast, adaptive stopping can lead the observer to converge before the time limit on a large fraction of trials (bright orange distributions). On these trials, the observer converges because it is sufficiently certain of its belief, which typically corresponds to a low error (left column) that is generated by an optimistic stimulus sequence (right column). On a small fraction trials, the observer is certain but of the wrong latent state; these trials have high error generated by pessimistic stimulus sequences. Once the observer's belief converges, the corresponding trial is removed from the distribution; inference continues on the remaining trials (brown distributions). Relative to a fixed stopping rule based on the SAT, adaptive stopping leads to a rapid convergence of those inference trajectories with the lowest error and lowest excess surprise. Pie charts show the fraction of all 10^5 trials that converged before the given time (white), at the given time (bright orange), or have not yet converged (brown). (b) Same as panel (a), but shown for $N = 100$ latent states.

distributions]. Consistent with the trial-to-trial structure observed in Fig. 3, the adaptive stopping rule—which is based only on the observer's uncertainty—tends to select those times and trials with the lowest excess surprise. These predominantly correspond to trials with low errors, but include a small fraction of trials in which the observer was certain but of the wrong latent state, and thus had high error. As time passes and the excess surprise tends toward zero, the adaptive stopping rule selects trials with higher (but still negative) excess surprise that nevertheless generate similar distributions of error. These results are qualitatively consistent across more complex inference scenarios [Fig. 5(b)]. Together, these results confirm that the adaptive stopping rule leverages trial-to-trial variability by differentiating optimistic versus pessimistic sequences, enabling the observer to make accurate inferences in far less time.

F. Adaptive stopping confers several key advantages

Adaptive stopping indirectly exploits the statistical structure of very brief stimulus sequences. This strategy has multiple advantages that increase the probability of accurate inferences and, as we will show, the actions that depend on them.

First, adaptive inference rules provide subjective information about the objective level of error. If an inference does not converge within the time limit, it is highly likely that the underlying stimulus sequence is pessimistic and the inference error is large. Alternatively, if an inference does converge, it is highly likely that the underlying stimulus sequence is optimistic and the inference error is small, regardless of *when* the inference converges [Figs. 4(f), 4(g) and 5(a), 5(b)]. This meta-knowledge is related to the notion of confidence [24] and can be exploited by the observer to plan appropriate actions. For example, in planning an escape, animals can freeze if

large bulk of low errors correspond to cases in which the observer is certain of the correct latent state. This distribution is similar for the set of early inference trajectories that terminated at different stopping times (orange distributions); in contrast, the set of late inference trajectories that did not drop below the uncertainty threshold during the time $t \leq T_{\max}$ exhibit a different distribution of errors (brown). For comparison, the distribution of errors generated by the SAT at $t = T_{\max}$ is shown in black. (i) Same as panels (g), (h), but for $N = 100$ latent states.

time has run out and they remain uncertain of the direction of an approaching predator; alternatively, if their inference rapidly converges, they can plan coordinated actions under the assumption of a reliably accurate estimate of the approach direction.

Second, those inferences that converge within the time limit are substantially faster, and have substantially lower error, compared to a fixed strategy based on the SAT [Fig. 4(g)]. This is because they exploit random fluctuations in short stimulus sequences—information that is not used if the observer merely waits a fixed amount of time.

Third, adaptive stopping improves the efficiency of the underlying sensory processing that supports inference. Because an adaptive observer terminates inference as soon as their uncertainty has dropped below threshold, they use far fewer stimuli on average to achieve the same accuracy of inference as an observer that waits for a predetermined amount of time. As a result, the underlying sensory system has fewer stimuli to encode, process, and transmit.

To concretely illustrate the advantages of adaptive stopping, we couple our ideal observer to a decision maker that can select and initiate actions, and we use this to model a scenario that mimics escape behavior in the fruit fly, *D. melanogaster* [Fig. 6(a)]. We model the fly as an agent that has limited time to infer the direction of a looming predator and execute an escape in the opposite direction. We assume that this takes place in two stages: first, the model fly has a maximum time limit of $T_{\max}$ to infer the approach direction from noisy stimuli that signal one of N latent states; here, $T_{\max}$ serves as a proxy for the looming speed of the predator, whereby faster looming speeds impose stronger time limitations on inference. Second, after inferring the direction of approach, the model fly must execute an escape in the opposite direction before colliding with the predator. Here, we assume that there are tradeoffs between the timing, precision, and accuracy of the escape: we assume that more precise actions require more time to execute, and that more accurate actions allow more time before a collision with the predator [heatmap in Fig. 6(a); Methods]. As we discuss in more detail below, the former assumption is motivated by observations that flies exhibit two different modes of escape: a slow escape that takes more time to coordinate and leads to more stable flight (and is thus more precise), and a fast escape that takes less time to coordinate and leads to less stable flight (and is thus less precise) [8,25]. Under these assumptions, the model fly can successfully avoid a collision by slowly coordinating a precise and accurate escape away from the predator, or by quickly making an imprecise and inaccurate escape.

