Location- and feature-based selection histories make independent, qualitatively distinct contributions to urgent visuomotor performance

Emily E Oor[1], **Emilio Salinas**[2] **and Terrence R Stanford**[2,*]

[1]Department of Psychology, Wake Forest University, Winston-Salem, North Carolina, United States of America

[2]Department of Translational Neuroscience, Wake Forest University School of Medicine, Winston-Salem, North Carolina, United States of America

*Corresponding author

e-mail: stanford@wakehealth.edu (TRS)

draft May 29, 2024, 19:22

Abstract

Attention mechanisms that guide visuomotor behaviors are classified into three broad types according to their reliance on stimulus salience, current goals, and selection histories (i.e., recent experience with events of many sorts). These forms of attentional control are clearly distinct and multifaceted, but what is largely unresolved is how they interact dynamically to determine impending visuomotor choices. To investigate this, we trained two macaque monkeys to perform an urgent version of an oddball search task in which a red target appears among three green distracters, or vice versa. By imposing urgency, performance can be tracked continuously as it transitions from uninformed guesses to informed choices, and this, in turn, permits assessment of attentional control as a function of time. We found that the probability of making a correct choice was strongly modulated by the histories of preceding target colors and target locations. Crucially, although both effects were gated by success (or reward), the two variables played dynamically distinct roles: whereas location history promoted an early motor bias, color history modulated the later perceptual evaluation. Furthermore, target color and location influenced performance independently of each other. The results show that, when combined, selection histories can give rise to enormous swings in visuomotor performance even in simple tasks with highly discriminable stimuli.

Introduction

The choice of where to look, which is integral to most daily human activities, reflects dynamic interplay between information sensed from the external environment and internal goals specific to the task at hand. This interaction is likely a key contributor to the variability in choice and associated outcome that typifies real-world behavior. But even in a laboratory setting, performance on visuomotor choice tasks varies not only between subjects but also for the same subject across nominally identical sessions and trials. Understanding the drivers of behavioral variance is essential to determining the factors that limit performance and, ultimately, to deducing their neural basis.

It is a given that attentional state contributes substantially to the conditions that determine success or failure on any given goal-oriented task. Endogenous mechanisms allocate attention volitionally, prioritizing choices according to internal goals, whereas exogenous mechanisms attract attention automatically in proportion to the physical saliencies of external features or events (Itti and Koch, 2001; Theeuwes, 2010; Carrasco, 2011; Wolfe and Horowitz, 2017). How strongly these mechanisms engage their respective neural substrates, and the degree to which these attentional pointers align or conflict, contribute to determining if, when, and how any particular choice is made (Busse et al., 2008; Markowitz et al., 2011; Chen et al., 2013; Scerra et al., 2019; Chen et al., 2020; Goldstein et al., 2022; Seideman et al., 2022; Oor et al., 2023).

Although exogenous and endogenous attentional mechanisms are potent influences, physical stimulus properties and explicit knowledge of task requirements do not solely determine choice performance in any given instance. It is now generally accepted that current choices are constrained by the recent history of past choices in ways that are neither consistent with task goals nor a straightforward reflection of the physical saliencies of available options (Fecteau and Munoz, 2003; Awh et al., 2012; Theeuwes, 2018, 2019; Anderson et al., 2021). "Selection history" is the umbrella term for attention mechanisms that are distinct from stimulus-driven and goal-directed forms in that they comprise a collection of implicit biases shaped by past experience. Potential sources of selection history bias include feature-based priming (Maljkovic and Nakayama, 1994; McPeek et al., 1999; Kristjánsson and Campana, 2010; Brascamp et al., 2011; Kristjánsson and Ásgeirsson, 2019), location bias (Maljkovic and Nakayama, 1996; Chelazzi et al., 2014; Hickey

et al., 2014), and stimulus- or action-reward associations (Barraclough et al., 2004; Hikosaka et al., 2006; Anderson et al., 2011; Abrahamyan et al., 2016; Hauser et al., 2018; Anderson, 2019; Hermoso-Mendizabal et al., 2020; Lak et al., 2020), among others. Given their heterogeneity, selection history biases are likely to impact the coding of both intrinsic and extrinsic decision variables — and do so on multiple time scales — to drive choice trends (Awh et al., 2012; Theeuwes, 2018, 2019; Anderson et al., 2021).

