Article

Current Biology

Sequential sampling from memory underlies action selection during abstract decision-making

Highlights

- Monkeys made perceptual decisions without knowledge of the action to report them
- Decision formation was postponed until the relevant motor actions were revealed
- Perceptual information was sequentially sampled from memory for decision formation
- Neurons in area LIP represent accumulating evidence sampled from short-term memory

Authors

S. Shushruth, Ariel Zylberberg, Michael N. Shadlen

Correspondence

shushruth@gmail.com (S.S.), ariel.zylberberg@gmail.com (A.Z.), shadlen@columbia.edu (M.N.S.)

In brief

Shushruth et al. investigate how monkeys make visual decisions before the actions to report their choice are specified. Monkeys do not make a decision while viewing the stimulus but wait until the actions are revealed. Neurons in parietal cortex accumulate visual information stored in short-term memory to decide which action to take.

Shushruth et al., 2022, Current Biology *32*, 1–12
May 9, 2022
https://doi.org/10.1016/j.cub.2022.03.014

CellPress

Current Biology

CellPress
OPEN ACCESS

Article

Sequential sampling from memory underlies action selection during abstract decision-making

S. Shushruth,[1,*] Ariel Zylberberg,[1,*] and Michael N. Shadlen[1,2,3,4,5,*]
[1]Zuckerman Mind Brain Behavior Institute, Department of Neuroscience, New York, NY 10027, USA
[2]Howard Hughes Medical Institute, New York, NY 10027, USA
[3]Kavli Institute, Columbia University, 612 West 130th Street, New York, NY 10027, USA
[4]Twitter: @shadlen
[5]Lead contact
*Correspondence: shushruth@gmail.com (S.S.), ariel.zylberberg@gmail.com (A.Z.), shadlen@columbia.edu (M.N.S.)
https://doi.org/10.1016/j.cub.2022.03.014

SUMMARY

The study of perceptual decision-making in monkeys has provided insights into the process by which sensory evidence is integrated toward a decision. When monkeys make decisions with the knowledge of the motor actions the decisions bear upon, the process of evidence integration is instantiated by neurons involved in the selection of said actions. It is less clear how monkeys make decisions when unaware of the actions required to communicate their choice—what we refer to as "abstract" decisions. We investigated this by training monkeys to associate the direction of motion of a noisy random-dot display with the color of two targets. Crucially, the targets were displayed at unpredictable locations after the motion stimulus was extinguished. We found that the monkeys postponed decision formation until the targets were revealed. Neurons in the parietal association area LIP represented the integration of evidence leading to a choice, but as the stimulus was no longer visible, the samples of evidence must have been retrieved from short-term memory. Our results imply that when decisions are temporally unyoked from the motor actions they bear upon, decision formation is protracted until they can be framed in terms of motor actions.

INTRODUCTION

A decision is a commitment to a proposition or plan of action based on evidence, prior knowledge, priorities, and value. Perceptual decision-making refers to the class of decisions in which the dominant source of evidence is derived from sensation and in which the decision is a provisional action or a mental assignment to a category. Viewed from the perspective of information processing, perceptual decision-making establishes a compressed distillation of sensory data into distinct categories. Viewed from the perspective of behavior, it effects an intention, satisfying policy objectives, such as obtaining reward. These perspectives are naturally connected because we decide about a perceptual category to make a choice. The study of decision-making in laboratory animals tends to conflate these depictions, perhaps by necessity.

There is recent interest in characterizing the neural processes that underlie decisions about category membership, independent of intention,[1–4] which we will refer to as abstract decisions. Categorical labels introduce flexibility to sensorimotor programs.[5,6] For example, one can assign the labels, "A" and "B" or "rightward" and "leftward" to consolidate motion perceived to the right or left, irrespective of its precise direction or motion strength. These abstract labels then allow for the implementation of flexible action plans, such as "press a red button if you see rightward motion."

The extent to which nonhuman primates can assign abstract labels to sensory percepts and exploit them to be flexible in their actions is unclear. The process of abstraction, by definition, unyokes the sensory evaluation processes from the process of acting on the sensory information. However, multiple lines of research in macaques suggest that the process of sensory evaluation is intimately coupled to the actions that can result from the evaluative process.[7–9] This framework, wherein cognitive processes are embodied in terms of the motor actions they afford, is supported by the patterns of neural activity found in association and premotor cortices of monkeys.[7,10,11] However, monkeys can be trained to decide on properties of sensory stimuli even when unaware of the motor action required to report their decision.[1,2,4,12–17] In these studies, monkeys were required to commit to a category assignment without committing to an action. It is unclear from these studies how the abstract representation is established and how it is ultimately translated to the response. That is what we set out to clarify.

We trained two monkeys to decide on the net direction of stochastic random-dot motion (RDM) and associate two possible directions with two colors. The monkeys reported the direction of motion by making an eye movement to the target of the associated color, but these targets were revealed at unpredictable locations after the motion stimulus had been extinguished. To perform the task well, monkeys needed to integrate motion information in the stimulus over time to make an abstract decision about the direction of motion. This imposition allowed us to investigate how an abstract perceptual decision is formed when the actions associated with the decision are yet to be specified. Furthermore, since the decision-making phase is unyoked from the motor

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

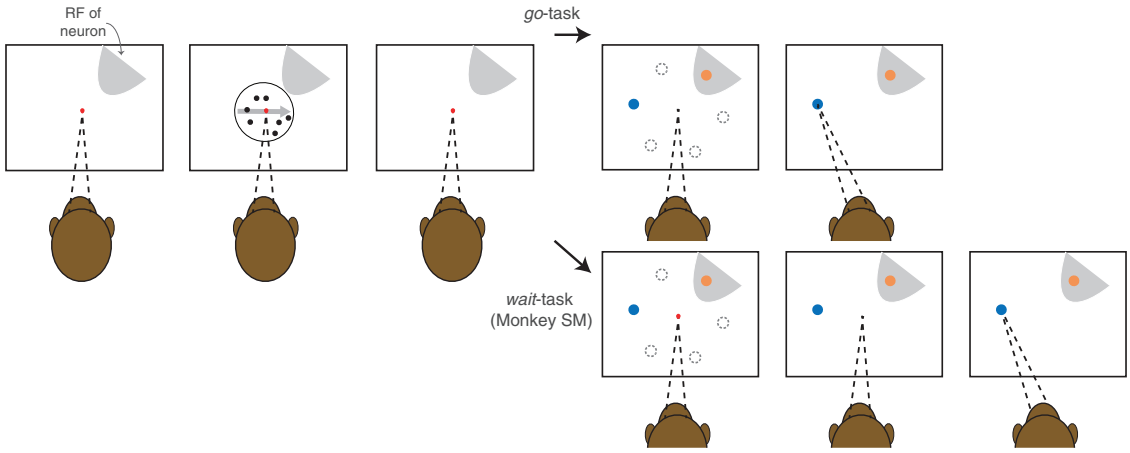

Figure 1. Behavioral task

The monkey fixates at an instructed location (red dot). After a delay, a random-dot motion stimulus appears around the fixation point. The stimulus terminates after a variable duration (350–800 ms). After another short delay (200–333 ms), a blue and a yellow target appear at unpredictable peripheral locations. In the "*go*" version of the task (top panel), the fixation is extinguished at the time the targets appear, and the monkey can report the decision by choosing one of the colored targets. In the "*wait*" version of the task (bottom panel), the monkey must wait until the fixation point is extinguished before choosing a target. During recording sessions, the target locations on each trial are pseudorandomly chosen from a restricted set of locations based on the receptive field of the neuron being recorded. The unchosen locations are illustrated by dashed gray circles (not shown to the monkey). During training sessions, the target locations were less constrained.

planning phase, the task also permits investigation of the conversion of an abstract decision to an action.

Surprisingly, we found that evidence evaluation and action selection—the two aspects of abstract decision-making that our task was supposed to unyoke—were, in fact, intimately coupled. The behavior of the monkeys showed that they based their decision on motion evidence integrated over time. However, this integration transpired during the action-selection epoch instead of the epoch when the evidence was presented. Furthermore, activity of neurons in the sensorimotor association area LIP represented decision formation during the action-selection epoch. Our results suggest that monkeys form abstract perceptual decisions by evaluating sensory information from iconic short-term memory[18] for action selection.

RESULTS

We trained two monkeys (monkey-AN and monkey-SM) to decide whether the net direction of a RDM stimulus was to the right or left. The monkeys reported their decision by making an eye movement to a blue or yellow target based on the association they had learned between the direction of motion and target colors (Figure 1). The two targets appeared after a short delay (200–333 ms) following the termination of the motion stimulus, and the locations of the two targets were randomized across trials. Thus, all the evidence bearing on the decision was supplied before the monkeys were instructed about the motor act that would be required to report the decision. Unlike previous studies,[12,14] both monkeys were naive to the RDM stimulus when they began training on the task. Since one of our goals was to investigate how decisions are converted to motor actions, the monkeys were allowed to report the decision as soon as the targets were presented (*go*-task). Monkey-SM was also trained on a variant of the task in which an additional waiting time was imposed after the appearance of the targets (*wait*-task).

The abstract decision-making task proved to be challenging for the monkeys to learn (Figure S1). Monkey-AN required 28 sessions to acquire the motion-color association and failed to improve beyond competency at the highest motion strengths for the next ~40 sessions (~50,000 trials). Only then did the monkey begin to exhibit gradual improvement, quantified by a reduction in psychophysical threshold—the motion strength required to support accuracy greater than 75% correct (Equation 1). Monkey-SM learned the motion-color association quickly but made little progress over months of training. After 127 sessions (983 trials per session on average), the thresholds still hovered around 25% coherence. This monkey was then trained on the *wait* variant of the task for an additional 58 sessions (740 trials per session) until the thresholds decreased and stabilized at ~11% coherence.

By the final training session, both monkeys performed the task above chance for all non-zero motion coherences (Figures 2A and 2B), although they made errors on the easiest motion strength. Such asymptotic performance is commonly interpreted as a sign that the decision-maker lapses (e.g., guesses) on a fraction of all trials. The lapse rates were 9% and 11% of trials for monkeys AN and SM, respectively. Monkeys performing the same direction discrimination task with a direct mapping between motion direction and actions typically exhibit lapse rates under 2%, further attesting to the challenging nature of the present task, even after extensive training. Nonetheless, for the vast majority of trials, both monkeys used evidence from the RDM to choose the appropriate color. This was confirmed using psychophysical reverse correlation, which identifies the times that random fluctuations of information in the random-dot stimulus influence the decision. The analysis reveals that monkeys AN and SM based their decisions on information acquired over 357 and 261 ms, respectively (Figures 2C and 2D). Importantly, the influence ceases at least 350 ms before the color targets appear.