In such a scenario, an adaptive model fly can exploit the timing and expected error in its inference to improve the probability of a successful escape [left and middle columns in Fig. 6(b)]. If its certainty drops below threshold at or before the time limit, we assume that the model fly can initiate a slow escape [left column in Fig. 6(b)]. In this situation, the model fly believes that it has a correct estimate of the approach direction, and it can use its remaining time to plan and execute a precise escape; the earlier the inference, the more time it has to execute a sequence of actions, and the more precise the escape will be [Fig. 6(b), left column; note that execution time is higher, and the execution noise is lower, when the animal

is further from the time limit]. Alternatively, if the model fly reaches the time limit before its estimate has converged, it can initiate a fast escape. In this situation, the model fly is uncertain about the direction of approach, its estimation error is likely high, and the time for inference has run out. Rather than using its erroneous estimate, we assume that it instead initiates a “last ditch” fast escape that is taken in a random direction but at a more rapid speed [Fig. 6(b), middle column; note that execution time is shorter and the execution noise is flat, compared to a slow escape taken at the same time].

In contrast, a decision based purely on the SAT does not have information to guide different patterns of escape based on timing or error [Fig. 6(b), right column]. This leads to lower probability of a successful escape compared to the adaptive strategy, regardless of the time constraint or the complexity of the inference task [Fig. 6(c)].

G. Escape behavior in *Drosophila* exhibits signatures of adaptive stopping

The scenario described above [Figs. 6(a)–6(c)] highlights some of the key advantages of using an adaptive stopping rule to inform decisions and actions. To explore whether these advantages might be borne out in real escape behavior, we compared the properties of the underlying inference to measurements of escape behaviors in the fly. Importantly, these properties were intrinsic only to the inference process, and were independent of the model of escape decisions described above. In the fly, escape behavior has typically been studied by exposing flies to an expanding visual disk that mimics an approaching predator [8,20,25,26]. Flies exhibit two different modes of escape in response to such stimuli [20]: either a slow, precise escape that is preceded by an extended sequence of preparatory movements, or a fast, imprecise escape without extended preparation. These two different modes are mediated by different neural pathways [25], and they require different amounts of time to execute [Fig. 6(e), bottom; reproduced from Ref. [25]]. Slow escapes are accompanied by a longer period of wing extension, which in turn leads to a more stable takeoff; fast escapes, by contrast, are accompanied by quicker wing extensions that can lead to tumbles in flight [8]. Recovering from these tumbles is likely to require additional time and energy, and we thus assume that the fly would benefit from initiating a slow takeoff if time permits. Consistent with this idea, the fly is more likely to initiate a slow escape if the looming speed is slow, which allows more time to make a decision and execute a series of preparatory actions; as the looming speed increase, the fly is increasingly more likely to initiate a fast and imprecise escape [Fig. 6(e), top; reproduced from Ref. [25]]. Distinct modes of escape behavior have also been observed in other species, including in fish, where they are controlled by the Mauthner cell circuit [27].

As illustrated above, our theory naturally produces two distinct types of inference that could underlie these different modes of escape: “certain” inferences that fall below the uncertainty threshold within the time limit and can be used to guide slow, precise escapes, and “uncertain” inferences that could be used to trigger a fast, imprecise escape. We find that the probability of initiating a fast escape scales linearly with




FIG. 6. Adaptive stopping improves successful escape and reproduces qualitative features of fruit fly behavior. (a) We consider a scenario in which a model fly (i.e., our ideal observer) must infer the direction of an approaching predator and use the inference to guide an escape in the opposite direction. The inference must be completed with T_{max} timesteps; the execution of the escape requires additional time and precision. We compute the probability that the escape be completed before the model fly collides with the predator; we assume that collisions happen later if the model fly escapes away from the predator (grayscale colormap), and that more precise actions with lower execution noise—which are more likely to avoid the predator—require more time to execute (orange shaded distribution). (b) The adaptive stopping rule (left and middle columns) can be exploited to improve the precision and type of escape (rows and columns, respectively); more rapid inferences can be used to coordinate more precise escapes, as indicated by the narrower distributions at top. Similarly, early versus late inferences can be used to trigger coordinated escapes that are centered on the estimated escape direction but are slower to execute (left column), versus last-ditch escapes that are executed in a random direction but are fast (middle column). The SAT (right column) does not permit such flexibility. (c) Adaptive stopping improves the probability of a successful escape across a wide range of time constraints and inference tasks (red versus blue respectively denote an increase versus decrease in the probability of successful escape). (d)–(g) Statistics of escape behavior compared between the fruit fly, the adaptive stopping model, and the probabilistic SAT model, measured as a function of time constraint (upper row) and escape duration (bottom row). (d) Schematics illustrating the different factors that are compared in panels (e)–(g). (e) Escape behavior of *D. melanogaster* in response to rapidly looming visual stimuli (reproduced from Ref. [25]). Upper: fraction of fast and imprecise escapes as a function of the visual speed of the looming stimulus. Faster looming corresponds to lower r/v values, where r is the stimulus size and v is the looming speed. Lower: distribution of escape durations. Modes of the distribution correspond to two escape modes—fast and imprecise escapes of short duration, and

the logarithm of the time limit in a manner that resembles fly behavior [compare the upper panels of Figs. 6(f) and 6(e)]. Furthermore, we observe a bimodal distribution of escape durations, which in our model corresponds to the remaining time between the decision to escape and the time limit $T_{\max}$ [compare the lower panels of Figs. 6(f) and 6(e)]. Both of these features are consistent across a range of α values. We note that the time remaining for inference and the time remaining for action are partially confounded in existing experimental data; future experiments could disentangle their individual impact on escape behavior.

