J Neurophysiol 132: 1608–1620, 2024.
First published October 9, 2024; doi:10.1152/jn.00181.2024

american physiological society

JNP JOURNAL OF NEUROPHYSIOLOGY

RESEARCH ARTICLE

Higher Neural Functions and Behavior

Rat movements reflect internal decision dynamics in an evidence accumulation task

Gary A. Kane,[1]* Ryan A. Senne,[2]* and Benjamin B. Scott[1]

[1]*Department of Psychological and Brain Sciences and Center for Systems Neuroscience, Boston University, Boston, Massachusetts, United States and* [2]*Graduate Program for Neuroscience, Boston University, Boston, Massachusetts, United States*

Abstract

Perceptual decision-making involves multiple cognitive processes, including accumulation of sensory evidence, planning, and executing a motor action. How these processes are intertwined is unclear; some models assume that decision-related processes precede motor execution, whereas others propose that movements reflecting ongoing decision processes occur before commitment to a choice. Here we combine two complementary methods to study the relationship between decision processes and the movements leading up to a choice. The first is a free-response pulse-based evidence accumulation task, in which stimuli continue until choice is reported, and the second is a motion-based drift diffusion model (mDDM), in which movement variables from video pose estimation constrain decision parameters on a trial-by-trial basis. We find that the mDDM provides a better fit to rats' decisions in the free-response accumulation task than traditional drift diffusion models. Interestingly, on each trial we observed a period, before choice, that was characterized by head immobility. The length of this period was positively correlated with the rats' decision bounds, and stimuli presented during this period had the greatest impact on choice. Together these results support a model in which internal decision dynamics are reflected in movements and demonstrate that inclusion of movement parameters improves the performance of diffusion-to-bound decision models.

NEW & NOTEWORTHY In this study we combine a novel pulse-based evidence accumulation task with a newly developed motion-based drift diffusion model (mDDM). In this model, we incorporate movement parameters derived from high-resolution video data to estimate parameters of the model on a trial-by-trial basis. We find that this new model is an improved description of animal choice behavior.

accumulation; decision-making; motion tracking; response time

INTRODUCTION

Perceptual decision-making is a complex and integral part of our interaction with the world. It involves the accumulation of sensory evidence, comparison with internal models, and subsequent decision commitment (1). This process allows us to select the appropriate motor response based on the evidence available. Traditional models conceptualize perceptual decision-making as a serial process in which evidence accumulation precedes decision commitment and motor execution (2). More recent studies suggest an alternative model in which evidence accumulation and motor actions occur in parallel (3–6). In this model, aspects of ongoing motor activity may reflect underlying cognitive dynamics (5–8). Reconciling these models has been challenging, and the relationship between motor actions and cognitive processing remains unclear (9, 10).

To better understand how movement is related to the decision-making process, we developed and applied two complementary methods. The first is a free-response evidence accumulation task for rodents. In this task, inspired by previous pulse-based accumulation tasks (11–14), agents observe a series of flashes from two choice ports. Selection of the choice port associated with the higher flash probability is rewarded with a drop of sugar water. This task has two key features: rats are free to move at any time of the cue period,

*G. A. Kane and R. A. Senne contributed equally to this work.
Correspondence: B. B. Scott (bbs@bu.edu).
Submitted 29 April 2024 / Revised 17 September 2024 / Accepted 2 October 2024

i.e., no head or nose fixation is involved, and evidence is presented from trial initiation until a choice is reported. This design allowed us to characterize animal movements and evidence accumulation throughout the duration of a trial and to study how movements and decision processes are interrelated. The second method is a drift diffusion model (DDM), which incorporates pose estimation to constrain decision parameters on a trial-by-trial basis. In this method, termed mDDM, neural network-based video pose estimation is used to track the movements of a decision-maker moment by moment. These movement parameters are then used to constrain decision parameters of the DDM on a trial-by-trial basis.

We trained male and female rats to perform the free-response task in an automated operant training facility. Rats learned the task quickly, exhibited high accuracy, and showed hallmarks of evidence accumulation, including response time (RT) distributions that are well fit by a traditional DDM. In parallel, we recorded behavioral videos of rats performing the task and then fit the mDDM. Using this approach, we find that the mDDM outperforms traditional DDMs, based on standard model comparison metrics, and provides a trial-by-trial estimate of decision parameters including starting point, nondecision time, and decision threshold. Interestingly, we also observed a period of relative immobility, before movement onset, and that evidence presented during this period was a better predictor of subsequent choice than later evidence. Together these results are consistent with a model of decision-making in which movements reflect the internal dynamics of the decision process but there exists a period of deliberation, before movements, when most evidence is accumulated. In addition, the task and modeling approach we describe is well suited for studying the decision-making process and could be paired with neural recordings in future studies to further characterize perceptual decision-making in rats and other species.

MATERIALS AND METHODS

All experiments and procedures were performed in accordance with protocols approved by the Boston University Animal Care and Use Committee. Long Evans adult rats ($N = 13$; aged 3 mo to 2 yr) were purchased from Taconic or bred in house. Both male ($N = 3$) and female ($N = 10$) rats were used and trained in the same room at the same time. Rats were food restricted to 90–100% of their body weight and fed once per day (typically 3 or 4 pellets of food per day). They received 0.025–0.04 mL (depending on the session and rat) of 10% sucrose (100 g/L) as rewards. Reward volume was consistent within a session. Daily trained rats were housed on a 14:10-h ON:OFF light schedule, with the ON phase corresponding to daylight hours in Boston, MA.

Behavioral Control System

The behavioral training was performed in a high-throughput semiautomated facility (15, 16). Rats were trained in custom acrylic chambers with three nose ports. Nose ports were three-dimensional (3-D) printed (Sanworks or custom made) and equipped with a visible LED for stimulus delivery (Sanworks), a peristaltic pump for reward delivery, and an infrared (IR) LED and photodetector as a beam break (Sanworks). Behavioral control software to implement the task and control individual boxes was written in MATLAB. Boxes were controlled through a Teensy-based microcontroller system (Bpod Sanworks). A custom-written Python application was used to control multiple Bpod instances from a single control computer, requiring an edited version of the Bpod MATLAB software library (edited Bpod library: https://github.com/RatAcad/Bpod_Gen2; custom Python application: https://github.com/RatAcad/BpodAcademy). Further details about the software implementation can be found in the respective code repositories. The floor of the chambers contained bedding that was changed after each session.

Daily Training

Rats were housed in pairs in an animal facility and moved to the training room in the laboratory each morning for training. Rats ran for 2-h shifts, 5 days per week in the late morning (10:00 AM–noon) or early afternoon (1:00 PM–3:00 PM). Feeding was conducted in the late afternoon after behavioral training (1:00 PM–4:00 PM). Each chamber also contained a video camera mounted above the chamber.

Training Pipeline

Rats were rewarded with 10% sucrose at all stages of training and testing. Training took 1–4 wk depending on the rat and progressed through three stages. In *stage 1* (1–3 days), rats were rewarded for inserting their nose into a side port with an illuminated LED. In *stage 2* (1–21 days), rats received reward for inserting their nose into the center port and then the side port with an illuminated LED. In *stage 3*, they received a reward for inserting their nose into the center port and then the side port with a flashing LED. Once rats reached criterion 80–100% correct, the probability of flashes on the incorrect side increased in the following way 100:0 → 90:10 → 80:20. Progression through this third stage took 1–2 days per condition. After completion of all stages and conditions, rats performed the task with a flash probability of 75:25.

Behavioral Task

At the start of a trial, the light in the center port turned on, indicating that the rat could initiate the trial at its leisure (Fig. 1*A*). Once the rat nose poked in the center port, a simultaneous flash would occur in the left and right ports, signaling the start of the trial. After this one light flash on both sides, the rat would see a series of light flashes according to a Bernoulli process: every 100 ms, the rat would see a flash on either the right or the left side with a probability of 75:25 that the flash would occur on the correct side versus the incorrect side; the correct side was drawn randomly on each trial. To record a response, the rat nose poked in either the left or right port, and the series of light flashes would terminate as soon as this decision was recorded. If the rat responded correctly, the light in the correct side port turned on for 2 s and a 30-μL sucrose water reward was delivered immediately. If the rat responded incorrectly, all lights turned off and no reward was delivered, and new trials would start after a 2-s delay. For daily testing, rats participated in 2-h sessions Monday–Friday.