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

CellPress
OPEN ACCESS

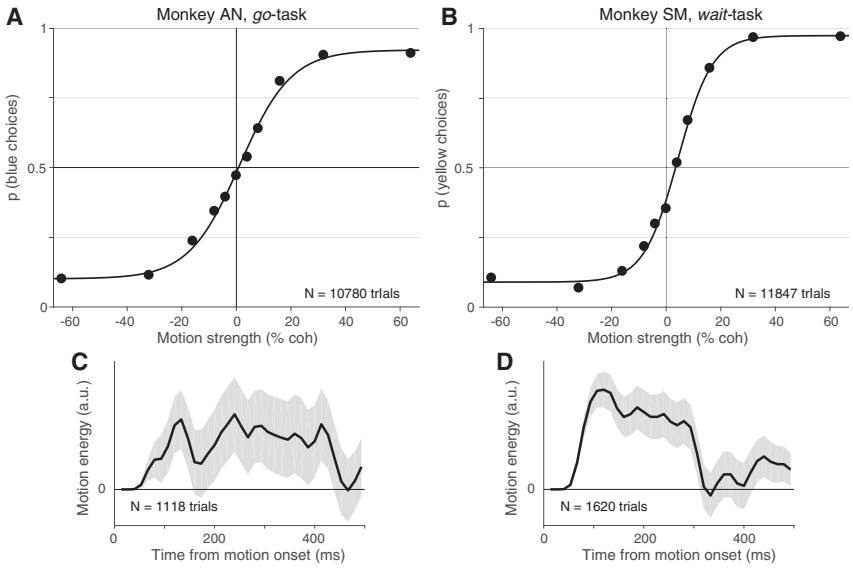

Figure 2. Color choices are governed by the strength and direction of motion

(A) Effect of motion strength on decisions for monkey-AN on the *go*-task. Monkey-AN was trained to associate blue with rightward and yellow with leftward motion. The proportion of blue (rightward) choices are plotted as a function of signed motion strength (rightward motion is positive signed). Curves are logistic regression fits to the data. Error bars are SE.

(B) Effect of motion strength on decisions for monkey-SM on the *wait*-task. Monkey-SM was trained to associate yellow with rightward and blue with leftward motion. The proportion of yellow (rightward) choices are plotted as a function of signed motion strength. Otherwise, same conventions as in (A).

(C and D) The influence of fluctuations in motion information on choices plotted as a function of time from motion onset. Curves represent the mean motion energy in support of the direction chosen by the monkey on 0% coherence trials (shading, ±1 SEM).

Action selection during abstract decision-making is a deliberative process

The behavioral task is structured to separate the decision-making epoch from the action-selection epoch. The natural expectation is that the monkey decides about the direction of motion in the random-dot display while it is visible and is thus prepared to make an eye movement to the blue or yellow target once they are displayed (Figure 3, Strategy 1). If so, the action-selection epoch would involve a simple translation of the decision into an eye movement. This should take as little as 200 ms—the amount of time required to search for a colored target accompanied by a single highly discriminable distractor.[19–21] Unexpectedly, we found that both monkeys required a prolonged action-selection epoch to integrate motion evidence toward a decision. Given that only the two-colored targets were visible during this epoch, we hypothesized that monkeys could be using this time to sample information stored in memory to render their decisions (Figure 3, Strategy 2). Behavioral results from monkeys AN and SM provided distinct but complementary insights into the deliberative process. We next proceed to describe them separately, followed by the data from neural recordings, which were strikingly similar.

Monkey-AN: Prolongation of go-RTs

Monkey-AN showed a natural inclination to deliberate after the appearance of the targets. Figure 4A displays the proportion of blue choices (bottom) and their associated response times (go-RT, top) as a function of motion strength and direction. The RTs, measured from onset of the colored choice targets (at least 300 ms after the motion stimulus was extinguished), were 2–4 times slower than expected[21] if the decision had been made during the motion-viewing epoch (Figure 3 [strategy 1]). The averages range from 440 ms for the easiest condition to 771 ms for the most difficult, and they exhibit a clear dependency on the strength and direction of motion. The range of go-RTs between easiest to most difficult blue choices and the range between easiest to most difficult yellow choices is an order of magnitude longer than the range associated with changes in reward expectation or confidence (typically <20 ms; e.g., Gold and Shadlen[12]). The dependency of go-RT on motion strength resembles the pattern of response times—relative to onset of the RDM—seen in earlier studies, where monkeys were free to indicate their saccadic choices to targets that were already present during motion viewing.[23]

We therefore considered the possibility that the pattern of go-RTs might result from sequential sampling of evidence experienced earlier in the trial, that is, from memory (Figure 3, Strategy 2). To evaluate this, we appropriated a bounded evidence-accumulation model (drift-diffusion) that is known to reconcile the choice proportions with the response times of subjects when they are allowed to answer whenever ready. In such "free response" tasks, the decision-maker knows how to answer while viewing the motion and simply stops the trial by pushing a button or making an eye movement to one of two visible targets. We wondered if the same type of model could reconcile the choices and go-RTs of monkey-AN.

The curves in Figure 4A are fits of a bounded drift-diffusion model to the proportion of blue choices and the mean go-RTs. The simplest version of this model assumes that the drift rate is proportional to signed motion coherence and the terminating bounds do not change as a function of time.[24] Any bias is accommodated by an offset to the drift rate.[25] The mean go-RT for each signed motion strength is predicted by the expectation of the bound termination times plus a constant nondecision time, which captures contributions to the RT that do not depend on the motion strength and bias. We used separate terms (t^{b}_{nd} and t^{y}_{nd}) to describe the faster blue and slower yellow choices. The model to this point uses only five degrees of freedom to explain the choice proportions and mean go-RT across 11 motion strengths (Equations 3 and 4).

We incorporated one additional feature to accommodate the failure of monkey-AN to achieve perfect performance on the easiest conditions (±64% coh). Such errors are typically attributed to lapses in which the subject ignores the evidence and guesses. However, we noticed that the go-RTs associated

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

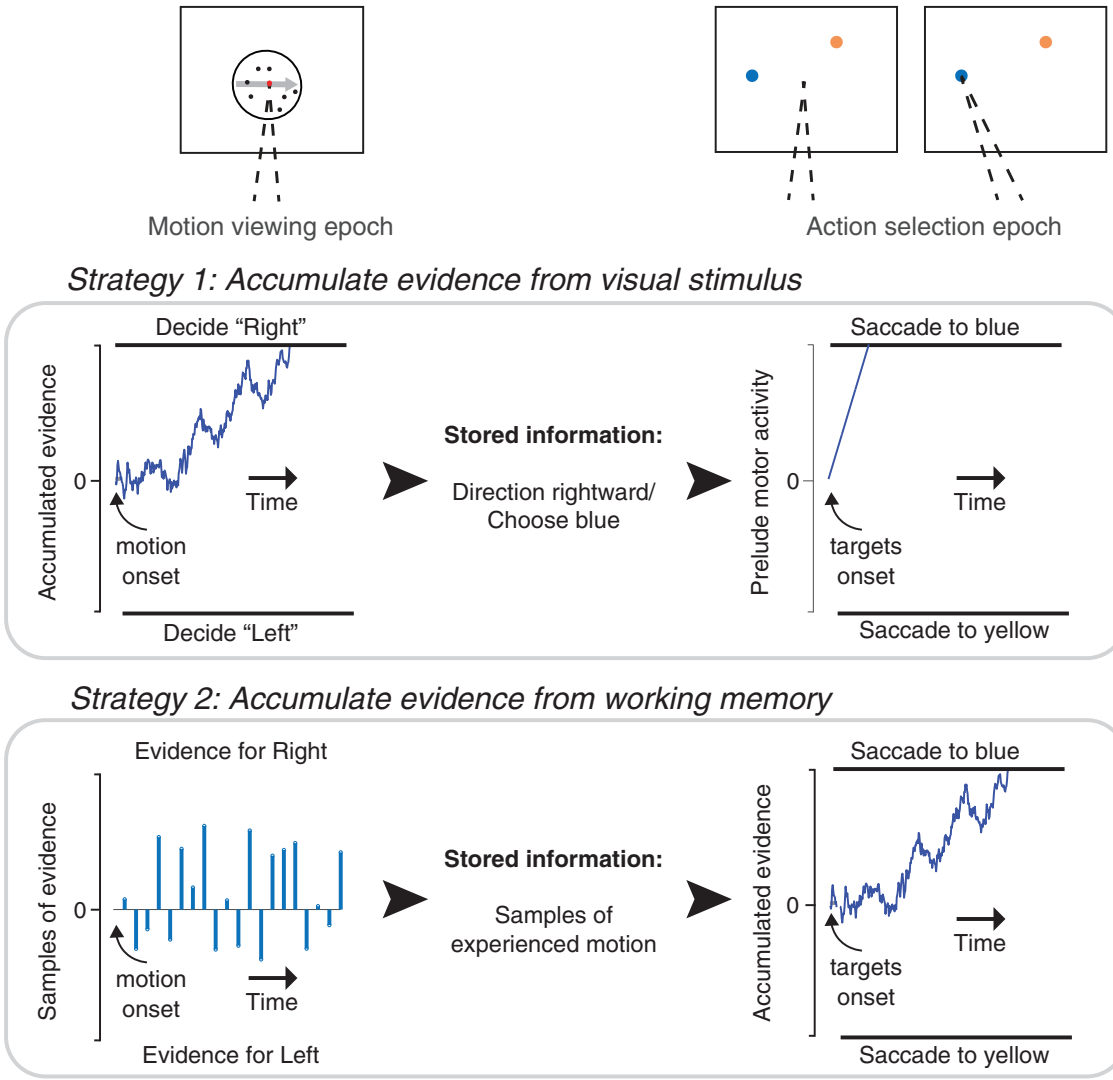

Figure 3. Putative strategies

Schematic of strategies that monkeys could adopt to solve the task. Strategy 1: during motion viewing, evidence for motion direction is accumulated to decide if the motion is to the right or left. The result of the decision about direction and/or its color association is stored. When the colored targets are presented, the previously made decision guides an immediate saccade to the target with the chosen color. The saccadic latency might vary by 10–20 ms as a function of confidence in the decision. Strategy 2: during motion viewing, the experienced evidence is stored in short-term memory. When the targets are shown, the stored evidence is evaluated during action-selection to decide which of the two-colored targets to choose. The drawing gives the impression of many samples, but the samples themselves might represent several tens of ms of motion information (as in Kang et al.[22]).

with errors on the strong leftward-motion condition (blue choices) had the slow go-RTs associated with correct leftward-motion (yellow) choices. Similarly, the errors associated with the strong rightward-motion condition (yellow choices) had the fast go-RTs associated with correct rightward-motion (blue) choices (Figure S2B). This indicates that the lapses were not guesses but an error in the association between direction and color (Data S1). We accommodated this feature in the model, assuming that this type of error occurred on a small fraction of trials, independently of motion strength (STAR Methods). The model captures the coherence dependence of the go-RTs on correct choices (R^2 = 0.99) while also accounting for the accuracy of the monkey's choices (Figure 4A). The fidelity of the fits supports the hypothesis that the prolonged go-RTs reflect a bounded accumulation of noisy evidence leading to the rendering of the decision. As this sampling began at least 300 ms after the motion stimulus was extinguished, the samples must be derived from memory.