In contrast, these two distinct modes of escape behavior cannot easily be explained as a manifestation of the SAT. The SAT suggests that later actions—which would occur after a longer collection of sensory stimuli—should in principle be more accurate than earlier actions. Thus, according to the SAT, the only rational strategy would be to wait to initiate an escape until the last possible moment, when the fly’s inferences are most likely to be accurate. This would give rise to highly stereotyped behavior, in contrast to what is observed in real flies. To permit a more rigorous comparison between fly behavior and the SAT, we derived a family of probabilistic stopping rules that directly exploit the SAT (Methods). At each timepoint up until $T_{\max}$, the model fly probabilistically decides whether to continue performing inference, or whether to stop the inference process and initiate an escape. We assume the decision to continue performing inference is informed solely by the SAT curve; the higher the average error at a given timepoint, the more likely the model fly is to continue performing inference. To capture a variety of inference scenarios, we model this SAT curve as an exponential function parameterized by $\lambda \leq 0$. When λ is close to 0, the inference task is difficult, and the animal favors performing inference; this tends to force a large fraction of fast escapes at the time limit. For λ large and negative, the inference task is easy, and the animal favors initiating a slow and early escape. In contrast to the adaptive stopping rule, we find that this probabilistic SAT rule does not reproduce the behavior of real flies [Figs. 6(e) and 6(g)]; the fraction of fast escapes does not scale linearly with the logarithm of the time limit for any value of λ [Fig. 6(g), upper], and the resulting distribution of escape durations is bimodal for only a very narrow range of λ values [Fig. 6(g), lower].

Together, these results suggest that escape decisions in the fly might exploit principles of rapid inference similar to the adaptive strategy described here. In constructing this adaptive model, we assumed that any increase in the overall speed of inference was used to coordinate a more precise escape [Fig. 6(h)]; as a result, slow escapes would be initiated at earlier times compared to fast escapes, but they would be terminated at approximately the same time. This assumption is borne out in the behavior of real flies: slow escapes take longer to execute than fast escapes, but the timing of the final takeoff is similar for both types of escape [Fig. 6(i), reproduced from Ref. [25]].

While we do not attempt to model this process at the level of neural circuits that mediate escape behavior, we note that the full posterior distribution over possible approach directions could be dynamically computed using a relatively simple circuit architecture [Fig. 6(k)]. This computation relies on a set of tuning curves that approximate the log-likelihood of an incoming stimulus given possible approach directions, as has been proposed for probabilistic computations in the cortex [28]. This proposed circuit could dynamically compute and encode the posterior in the activity of neurons specific to each approach direction, which could occur through a sequence of neurally plausible operations. First, the output of these tuning curves is additively combined with the population firing rates transformed with a logarithmic nonlinearity. Second, the resulting approximation of the unnormalized log-posterior is transformed with exponential nonlinearities to encode unnormalized probabilities of each approach direction. In the final step, additive pooling and divisive normalization generate activity that encodes the normalized posterior probability of approach at each direction. Maintaining the full posterior over the approach directions enables the computation of all quantities relevant for our proposed theory, such as the estimated approach direction and the uncertainty in that estimate, which can easily be read-out from this population. The specific implementation of rapid inference in the fly brain remains an interesting subject for future work.

III. DISCUSSION

To survive, animals are faced with the daunting task of making life-threatening inferences based on a limited set of

slow, deliberate escapes of long duration. (f), (g) Escape behavior predicted by the adaptive stopping model (f) and the probabilistic SAT model (g). Upper: fraction of fast escapes as a function of the time constraint $T_{\max}$, which limits the maximum duration of inference and serves as a proxy for the looming speed. Colors correspond to different values of the trade-off parameter α (f) or the difficulty parameter λ (g). Lower: probability distribution of escape durations, measured as the time between the initiation of the escape and the time limit $T_{\max}=20$ samples. Insets depict probability distributions for different parameter values. (h) Our adaptive stopping model assumes that slow escapes are initiated when the inference converges before or at the time limit; all remaining time is used to coordinate a precise escape. Thus, the final takeoff times are the same for slow and fast escapes. (i) Flies show similar distributions of takeoff times for slow and fast escapes across a range of looming speeds (reproduced and modified from Ref. [25]). In the limit of very fast looms (small r/v), takeoffs from slow escapes lag slightly behind those of fast escapes. The inset shows the distribution of escape durations, analogous to the lower panel of e and colored by slow (gray) versus fast (black) escapes. (j) The assumptions of the adaptive stopping model, visualized in the same way as the data in panel (i). (k) Schematic diagram of a circuit that computes the posterior over approach directions. From the left: tuning curves instantiate the log-likelihood of the stimulus, $\log p(s_t|\theta)$; temporal feedback adds the log prior $p(\theta|s_{t<})$; an exponential nonlinearity generates a non-normalized posterior distribution $p(s_t|\theta)p(\theta|s_{t<})$; a sum and divisive normalization together generate the normalized posterior $\frac{1}{Z} p(s_t|\theta)p(\theta|s_{t<})$; each output neuron encodes the belief about each possible approach direction given the stimulus history $p(\theta|s_{t\leq})$. The posterior, transformed by a log-nonlinearity, is then fed-back to update the prior at the next time-step.

noisy stimuli. In this work, we identified statistical regularities that could be exploited in such dire situations. Our key observation is that on very short timescales, inference errors can strongly deviate from their average behavior. We demonstrated that the structure of such deviations could be used by an ideal observer to decide when to stop collecting stimuli to increase the speed of inference and bound the resulting error. We further showed that the behavior of such an adaptive observer can locally deviate from the classical SAT—a pattern that emerges on average. Finally, we showed how this adaptive strategy could be used to guide more effective escape behavior, and identified similarities with escape behavior in the fruit fly *D. melanogaster*.