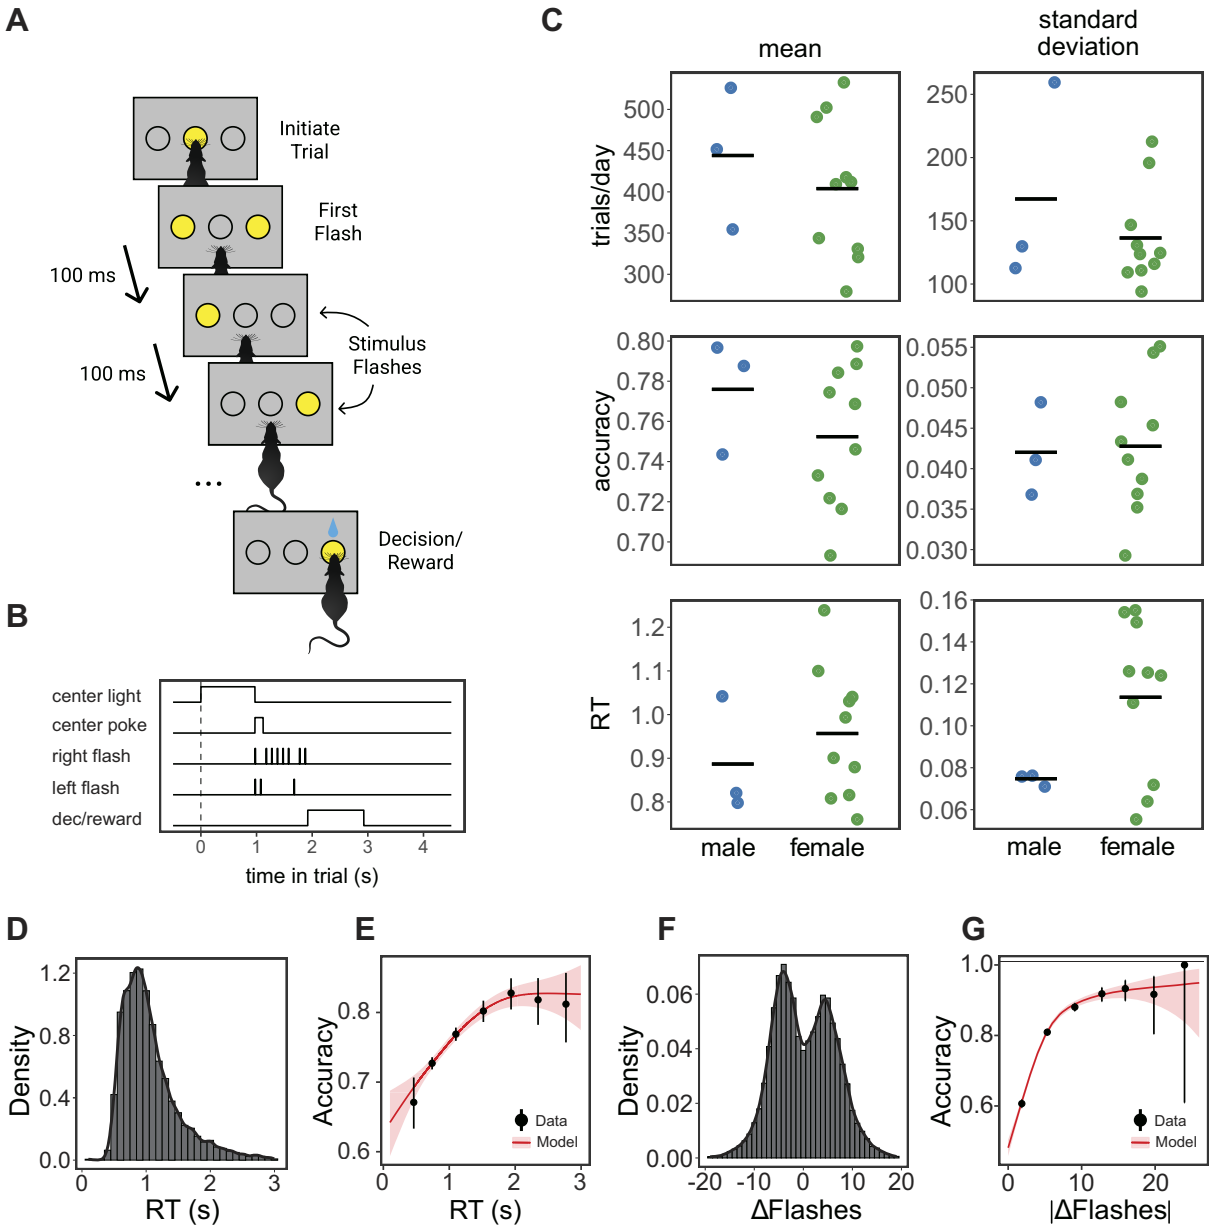

Figure 1. A free-response pulse-based accumulation task for rats. *A*: schematic of the behavioral task. Rats initiate a trial by inserting their nose into the center port of a 3-port operant training chamber. After initiation, rats are presented with a series of flashes from the left and right ports according to a Bernoulli process. The trial ends when rats insert their nose into 1 of the 2 side ports. Rats obtain reward if they choose the side with the higher flash probability. *B*: timing of events in an example behavioral trial. *C*: behavioral performance of rats performing the task. Colored dots represent individual animals ($n = 3$ male, blue dots; 10 female, green dots) Rats exhibit similar trial numbers [1-way repeated-measures (RM) ANOVA, $P = 0.458$, $F = 0.592$] (*top*) and accuracy (1-way RM ANOVA, $P = 0.469$, $F = 0.564$) and response time (RT) (1-way RM ANOVA, $P = 0.393$, $F = 0.789$) (*bottom*) across sexes. *Left*: the mean across days for individual rats. *Right*: the standard deviation across days for individual rats. *D*: histogram of RTs across all rats. *E*: on trials with greater RTs, rats exhibited greater accuracy. Black dots indicate data; red line indicates best fit from a generalized additive mixed model. *F*: histogram of flash differences (#R − #L) across the cue period of all trials. Positive values represent greater rightward flashes. *G*: accuracy increases with difference in flashes (absolute value of #R − #L). Black dots indicate data; red line indicates best fit from a generalized additive mixed model.

Diffusion Models of Decision-Making

Two diffusion-to-bound models were applied to behavioral data: the extended drift diffusion model (17) and a novel pulsed evidence drift diffusion model based on Ref. 11.

Both modes assume that evidence, represented by the decision variable (x), changes over time until it reaches a decision bound. The magnitude of the decision bound (a), which represents the total evidence necessary for an agent to come to a decision, is a free parameter. In addition to the decision bound (x), both models share two additional parameters that govern the how x changes over time:

1) starting point, x_0, which represents the initial value of the decision variable (x) and captures the side bias of the animal
2) drift-rate, v, which represents the strength of the evidence

The change in the decision variable (dx) over time is governed by the following equation:

$$\mathrm{d}x = v_i \cdot \mathrm{d}t + cdW$$

where cdW is the Gaussian noise term that represents the diffusion of the decision variable over time. The extended drift diffusion model also allows for random trial-to-trial variability in the drift rate and the starting point:

$$v_i \sim N(v, \theta_v)$$

$$x_{0_i} \sim \mathrm{Unif}(x_0 - \theta_{x_0}, x_0 + \theta_{x_0})$$

Once x reaches the decision bound, a decision is reached (a for right choices, $-a$ for left choices). The time taken for x to reach the decision bound is the decision time. The reaction time of the animals is the decision time plus another free parameter, the nondecision time (ndt). The ndt is typically interpreted to include the time it takes for processes that are not related to evidence accumulation, such as motor movements to the decision target (18). The ndt can vary from trial to trial based on the following equation:

$$\mathrm{ndt}_i \sim \mathrm{Unif}(\mathrm{ndt} - \theta_{\mathrm{ndt}}, \mathrm{ndt} + \theta_{\mathrm{ndt}})$$

In the pulse diffusion model, evidence is accumulated per unit of stimulus, where s_t is the stimulus at time t: $s_t = 1$ during a right pulse, $s_t = -1$ during a left pulse, and $s_t = 0$ otherwise, as opposed to continuously in time.

$$\mathrm{d}x = v_t \cdot s_t \cdot \mathrm{d}t + cdW$$

Additionally, consistent with the pulsed evidence accumulation model in Ref. 11, the drift rate does not vary trial by trial, but it drifts over the course of the trial:

$$v_t \sim N(v, \theta_v)$$

Furthermore, in this model, variability in the starting point across trials follows a normal distribution rather than a uniform distribution:

$$x_{0_i} \sim N(x_0, \theta_{x_0})$$

For model comparison, we used the Bayesian information criterion (BIC) calculated as

$$\mathrm{BIC} = k \cdot \ln(n) - 2 \cdot \ln(\hat{L})$$

where k is the number of free parameters, n is the number of data points, and $\hat{L}$ is the likelihood of the fitted model.

Video Recording and Image Processing

Video was recorded during a subset of daily testing sessions with USB webcams. Videos were recorded at 30 fps with the same custom Python application used to control behavioral boxes (https://github.com/RatAcad/BpodAcademy). To synchronize video data with behavioral data, the time of each camera frame and a transistor-transistor logic (TTL) signal indicating the start of the trial (center light turned on) and trial initiation (nose poke in center port) were both recorded with the Python *time* package. TTL signals were sent from the Bpod State Machine and recorded using serial input from a Teensy 3.2 microcontroller.

The position of the rat during a behavioral session was extracted from the video data with DeepLabCut [Mathis et al. (19)]. Experimenters identified the position of the ears, nose, and back of the head (see Fig. 3*A*) on 620 frames across five animals, which were used as training data (95% of frames) and validation (5% of frames). A ResNet-50 base neural network was used for 160,000 training iterations with a test error of 6 pixels (640 × 480-pixel images). A p-cutoff of 0.75 was used to condition (X, Y) coordinates for future analysis.

To standardize rats' positions across sessions despite slight differences in camera position and angle in different boxes, keypoints were first corrected for any translation and rotation with the following procedure:

1) The (X, Y) coordinates of each nose poke were defined as the average coordinate of the nose at the time of each nose poke (e.g., to find the center nose poke, the average of the nose position on frames that aligned in time with rats' nose poke in the center port).
2) All keypoints were translated such that the center nose poke was at coordinate (0,0).
3) Keypoints were rotated and scaled such that the left and right nose poke positions were at coordinates (−1, 0) and (1, 0), respectively.

Finally, two keypoints of interest were computed: *1*) the center of the head as the center point in between the nose, back of head, and between the ears and *2*) the angle of the head relative to the nose poke wall as the angle of the intersection between a line drawn from the nose and back of head and a line drawn between the nose pokes.