Because drift-diffusion models with time-independent (i.e., "flat") bounds cannot explain the difference in response times between correct and error choices at a given motion strength, we considered a more elaborate version of the model to explain the go-RTs on errors. The model incorporates decision-termination bounds that can change with elapsed time (Equation 11; Figure S2A). We fit the extended model to the choice and go-RT data, including the go-RTs on errors. The best-fitting model

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

CellPress
OPEN ACCESS

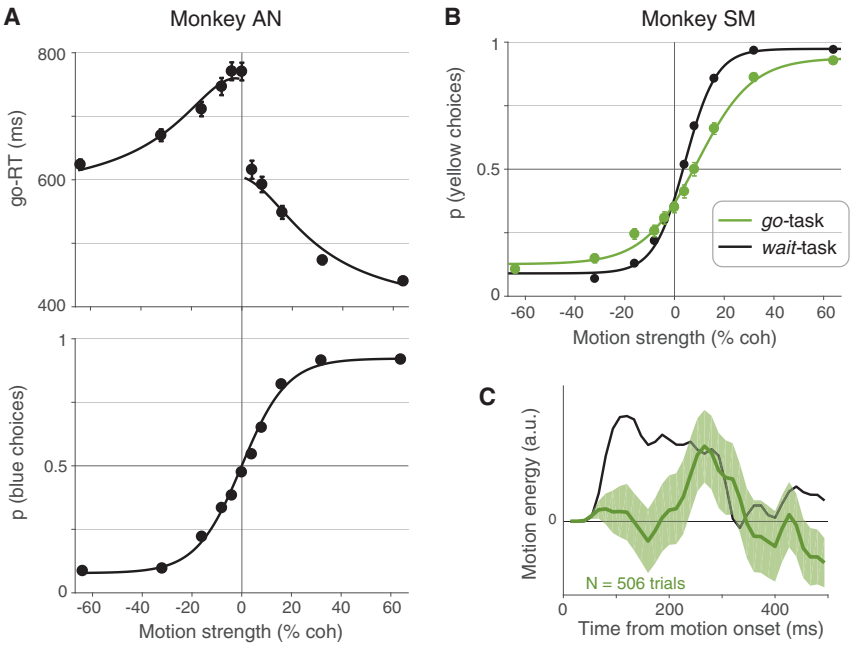

Figure 4. Deliberation during action selection

(A) (Top) Go-RTs of monkey-AN plotted as a function of signed motion coherence. Error bars are SEM. Curves are fits to a bounded drift-diffusion model. The model is also constrained by the choice proportions. (Bottom) Same data as in Figure 2A. Curve is the fit of the bounded diffusion model, which accounts for both the choice proportions and the go-RTs. Table S2 shows the best-fitting model parameters. See also Figure S2 and Data S1.

(B) Proportion of rightward (yellow) choices as a function of motion strength for monkey-SM from the last four training sessions on the *go*-task (green; see Figure S1). The same data as Figure 2B are shown for comparison (black). Lines are logistic regression fits to the data.

(C) Influence of motion fluctuations on choice in the last four training sessions of the *go*-task for monkey-SM (green curve). Same conventions as Figure 2D. Data from the *wait*-task (black curve, same as in Figure 2D) are shown for comparison.

See also Figure S3.

(Figure S2) yields an expectation of the integration time for each motion strength. For 0% coherence, the expectation is 243 ms, which is consistent with the psychophysical reverse-correlation analysis above (Figure 2C). Note that the reverse-correlation analysis also shows that the monkey uses the earliest epochs of motion evidence to inform its decision. Taken together, the analyses of go-RT and reverse correlation suggest that monkey-AN stores at least 300 ms of information about the motion in some form. The estimate is longer than the expectations because the duration of stimulus information needed for decision termination is not known before the accumulation process transpires.

Monkey-SM: Improved integration after enforced wait

Unlike monkey-AN, monkey-SM did not show a tendency to deliberate after target onset in the *go*-version of the task. However, although monkey-SM learned the direction-color association and performed better than monkey-AN at the strongest motion conditions, it failed to achieve proficiency on the more difficult conditions (Figure S1). Even after extensive training, sensitivity plateaued at an unacceptable level (Figure 4B [green]), and psychophysical reverse correlation revealed only a weak, transient impact of motion information on choice (Figure 4C [green]). We therefore suspected that this monkey based its decisions on a brief sample of information from the first, last, or random glimpse of the display (e.g., see Stine et al.[26]). We confirmed this using a variant of the *go*-task in which the strength of motion was modulated as a function of time within a trial (Figure S3). The coherence started at 0% and either stepped or changed gradually to a large positive or negative value. The time of the step or the rate of change varied across trials. The monkey's performance deteriorated to chance when the strong motion was concentrated at the end of the trial (Figure S3). We deduced that the monkey based its decisions on motion information sampled over a short time window at the beginning of the trial. Not surprisingly, the go-RTs exhibited no sign of deliberation. They were nearly as fast as a simple saccadic reaction time to a single target (192 ± 0.4 ms) and showed no influence of the previously experienced motion strength.

Based on our experience with monkey-AN, we wondered if monkey-SM failed to integrate after the color-choice targets appeared. We therefore introduced a wait time after the onset of the targets. This simple modification led to a 2-fold improvement in sensitivity (Figures 4B [green versus black traces] and Figure S1). This degree of improvement would require at least a 4-fold increase in the number of independent samples of evidence the monkey used to form its decision. Indeed, psychophysical reverse correlation revealed a longer time window over which motion information influenced decisions: from 40 ms, before the introduction of the enforced wait, to 261 ms, after ~40 sessions of training (Figure 4C). Thus, the imposition of a wait after the onset of the targets encouraged monkey-SM to use more information to inform its decision—information that was acquired earlier in the motion-viewing epoch.

The behavioral data from both the monkeys therefore provide complementary evidence that deliberation during the action-selection epoch is necessary for integrating previously observed motion information. During motion viewing, both monkeys must store some representation of the motion in short-term memory. The go-RTs from monkey-AN indicate that the stored information is sampled sequentially in the action-selection period. Owing to the enforced wait, we lack meaningful go-RTs for monkey-SM. However, as we next show, the neural recordings demonstrate that monkey-SM also samples sequentially from memory during the action-selection epoch.

LIP neurons represent the accumulation of evidence from memory

We recorded from single units with spatially selective persistent activity in area LIP.[27,28] Such neurons are known to represent an evolving decision variable—the accumulated evidence for and against a motion direction—when one of the choice targets is

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

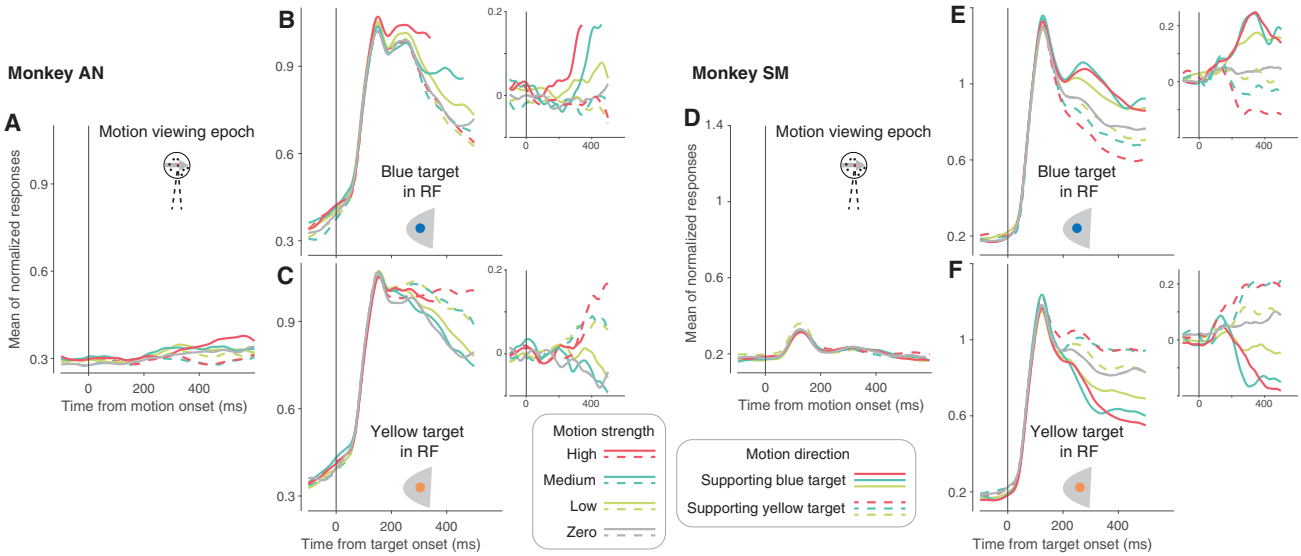

Figure 5. LIP activity during motion viewing and target selection

The graphs show average normalized responses as a function of time aligned to motion onset or target onset. Data from the two monkeys are shown separately (left, AN, 29 neurons; right, SM, 31 neurons).

(A and D) Responses aligned to motion onset. All trials are included.

(B, C, E, and F) Data from trials in which the blue (B and E) or yellow (C and F) target were in the neuronal response field. The responses are aligned to the onset of the target. In (B) and (C), traces extend until at least 33% of the trials have not terminated. Insets show residual responses after removal of the large visual response to the target. They isolate the component of the response that is controlled by the strength and direction of motion. In all panels, coherences are grouped as high (±64% and ±32%), medium (±16%), low (±8% and ±4%), and 0%. Grouping of the direction of motion (for coherences ≠0%) is based on the preferred color-motion association for each neuron. This was consistent with the association the monkey had learned between motion direction and target color, except for six neurons in monkey-SM for which the association was reversed (see STAR Methods). The responses aligned to saccade are shown in Figure S4.

in the neural response field (RF).[23,29] The present study differs from previous reports in two critical aspects: (1) the choice targets were not visible during motion viewing, and (2) the locations of the choice targets were unpredictable. Under these conditions, the neural responses accompanying motion viewing were only weakly modulated by motion strength in monkey-AN (Figure 5A) and unmodulated in monkey-SM (Figure 5D).

The action-selection epoch begins with the appearance of the color-choice targets. When a target was in the neural RF, it elicited a strong visual response beginning ~50 ms after onset (Figures 5B, 5C, 5E, and 5F), consistent with previous reports.[30] The subsequent evolution of the response reflected both the strength and direction of the RDM stimulus that had been presented in the previous epoch. To better visualize the relationship between the neuronal response and the previously presented RDM stimulus, we removed the visual response (STAR Methods). The residual responses (Figures 5B, 5C, 5E, and 5F [insets]) were effectively detrended with respect to any influences that were unaffected by the strength and direction of motion. The residual responses exhibit a clear dependency on the strength and direction of the RDM.