Although selection history influences are measurable as effects on performance accuracy, reaction time (**RT**), or both, discerning how multiple sources of variance contribute to determining choice outcome is not straightforward. Just as these implicit biases need not align with goal-directed or stimulus-based drivers of attention, different sources of selection history bias may reinforce or conflict with each other to varying degrees according to their respective history dependence profiles. Thus, via internal tradeoffs, many different combinations of selection history bias could yield very similar behavioral outcomes as measured by accuracy and RT. Moreover, because history induced biases are not explicitly strategic — they develop even when history offers no predictive value to the current trial — their expression in measures of overt behavior will be attenuated in laboratory tasks that prioritize accuracy and allow sufficient time for choices to be guided predominantly by goal-driven mechanisms.

The current study uses an urgent task, i.e., one wherein the time available to respond is limited (Salinas and Stanford, 2021; Stanford and Salinas, 2021), to isolate and quantify how selection history biases deriving from target feature (color) and target location contribute to performance during the search for a color singleton (i.e., a stimulus with a unique color). With this novel task design the goal was to differentiate sources of history bias on the basis of both inter- and intra-trial temporal dynamics and, having done so, determine the degree to which they act independently or interact in biasing each choice. Urgency is the key, first, for maximizing the measurable impact of bias, and second, for distinguishing the influences of biases that act on stimulus features (e.g., color priming) from those that do not (e.g., location priming). This is because, under time pressure, individual choices can be precisely identified via their unique temporal signatures as either uninformed guesses or perceptually informed decisions (Salinas and Stanford, 2021; Stanford and Salinas, 2021), and crucially, spatial and perceptual biases are predicted to manifest in very different ways in each case: guesses can be strongly biased toward an arbitrary location early on, even before presentation of the stimulus array, whereas informed decisions can only be influenced by perceptual biases later, once the features of the stimulus array have been at least partially processed.

By identifying these different sources of bias, it becomes clear that, at least in monkeys, their combined impact on behavior can be extraordinarily large, producing major swings in performance even in visuomotor tasks that involve simple rules and highly discriminable stimuli, and which, in principle, should be very easy.

Methods

Subjects

All surgical and behavioral recording procedures were conducted in compliance with the Wake Forest University School of Medicine (WFUSM) Institutional Animal Care and Use Committee (IACUC), the National Institutes of Health Guide for the Care and Use of Laboratory Animals, and all relevant USDA regulations. Data were obtained from two purpose-bred adult male rhesus monkeys (Macaca mulatta), both of which met WFUSM veterinary standards for health quality upon procurement. Daily health status was monitored by veterinary and lab staff throughout the duration of the study. Animals were housed in Allentown quad format cages, which met

all regulatory requirements, in rooms equipped with a controlled light-dark cycle. Behavioral enrichment was provided by laboratory staff as per a WFUSM environmental enrichment policy. Feeding was ad libitum and provided to each animal's home cage by WFUSM Animal Resources Program staff. Supplemental food enrichment was provided daily by laboratory staff.

Monkeys N and C participated in a previous study performing a similar urgent choice paradigm paired with recordings from the superior colliculus, thus the current experiment utilized existing surgical preparations. Standard sterile surgical techniques under general anesthesia were employed for securing an MRI compatible head post and an MRI-compatible recording cylinder to the skull (Wyder et al., 2003). Post-surgical pain was managed with the opioid analgesic Buprenorphine hydrochloride delivered intramuscularly in 0.01–0.02 mg/kg doses, once perioperatively, followed by doses in 8–12 hrs increments for the following 24–48 hrs (unless otherwise specified by veterinary staff). Post-surgical inflammation was managed with the non-steroidal anti-inflammatory drug Ketoprofen delivered intramuscularly in 2.5–5.0 mg/kg doses, once perioperatively, followed by once daily administration for two days (unless otherwise specified by veterinary staff). After a minimum recovery period of two weeks, positive reinforcement was used to train willingness to tolerate head restraint and, subsequently, to perform behavioral tasks involving saccadic eye movements to visual stimuli appearing on a computer monitor.