Sinusoidal Movement Model

To simplify the complex movement trajectories of rats (typically 100–300 observations per trial), we fit a sinusoidal model to each rat's head position for every trial. This model reduces the dimensionality of the data while capturing key aspects of the rat's movement. In this model we define the following five free parameters:

1) m_delay (movement delay): the starting point of the sinusoidal curve, i.e. how much time elapses before the animal starts moving
2) p_offset (position_offset): the offset of the curve from the horizontal center
3) A (amplitude): the amplitude of the sinusoidal curve
4) φ (starting phase): the initial phase of the sinusoidal curve; this parameter controls "where" on the sine wave the animal begins its movement
5) ω (sinusoidal period): how quickly the animal moves (i.e., did the rat accelerate more quickly to the port, indicative of a high value for ω)

Before movement (i.e., when the time in the trial is less than m_delay) the position of the animal's head is defined by p_offset. After movement the position of the animal's head is defined by

$$\mathrm{predicted\_trajectory} = \mathrm{choice} \times A \times \sin\left(\omega \times \mathrm{time} + \frac{\varphi}{\mathrm{d}t}\right) + \mathrm{p\_offset}$$

where "choice" is the rat's choice on the trial (coded as −1 for left, 1 for right), $\mathrm{d}t = 0.01$ is the resolution of the trajectories, and "time" is the elapsed time since m_delay. Thus, the free parameters of the model control the overall characteristics of the sinusoidal wave [e.g., how tall is the wave (A), how

shifted is the wave relative to center (p_offset), where in the sine wave do we start (φ), etc.).

Free parameters were estimated by maximizing the R^2 between the predicted and observed trajectories using generalized simulated annealing (GenSA R package; Ref. 20). Letting y_i be the ith datapoint in a dataset, $\hat{y}_i$ the estimated value for the ith datapoint, and $\bar{y}$ the mean of the dataset, R^2 is typically defined as

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2}$$

More intuitively, R^2 represents the proportion of variance in the dataset that can be explained by the model.

Lapse Logistic Regression and Hierarchical Bayesian Logistic Regression

The choice behavior was modeled with a modified logistic regression with a lapse parameter. The probability of choosing right, Pr(R), is given by

$$\Pr(\mathrm{R}) = \frac{\psi}{2} + \frac{1-\psi}{1+\exp(-(\beta_1 \Delta\mathrm{Flashes} + \beta_0))}$$

where ψ is the lapse parameter such that $0 \leq \psi \leq 1$, β_1 is the slope parameter controlling the sensitivity to the flash difference, and β_0 is a bias term. The model can be interpreted as follows:

1) With probability ψ, the animal "lapses," meaning that despite strong sensory evidence the animal makes an incorrect decision. This then implies that animals choose right on a lapse trial with probability $\frac{\psi}{2}$.
2) With probability $1-\psi$, the animal makes a choice based on a standard logistic function.

To then model the choice behavior of animals when the number of flash bins in the pre and post periods are identical (i.e., the same number of flashes occurred from trial initiation to when an animal moved and after movement until choice) we fit a hierarchical Bayesian logistic regression. Thus, we modeled the probability of an animal choosing right as

$$\mathrm{logit}(\Pr(\mathrm{R})) = (\alpha_0 + \beta_0) + (\alpha_1 + \beta_1) \times \Delta\mathrm{Flashes}$$

where β_0 and β_1 are the fixed effects (i.e., the population-level intercept and slope respectively) and α_0 and α_1 are the random effects (i.e., subject-specific deviations from the population-level effects). We used weakly informative priors for all parameters:

$$\beta_0 \sim N(0, 10)$$

$$\beta_1 \sim N(0, 10)$$

$$\alpha_0 \sim N(0, \sigma)$$

$$\alpha_1 \sim N(0, \sigma)$$

$$\sigma \sim \mathrm{HalfCauchy}(2)$$

The model was fit with the Bambi package (21) in Python, which utilizes the No-U-Turn Sampler (NUTS), a variant of Hamiltonian Monte Carlo. We ran 4 chains with 2,000 iterations each, including a 1,000-iteration warm-up period, resulting in 4,000 postwarm-up samples for each parameter. Convergence was assessed with the Gelman–Rubin statistic ($\hat{\mathrm{R}}$), ensuring all values were below 1.1, and by visual inspection of trace plots for each parameter.

Data Analysis and Model Fitting

Data analysis was performed in R version 4.3.1. Calculations of standard statistical tests [e.g., Kolmogorov–Smirnov (KS) tests, ANOVAs] were performed with the stats and aov packages. To fit the accuracy versus response time curve in Fig. 1*E*, we used the mgcv (22) package to fit a generalized additive mixed model. The beta regression was fit with the glmmTMB (23) package.

DDM and pulse DDM models were fit with a custom R package (rddm; https://github.com/gkane26/rddm). For comparison of the DDM and the pulse DDM, parameters were estimated by maximizing the quantile maximum probability estimate (24, 25) with differential evolution as the optimization routine (26). To estimate the influence of movement parameters from the sinusoidal model on trial-by-trial changes in DDM parameters, DDM model parameters were estimated by maximum likelihood estimation. Trial-by-trial likelihoods (as implemented in the rddm package) were calculated with the rtdists R package (27).

To fit our lapse parameter binomial generalized linear model (GLM), we wrote a custom link function in R and performed constrained optimization of the negative log-likelihood function using the L-BFGS-B algorithm in the optim function in R such that the lapse parameter was bounded between 0 and 1. For the optimization we used 10-fold cross-validation. To do this, we fit the model to each rat and did cross-validation (CV) within each animal's trial data.

All plots were generated with ggplot2 and matplotlib.

RESULTS

Rats Accumulate Evidence in a Free-Response Perceptual Decision-Making Task

We developed a free-response version of a visual pulse-based evidence accumulation task (12, 28, 29), which rats performed in a three-port operant chamber (Fig. 1*A*). To initiate the trial, the rat would nose poke in the center port. After trial initiation, lights on the left- and right-side reward ports flashed bilaterally, indicating the start of the cue period. During the cue period, the left- and right-side light ports flickered based on a Bernoulli process (Fig. 1*B*). The odds were set so that, in each 100-ms time bin of the cue period, the correct reward port would illuminate briefly (10 ms) with a 75% probability, whereas the incorrect port had a 25% chance of illuminating. During the cue period the animal was free to respond at any time. In contrast to previous versions of the pulse-based accumulation tasks, rats were not required to maintain fixation at the center poke and could move their head freely during the cue period (11–13). The trial ended when the rat poked its nose into one of the two side light pokes. If the poke with the higher generative probability was selected, the animal received a drop of sugar water (10% sucrose); if the opposite side was selected, the animal received no reward.

Rats were trained in a 2-h session once per day, 5 days a week. After progressing through a series of training stages

(MATERIALS AND METHODS), rats performed an average of 409 trials per session, reached high accuracy (76% correct on average), and exhibited long RTs (mean = 0.956 s), defined here as the time difference between the initiation of the trial and the animal reporting its choice (Fig. 1, *C* and *D*). We observed no differences between males or females in trials per session [1-way repeated-measures (RM) ANOVA, $P = 0.458$, $F = 0.592$], accuracy (1-way RM ANOVA, $P = 0.469$, $F = 0.564$), or RT (1-way RM ANOVA, $P = 0.393$, $F = 0.789$). Consistent with an evidence accumulation strategy, rats exhibited increased accuracy on trials with longer RTs (Fig. 1, *D* and *E*) and greater evidence favoring one side (Fig. 1, *F* and *G*).

Next, we sought to evaluate how well the rats' RT distributions and accuracy were fit by drift diffusion models (Fig. 2). To do this we fit two variants of the DDM, an extended DDM and a pulse-based DDM (11, 17). The extended DDM provided a good description of the rats' RT distributions (Fig. 2*B*) and indicated relatively high decision bounds, consistent with evidence accumulation. Next, to determine whether the model fits could be improved by considering the discrete nature of the stimulus, we used a pulse-based DDM (pDDM). The main difference in the pDDM is that evidence is only accumulated at the time of stimulus presentation and not continuously (MATERIALS AND METHODS). Both models explained the full distribution of choices and RTs equally, determined by comparing the negative log-likelihoods of the models (exact 2-sided KS test, $P = 0.9992$; Fig. 2*C*). Moreover, we did not observe significant differences in the magnitude of side bias or nondecision time between the DDM and pDDM (Fig. 2*D*). The pDDM did consistently predict larger drift rates and boundary parameter values (exact 2-sided KS test, $P_{\text{drift rate}} = 1.923\text{e}{-7}$, $P_{\text{boundary}} = 0.00275$; Fig. 2*D*); however, these increases were correlated such that both models predicted similar RT distributions (Fig. 2*B*). Taken together, these results suggest that rats use an evidence accumulation strategy to make decisions and that the pDDM and DDM provide an equally good description of rat RTs during the task.

Rats' Head Movements during Decision-Making Are Well Fit by a Sinusoidal Model

In our task rats are free to move their heads during the cue period. This provided an opportunity to study head movements during evidence accumulation. We recorded video from a subset of rats ($N = 6$), extracted their movement trajectories, and fit a five-parameter sinusoidal model to their head movements on each trial.

A ResNet50-based DeepLabCut network (19) was trained to track the nose, ears, and back of the head of rats as they performed the perceptual decision-making task (Fig. 3, *A* and *B*). The position of the center of the head and the angle of the head relative to the nose poke wall were derived from these keypoints. To characterize movement trajectories with a small number of interpretable parameters, we utilized a five-parameter sinusoidal model (see MATERIALS AND METHODS; Fig. 3*C*). The sinusoidal model was fit to the head movements of each rat on each trial and explained most of the data variance, with a median R^2 value of 0.9739 (Fig. 3*D*). Model fits declined slightly but significantly for longer-RT trials [mixed-effects beta regression; $P(\beta_{\text{RT}}) < 2\text{e}{-16}$) (Fig. 3*E*). However, the model provided a good description of both correct and error trials for both left and right choices (Fig. 3*E*). These results suggest that rats' head movements during perceptual decision-making in a three-port operant chamber are well captured by a five-parameter sinusoid model.