This work highlights the relevance of maintaining not only point estimates of behaviorally relevant quantities, but also the “meta knowledge” captured by the observer’s perceptual uncertainty about the accuracy of those point estimates. The dynamics of the observer’s belief can provide useful information beyond the observer’s “best guess” about the state of the environment. Our observer uses this meta knowledge to decide when to stop gathering sensory stimuli and commit to an action. The problem of deciding when to stop collecting data during dynamic inference is often studied in the context of Bayesian optimal stopping [29] or sequential decision making [30]. These problems are typically solved through backwards induction and dynamic programming [31], and often require specifying a cost to taking an action [32]. Here, we consider an alternative approach that exploits the idiosyncratic properties of very brief sequences of observations, without invoking the downstream cost of taking actions based on those observations. Deviations of random variables from their asymptotic behavior fall within the purview of large deviation theory [23], but are typically analyzed using objective measures that are not available to an animal. Here, we contend with an animal’s need to make subjective inferences of such objective quantities. Moving forward, a synthesis of dynamic Bayesian inference and large deviation theory could be used to devise more sophisticated approaches for performing rapid inferences with very limited samples.

Our inference model captured only very basic aspects of rapid inference and escape planning in animals. Here, we assumed that the observer knows how much time is available for performing inference, and it can optimize its inference subject to this time limit. In real-world settings, the observer must infer this time limit from the same incoming stimuli that it uses to estimate other environmental properties, such as the direction of an approaching predator. Moreover, real-world inference must be performed using complex stimuli, such as images [33], sounds [34], or smells [16], that evolve over time and according to multiple interacting latent states. For example, an animal might need to infer the shortest path toward its shelter, together with the evolving trajectory of the predator [35]. However, despite the simplicity of our scenario, we identified short-timescale structure that could be exploited to increase inference speed; we therefore expect complex scenarios to exhibit even richer structure that could be exploited to further optimize rapid inference and planning. This simple scenario is the first normative perspective to capture several qualitative features of escape behavior in the fruit fly, including (i) the existence of two different modes

of escape based on two qualitatively different outcomes of inference, (ii) the relative propensity of each mode based on the time limitations of inference, and (iii) the distribution of escape durations made within each mode. If flies indeed exploit the short-timescale structure of incoming stimuli, we would additionally expect that the error of their inferences would be low for all long-mode escapes, regardless of whether these escapes were initiated earlier versus later in time. This would suggest that observed variability in the accuracy of the escape arises from the extent of preparatory movement, and not from the accuracy of the inference that preceded the movement. This highlights the challenge inherent in studying the dynamics of rapid inference and escape planning, namely that the time for performing inferences and for executing actions are both strongly limited. Under our interpretation, flies perform fast and inaccurate escapes because their perceptual estimate never converged, and thus they were left without any remaining time to execute a long, deliberate escape. One possibility is that the time allotted by the animal to plan the escape is shorter than the time to infer and act together. New experiments will be necessary to dissociate these and other factors. For example, it has been demonstrated that when planning escape routes, mice use heuristics that modularize the space of parameters that describe the escape path [36]. The theory presented here could be broadened to make predictions about specific features of actions that could result from rapid inferences.

Our approach is complementary to a long tradition in computational neuroscience of using drift-diffusion models (DDMs) to study the dynamics of decision making. Within this class of models, decision making can be understood as a specific form of categorical inference in which an observer must discern one of several (typically two) alternative distributions of the data [37]. These models reveal rich and diverse phenomenology (see, e.g., Ref. [10]) and can easily be augmented with features such as dynamic boundaries or urgency signals that allow the observer to control the expected duration of inference [32,38]. Moreover, the parameters of DDMs can be mapped onto neural quantities that can be directly observed in experiments [39]. In this work, we take a more general theoretical approach that can be more naturally extended to a broader range of scenarios, but comes at the cost of immediate points of comparison to neural data. The principles described here apply to any Bayesian inference scenario and are therefore not limited to the categorical inferences solved by DDMs. As such, they can easily be extended to any problem in which an observer must make a rapid decision based on an evolving posterior belief about a behaviorally relevant state of the environment, including multiple continuous latent states and stimuli. Moreover, we do not need to augment the ideal observer to control the duration of inference; rather, we ask how the observer could use its own posterior belief to exploit structure *intrinsic* to individual stimulus sequences, and we analyze the patterns of timing and error that would result from such a strategy.

The cost of this generality is that this work cannot, in its current form, be readily and directly mapped onto specific neural quantities, such as the firing rates of neurons during decision making under uncertainty [10,14,38]. Several lines of existing work have proposed neural implementations of

decision making that follow the SAT [9,10]. Yet other approaches consider neural representations of uncertainty that could underlie such probabilistic computations [40]. In flies, the neural basis of escape behavior has been mapped onto different descending pathways that mediate short- versus long-mode escapes [25]. Neural pathways dedicated to the control of escape behaviors are also known in mice [33] and fish [27,41]. Whether and how these pathways might be shaped by ongoing inferences in a manner that explains individual variability in behavior remains an interesting direction for future research.