To quantify the rate of change of residual responses (buildup rate), we identified the time at which the raw responses first diverge. For each neuron, we then computed the buildup rate for each coherence as the slope of a line fit to the average of the residual firing rates. Each point displayed in Figure 6 is the mean buildup rate across neurons. These buildup rates exhibited a linear dependence on motion strength. For monkey-SM, the linear dependence was statistically significant in all four conditions (i.e., for all combinations of direction of motion and color of target in the RF; Table S1). For monkey-AN, the linear dependence was statistically significant in three of the four combinations of motion and direction ($p<0.05$, Table S1). The buildup rates are comparable with those obtained when the motion is viewed in the presence of saccadic choice targets (e.g., see Figure 3G in Shushruth et al.[31]).

Thus, for both monkeys, the neural responses during action selection exhibit the hallmark of a decision variable, which must be informed by evidence acquired earlier. This is consistent with the pattern of go-RTs in monkey-AN, which also supports sequential sampling of evidence during the action-selection epoch. Analyses of the time-dependent changes in response variance and autocorrelation lend additional support for sampling of noisy evidence from memory in both monkeys.

The coherence dependent ramping evident in trial-averaged response residuals could reflect the deterministic component of evidence accumulation—a ramp with slope equal to the statistical expectation of the momentary evidence for yellow or blue. On single trials, theoretically, the decision variable also includes an accumulation of noise. This is the diffusion component of drift-diffusion, which is thought to explain stochastic choice and variable decision times. Although suppressed in the trial averages, the diffusion component can be detected in the evolution of the variance and autocorrelation of the neural firing rates (variance and correlation of the conditional expectation: VarCE and CorCE, respectively[31–33]; STAR Methods; Equations 12, 13, 14, 15, and 16).

We computed the VarCE across trials in the epoch of putative accumulation (coincident with the time of the buildup). The

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

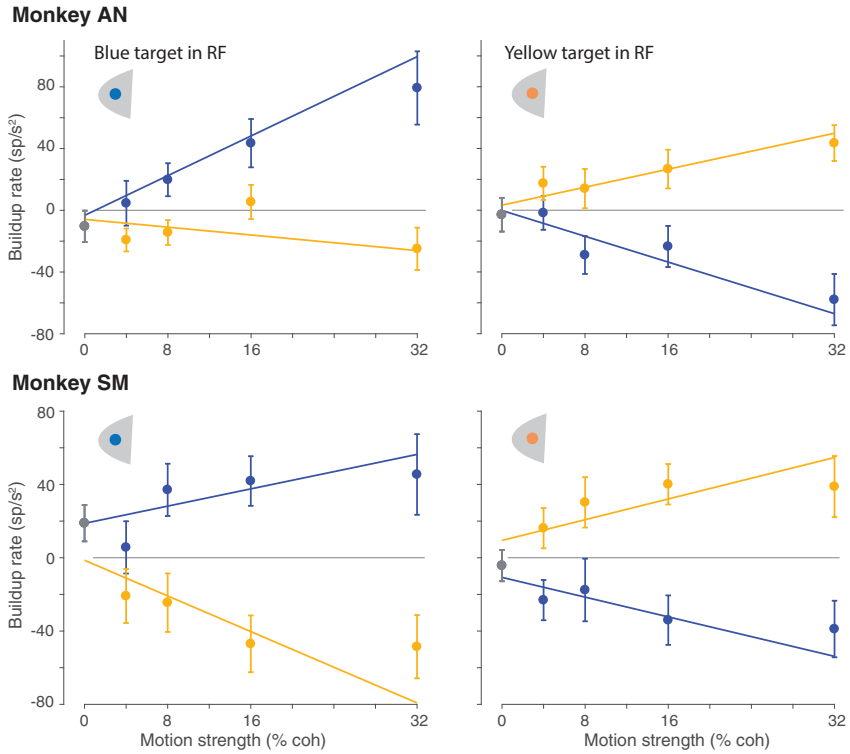

Figure 6. Buildup of neural activity depends on the strength and direction of motion

Buildup rates were estimated for each neuron, using trials with the same motion strength, direction, and color-target in the response field (top, monkey-AN; bottom, monkey-SM). Symbols are averages across neurons (error bars are SEM) The lines in the graph are weighted least square fits to the average buildup rates, grouped by motion direction. The 0% coherence point (gray) is included in both weighted regressions in each panel.

See also Table S1.

VarCE underwent a linear increase as a function of time over most of this epoch (Figures 7A and 7F). This is the pattern expected for partial sums (i.e., accumulation up to time, *t*) of independent samples of noise. The autocorrelation of the responses (CorCE) also showed signatures of a diffusion process: a decrease in autocorrelation as a function of the time separation between the bins (i.e., lag), and an increase in autocorrelation between adjacent bins as a function of time (Figures 7C, 7E, 7H, and 7J). The estimated autocorrelation pattern for both monkeys hewed closely to the theoretical predictions (Figures 7D [R^2 = 0.84, monkey-AN] and 7I [R^2 = 0.89, monkey-SM]). Such conformance lends further support for the conclusion that evidence integration in both monkeys occurs in the action-selection epoch. Indeed, the same analyses applied to the neural responses in the motion-viewing epoch fail to conform to the theoretical predictions of diffusion (i.e., integration of noisy evidence) (R^2 = 0.45, monkey-AN; R^2 = 0.26, monkey-SM; Figure S5). This was already obvious from the response averages in Figures 5A and 5D. The demonstration of a dynamic process of evidence accumulation during the action-selection epoch rules out the possibility that the LIP activity in the action-selection epoch is simply a reflection of confidence in a decision that was made during motion viewing. If so, the neural activity in the action-selection epoch should not represent the accumulation of noisy evidence.

Importantly, analyses of the neural responses support the thesis that both monkeys form their decisions in the action selection epoch. Moreover, they do so through the accumulation of noisy samples of evidence to a threshold. The strategy is strikingly different from the previous studies[12,14] and it is all the more remarkable because it holds across two variants of the task. Below, we consider the reasons why our monkeys postpone decision-making until the action selection. The fact that they do implies that their decisions are guided by evidence stored in short-term memory. This conclusion is at least equal in importance to the deferment of the decision process until action selection.

DISCUSSION

The study of perceptual decision-making in monkeys has provided insights into the process by which sequential samples of sensory evidence are accumulated over time.[29,34] A peculiar observation in these studies is that the accumulation of evidence is instantiated by neurons associated with motor planning.[23,33,35–37] This observation has led to the proposal that perceptual decision-making is embodied as a choice between potential actions.[7,8] Yet, monkeys can make perceptual decisions when they are unsure of the action that will be required of them to report their decision.[12,14–17] We investigated how monkeys accumulate sensory evidence under these circumstances, using monkeys that had never learned an association between the decision and the action to report it. The monkeys learned to associate leftward and rightward motion with the color of choice targets, which appeared at unpredictable locations after motion viewing.

We had anticipated that they would not represent evidence in the form of an oculomotor plan but rather as a plan to choose the appropriate color. Instead, we found that the monkeys formed their decisions after the color-choice targets appeared—that is, during the period of action selection—after the source of sensory evidence had been extinguished. Both monkeys based their decisions on samples of evidence that must have been retrieved from short-term memory. Monkey-AN developed this strategy spontaneously; monkey-SM did not, but it appears to have adopted this strategy once we imposed a second waiting period during the action-selection epoch. The striking change was evident in the longer time span of stimulus information used to inform decisions (Figure 4C) to achieve a level of proficiency comparable with monkey-AN and many others we have trained on direction discrimination tasks. The go-RT from monkey-AN exhibited one peculiar feature. The difference in nondecision times for blue and yellow choices was nearly as long as the entire range of go-RT for either choice. In Data S1, we show that the