Rats' Movements Provide Information about Internal Decision Variables

A long-standing question in the field is whether movements before a decision are reflective of the latent decision-making process (7, 8). To address this question, we tested whether DDM models with movement information (i.e., mDDMs) could explain rat choices and RTs better than a DDM without movement information (i.e., a traditional DDM). We developed three different mDDM variants to test three hypotheses about the relationship between movement parameters and decision parameters (Fig. 4). Our first hypothesis was that, on each trial, the time spent moving to the side reward port positively correlates with nondecision time. This hypothesis was evaluated with a model where the nondecision time was a function of the movement time (defined as response time minus the delay in the sinusoidal model). We referred to this variant as the movement time DDM (mtDDM) (Fig. 4, *A* and *D*).

Our second hypothesis postulated that the delay before a movement onset predicts the decision boundary parameter. This hypothesis reflects the intuitive idea that if a rat spends more time waiting before starting to move, it spends more time accumulating evidence. To test this hypothesis, we used a model where the boundary was a linear function of the delay in the sinusoidal movement model. This variant was known as the movement delay DDM (mdDDM) (Fig. 4, *B* and *E*).

The third hypothesis proposed a correlation between the initial position of a rat's head and its side bias. For example, if a rat's head is slightly angled to the right at the start of a trial, it would be more likely to go right than if it were angled slightly to the left. We tested this hypothesis by using a DDM variant where the side bias parameter was a function of the side bias as measured by the sinusoidal movement model. We termed this model the movement orientation DDM (moDDM) (Fig. 4, *C* and *F*).

We compared the performance of each of these new models to the standard DDM. After fitting to rat behavioral data, all three movement-informed models (mtDDM, mdDDM, and moDDM) resulted in a lower Bayes information criterion (BIC) value compared to the standard DDM (Fig. 4*G*). However, we found that the mtDDM, which estimates the nondecision time parameter, resulted in the lowest BIC value. This suggests that the nondecision time variable is the most crucial for improving model fits, meaning that rats have significant variability in how much time they take to report their choices after movement onset, which could be indicative of other latent dynamics of the decision process. Overall, these results show that incorporating movement parameters into decision models can improve model fits.

Rats Weight Evidence before Movement Onset Most Heavily

Our model comparison approach suggests that movement time is correlated with nondecision time and that the delay

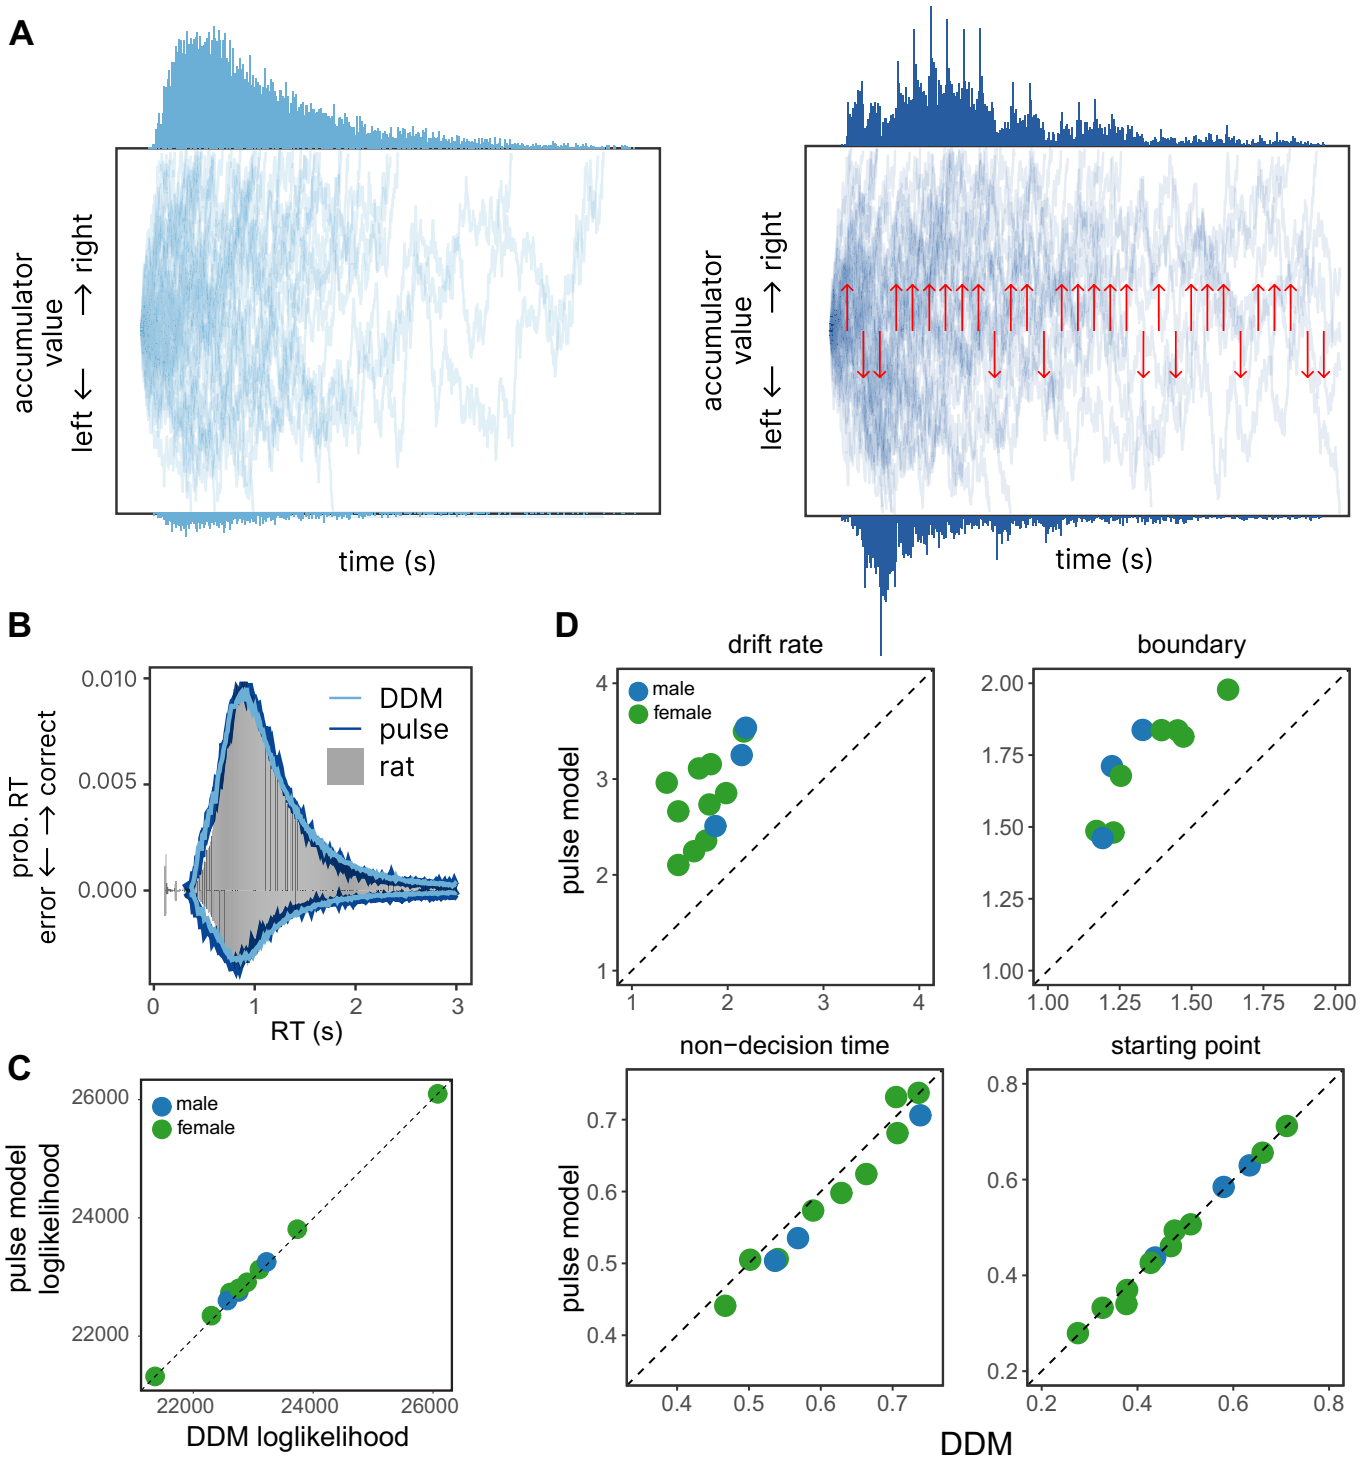

Figure 2. Response times in the accumulation task are well described by diffusion models. *A*: schematic of the accumulator value across 50 example trials as predicted by the drift diffusion model (DDM, *left*) and pulse-based model (*right*). The *top* histogram represents model-predicted response times to the upper boundary (i.e., decisions to go right); the *bottom* histogram represents model-predicted response times to the bottom boundary (i.e., decisions to go left). *B*: histogram of response times (RTs) for a single rat showing correspondence between the data (gray), DDM (light blue line), and pulse model (dark blue line). *C*: comparison of the goodness of fit (log-likelihood) for the pulse-based model (*y*-axis) and DDM (*x*-axis). Each dot represents a different animal. *D*: comparison of 4 parameters, drift rate, boundary, nondecision time, and starting point, between the pulse-based model (*y*-axis) and DDM (*x*-axis). Each dot represents an individual rat. Both models predict similar nondecision times and starting points, whereas the pulse model predicts a larger drift rate and boundary value.