Outlook. To survive, organisms must often exploit all sources of available information. In this work, we demonstrate that brief stimulus sequences contain enough information to allow organisms to tip the scales in favor of increasing speed during critical inferences that may decide life or death in a fraction of a second.

IV. METHODS

A. Time-constrained inference of the Gaussian mean

To illustrate structured variability on individual inference trials, we consider a simple Bayesian model of Gaussian mean inference, fully described in Ref. [42]. In brief, the observer receives i.i.d. stimuli s_t that follow a Gaussian distribution with an unobserved mean $\theta = 0$ and known variance $\sigma^2 = 9$. The posterior over the latent mean $p(\theta|s_{t \leq t})$ is also a Gaussian distribution, whose mean $\hat{\theta}$ is the minimal squared-error estimate of θ . We simulated $K = 10^6$ inference trials; on each trial, the observer begins with a Gaussian prior with mean $\mu_\theta = 0$ and variance $\sigma_\theta^2 = 1$. We then clustered the error trajectories using the standard k-means algorithm. We determined the convergence time of a trajectory to be the time point where the error is within 20% of the final error of that trajectory. To consider environments with different variability, we considered three different values of the variance σ_θ^2 of the latent mean θ : $\sigma_\theta^2 = 1, 2, 4$. We computed the probability of low-error sequences p_{le} as: $p_{le} = \prod_{t=1}^{T_{\max}} p(\mu - \epsilon \leq s_t \leq \mu + \epsilon)$, where ϵ is the distance from the mean μ . The probability ratio is displayed in Fig. 1(g) and was computed as: $\log(p_{le}/(1 - p_{le}))$. We considered 100 values of ϵ spaced linearly between 0.1σ and the radius of equal probability (i.e., the value of ϵ for which the probability of the bulk and the tails of the distribution were equal).

B. Time-constrained inference of the direction of an approaching predator

We consider a scenario in which an ideal observer must infer the direction of an approaching predator from one of N possible angles of approach $\theta_n \in \{\theta_1, \dots, \theta_N\}$ [Figs. 2(a) and 2(b)]. To infer this direction, the observer relies on a brief sequence of noisy sensory stimuli $(s_1, s_2, \dots, s_{T_{\max}})$, where $T_{\max}$ limits the total inference time. We assume that the observer begins with a uniform prior over possible approach angles. On a given trial, i.i.d. sensory stimuli are drawn from a von Mises distribution with a fixed latent state θ_n ; i.e., $p(s_t|\theta_n) = \text{vonMises}(\theta_n, \kappa)$, where $\kappa = 1$ is the noise level, and where $\theta_n \in \{\theta\}$ is chosen randomly on each trial [Fig. 2(c)]. At each time step t within a trial, the observer updates the

posterior distribution $p(\theta|s_1, \dots, s_t)$ and computes the estimated direction $\hat{\theta}_t = \bar{p}(\theta_n|s_1, \dots, s_t)$, where $\bar{p}$ denotes the circular mean of the posterior p , and is computed using the function *circmean.m* from the MATLAB Circular Statistics Toolbox [43]. We compute the inference error at time t as $E(t) = \min (\theta + 2\pi n - \hat{\theta}_t)^2$, where minimizing over $n \in [-1, 0, 1]$ accounts for possible wrapping errors. To compare across different inference scenarios, we compute the normalized error $\bar{E}(t) = E(t)/\langle E(1) \rangle_K$, where the normalization is performed relative to the average initial error at time $t = 1$, computed over a set of K inference trajectories.

1. Quantification of error dynamics

We performed inference using $N = [3, 4, 5, 6, 7, 8, 100]$ latent states distributed uniformly on a circle. For each number of latent states, we simulated $K = 10^5$ trials of inference over $T = 30$ timesteps. Unless otherwise stated, all results in the main text were shown for $N = 5$ latent states using a maximum inference time of $T_{\max} = 10$. In Fig. 3, we used k-means clustering to cluster error trajectories into 6 clusters; we chose this number for illustration, but our results were qualitatively consistent for different numbers of clusters. We used Eq. (1) to compute the excess surprise of a given stimulus sequence up to and including time t . To compute uncertainty, we used the entropy of the posterior at time t :

$$H_t(\theta|s_{t \leq t}) = - \sum_{n=1}^N p(\theta_n|s_{t \leq t}) \log p(\theta_n|s_{t \leq t}), \quad (3)$$

where $s_{t \leq t} = s_1, \dots, s_t$ is the specific sequence of sensory stimuli up to and including time t . Note that $H_t(\theta|s_{t \leq t})$ refers to entropy of the posterior at a time-point t , and not to the conditional entropy of θ given a t -samples-long sequence of stimuli. When analyzing distributions of uncertainty and excess surprise, as shown in Fig. 3(e), we split the inference trajectories into the bottom and top quartiles of uncertainty [Fig. 3(e), lower panel] or excess surprise [Fig. 3(e), upper panel], computed at time $t = T_{\max} = 10$.