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

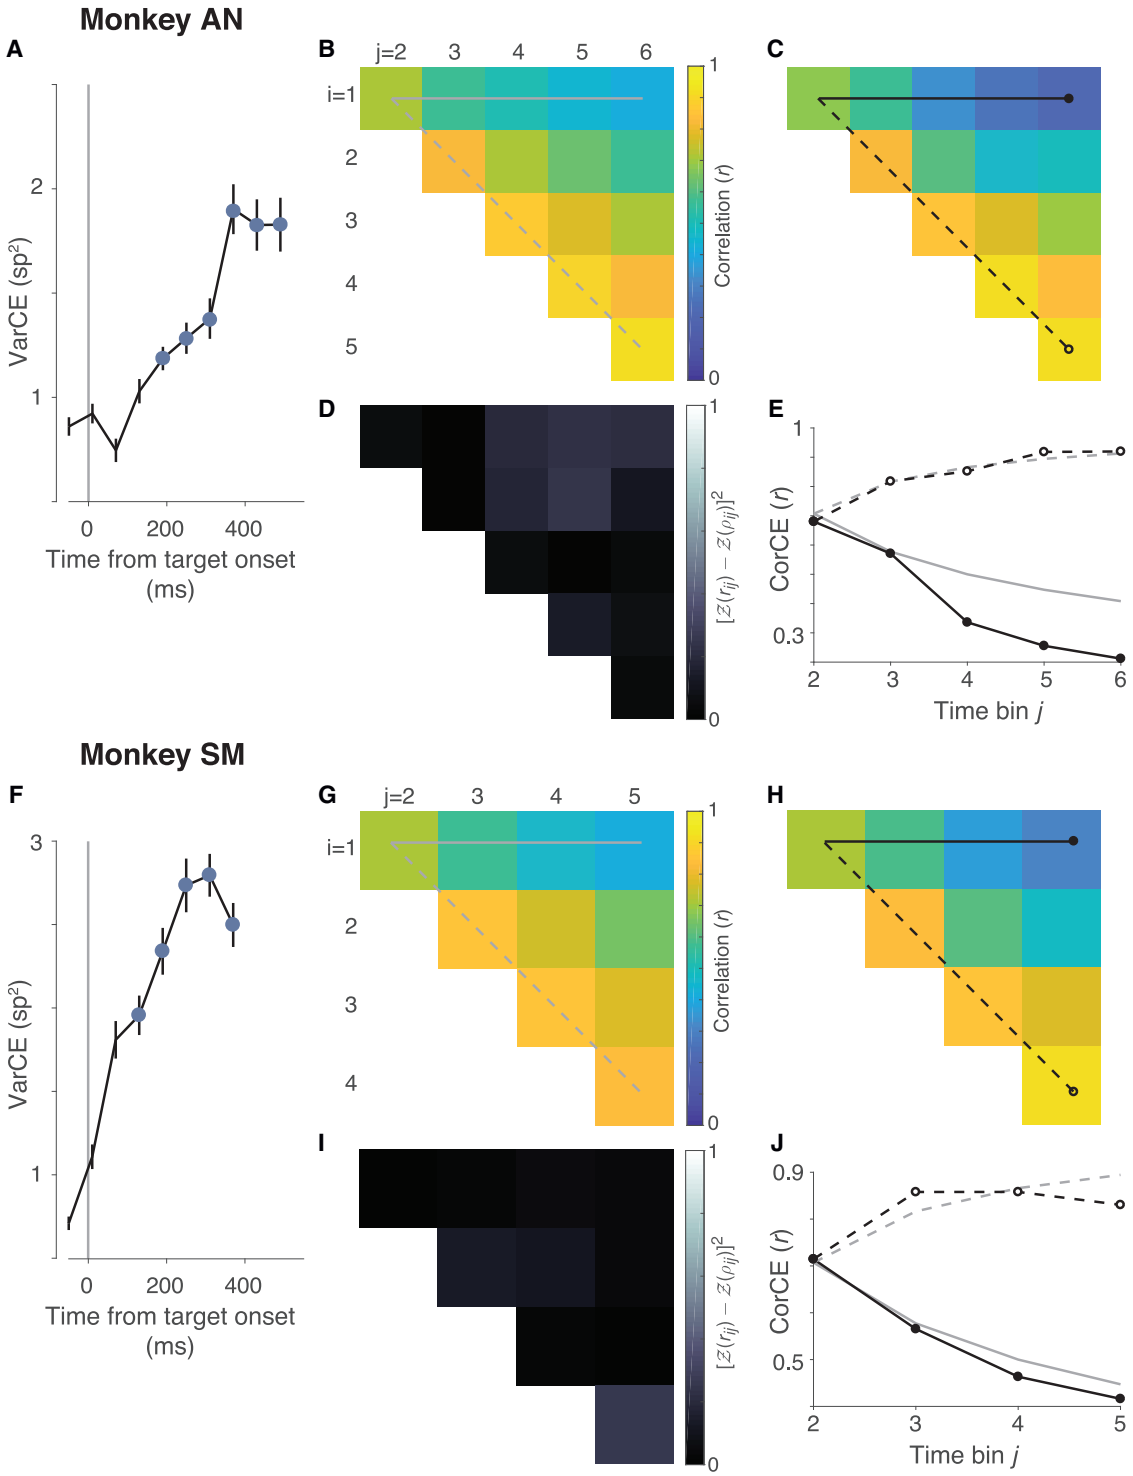

Figure 7. Variance and autocorrelation of decision-related neural responses during action selection

The analyses depicted here evaluate predictions that the neural activity during action-selection epoch on single trials includes a representation of accumulated noise.

(A and F) Variance of neural responses aligned to target onset. Filled symbols are estimates of the variance of the conditional expectations (VarCEs) of the spike counts in 60 ms bins spanning the putative integration epoch. Error bars are SE.

(B and G) Theoretical correlations between the cumulative sums of independent, identically distributed random numbers from the 1st to i^{th} and from 1st to j^{th} samples. The unique values of the correlation matrix are displayed as an upper triangular matrix. The horizontal solid line shows the correlation between the first

(legend continued on next page)

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

asymmetric go-RTs are explained by a model in which the monkey makes two decisions in series[22] but is willing to terminate with a blue choice if there is sufficient evidence.

For both monkeys, however, neural recordings from area LIP provided further confirmation that a sampling process transpired during the action-selection epoch. On trials when one of the color-choice targets appeared in the neural RF, it produced a visual response plus a signal reflecting the direction and strength of the previously presented motion. The time course of the evolution was characteristic of an integration process—more specifically, the integration of noisy evidence acquired from the stimulus. Although the memory requirements for the protracted integration of evidence may seem daunting, in Data S1, we show that it is not necessary to "replay" the entire sequence of evidence samples during the action-selection epoch; instead, storing a few samples of evidence is sufficient to achieve high levels of accuracy in the task (see also Kang et al.[22]).

A common strategy to dissociate a decision from a plan of action exploits the delayed match-to-sample design,[1] wherein a subject evaluates a sample stimulus and then, after a short delay, is presented with a second stimulus, which is compared with the first and classified as the same or different. It is assumed that the subject forms a decision about the identity or category membership of the sample before the test stimulus is presented and holds the outcome of this categorical decision in working memory. Using this approach, it has been shown that monkeys can report whether the test and sample belong to the same category[1,2,4,38] or share similar properties, such as magnitude,[39] numerosity,[40] or speed/direction.[41] These studies focus mainly on neural activity in association cortex during the sample and delay periods. This activity often varies systematically with the relevant properties of the sample stimulus and is thus interpreted as a decision that is independent of any planned action. Our results suggest an alternative interpretation. Instead of processing the sample stimulus to make a decision about category, it is processed as an instruction to brain circuits that organize the response to the test stimulus. The instruction might establish a criterion to classify the test, or it might establish the appropriate sensory-response mapping. Such a mechanism has been documented in a simple olfactory delayed match-to-sample task in mice.[42] It does not require deciding about the sample; it requires enacting a memory, cued by the sample, of the appropriate sensory-response mapping between test stimulus and behavioral response.

Two earlier studies of abstract perceptual decision-making used tasks similar to ours but reached the opposite conclusion. The task in the Gold and Shadlen study[12] was nearly identical to ours, but their monkeys failed to exhibit any signs of deliberation in the action-selection epoch. The saccadic latencies were ~200 ms from color-target appearance, suggesting the monkeys had formed their decision about the color rule before the targets appeared. The only salient difference with the present study is that their monkeys had been trained previously to associate motion with eye movements to targets. We suspect that having learned to accumulate evidence for motion as an evolving plan to make a saccade, they were able to form a decision in another intentional way—for color rule instead of target location. A similar explanation applies to the study by Bennur and Gold.[14] Their monkeys made decisions in the presence of saccadic choice targets. In the version of their task that resembles ours (their version 3), the monkeys were required to associate left and right motion with up and down targets or with down and up targets, depending on the target color cued after the motion had been shown. Decision-related changes in LIP activity were apparent during motion viewing and before the action-selection epoch. As in the Gold and Shadlen study,[12] the monkeys had been trained on a direct association between direction and an action and thus required only a slight elaboration: to switch the stimulus-response associations in accordance with the color cue.

Our result was anticipated by Wang and colleagues,[15] who used a spatial integration task with separate evaluation and action-selection epochs. The task structure used for one monkey resembles our *go*-task. It imposes a delay between the extinction of a static discriminandum and the presentation of the choice options. Similar to our monkey-AN, their monkey-T exhibited go-RTs that depended on the strength of the evidence experienced beforehand. They also report that the rate of rise of neuronal responses in Area PMd during action selection was dependent on stimulus strength. The Wang study also supports the hypothesis that sampling of evidence from memory may be necessary to form a perceptual decision when the evidence is provided before it is possible to accommodate it in an intentional context.

The near limitless capacity for abstraction in humans gives an impression of disembodied ideation. Humans can evaluate propositions about the world—what things are and what categories they belong to—without using them as objects of possible action. An alternative formulation, rooted in ecological perception,[43] suggests that knowledge of the environment is in the service of what we might do in the form of considerations and intentions.[44,45] One activity humans pursue is reporting to other humans. The conversion of a provisional report to an action, such as "look at the blue spot if the motion is rightward," permits humans to form a decision before the action is specified. The same logical structure applies to what monkeys—previously

sample and the cumulative sum to the j^{th} sample (lag = $j - i$). It shows decreasing correlation as a function of lag. The dashed line identifies the first juxtadiagonal set of correlations between pairs with the same lag = 1. It shows an increase in correlation as a function of time of the pairs of samples.

(C and H) Correlations estimated from the neural response. These are the correlations between the conditional expectation of the spike counts (CorCEs) in time bins i and j. If the rates on single trials are determined by unbounded drift-diffusion, these correlations should match the values in (B) and (G). The top row and first juxtadiagonal are identified as in (B) and (G).

(D and I) Deviance of the estimated correlations from theoretical correlations (sum of squares measure).

(E and J) Comparison of theoretical and estimated correlations in the top row and first juxtadiagonal of the matrices in (B), (G), (C), and (H). Gray traces show the theoretical values in (B) and (G). Black lines connect the CorCE values in (C) and (H). Line and symbol styles distinguish the top row (correlation as a function of lag) and first juxtadiagonal (correlations of neighboring bins as a function of time).

See also Figure S5.

trained to associate rightward/leftward motion with an eye movement to the right/left—can achieve in abstract decision tasks similar to ours. If humans are not informed about the axis of discrimination until after the motion has been viewed, then similar to the monkey, humans too must rely on memory.[46] Further, studies of iconic short-term memory demonstrate that such memory can be formed strategically to anticipate knowledge of the operations that may be required.[18,47,48] Thus, both monkeys in our task must have learned to store the appropriate motion information in short-term memory buffers to enable action selection based on the colors of the choice targets. The site of such buffers is unknown, but there is literature to suggest sensory areas, such as MST[49] and prefrontal cortex,[41,50] as possible candidates. It is also possible that LIP neurons with RFs that overlap the RDM play a role, as they are known to exhibit direction selectivity.[2,51]

The difficulty that our abstract decision task poses for naive monkeys might raise concerns about the relevance of our finding to human cognitive function. The abstract decision in our task requires the animal to either (1) build a hierarchical decision in which the outcome of the motion decision substitutes for the colored object to instruct the blue-yellow choice or (2) store evidence from motion to resolve the subsequent color choice. The hierarchical strategy is the one humans appear to exercise, as the effect of the strength of evidence on go-RTs is minimal in human subjects.[52,53]

At first glance, the hierarchical strategy might appear to be the more sophisticated of the two. It is more complex, and the nested structure seems similar to a building block for language. However, the second strategy also connects to a sophisticated element of cognition: the capacity to use recent, but temporally non-adjacent, information to guide a decision. This is critical for learning causal relations, and it too plays a role in language. We make strategic use of short-term memory to store semantic content (analogous to samples of evidence), which we incorporate in locution later—analogous to action selection—in accordance with syntactic demands. The process is embarrassingly vivid when we lose the train of our thought. Such embarrassment is mitigated by the strategic use of short-term memory, adhering to the old adage, "put your mind in gear before you put your mouth in motion" (A. Shadlen, personal communication).

Clearly, expressions of perceptual decisions through eye movements and expressions of ideas through language invite more contrast than comparison, but the structural similarity may prove useful for neurobiology. Thus, it is that the monkey's crude approximation to abstract decision-making elucidates a critical building block of our own ideation.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.cub.2022.03.014.

ACKNOWLEDGMENTS

The research was supported by Howard Hughes Medical Institute, an R01 from NEI (M.N.S., R01EY11378), an R01 from NIMH (M.N.S., R01MH122513), and an Young Investigator grant from BBRF (S.S., #23556). We thank Brian Madeira and Cornel Duhaney for technical support and animal care. We thank Danique Jeurissen, Gabriel Stine, Natalie Steinemann, Simon Kelly, and Redmond O'Connell for comments on an earlier draft of the manuscript and Arthur Shadlen (grandpa) for advice and/or reprimand. We are especially grateful to animal technicians, veterinary staff, and other essential workers at the Zuckerman Institute who made it possible to collect the data set from monkey-SM during the SARS-CoV-2 pandemic.

AUTHOR CONTRIBUTIONS

Conceptualization, S.S. and M.N.S.; software, S.S., A.Z., and M.N.S.; investigation, S.S.; formal analysis, S.S., A.Z., and M.N.S.; writing—original draft, S.S. and M.N.S.; writing—review & editing, S.S., A.Z., and M.N.S.; funding acquisition, S.S. and M.N.S.