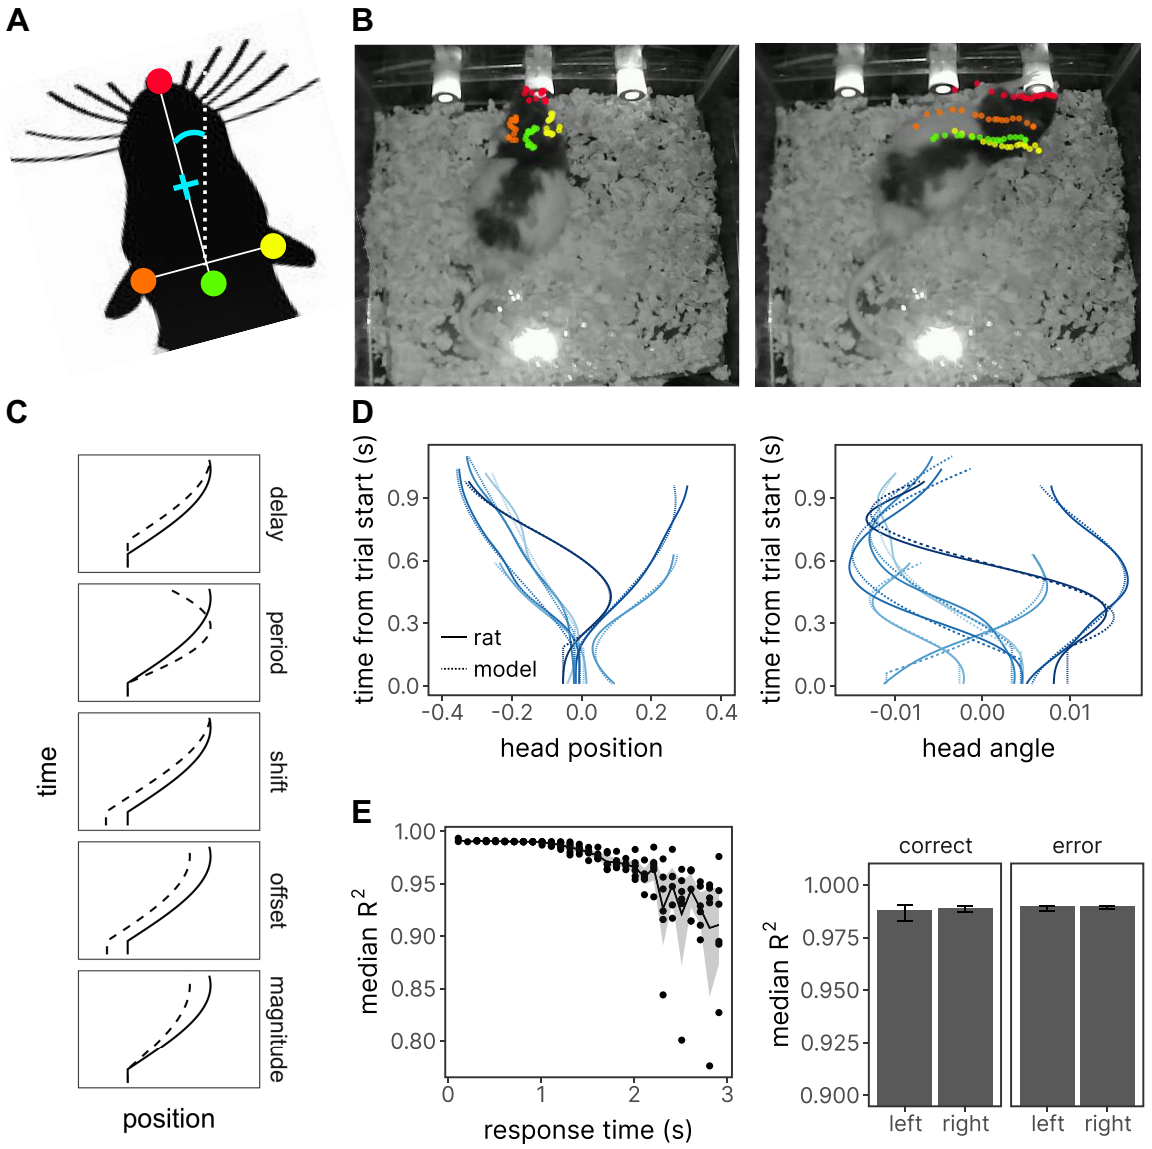

Figure 3. Rats' head movement trajectories during decision-making are well described by a 5-parameter model. *A*: schematic of the key tracking points on the rat head. Colored dots represent the 4 target features tracked by DeepLabCut: nose (red), left ear (orange), right ear (yellow), and back of the head (green). The teal plus sign and angle indicate the center of head and head angle, respectively, derived from these target features. *B*: example images from a single trial overlaid with the position of the target features described in *A*. *Left*: the rat entering the center port, with target features for the previous 10 frames (~300 ms) overlaid. *Right*: the rat entering the right port, with the target features for the previous 10 frames overlaid. Note the rightward trajectory of the animal indicated by the trail of dots reflecting the rightward choice of the animal. *C*: schematic showing an example trajectory of the rat's head position and how each of the 5 parameters alters the fit. Solid line indicates the rat's head movement on a single trial, and dashed line indicates how changes in each of the parameters alter the fit. *D*: trajectories for head position (*left*) and angle (*right*) position on 8 example trials. Solid blue lines indicate the data, and dashed blue lines indicate model fits. *E*: goodness of fit of the model (R^2) across trials with different response times (*left*) and across correct and error trials (*right*).

in movement onset is correlated with decision bounds. One intuitive explanation for this result is that the premovement period reflects a period of deliberation when the animal weights sensory information more strongly. To assess this hypothesis, we took advantage of the known timing of sensory pulses and measured the influence of stimuli before and after movement on the animal's decision using a binomial generalized linear model (GLM) (Fig. 5, *A* and *B*). We fit three distinct models: "All," "Before," and "After." The All model represents the differences in right versus left flashes across the entire trial period. The Before and After models consider only the differences in flashes before and after movement onset, respectively. To assess which model fit best, we used 10-fold cross-validation and employed negative log-likelihood as the objective measurement (Fig. 5*C*). Across all rats, we found that the All and Before models were similar and performed better than the After model.

Despite these results, it is possible that the After model performs worse simply because more trials had greater sensory evidence in the premovement period (Fig. 6*A*). Therefore, to further validate these results, we restricted the available data to trials with an equivalent number of flashes presented to the rat before movement initiation and after movement until choice (Fig. 6, *A* and *B*). First, we