2. Optimization of adaptive stopping rules

We parameterized uncertainty thresholds using the following functional form: $U(t; \bar{p}) = p(1) + p(2)\exp(p(3)t)$. We optimized the parameters $\bar{p}$ to minimize the cost function given in Eq. (2), using 21 different values of α evenly spaced between and including 0 and 1. Given a fixed ensemble of K inference trajectories, each parameter setting will cause a fraction $K_{t_{\text{stop}}}$ of these trajectories to fall below threshold at a time $t_{\text{stop}} \leq T_{\max}$. Any trajectories that did not fall below threshold before or at time $T_{\max}$ were stopped at $T_{\max}$. Figure 4(d) shows the average normalized error e_{avg} and stopping time t_{avg} for different α ; these are the values that are optimized in Eq. (2) for a given α .

To examine violations in the speed-accuracy tradeoff, we used the distribution of trajectories that stopped at each value of t_{stop} to compute the average normalized error as a function of time, $\langle \bar{E}(t_{\text{stop}}) \rangle_{K_{t_{\text{stop}}}}$. These curves are shown in Fig. 4(f), together with the corresponding distributions of stopping times [Fig. 4(e)]. We then decomposed inference trajectories into those trajectories that converged before or at $t = T_{\max}$ ("early

stops”), and those that did not (“late stops”). Within each group, we computed the median, 20%, and 80% quantiles of the distribution of normalized error at each stopping time [Fig. 4(g)]; the full error distributions are shown in Fig. 4(h).

C. Simplified model of escape via adaptive stopping

We built a simple model to illustrate the computational advantages of adaptive stopping. We note that due to multiple factors that determine successful escape (the duration of the escape, the timing of inference, and the accuracy of the latent state estimate), this scenario would be difficult to model using standard methods of Bayesian decision theory. Analogously to the results presented in Fig. 4, we simulated an ideal observer that estimates the direction of an approaching predator, and uses this estimate to initiate an escape in the opposite direction. The ideal observer uses an adaptive uncertainty threshold to determine when to stop performing inference and initiate an escape; if the inference converges before or at the time limit $t = T_{\max}$ (“early stops”), the observer initiates a slow escape; otherwise (“late stops”), the observer initiates a fast escape. We assume that slow escapes are more precise but require longer time to execute; in contrast, fast escapes are imprecise but are fast to execute. We model this by assuming that fast escapes are executed in a random direction and cost a time $\Delta t_{\text{fast}} \sim P(\Delta t; \lambda_{\text{fast}})$, where we take $P(x; \lambda) = \exp(-x/\lambda)/\lambda$, and $\lambda_{\text{fast}} = 0.02$. We assume that slow escapes cost a time $\Delta t_{\text{slow}} \sim \exp(\Delta t; \lambda_{\text{slow}}(t_{\text{stop}}))$. We assume that $\lambda_{\text{slow}}(t_{\text{stop}}) = T_{\max} - t_{\text{stop}} + C_{\text{slow}}$, where C_{slow} is a constant; i.e., the average escape timescales with the duration between the time at which the inference converged and the maximum time limit. We take $C_{\text{slow}} = 0.1$, such that slow escapes at $t = T_{\max}$ cost more time than fast escapes.

In addition to differences in duration, we assume that slow and fast escapes differ in their precision. In contrast to fast escapes, which are executed in a random direction, we assume that slow escapes are initiated at the direction specified by the inference process, but that they are corrupted by execution noise whose variance scales with $T_{\max} - t_{\text{stop}}$. In other words, for early inferences that stop well before the time limit, the observer has a long amount of time to coordinate an escape, and the execution noise is low; for late inferences that stop near the time limit, the observer has little time to coordinate an escape, and the execution noise is high. We use this to assume that the executed escape direction is $\theta_{\text{exec}} \sim \text{vonMises}(\pi - \hat{\theta}, \kappa(t_{\text{stop}}))$, where $\hat{\theta}$ is the observer’s estimate of the direction of predator approach. We take $\kappa(t_{\text{stop}}) = 1 + T_{\max} - t_{\text{stop}}$; i.e., the variance increases as t_{stop} approaches $T_{\max}$.

Given an executed escape direction θ_{exec} and duration t_{exec} , we compute the probability that the escape is executed within a time limit that depends on orientation from the predator: $P(\text{successful escape}) \equiv P(t_{\text{exec}} < t_{\text{thresh}}(\theta_{\text{exec}}))$, where $t_{\text{thresh}}(\theta_{\text{exec}}) = \text{vonMises}_s(\pi - \theta - \theta_{\text{exec}}, \kappa_{\text{thresh}})$; in other words, t_{thresh} is largest (and thus there is the highest probability of escape) when the escape is executed in the opposite direction as the predator (i.e., when $\theta_{\text{exec}} = \pi - \theta$). Here, vonMises_s denotes that the von Mises function is scaled by its maximum value, such that $t_{\text{thresh}}(\pi - \theta) = 1$. In practice, we use the joint distribution of escape directions and escape times to compute $P(\text{successful escape})$

for an individual inference trajectory; we then average this probability across $K = 100\,000$ inference trajectories for a given time limit and inference task. Figure 6(c) compares the performance of this model to an analogous SAT version that only uses slow escapes at time $t = T_{\max}$.

D. Comparisons to data

We compared models of dynamic inference to the behavior of *D. melanogaster* published in Ref. [25]. We used two behavioral features as a basis for comparison: (i) the probability of eliciting a fast escape as a function of the time constraint and (ii) the distribution of escape durations for a fixed time constraint.