DECLARATION OF INTERESTS

MNS is a member of the advisory board of Current Biology. The other authors declare no competing interests.

Received: December 17, 2021
Revised: March 2, 2022
Accepted: March 3, 2022
Published: March 29, 2022

REFERENCES

1. Freedman, D.J., Riesenhuber, M., Poggio, T., and Miller, E.K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. Science *291*, 312–316. https://www.ncbi.nlm.nih.gov/pubmed/11209083.
2. Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85–88. http://www.ncbi.nlm.nih.gov/pubmed/16936716.
3. Seger, C.A., and Miller, E.K. (2010). Category learning in the brain. Annu. Rev. Neurosci. *33*, 203–219. https://www.ncbi.nlm.nih.gov/pubmed/20572771.
4. Goodwin, S.J., Blackman, R.K., Sakellaridi, S., and Chafee, M.V. (2012). Executive control over cognition: stronger and earlier rule-based modulation of spatial category signals in prefrontal cortex relative to parietal cortex. J. Neurosci. *32*, 3499–3515. https://www.ncbi.nlm.nih.gov/pubmed/22399773.
5. Tenenbaum, J.B., Kemp, C., Griffiths, T.L., and Goodman, N.D. (2011). How to grow a mind: statistics, structure, and abstraction. Science *331*, 1279–1285. https://www.ncbi.nlm.nih.gov/pubmed/21393536.
6. Chafee, M.V., and Crowe, D.A. (2012). Thinking in spatial terms: decoupling spatial representation from sensorimotor control in monkey posterior

parietal areas 7a and LIP. Front. Integr. Neurosci. *6*, 112. https://www.ncbi.nlm.nih.gov/pubmed/23355813.

7. Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. Philos. Trans. R. Soc. Lond. B Biol. Sci. *362*, 1585–1599. http://www.ncbi.nlm.nih.gov/pubmed/17428779.
8. Shadlen, M.N., Kiani, R., Hanks, T.D., and Churchland, A.K. (2008). Neurobiology of decision making: an intentional framework. In Better than Conscious? Decision Making, the Human Mind, and Implications for Institutions, C. Engel, and W. Singer, eds. (MIT Press), pp. 71–101.
9. Cisek, P., and Kalaska, J.F. (2010). Neural mechanisms for interacting with a world full of action choices. Annu. Rev. Neurosci. *33*, 269–298. https://www.ncbi.nlm.nih.gov/pubmed/20345247.
10. Klaes, C., Westendorff, S., Chakrabarti, S., and Gail, A. (2011). Choosing goals, not rules: deciding among rule-based action plans. Neuron *70*, 536–548. http://www.ncbi.nlm.nih.gov/pubmed/21555078.
11. Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791–806. http://www.ncbi.nlm.nih.gov/pubmed/24183028.
12. Gold, J.I., and Shadlen, M.N. (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. J. Neurosci. *23*, 632–651. http://www.ncbi.nlm.nih.gov/pubmed/12533623.
13. Genovesio, A., Tsujimoto, S., and Wise, S.P. (2009). Feature- and order-based timing representations in the frontal cortex. Neuron *63*, 254–266. http://www.ncbi.nlm.nih.gov/pubmed/19640483.
14. Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J. Neurosci. *31*, 913–921. http://www.ncbi.nlm.nih.gov/pubmed/21248116.
15. Wang, M., Montanède, C., Chandrasekaran, C., Peixoto, D., Shenoy, K.V., and Kalaska, J.F. (2019). Macaque dorsal premotor cortex exhibits decision-related activity only when specific stimulus-response associations are known. Nat. Commun. *10*, 1793. https://www.ncbi.nlm.nih.gov/pubmed/30996222.
16. Quinn, K.R., Seillier, L., Butts, D.A., and Nienborg, H. (2021). Decision-related feedback in visual cortex lacks spatial selectivity. Nat. Commun. *12*, 4473.
17. Horwitz, G.D., Batista, A.P., and Newsome, W.T. (2004). Representation of an abstract perceptual decision in macaque superior colliculus. J. Neurophysiol. *91*, 2281–2296.
18. Sperling, G. (1960). The information available in brief visual presentations. Psychol. Monogr.: Gen. Appl. *74*, 1–29.
19. Schall, J.D., and Thompson, K.G. (1999). Neural selection and control of visually guided eye movements. Annu. Rev. Neurosci. *22*, 241–259. https://www.ncbi.nlm.nih.gov/pubmed/10202539.
20. Seideman, J.A., Stanford, T.R., and Salinas, E. (2018). Saccade metrics reflect decision-making dynamics during urgent choices. Nat. Commun. *9*, 2907. https://www.ncbi.nlm.nih.gov/pubmed/30046066.
21. Tanaka, T., Nishida, S., and Ogawa, T. (2015). Different target-discrimination times can be followed by the same saccade-initiation timing in different stimulus conditions during visual searches. J. Neurophysiol. *114*, 366–380. https://www.ncbi.nlm.nih.gov/pubmed/25995344.
22. Kang, Y.H., Löffler, A., Jeurissen, D., Zylberberg, A., Wolpert, D.M., and Shadlen, M.N. (2021). Multiple decisions about one object involve parallel sensory acquisition but time-multiplexed evidence incorporation. Elife *10*, e63721. https://www.ncbi.nlm.nih.gov/pubmed/33688829.
23. Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. *22*, 9475–9489. http://www.ncbi.nlm.nih.gov/pubmed/12417672.
24. Palmer, J., Huk, A.C., and Shadlen, M.N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. J. Vis. *5*, 376–404.
25. Hanks, T.D., Mazurek, M.E., Kiani, R., Hopp, E., and Shadlen, M.N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. J. Neurosci. *31*, 6339–6352.
26. Stine, G.M., Zylberberg, A., Ditterich, J., and Shadlen, M.N. (2020). Differentiating between integration and non-integration strategies in perceptual decision making. Elife *9*, e55365. https://www.ncbi.nlm.nih.gov/pubmed/32338595.
27. Gnadt, J.W., and Andersen, R.A. (1988). Memory related motor planning activity in posterior parietal cortex of macaque. Exp. Brain Res. *70*, 216–220. http://www.ncbi.nlm.nih.gov/pubmed/3402565.
28. Barash, S., Bracewell, R.M., Fogassi, L., Gnadt, J.W., and Andersen, R.A. (1991). Saccade-related activity in the lateral intraparietal area. II. Spatial properties. J. Neurophysiol. *66*, 1109–1124. https://www.ncbi.nlm.nih.gov/pubmed/1753277.
29. Shadlen, M.N., and Newsome, W.T. (1996). Motion perception: seeing and deciding. Proc. Natl. Acad. Sci. USA *93*, 628–633. http://www.ncbi.nlm.nih.gov/pubmed/8570606.
30. Bisley, J.W., and Goldberg, M.E. (2003). Neuronal activity in the lateral intraparietal area and spatial attention. Science *299*, 81–86. https://www.ncbi.nlm.nih.gov/pubmed/12511644.
31. Shushruth, S., Mazurek, M., and Shadlen, M.N. (2018). Comparison of decision-related signals in sensory and motor preparatory responses of neurons in area LIP. J. Neurosci. *38*, 6350–6365. https://www.ncbi.nlm.nih.gov/pubmed/29899029.
32. Churchland, A.K., Kiani, R., Chaudhuri, R., Wang, X.J., Pouget, A., and Shadlen, M.N. (2011). Variance as a signature of neural computations during decision making. Neuron *69*, 818–831. https://www.ncbi.nlm.nih.gov/pubmed/21338889.
33. de Lafuente, V., Jazayeri, M., and Shadlen, M.N. (2015). Representation of accumulating evidence for a decision in two parietal areas. J. Neurosci. *35*, 4306–4318. https://www.ncbi.nlm.nih.gov/pubmed/25762677.
34. Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. Nature *341*, 52–54. https://www.ncbi.nlm.nih.gov/pubmed/2770878.
35. Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. *2*, 176–185. http://www.ncbi.nlm.nih.gov/pubmed/10195203.
36. Ding, L., and Gold, J.I. (2010). Caudate encodes multiple computations for perceptual decisions. J. Neurosci. *30*, 15747–15759. https://www.ncbi.nlm.nih.gov/pubmed/21106814.
37. Ding, L., and Gold, J.I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. Cereb. Cortex *22*, 1052–1067. http://www.ncbi.nlm.nih.gov/pubmed/21765183.
38. Fitzgerald, J.K., Freedman, D.J., and Assad, J.A. (2011). Generalized associative representations in parietal cortex. Nat. Neurosci. *14*, 1075–1079. https://www.ncbi.nlm.nih.gov/pubmed/21765425.
39. Genovesio, A., Tsujimoto, S., and Wise, S.P. (2011). Prefrontal cortex activity during the discrimination of relative distance. J. Neurosci. *31*, 3968–3980. https://www.ncbi.nlm.nih.gov/pubmed/21411640.
40. Nieder, A., Freedman, D.J., and Miller, E.K. (2002). Representation of the quantity of visual items in the primate prefrontal cortex. Science *297*, 1708–1711. https://www.ncbi.nlm.nih.gov/pubmed/12215649.
41. Hussar, C.R., and Pasternak, T. (2009). Flexibility of sensory representations in prefrontal cortex depends on cell type. Neuron *64*, 730–743. https://www.ncbi.nlm.nih.gov/pubmed/20005828.
42. Wu, Z., Litwin-Kumar, A., Shamash, P., Taylor, A., Axel, R., and Shadlen, M.N. (2020). Context-dependent decision making in a premotor circuit. Neuron *106*, 316–328.e6. https://www.ncbi.nlm.nih.gov/pubmed/32105611.
43. Gibson, J.J. (1979). The Ecological Approach to Visual Perception (Psychology Press & Routledge Classic Editions).
44. Merleau-Ponty, M. (1962). Phenomenology of Perception (Routledge).