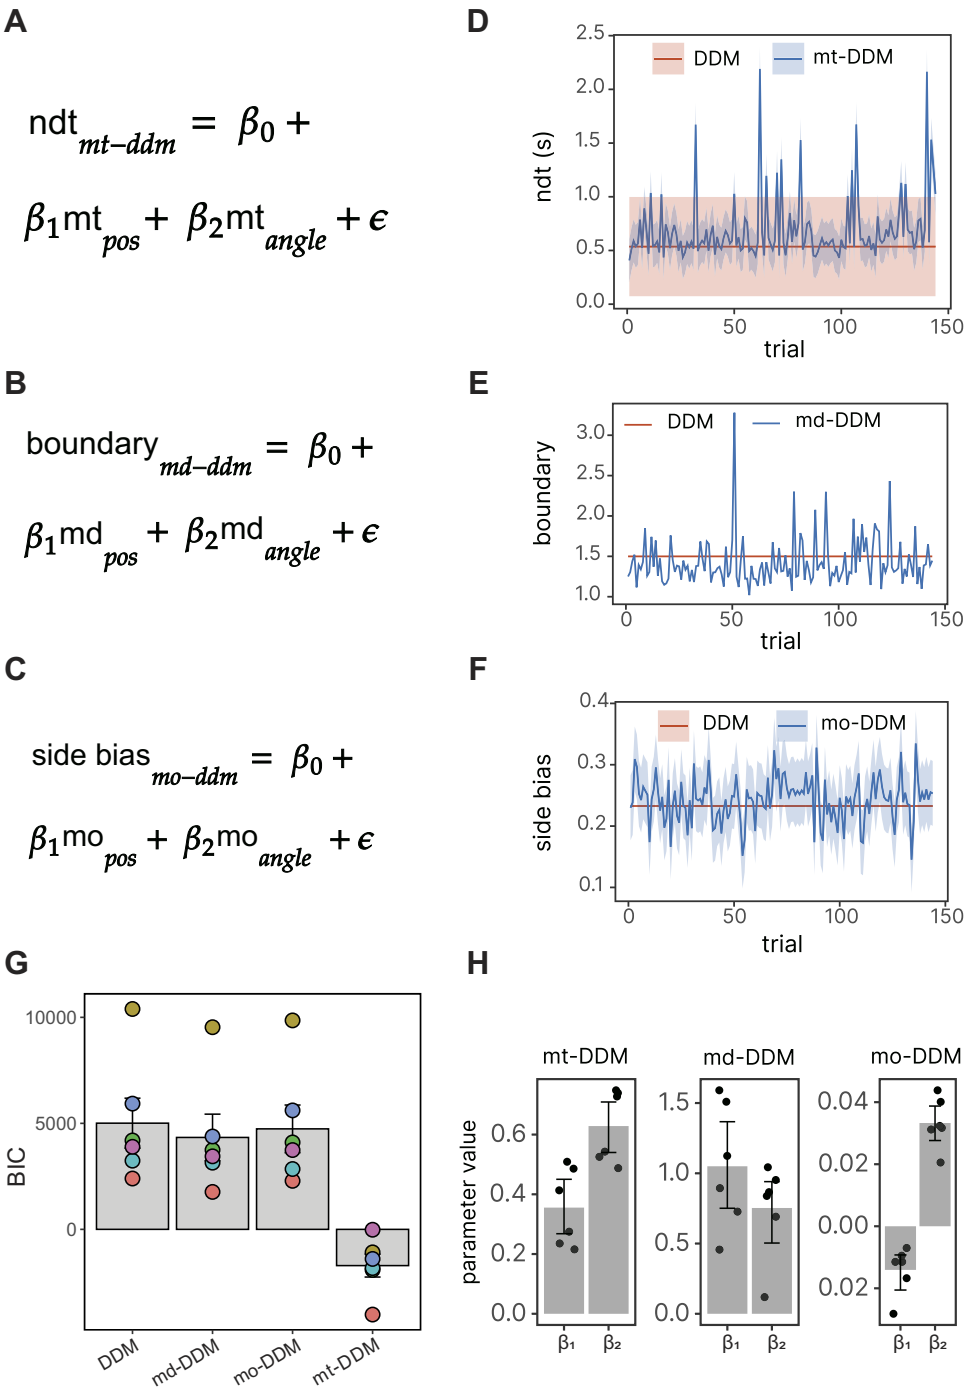

Figure 4. Integration of movement variables improves diffusion model fits. *A*: equation for the nondecision time (ndt) in the movement drift diffusion model (DDM) (mt-DDM). *D*: comparison of the ndt predicted by the DDM and mt-DDM during an example behavioral session. *B*: equation for the decision boundary in the movement delay DDM (md-DDM). *E*: comparison of the boundary predicted by the DDM and md-DDM during an example behavioral session. *C*: equation for the side bias in the movement offset DDM (mo-DDM). *F*: comparison of the side bias predicted by the DDM and mo-DDM during an example behavioral session. *G*: model comparison [Bayes information criterion (BIC)] between the DDM, mt-DDM, md-DDM, and mo-DDM. *H*: values for each of the 2 movement parameters [position (β_1) and angle (β_2)] in each of the 3 models, mt-DDM, md-DDM, and mo-DDM.

evaluated choices on trials in which the sign of accumulated evidence was different before and after movement (Fig. 6*C*). We found that, on these trials, rats were strongly biased to choose the reward port with greater evidence before movement. Next, we fit the same three models (Before, After, and All) as in our previous logistic regression to data from these trials. To maximize statistical power, we fit a hierarchical Bayesian logistic regression (see MATERIALS AND METHODS). This approach allowed us to fit population-level parameters, as well as random variations at the subject level (see MATERIALS AND METHODS). To evaluate model performance first, we calculated and compared the expected log pointwise predictive density (ELPD), using Bayesian leave-one-out cross-validation (30). We found that the Before model performed best ($Before_{ELPD} = -380.40$, $All_{ELPD} = -420.19$, $After_{ELPD} = -595.40$). Next, to estimate the accuracy (ability to predict the choice of the animal) of each model we sampled the posterior distribution and found that the Before model performed best, albeit similarly to the All model, and better than the After model (Fig. 6, *D* and *E*). Finally, we compared the receiver operating characteristic (ROCs) for each model to assess the quality of the binary classifiers; these plot the true positive rate (i.e., the model predicts the

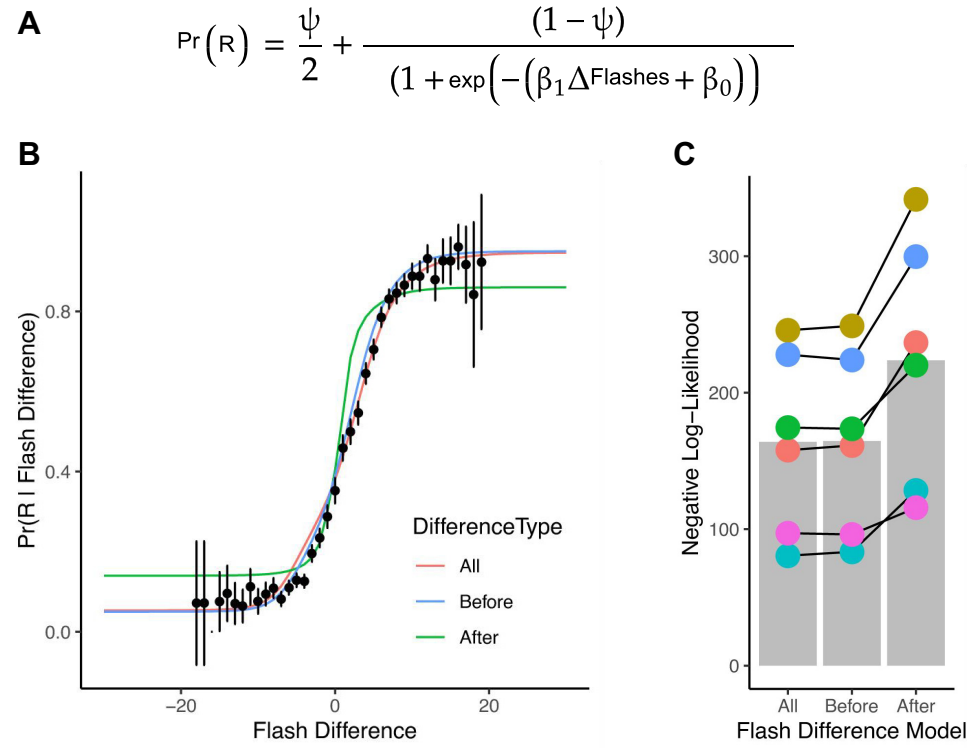

Figure 5. Rats' choices are better predicted by stimuli before movement onset than stimuli presented during movement execution. *A*: functional form of the model for the movement flash regression analysis; ψ represents the lapse parameter. *B*: psychometric curves comparing model predictions and rat behavior. Black dots represent rat data; error bars reflect the 95% confidence interval. Colored lines represent the predictions of the model shown in *A*. In the "All" model, flash difference is computed using the entire stimuli presented during the trial. In the "Before" model, only the stimuli presented before movement onset are used. In the "After" model, only stimuli presented during movements are used. *C*: bar plot comparing the negative log-likelihood plots for all fitted models. Colored dots represent model fits for each of the 6 rats. Before and All are the best fits to the behavioral data.

animal went to the right and the animal went to the right) against the false positive rate (the model predicts the animal went to the right but the animal went to the left) (Fig. 6*F*). Again, we found that the Before model was the best, followed by All, and finally After ($\text{Before}_{\text{AUC}} = 0.89$, $\text{All}_{\text{AUC}} = 0.87$, $\text{After}_{\text{AUC}} = 0.73$).

Together, these results are consistent with a model in which rats base their decisions on an accumulation process where sensory information received before movement onset is weighted most heavily.

DISCUSSION

In this study, we used a free-response version of a pulse-based perceptual integration task to study head movements during evidence accumulation. A key feature of this task is that visual stimuli are presented until an animal reports its decision. Behavioral analysis suggests that rats solve this pulse-based task by accumulating sensory evidence: high accuracy with a positively skewed response time distribution, accuracy increased with response time, and choices and response times were well fit by DDMs. Furthermore, we show that the use of head-related positional data to estimate DDM parameters on a trial-by-trial basis (mDDM) provided better fits than a traditional DDM. We also found that even though rats are free to move during the entire trial, rats tend to exhibit a period of fixation before movement to their choice. Interestingly, we find that sensory information presented during this interval before movement initiation was as effective in predicting choice as stimuli presented over the whole trial and significantly better than stimuli presented after movement initiation. Together these findings suggest that this task is useful for studying the decision-making process, that rats are capable of performing a free-response task with high accuracy and long response times, and that movement-related information can be used to infer the latent parameters of a DDM.

These results also offer insight into the relationship between sensorimotor and deliberative processes during decision-making. A traditional class of models posits a serial relationship between deliberation and movement, such that motor output is entirely decoupled from evidence accumulation and choice (2, 10). However, the data presented here are inconsistent with these models. For example, we show that that head position and angle before sensory cues predict the starting point of accumulation in a drift diffusion process. In contrast, our data could be explained by models of embodied cognition, in which decision processing is coupled with ongoing movement (5, 6, 10, 31). We believe the behavioral task and analytic approach we describe here could provide a foundation for further investigation into the relationship between decision processes and movement. For example, the free-response task could be combined with perturbation experiments to alter head or bodily movements and systematically evaluate how this potentially changes subsequent decisions (32). As video recording can easily be implemented in our task, closed-loop approaches could be used to introduce precise stimulus patterns that depend on the animals' ongoing movement trajectory in an online fashion. These experiments could be used to precisely map the relative timing of movement, evidence accumulation, and choice and could be used to detect and assess "changes of mind" mid-trial (33).