1. Adaptive stopping model

Since neither of these two features was explicitly represented in our model, we assumed that our ideal observer initiated a fast and inaccurate escape if its inference did not converge within the time limit. As above, we simulated the inference process for $K = 10^5$ trials, and we took the fraction of trials that did not converge at or before $t = T_{\max}$. To estimate the “duration” of escape, we assumed that both early and late stops could be used to initiate an escape, and that the duration of this escape would last the remainder of the available inference time: i.e., duration = $T_{\max} - t_{\text{stop}}$, where t_{stop} is the time at which the inference converged. This means that the final execution time is taken to be the same for slow and fast escapes [Fig. 6(j)]. To validate this assumption with data, we took the results shown in the righthand panel of Fig. 5(a) in Ref. [25], and we vertically aligned each pair of errorbars (corresponding to different values of r/v) to the mean takeoff time of the fast escapes.

2. Probabilistic speed-accuracy tradeoff model

To construct a model for escape behavior based on the classic SAT, we assumed that the observer makes a decision at each time step t to either continue gathering data, made with probability $p_{\text{sample}}(t)$, or to stop the inference and initiate an escape, made with probability $p_{\text{escape}}(t) = 1 - p_{\text{sample}}(t)$. We used the SAT curve Fig. 2(d) to specify the probability $p_{\text{sample}}(t)$; the larger the average error at time t , the higher the probability that the agent will continue to sample. Because the average error is roughly exponential as a function of time, we assumed that $p_{\text{sample}}(t) = \exp(\lambda t)$, where λ is a “difficulty parameter.” For $\lambda = 0$, the inference task is difficult, and inference always continues until the time limit $T_{\max}$. For $\lambda \rightarrow -\infty$, the inference task is easy, and the agent stops sampling immediately after the first sample is observed. For a given value of λ , one can analytically calculate the total fraction of converged inference trials for which the agent stopped sampling and initiated an escape within the time limit $T_{\max}$:

$$P_{\text{conv}}(T_{\max}) = \sum_{i=1}^{T_{\max}} (1 - \exp(\lambda i)) \prod_{j=0}^{i-1} \exp(\lambda j), \quad (4)$$

and conversely, the fraction of nonconverged trials is given by

$$P_{\text{nonconv}}(T_{\max}) = 1 - P_{\text{conv}}(T_{\max}). \quad (5)$$

We used these calculations to plot the results in Fig. 6(g).

ACKNOWLEDGMENTS

We thank Gwyneth Card for useful discussions about escape behavior in flies. A.M.H. was funded by the Howard Hughes Medical Institute.

DATA AVAILABILITY

The code that supports the findings of this article is available from the authors upon reasonable request.