45. Clark, A. (1997). Being There: Putting Brain, Body, and World Together Again (MIT Press).
46. Bang, D., and Fleming, S.M. (2018). Distinct encoding of decision confidence in human medial prefrontal cortex. Proc. Natl. Acad. Sci. USA *115*, 6082–6087.
47. Gegenfurtner, K.R., and Sperling, G. (1993). Information transfer in iconic memory experiments. J. Exp. Psychol. Hum. Percept. Perform. *19*, 845–866.
48. Sperling, G., and Weichselgartner, E. (1995). Episodic theory of the dynamics of spatial attention. Psychol. Rev. *102*, 503–532.
49. Teeuwen, R.R.M., Wacongne, C., Schnabel, U.H., Self, M.W., and Roelfsema, P.R. (2021). A neuronal basis of iconic memory in macaque primary visual cortex. Curr. Biol. *31*, 5401–5414.e4. https://www.sciencedirect.com/science/article/pii/S0960982221012926.
50. Zaksas, D., and Pasternak, T. (2006). Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. J. Neurosci. *26*, 11726–11742.
51. Fanini, A., and Assad, J.A. (2009). Direction selectivity of neurons in the macaque lateral intraparietal area. J. Neurophysiol. *101*, 289–305.
52. Twomey, D.M., Kelly, S.P., and O'Connell, R.G. (2016). Abstract and effector-selective decision signals exhibit qualitatively distinct dynamics before delayed perceptual reports. J. Neurosci. *36*, 7346–7352. https://www.ncbi.nlm.nih.gov/pubmed/27413146.
53. Coallier, É., and Kalaska, J.F. (2014). Reach target selection in humans using ambiguous decision cues containing variable amounts of conflicting sensory evidence supporting each target choice. J. Neurophysiol. *112*, 2916–2938.
54. Lewis, J.W., and Van Essen, D.C. (2000). Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. J. Comp. Neurol. *428*, 112–137. https://www.ncbi.nlm.nih.gov/pubmed/11058227.
55. Quick, R.F. (1974). A vector-magnitude model of contrast detection. Kybernetik *16*, 65–67.
56. Shadlen, M.N., Hanks, T.D., Churchland, A.K., Kiani, R., and Yang, T. (2006). The speed and accuracy of a simple perceptual decision: a mathematical primer. In Bayesian Brain: Probabilistic Approaches to Neural Coding, K. Doya, et al., eds. (MIT Press), pp. 209–237.
57. Ratcliff, R., and Rouder, J.N. (1998). Modeling response times for two-choice decisions. Psychol. Sci. *9*, 347–356.
58. Chang, J.S., and Cooper, G. (1970). A practical difference scheme for Fokker-Planck equations. J. Comp. Phys. *6*, 1–16.
59. Kiani, R., and Shadlen, M.N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. Science *324*, 759–764.
60. Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. J. Neurosci. *28*, 3017–3029. http://www.ncbi.nlm.nih.gov/pubmed/18354005.
61. Nawrot, M.P., Boucsein, C., Rodriguez Molina, V.R., Riehle, A., Aertsen, A., and Rotter, S. (2008). Measurement of variability dynamics in cortical spike trains. J. Neurosci. Methods *169*, 374–390.
62. Efron, B. (1987). The Jackknife, the Bootstrap, and Other Resampling Plans (CBMS-NSF Regional Conference Series in Applied Mathematics) (Society for Industrial Mathematics)).

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Behavioral and neural data	This paper	Mendeley Data: https://doi.org/10.17632/kvh98g9js6.1
Experimental models: Organisms/strains		
Macaca mulatta	California National Primate Research Center (Monkey AN) and Primate Products Inc. (Monkey SM)	N/A
Software and algorithms		
Matlab	Mathworks Inc.	N/A
Analysis code	This paper	Github: https://github.com/sshruth/Shushruth_etal_CurrBiol2022

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Michael N Shadlen (shadlen@columbia.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Behavioral and neural data have been deposited to Mendeley Data and are publicly available as of the date of publication. DOI is listed in the key resources table.
- Original code to analyze the data has been deposited in a Github repository and is publicly available as of the date of publication. URL is listed in the key resources table.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All training, surgery, and experimental procedures were conducted in accordance with the Public Health Service Policy on Humane Care and Use of Laboratory Animals. Experiments were approved by the Columbia University Institutional Animal Care and Use Committee (IACUC) under protocol number AC-AAAW4454. Two adult macaque monkeys (one 12 year old female, AN; one 9 year old male, SM) were used as subjects in this study.

METHOD DETAILS

Behavioral task

The monkeys performed a behavioral task in which they decided whether the net direction of a stochastic random-dot motion (RDM) stimulus was to the left or right. The animals initiated trials by fixating on a point (fixation point; FP) presented on an otherwise black screen. The RDM stimulus was then presented within a circular aperture (radius 2.5° or 3°) centered on the FP. The first three frames of the stimulus consist of white dots randomly plotted at a density of 16.7 dots$\cdot$deg$^{-2}\cdot$s^{-1}. From the fourth frame, each dot from three frames before is replotted—either displaced to the right or left, or at a random location. The probability with which a dot is displaced to the right or left determines the stimulus strength (coherence; C) and on each trial, C was randomly chosen from the set $\{0, \pm 0.04, \pm 0.08, \pm 0.16, \pm 0.32, \pm 0.64\}$, the positive sign indicating rightward motion. The motion strengths and the two directions were randomly interleaved. The stimulus was presented for a variable duration drawn from a truncated exponential distribution (range 350–800 ms, mean 500 ms). Two targets, one blue and one yellow, were presented after a short delay (333 ms, monkey-AN; 200 ms, monkey-SM) at eccentric locations that varied across trials. The monkeys had to report the perceived direction of motion by choosing the target of the associated color (blue for rightward and yellow for leftward, monkey-AN; vice-versa for monkey-SM). In the *go*-task (Figure 1, top), the FP was extinguished simultaneously with the onset of the colored targets. In the

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

wait-task (Figure 1, bottom), the FP stayed on for a variable duration (drawn from an inverted truncated exponential distribution, range 400–1200 ms, mean 900 ms).

Electrophysiology

We recorded spikes from 60 well-isolated single units (29 monkey-AN; 31 monkey-SM) in area LIP_v.[54] The sample size for monkey-AN was limited by a serious illness, leading to euthanasia. Monkey-SM was just ready for recording when New York entered lockdown owing to the SARS-CoV2 pandemic. We justified as mission-critical the need to obtain a neural data set of power equivalent to the first monkey. The neural data were analyzed separately for each monkeys, and all of the central findings are statistically significant for each separately.

MRI was used to localize LIP_v and to guide the placement of recording electrodes. We screened for neurons that exhibited spatially selective persistent activity using a memory-guided saccade task.[27] In the screening task, a target is flashed in the periphery while the monkey fixates on a central spot. The monkey has to remember the location of the target and execute a saccade to that location when instructed. The response field (RF) of each neuron was identified as the region of visual space that elicited the highest activity during the interval between the target flash and the eventual saccade.

During recording experiments, the locations for target presentation were chosen based on the location of the neuronal RF. For monkey-AN, six locations (including the RF) were chosen, equally spaced on an imaginary circle. On each trial, pairs of locations $2\pi/3$ rad apart were pseudorandomly selected to display the targets. The RF location was oversampled to increase the concentration of trials from which we could analyze neural data. A similar approach was taken in monkey-SM except that the number of possible locations were restricted to four and the target pairs were situated $\pi/2$ rad apart. Each colored target appeared in the RF on 33% and 28% of the trials for monkey-AN and monkey-SM, respectively. Note that the monkeys were trained to generalize across a larger set of locations and these spatial restrictions on target locations were implemented during recording sessions.

QUANTIFICATION AND STATISTICAL ANALYSIS

Analyses of behavioral data

Both monkeys were taught the association between the color of the target and the direction of motion using only the strongest motion strength (±64% coh). We then introduced the next easiest stimulus strength (±32% coh) and continued to add more coherences until we reached 0%. To assess the improvement of sensitivity across training sessions, we fit the choice-accuracy, $P_{correct}$, as a function of motion strength, $|C|$, for each session with a Weibull function[55] of the following form:

$$P_{correct} = 0.5 + (0.5 - \lambda)\left[1 - e^{-\left(\frac{|C|}{\alpha}\right)^{\beta}}\right] \quad \text{(Equation 1)}$$

where λ is the lapse rate, β is the shape parameter, and α is the threshold if $\lambda = 0$. We interpolated from these fits the $|C|$ that supports 75% accuracy and report that as the threshold (e.g., Figure S1).

The quantification of learning rate is from the introduction of the ±32% coh. The rates (e.g., Figure S1) are based on approximate number of sessions (and trials), because both monkeys experienced interruptions to training. For interruptions lasting more than a month, we excluded sessions after resumption until the monkey re-established thresholds similar to those prior to the interruption. This was also the case for monkey-SM when we switched from the *go*-task to the *wait*-task.

In Figures 2A, 2B, and 4B, we fit the choices of the monkeys with a logistic model of the following form:

$$P_{right} = \lambda + (1 - 2\lambda)[1 + \exp(-(\beta_0 + \beta_1 C)]^{-1} \quad \text{(Equation 2)}$$

where λ, β_0, β_1 are fit parameters. This is also the analytic solution to symmetric diffusion (when $\lambda = 0$), and thus comparable to the fits of the models which are constrained to explain both choice and go-RT.

The go-reaction times (go-RT) of monkey-AN were fit with a bounded evidence accumulation model,[56] modified to account for errors at the highest motion strength. In this model, the instantaneous evidence about motion at each time step is assumed to arise from a normal distribution with variance Δt and mean $\kappa(C + C_0)\Delta t$, where C is the signed motion coherence, C_0 is bias (expressed in units of signed coherence), and κ is a scaling parameter. The samples of instantaneous evidence are assumed to be independent and accumulated over time until the decision terminates, which occurs when the accumulated evidence reaches one of the bounds $\pm B$ leading to the choice of one of the targets. The mean go-RT is the expectation of the time taken for the accumulated evidence to reach the bound plus a constant—the non-decision time t_{nd} comprising all contributions to the go-RT that do not depend on motion strength/direction and bias (e.g., sensory and motor delays). To account for asymmetric go-RTs in some configurations, we used two different non-decision times (t^b_{nd} and t^y_{nd}) for blue and yellow target choices respectively.

In this framework, the mean go-RT for correct choices (i.e. choices consistent with the sign of the drift rate, $\kappa[C + C_0]$) is described by

$$\tilde{T}^x(C|\theta) = \frac{B}{\kappa(C + C_0)} \tanh[\kappa(C + C_0)B] + t^x_{nd} \quad \text{(Equation 3)}$$

where $x \in \{b, y\}$ and θ are the fitted parameters $\{B, \kappa, C_0, t^b_{nd}, t^y_{nd}\}$. The proportion of blue choices is determined by three of these parameters:

$$\tilde{P}^b(C|B,\kappa,C_0) = [1+\exp(-2\kappa(C+C_0)B)]^{-1} \quad \text{(Equation 4)}$$

where $\tilde{P}^b$ is the probability of the diffusion process terminating at the bound for blue choices. We first established an estimate of the bias from a logistic fit to the choices (Equation 2), expressing the bias in units of coherence ($\zeta = \beta_0/\beta_1$). Because the model explains the go-RT only when the choice is consistent with the sign of the drift rate,[57] we used the mean go-RT for positive choices at $C+\zeta>0$ and negative choices for $C+\zeta<0$.