Identification of the circuit mechanism that links deliberation and movements may require neural recordings

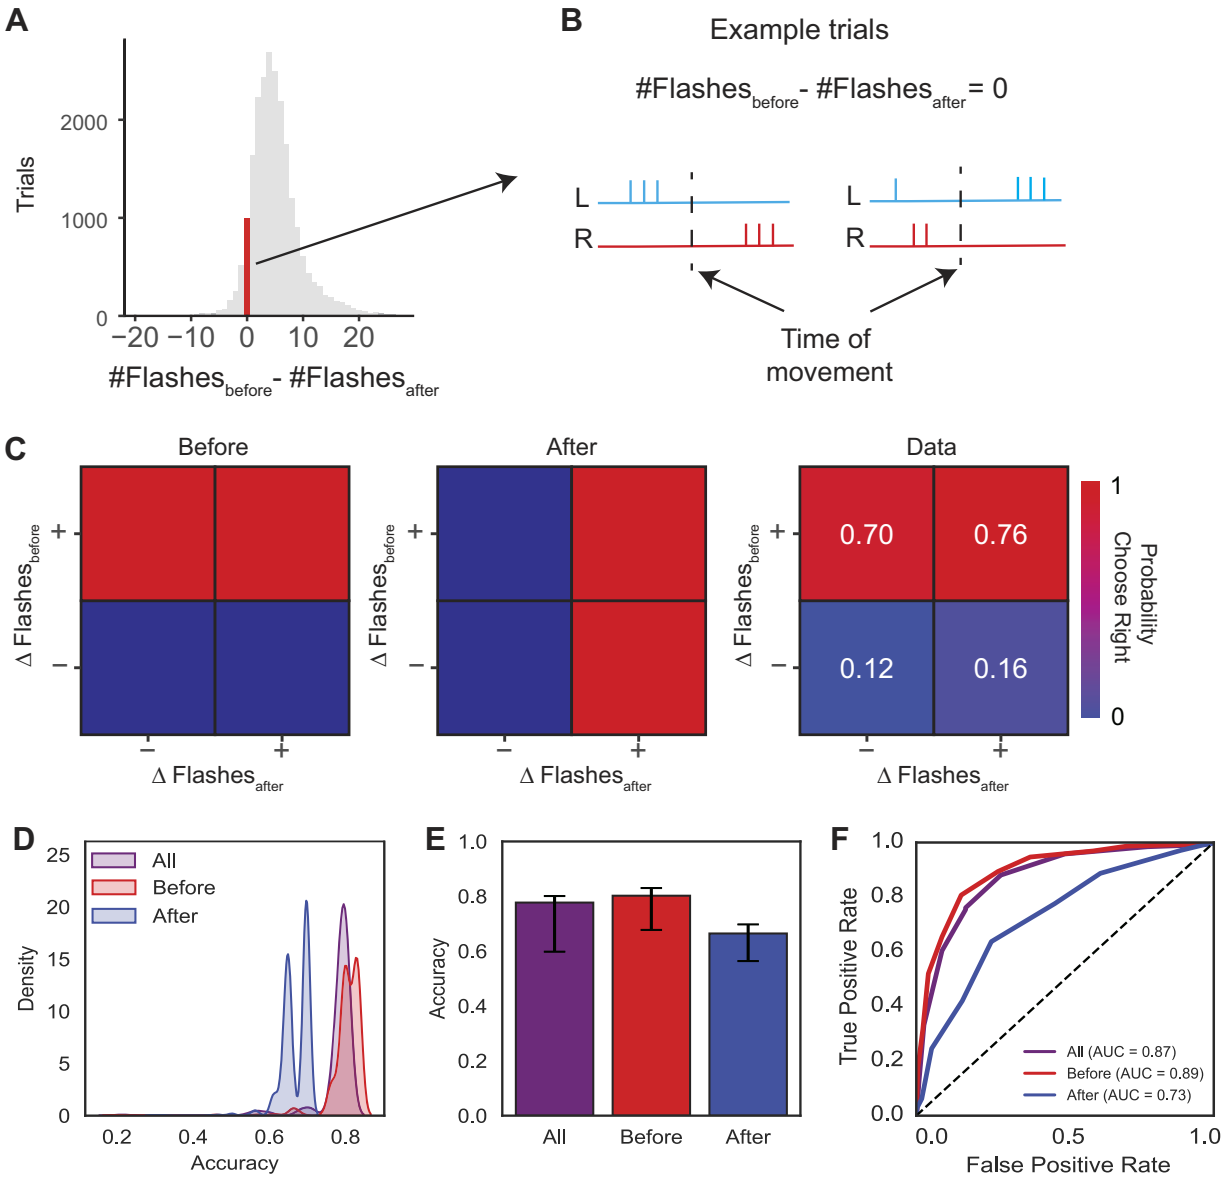

Figure 6. Analysis of choices for trials with equal numbers of flashes before and after moment. *A*: histogram showing the distribution of trials with different numbers of flashes before and after movement. Red bar indicates the number of trials in which the number of flashes before movement equals the number of flashes after movement ($n = 995$ trials). *B*: schematic of 2 example trials with equal number of flashes before and after movement. The timing of left (blue) and right (red) pulses is illustrated as vertical ticks. *C*: analysis of choices based on evidence before and after movement. Color indicates the probability of choosing the right side. *Left*: schematic of performance when choices depend exclusively on stimuli before movement onset. *Center*: schematic of performance when choices depend exclusively on stimuli after movement onset. *Right*: performance of rats. *D*: kernel density estimates of the accuracy of the "Before" (red), "All" (purple), and "After" (blue) models. Data for accuracy were obtained by sampling the posterior distribution of each model 1,000 times. *E*: comparison of the accuracy of each model. Error bars are the 94% credibility intervals. *F*: receiver operating characteristic (ROC) curves for choice decoder based on Before (red), All (purple), and After (blue) models. Before model outperformed All and After models ($\text{Before}_{AUC} = 0.89$, $\text{All}_{AUC} = 0.87$, $\text{After}_{AUC} = 0.73$). AUC, area under the curve.

combined with advanced statistical approaches to dissociate motor and cognitive factors. Prior studies have performed electrophysiological or calcium recording while rodents and primates performed continuous and pulsed-based accumulation tasks (3, 28, 34–36). These studies have identified cells that exhibit changes in firing rate that correlated with decision and task variables, including sensory evidence. Some have interpreted these dynamics as evidence in favor of a neural implementation of the drift diffusion-like process (1). However, despite technological advancements in neural measurement and analysis, concrete interpretations of dynamics recorded during these decision tasks have remained elusive (4, 37). Moreover, it is evident that movement is widely represented across the cortex and therefore should be considered as a factor driving neural dynamics during decision-making (4, 38). This presents a difficulty in interpreting neural dynamics recorded during these tasks. However, recent approaches have pioneered the use of subspace identification approaches that convincingly find lower-dimensional representations in which motor and cognitive variables are separated (39).

Comparison of Diffusion Models of Decision-Making

One of the advantages of the task we report here is that we can measure response times, which enables fitting of the traditional DDM and a more direct comparison between accumulation models (17, 40). The data from this task allowed us to fit both traditional DDMs, in which the drift occurs continuously throughout the cue period, as well as the pDDM, in which the drift is coupled to the exact time of stimulus presentations. For the dataset described here, the continuous and pulse models perform equivalently. However, we point out that the pDDM we fit is just one example of a much larger range of accumulation models designed for pulse-based tasks. For example, Brunton and colleagues (11) designed a nine-parameter accumulation model that incorporated estimates of noise from various sources. Alternative parameterizations of the pDDM may outperform the DDM in fitting to data from this task.

We chose to use the continuous DDM as a foundation for our movement DDM for several practical reasons. First and most importantly, it provided a good description of choices and response times in this task. Second, it was significantly less computationally intensive to calculate the trial-by-trial likelihood than the pulse DDM. This approach revealed that information from rats' movement trajectories improved inference of DDM parameters (i.e., that the parameters better described rat behavior) on a trial-by-trial basis, as indicated by an improvement in the Bayes information criterion (BIC). In future studies it may be possible to incorporate the movement variables into the pDDM, which may improve model performance further.

Limitations of the Present Study

Our results suggest that there is a separate stage in the evidence accumulation process, before movement, when the rat is heavily integrating sensory evidence. It is worth noting that we cannot effectively rule out the possibility that evidence accumulation occurs over the entire trial. One reason for this could be that across most of the trials the flash difference across the premovement period of the trial (i.e., the period in which the animal is initially immobile while gathering sensory evidence) and the entire trial are highly correlated (Pearson's $r = 0.96$). This suggests that we cannot fully distinguish between the premovement and whole trial flashes in terms of predicting choice. However, since the premovement is largely indistinguishable from the model fit of the whole trial flashes, this would still suggest that stimuli integrated before movement onset are most important for determining choice.

A second limitation is that our identification of movement onset is based on the sinusoidal model of head movements. Although we show this model to be an accurate description of the rat's head trajectory during our task, we did not evaluate other forms of movement, such as limb or trunk position. It is possible that incorporation of other forms of movement into the DDM would further improve model fits. Body movements and changes in posture can drive changes in circuits that are of interest in evidence accumulation studies, such as the posterior parietal cortex (41, 42). Therefore, a more global picture of movements during evidence accumulation may be valuable to better interpret neural dynamics associated with decision-making.

DATA AVAILABILITY

Data will be made available upon reasonable request.

ACKNOWLEDGMENTS

We thank Bence Ölveczky (Harvard University) for sharing details of training chambers and Josh Saunders (Sanworks), Dustin Clark (BU Research Computing Services), and Quan Do for technical assistance with implementing data collection. We thank Sinead O'Brien-Wernig and Amber Hickey and the BU Animal Science Center staff for care of the animals. We thank Hongjie Xia, Arula Ratnakar, Shiping Li, Rifqi Affan, and other members of the lab for help with data collection and maintenance of the behavioral training facility. Brian Depasquale, Chandramouli Chandrasekaran, and two anonymous reviewers provided valuable feedback on the manuscript.

GRANTS

This work was supported by a NIH R56MH132732 Award and a Whitehall Foundation Research Grant. G.A.K. was supported by a Boston University Center for Systems Neuroscience Distinguished Fellow award.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

G.A.K. and B.B.S. conceived and designed research; G.A.K. performed experiments; G.A.K. and R.A.S. analyzed data; G.A.K., R.A.S., and B.B.S. interpreted results of experiments; G.A.K. and R.A.S. prepared figures; G.A.K. and R.A.S. drafted manuscript; G.A.K., R.A.S., and B.B.S. edited and revised manuscript; G.A.K., R.A.S., and B.B.S. approved final version of manuscript.