- [1] A. A. Stocker and E. P. Simoncelli, Noise characteristics and prior expectations in human visual speed perception, *Nat. Neurosci.* **9**, 578 (2006).
- [2] W. F. Młynarski and A. M. Hermundstad, Adaptive coding for dynamic sensory inference, *eLife* **7**, e32055 (2018).
- [3] G. Tavoni, T. Doi, C. Pizzica, V. Balasubramanian, and J. I. Gold, Human inference reflects a normative balance of complexity and accuracy, *Nat. Hum. Behav.* **6**, 1153 (2022).
- [4] G. Tavoni, V. Balasubramanian, and J. I. Gold, What is optimal in optimal inference? *Curr. Opin. Behav. Sci.* **29**, 117 (2019).
- [5] D. A. Evans, A. V. Stempel, R. Vale, and T. Branco, Cognitive control of escape behaviour, *Trends Cognit. Sci.* **23**, 334 (2019).
- [6] W. E. Cooper and D. T. Blumstein, *Escaping from Predators: An Integrative View of Escape Decisions* (Cambridge University Press, Cambridge, UK, 2015).
- [7] P. Domenici, J. M. Blagburn, and J. P. Bacon, Animal escapology I: Theoretical issues and emerging trends in escape trajectories, *J. Exp. Biol.* **214**, 2463 (2011).
- [8] G. Card and M. Dickinson, Performance trade-offs in the flight initiation of *Drosophila*, *J. Exp. Biol.* **211**, 341 (2008).
- [9] R. P. Heitz, The speed-accuracy tradeoff: History, physiology, methodology, and behavior, *Front. Neurosci.* **8**, 86875 (2014).
- [10] R. Bogacz, E.-J. Wagenmakers, B. U. Forstmann, and S. Nieuwenhuis, The neural basis of the speed-accuracy tradeoff, *Trends Neurosci.* **33**, 10 (2010).
- [11] M. Dambacher and R. Hübner, Time pressure affects the efficiency of perceptual processing in decisions under conflict, *Psychol. Res.* **79**, 83 (2015).
- [12] B. U. Forstmann, R. Ratcliff, and E.-J. Wagenmakers, Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions, *Annu. Rev. Psychol.* **67**, 641 (2016).
- [13] R. Plamondon and A. M. Alimi, Speed/accuracy trade-offs in target-directed movements, *Behav. Brain Sci.* **20**, 279 (1997).
- [14] J. Drugowitsch, G. C. DeAngelis, D. E. Angelaki, and A. Pouget, Tuning the speed-accuracy trade-off to maximize reward rate in multisensory decision-making, *eLife* **4**, e06678 (2015).
- [15] A. G. Mendonça, J. Drugowitsch, M. I. Vicente, E. E. DeWitt, A. Pouget, and Z. F. Mainen, The impact of learning on perceptual decisions and its implication for speed-accuracy tradeoffs, *Nat. Commun.* **11**, 2757 (2020).
- [16] D. Rinberg, A. Koulakov, and A. Gelperin, Speed-accuracy tradeoff in olfaction, *Neuron* **51**, 351 (2006).
- [17] L. Chittka, A. G. Dyer, F. Bock, and A. Dornhaus, Bees trade off foraging speed for accuracy, *Nature (London)* **424**, 388 (2003).
- [18] T. Latty and M. Beekman, Speed-accuracy trade-offs during foraging decisions in the acellular slime mould *Physarum polycephalum*, *Proc. R. Soc. B* **278**, 539 (2011).
- [19] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, UK, 2003).
- [20] G. M. Card, Escape behaviors in insects, *Curr. Opin. Neurobiol.* **22**, 180 (2012).
- [21] E. P. Simoncelli, Optimal estimation in sensory systems, *Chapter 36 of The New Cognitive Neurosciences*, 4th ed., edited by M. Gazzaniga (MIT Press, 2009), pp. 525.
- [22] T. M. Cover, *Elements of Information Theory* (John Wiley & Sons, New York, NY, 1999).
- [23] H. Touchette, A basic introduction to large deviations: Theory, applications, simulations, [arXiv:1106.4146](https://arxiv.org/abs/1106.4146).
- [24] F. Meyniel, M. Sigman, and Z. F. Mainen, Confidence as Bayesian probability: From neural origins to behavior, *Neuron* **88**, 78 (2015).
- [25] C. R. Von Reyn, P. Breads, M. Y. Peek, G. Z. Zheng, W. R. Williamson, A. L. Yee, A. Leonardo, and G. M. Card, A spike-timing mechanism for action selection, *Nat. Neurosci.* **17**, 962 (2014).
- [26] C. R. Von Reyn, A. Nern, W. R. Williamson, P. Breads, M. Wu, S. Namiki, and G. M. Card, Feature integration drives probabilistic behavior in the drosophila escape response, *Neuron* **94**, 1190 (2017).
- [27] R. Eaton, R. Lee, and M. Foreman, The mauthner cell and other identified neurons of the brainstem escape network of fish, *Prog. Neurobiol.* **63**, 467 (2001).
- [28] M. Jazayeri and J. A. Movshon, Optimal representation of sensory information by neural populations, *Nat. Neurosci.* **9**, 690 (2006).
- [29] Z. Dai, H. Yu, B. K. H. Low, and P. Jaillet, in *Proceedings of the International Conference on Machine Learning* (PMLR, 2019), pp. 1496–1506.
- [30] K. J. Arrow, D. Blackwell, and M. A. Girshick, Bayes and minimax solutions of sequential decision problems, *Econometrica*, *J. Econ. Soc.* **17**, 213 (1949).
- [31] A. E. Brockwell and J. B. Kadane, A gridding method for Bayesian sequential decision problems, *J. Comput. Graphical Stat.* **12**, 566 (2003).
- [32] J. Drugowitsch, R. Moreno-Bote, A. K. Churchland, M. N. Shadlen, and A. Pouget, The cost of accumulating evidence in perceptual decision making, *J. Neurosci.* **32**, 3612 (2012).
- [33] D. A. Evans, A. V. Stempel, R. Vale, S. Ruehle, Y. Lefler, and T. Branco, A synaptic threshold mechanism for computing escape decisions, *Nature (London)* **558**, 590 (2018).
- [34] W. Młynarski, Efficient coding of spectrotemporal binaural sounds leads to emergence of the auditory space representation, *Front. Comput. Neurosci.* **8**, 26 (2014).
- [35] U. Mugan and M. A. MacIver, Spatial planning with long visual range benefits escape from visual predators in complex naturalistic environments, *Nat. Commun.* **11**, 3057 (2020).
- [36] P. Shamash, S. F. Olesen, P. Iordanidou, D. Campagner, N. Banerjee, and T. Branco, Mice learn multi-step routes

- by memorizing subgoal locations, *Nat. Neurosci.* **24**, 1270 (2021).
- [37] S. Bitzer, H. Park, F. Blankenburg, and S. J. Kiebel, Perceptual decision making: Drift-diffusion model is equivalent to a Bayesian model, *Front. Hum. Neurosci.* **8**, 102 (2014).
- [38] S. Tajima, J. Drugowitsch, N. Patel, and A. Pouget, Optimal policy for multi-alternative decisions, *Nat. Neurosci.* **22**, 1503 (2019).
- [39] M. N. Shadlen and W. T. Newsome, Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey, *J. Neurophysiol.* **86**, 1916 (2001).
- [40] Á. Koblinger, J. Fiser, and M. Lengyel, Representations of uncertainty: Where art thou? *Curr. Opin. Behav. Sci.* **38**, 150 (2021).
- [41] I. Temizer, J. C. Donovan, H. Baier, and J. L. Semmelhack, A visual pathway for looming-evoked escape in larval zebrafish, *Curr. Biol.* **25**, 1823 (2015).
- [42] K. P. Murphy, Conjugate Bayesian analysis of the Gaussian distribution (2007), <https://www.cs.ubc.ca/~murphyk/Papers/bayesGauss.pdf>.
- [43] P. Berens, Circstat: A MATLAB toolbox for circular statistics, *J. Stat. Software* **31**, 1 (2009).