Informed by the patterns of error go-RTs observed at the highest coherence (Figure S2B), we attribute the errors at the highest motion strength (lapse rate, λ) to a mistaken association between the sign of the terminating bound and its corresponding color-target ("direction-color confusion"). For weaker motion strengths the same confusion converts a fraction of correct terminations to erroneous color choices and the same fraction of incorrect terminations to correct color choices. We estimated λ from Equation 2, thereby enabling conversion of $\tilde{P}^b$ to the observed proportion of blue choices (P^b):

$$P^b = \tilde{P}^b - (\lambda\tilde{P}^b) + \lambda(1-\tilde{P}^b)\,. \quad \text{(Equation 5)}$$

In our formulation, the trials with direction-color confusion inherit the t_{nd} of the motion decision (not the chosen color) and the mean observed go-RT would include contributions from the trials lost and gained from that process. The fraction of confusion trials for blue choices at coherence C is

$$f_\lambda(C) = \lambda(1-\tilde{P}^b)/P^b \quad \text{(Equation 6)}$$

and the mean go-RT for blue choices observed to be correct would be

$$T^b(C) = \tilde{T}^b(C)[1-f_\lambda(C)] + \tilde{T}^y(C)f_\lambda(C)\,. \quad \text{(Equation 7)}$$

We used a maximum likelihood procedure to fit this model to the choice and mean go-RTs on the correct (relative to ζ) choices (Figure 4A). Table S2 shows the best-fitting model parameters. For each motion coherence, we calculate the average response time on correct trials ($RT^c(C)$), its standard error ($s^c(C)$), and the number of blue and yellow choices ($n^b(C)$ and $n^y(C)$ respectively). The parameters (Φ) are fit to maximize the function,

$$\hat{\Phi} = \underset{\Phi}{\operatorname{argmax}}\left(\mathcal{L}_c^{RT}(\Phi) + \mathcal{L}^{\text{choice}}(\Phi)\right) \quad \text{(Equation 8)}$$

The first term of the right hand side of Equation 8 is defined as:

$$\mathcal{L}_c^{RT}(\Phi) = \sum_C \log[\mathcal{N}(RT^c(C); T^c(C|\Phi), s^c(C))] \quad \text{(Equation 9)}$$

where $T^c(C|\Phi)$ is the mean go-RT predicted by the model for motion strength C (correct trials only), and $\mathcal{N}(\bullet;\mu,\sigma)$ is the normal probability-density function with mean μ and standard deviation σ.

The second term of the right hand of Equation 8 aims to maximize the probability of the observed choices given binomially-distributed errors,

$$\mathcal{L}^{\text{choice}}(\Phi) = \sum_C \log\left[P\left(n^b(C); n^b(C)+n^y(C), T^b(C|\Phi)\right)\right] \quad \text{(Equation 10)}$$

where $P(k;n,p)$ is the binomial probability of observing k blue choices out of n trials, given that blue choices occur with probability p, which we take to be equal to the model's predicted proportion of blue choices for parameters Φ.

We also fit an elaborated version of the bounded evidence accumulation model to include both correct and error trials (Figure S2; Data S1). In this model, the decision bounds (B) collapse over time:

$$B = \begin{cases} B_0 \times \exp(-a(t-B_{\text{del}})) & \text{if } t \geq B_{\text{del}} \\ B_0 & \text{otherwise} \end{cases} \quad \text{(Equation 11)}$$

where B_0 is the initial bound height, a is the rate of collapse and B_{del} is the delay to onset of collapse. The non-decision time was assumed constant and equal to t_{nd}. Instead of using Equations 3 and 4, $\tilde{T}^x(C)$ and $\tilde{P}^b(C)$ are obtained by numerical solution of Fokker-Planck equations.[58,59] Again, separate non-decision times were used for decisions terminating at each of the two bounds and errors at the highest coherence were modeled as 'direction-color confusion' using the approach described above.

We augmented these analyses with psychophysical reverse correlation, to provide an empirical estimate of the epoch in which the RDM stimulus affected the choice. The motion energy on individual trials (0% coherence only) was computed using spatiotemporal filters as described in Kiani et al.[60] The sign, right minus left or vice versa, was chosen such that positive indicates stimulus evidence in support of the monkey's choice on that trial (Figures 2C, 2D, and 4C). To determine the actual duration of motion that had a significant influence on choices, we recalculated kernels using different lengths of the random dot movie shown in each trial. We report the length of time that the stimulus affects choice as the shortest movie-length that accounts for all the statistically significant bins obtained using the full-length movie.

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

Analyses of neural data

For visualization of population average firing rates (Figure 5), spike times from single trials, $s_{i=1\ldots n}$, were represented as delta functions $\delta(s_i - t)$ and convolved with an 80 ms boxcar filter. For each neuron we grouped trials based on what was presented in its RF: blue target, yellow target or neither. We averaged across trials for each group and determined the maximum of the average responses across the three groups. The responses on all individual trials were divided by this maximum to obtain normalized firing rates. The population responses shown in Figure 5 were then computed from these normalized responses using relevant subsets of trials. For the motion viewing epoch, trials were grouped based on motion direction (0° or 180°) and coherence (High: 64% & 32%; Medium: 16%; Low: 8% & 4%; and 0%). In the target onset and saccade epochs, the grouping was based on which target was shown in the neuron's RF (blue or yellow), coherence (same coherence groups as in the motion viewing epoch) and the direction of motion (preferred vs. nonpreferred). For the majority of neurons, on trials in which a target appeared in the RF, a higher response was recorded when target appearance was preceded by the associated motion direction. For six neurons in monkey-SM, the non-associated direction elicited the higher response and was designated the preferred direction. To visualize the coherence dependent buildup of activity (insets of Figures 5B, 5C, 5E, and 5F), we detrended the population responses by subtracting the average responses to the 0% and ±4% coherence conditions. This detrending was done separately for trials with each colored target in the RF.

We pursued several analyses to characterize the neural responses during the epoch of action selection, after the onset of the color-choice targets. We defined the beginning of this epoch, t_∇, as the first of three consecutive 40 ms time bins, beginning at least 50 ms after target onset, in which the average responses associated with correct choices at the strongest motion diverged ($p<0.05$, Wilcoxon rank sum test). For monkeys AN and SM $t_\nabla = 170$ and 100 ms, respectively. Our analyses focus on early decision formation, before many decisions would be expected to terminate on the more difficult conditions. For monkey-AN, we set the end of the epoch as $t_\nabla + 300$ ms or 200 ms before saccade initiation, whichever occurred first. There are no overt terminating events for monkey-SM. We therefore chose $t_\nabla + 250$ ms.

The effect of signed motion strength on build-up rate (Figure 6) was established as follows. In the epoch defined above, we computed firing rates in 20 ms bins for each trial. For each neuron we grouped trials based on the target that appeared in the RF (blue or yellow). We removed the sensory component of the responses for each group by subtracting the average responses to the 0% and ±4% coherence conditions and computed the buildup rate for each coherence (the slope across bins). We excluded the ±64% coherence conditions from this analysis because there were too few time bins for monkey-AN, owing to fast go-RT, and an early plateau in monkey-SM, owing, we suspect, to fast decision terminations. We report the population mean and SE of the buildup rates and the fit to a linear model regressing these buildup rates against signed coherence in Figure 6.

The analyses summarized in Figure 7 compare the evolution of the variance and autocorrelation of the firing rate during the epoch of putative decision formation to the expected time course of these statistics under diffusion — if the spikes are associated with latent firing rates that represent the the sum of independent, identically distributed (*iid*) random numbers. The theory and algorithm are described in previous publications.[31–33] We used the spike counts in 60 ms bins in the epoch described above. This analysis focused on trials with the three weakest motion strengths (0%, ±4% and ±8% coh) to exploit the longer duration over which the decision process unfolds in these trials. The trials are initially grouped by neuron, the 5 unique signed coherences, and the target in the RF. We used the residuals of responses for each group to remove the contribution of motion strength and direction.

Consider, for the moment, trials from one neuron and one time bin. For each trial, i, we measure the raw spike count and compute the residual count by removing the mean count for all trials of the same combination of signed coherence and the color of the target in the RF, j,

$$\nu_{ij} = n_{ij} - \overline{n}_j \qquad \text{(Equation 12)}$$

The total variance across trials, is

$$\mathrm{Var}[\nu_{ij}]_i = \mathrm{Var}[n_{ij}]_i \qquad \text{(Equation 13)}$$

because variance is a central moment. We assume the noise component of evidence samples is the same for all the motion strengths. Therefore the variance across all combinations of signed coherence and the color of the target in the RF is $\mathrm{Var}[\nu], \forall\, j$. This is the total variance of the counts in the time bin under consideration. We are interested in the variance of the latent rate that gives rise to the spike counts on each trial. This is obtained by subtracting off the component of the variance attributed to the variable spike counts that would be observed even if the latent rate were fixed. For a Poisson point process this would be $\overline{n}_j$, but we assume the point process is a generalized renewal[61] and is thus approximated by the *point process variance*,

$$V_j^{PP}[\nu_j] = \varphi \overline{n}_j \qquad \text{(Equation 14)}$$

where the Fano factor, φ, is unknown. Note that the point process variance depends on the signed coherence. From the law of total variance, subtraction of this component from the total variance leaves the variance of the conditional expectation, $\mathrm{VarCE} \equiv \mathrm{Var}[\mathbb{E}(n)]$. There is a bookkeeping step that respects the dependence of V^{PP} on signed coherence and neuron (see previous citations), but using the residuals, we can obtain an estimate of the VarCE across all neurons at one time bin. Dividing by 0.06^2 yields the variance of the latent rates ($\mathrm{spikes}^2/\mathrm{s}^2$) across trials (in the time bin under consideration), although it depends on the unknown φ. For unbounded diffusion the VarCE should increase linearly as a function of time, because it is a cumulative sum of *iid* random numbers.

Please cite this article in press as: Shushruth et al., Sequential sampling from memory underlies action selection during abstract decision-making, Current Biology (2022), https://doi.org/10.1016/j.cub.2022.03.014

Diffusion also specifies the autocorrelation, between the cumulative sum of the first i samples and the cumulative sum of the first $j \geq i$ values:

$$\rho_{ij} = \sqrt{\frac{i}{j}} \quad \text{(Equation 15)}$$

This implies a decay of correlation as function of lag, $j - i$, and an increase in correlation for fixed lag, as a function of time. We obtain the estimates, r_{ij}, from data by forming the autocovariance matrix on residuals from all neurons. Note that the $\mathrm{Cov}^{PP} = 0$ for $i \neq j$, because by construction, given the rate in time bin j the stochastic realization of spike count does not depend on the rate or realization of spike count in bin i. Therefore the covariance of the conditional expectation (CovCE) is the raw covariance for $i \neq j$. Its diagonal ($i = j$) is the VarCE. This matrix is normalized in the usual way to produce a correlation matrix of conditional expectation (CorCE).

The CorCE depends on the VarCE which depends on φ, which is unknown. We chose the value that minimized the sum of squares:

$$\sum_{i=1}^{n-1} \sum_{j=2}^{n} \left[\mathcal{Z}(r_{ij}) - \mathcal{Z}(\rho_{ij})\right]^2 \quad \text{(Equation 16)}$$

where $\mathcal{Z}$ denotes standardization (Fisher-z transform).

The values of the variance plotted in Figures 7A and 7F are VarCE, using the fitted φ. The standard errors are estimated from a bootstrap procedure[62] in which trials were sampled (with replacement) while maintaining their grouping (same neuron, dot direction, coherence and color of target in RF). We also performed the same analysis using neural responses in the epoch between 190 to 550 ms after RDM onset (Figure S5).