REFERENCES

1. **Gold JI**, **Shadlen MN.** The neural basis of decision making. *Annu Rev Neurosci* 30: 535–574, 2007. doi:10.1146/annurev.neuro.29.051605.113038.
2. **Hanks TD**, **Summerfield C.** Perceptual decision making in rodents, monkeys, and humans. *Neuron* 93: 15–31, 2017. doi:10.1016/j.neuron.2016.12.003.
3. **Pinto L**, **Koay SA**, **Engelhard B**, **Yoon AM**, **Deverett B**, **Thiberge SY**, **Witten IB**, **Tank DW**, **Brody CD.** An accumulation-of-evidence task using visual pulses for mice navigating in virtual reality. *Front Behav Neurosci* 12: 36, 2018 [Erratum in *Front Behav Neurosci* 16: 1079746, 2022]. doi:10.3389/fnbeh.2018.00036.
4. **Musall S**, **Kaufman MT**, **Juavinett AL**, **Gluf S**, **Churchland AK.** Single-trial neural dynamics are dominated by richly varied movements. *Nat Neurosci* 22: 1677–1686, 2019. doi:10.1038/s41593-019-0502-4.
5. **Cisek P**, **Kalaska JF.** Neural mechanisms for interacting with a world full of action choices. *Annu Rev Neurosci* 33: 269–298, 2010. doi:10.1146/annurev.neuro.051508.135409.
6. **Cisek P**, **Pastor-Bernier A.** On the challenges and mechanisms of embodied decisions. *Philos Trans R Soc Lond B Biol Sci* 369: 20130479, 2014. doi:10.1098/rstb.2013.0479.
7. **Shadlen MN**, **Kiani R.** Decision making as a window on cognition. *Neuron* 80: 791–806, 2013. doi:10.1016/j.neuron.2013.10.047.
8. **Wispinski NJ**, **Gallivan JP**, **Chapman CS.** Models, movements, and minds: bridging the gap between decision making and action. *Ann NY Acad Sci* 1464: 30–51, 2020. doi:10.1111/nyas.13973.
9. **Selen LP**, **Shadlen MN**, **Wolpert DM.** Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J Neurosci* 32: 2276–2286, 2012. doi:10.1523/JNEUROSCI.5273-11.2012.

10. **Lepora NF**, **Pezzulo G.** Embodied choice: how action influences perceptual decision making. *PLoS Comput Biol* 11: e1004110, 2015. doi:10.1371/journal.pcbi.1004110.
11. **Brunton BW**, **Botvinick MM**, **Brody CD.** Rats and humans can optimally accumulate evidence for decision-making. *Science* 340: 95–98, 2013. doi:10.1126/science.1233912.
12. **Scott BB**, **Constantinople CM**, **Erlich JC**, **Tank DW**, **Brody CD.** Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *eLife* 4: e11308, 2015. doi:10.7554/eLife.11308.
13. **Gupta D**, **DePasquale B**, **Kopec CD**, **Brody CD.** Trial-history biases in evidence accumulation can give rise to apparent lapses in decision-making. *Nat Commun* 15: 662, 2024. doi:10.1038/s41467-024-44880-5.
14. **Chakravarty S**, **Delgado-Sallent C**, **Kane GA**, **Xia H**, **Do QH**, **Senne RA**, **Scott BB.** A cross-species framework for investigating perceptual evidence accumulation (Preprint). *bioRxiv* 2024.04.17.589945, 2024. doi:10.1101/2024.04.17.589945.
15. **Dhawale AK**, **Poddar R**, **Wolff SB**, **Normand VA**, **Kopelowitz E**, **Ölveczky BP.** Automated long-term recording and analysis of neural activity in behaving animals. *eLife* 6: e27702, 2017. doi:10.7554/eLife.27702.
16. **Poddar R**, **Kawai R**, **Ölveczky BP.** A Fully automated high-throughput training system for rodents. *PLoS One* 8: e83171, 2013. doi:10.1371/journal.pone.0083171.
17. **Ratcliff R.** A theory of memory retrieval. *Psychol Rev* 85: 59–108, 1978. doi:10.1037/0033-295X.85.2.59.
18. **Myers CE**, **Interian A**, **Moustafa AA.** A practical introduction to using the drift diffusion model of decision-making in cognitive psychology, neuroscience, and health sciences. *Front Psychol* 13: 1039172, 2022. doi:10.3389/fpsyg.2022.1039172.
19. **Mathis A**, **Mamidanna P**, **Cury KM**, **Abe T**, **Murthy V**, **Mathis MW**, **Bethge M.** DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci* 21: 1281–1289, 2018. doi:10.1038/s41593-018-0209-y.
20. **Xiang Y**, **Gubian S**, **Suomela B**, **Hoeng J.** Generalized simulated annealing for global optimization: the GenSA package. *R J* 5: 13, 2013. doi:10.32614/RJ-2013-002.
21. **Capretto T**, **Piho C**, **Kumar R**, **Westfall J**, **Yarkoni T**, **Martin OA.** Bambi: a simple interface for fitting Bayesian linear models in Python (Preprint). *arXiv* 2012.10754[stat.CO], 2020. doi:10.48550/arXiv.2012.10754.
22. **Wood SN.** *Generalized Additive Models: an Introduction with R* (2nd ed.). Boca Raton, FL: CRC Press/Taylor & Francis Group, 2017.
23. **Brooks ME**, **Kristensen K**, **van Benthem KJ**, **Magnusson A**, **Berg CW**, **Nielsen A**, **Skaug HJ**, **Mächler M**, **Bolker BM.** glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *R J* 9: 378, 2017. doi:10.32614/RJ-2017-066.
24. **Brown S**, **Heathcote A.** QMLE: fast, robust, and efficient estimation of distribution functions based on quantiles. *Behav Res Methods Instrum Comput* 35: 485–492, 2003. doi:10.3758/bf03195527.
25. **Heathcote A**, **Brown S**, **Cousineau D.** QMPE: estimating Lognormal, Wald, and Weibull RT distributions with a parameter-dependent lower bound. *Behav Res Methods Instrum Comput* 36: 277–290, 2004. doi:10.3758/bf03195574.
26. **Mullen KM**, **Ardia D**, **Gil DL**, **Windover D**, **Cline J.** DEoptim: an R package for global optimization by differential evolution. *J Stat Soft* 40: 1–26, 2011. doi:10.18637/jss.v040.i06.
27. **Voss A**, **Voss J.** A fast numerical algorithm for the estimation of diffusion model parameters. *J Math Psychol* 52: 1–9, 2008. doi:10.1016/j.jmp.2007.09.005.
28. **Scott BB**, **Constantinople CM**, **Akrami A**, **Hanks TD**, **Brody CD**, **Tank DW.** Fronto-parietal cortical circuits encode accumulated evidence with a diversity of timescales. *Neuron* 95: 385–398.e5, 2017. doi:10.1016/j.neuron.2017.06.013.
29. **Scott BB**, **Thiberge SY**, **Guo C**, **Tervo DG**, **Brody CD**, **Karpova AY**, **Tank DW.** Imaging cortical dynamics in GCaMP transgenic rats with a head-mounted widefield macroscope. *Neuron* 100: 1045–1058.e5, 2018. doi:10.1016/j.neuron.2018.09.050.
30. **Vehtari A**, **Gelman A**, **Gabry J.** Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC (Preprint). *arXiv* 1507.04544[stat.CO], 2015. doi:10.48550/arXiv.1507.04544.
31. **Redish AD.** Vicarious trial and error. *Nat Rev Neurosci* 17: 147–159, 2016. doi:10.1038/nrn.2015.30.
32. **Gallivan JP**, **Chapman CS**, **Wolpert DM**, **Flanagan JR.** Decision-making in sensorimotor control. *Nat Rev Neurosci* 19: 519–534, 2018. doi:10.1038/s41583-018-0045-9.
33. **Resulaj A**, **Kiani R**, **Wolpert DM**, **Shadlen MN.** Changes of mind in decision-making. *Nature* 461: 263–266, 2009. doi:10.1038/nature08275.
34. **Roitman JD**, **Shadlen MN.** Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22: 9475–9489, 2002. doi:10.1523/JNEUROSCI.22-21-09475.2002.
35. **Hanks TD**, **Kopec CD**, **Brunton BW**, **Duan CA**, **Erlich JC**, **Brody CD.** Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* 520: 220–223, 2015. doi:10.1038/nature14066.
36. **Boyd-Meredith JT**, **Piet AT**, **Dennis EJ**, **El Hady A**, **Brody CD.** Stable choice coding in rat frontal orienting fields across model-predicted changes of mind. *Nat Commun* 13: 3235, 2022. doi:10.1038/s41467-022-30736-3.
37. **Costello MG**, **Zhu D**, **Salinas E**, **Stanford TR.** Perceptual modulation of motor–but not visual–responses in the frontal eye field during an urgent-decision task. *J Neurosci* 33: 16394–16408, 2013. doi:10.1523/JNEUROSCI.1899-13.2013.
38. **Umeda T**, **Isa T**, **Nishimura Y.** The somatosensory cortex receives information about motor output. *Sci Adv* 5: eaaw5388, 2019. doi:10.1126/sciadv.aaw5388.
39. **Hasnain MA**, **Birnbaum JE**, **Nunez JL**, **Hartman EK**, **Chandrasekaran C**, **Economo MN.** Separating cognitive and motor processes in the behaving mouse (Preprint). *bioRxiv* 2023.08.23.554474, 2024. doi:10.1101/2023.08.23.554474.
40. **Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD.** The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113: 700–765, 2006. doi:10.1037/0033-295X.113.4.700.
41. **Mimica B**, **Dunn BA**, **Tombaz T**, **Bojja V**, **Whitlock JR.** Efficient cortical coding of 3D posture in freely behaving rats. *Science* 362: 584–589, 2018. doi:10.1126/science.aau2013.
42. **Tombaz T**, **Dunn BA**, **Hovde K**, **Cubero RJ**, **Mimica B**, **Mamidanna P**, **Roudi Y**, **Whitlock JR.** Action representation in the mouse parieto-frontal network. *Sci Rep* 10: 5559, 2020. doi:10.1038/s41598-020-62089-6.