Behavioral recording techniques

Standard operant methods with positive reinforcement were used to train monkey subjects on the behavioral task. For training and experimental sessions, liquid rewards for correct performance on visual tracking tasks served as the primary source of hydration. Animals were given the opportunity to work to satiety before returning to their home cage (satiety indicated by a decrease in motivation toward the later stages of a behavioral session). If necessary, supplementary liquid was provided in the animal's home cage to ensure adequate hydration as specified by WFUSM policy.

Behavioral tasks were performed with monkeys seated in a quiet, dimly lit room. Subjects maintained a comfortable upright seating position in a purpose-designed primate chair (Crist Instrument Co.) with head restrained to maintain a stable, straight ahead orientation with respect to the visual display. Visual stimulus arrays consisted of gray, red, and green filled circles measuring 2.3° of visual angle in diameter and were presented on a VIEWPixx LED monitor (VPixx Technologies Inc, Saint Bruno, Quebec, Canada; 1920 $\times$ 1200 screen resolution, 120 Hz refresh rate, 12 bit color) at a viewing distance of 57 cm. Eye movements were sampled at 1000 Hz using an infrared camera-based eye tracking system (EyeLink 1000, SR Research, Ottawa, Canada). Stimulus presentation, behavioral task implementation, and data acquisition were controlled and coordinated using a custom-designed data acquisition system (Ryklin Software Inc.).

Behavioral task

The compelled oddball (**CO**) task (Fig. 1a) is an urgent 4-alternative choice paradigm in which subjects must make a saccade to a color singleton presented among three uniformly colored distracters (Scerra et al., 2019). Subjects begin by fixating a central grey spot on a black background. After 200 ms, a stimulus array consisting of 4 gray placeholders spaced 90° apart is presented to indicate the locations at which the singleton target and 3 distracters will appear (Fig. 1a, Targets On). After a delay of 500–1000 ms, the fixation spot disappears (Fig. 1a, Go), instructing the subject to make a saccadic choice within 450 ms for a chance at juice reward. Crucially, the relevant color information is withheld for a period of time ranging from 0–225 ms (Fig. 1a, Gap), after which the gray placeholders change color to reveal the location of either a green target among 3

red distracters (Fig. 1a, Cue) or a red target among 3 green distracters (not shown). Target identities and locations are randomly and independently assigned from trial to trial such that the prior probability of a color repetition is 0.5, and the prior probability of a location repetition is 0.25. A trial is correct if the subject makes an eye movement to the location of the oddball target within the allowable time window of 450 ms.

Quantification and statistical analysis

All data analyses were performed in the Matlab programming environment (R2020a; The MathWorks, Natick, MA). Behavioral performance was quantified as detailed in prior reports (Stanford et al., 2010; Scerra et al., 2019; Goldstein et al., 2022; Seideman et al., 2022; Oor et al., 2023). Reaction time (**RT**) was measured from the time of the Go signal to the onset of the saccade (when eye velocity had surpassed 50°/s). For each trial, the processing time (**PT**) — the interval between Cue and saccade onset — was calculated by subtracting the duration of the Gap interval from the RT on that trial.

The psychophysical performance of each subject was evaluated using the tachometric curve (Fig. 1b), a psychometric function that plots the probability of a correct choice as a function of PT, or cue viewing time (Stanford et al., 2010; Scerra et al., 2019; Goldstein et al., 2022; Seideman et al., 2022; Oor et al., 2023). Tachometric curves were generated by sorting the trials by PT and computing the fraction correct for all trials within each PT bin, with bins (50 ms width) shifting every 1 ms. The tachometric curves of both subjects (Fig. 1b) show that performance varies from chance (probability correct $= 0.25$) to asymptotic with increasing PT. For PT < 100 ms, behavior is at chance and such trials are deemed "uninformed". For PT > 150 ms, performance rises rapidly to reflect the increasing likelihood that the choice is guided by color information and these choices are considered "informed".

Separate tachometric curves were generated for each history sequence, and for each resulting curve, two key parameters were calculated: the percent correct for uniformed choices (or low asymptote) and the percent correct for informed choices (or high asymptote). The calculation of these quantities was independent of the bin size used to depict the tachometric curves: the low asymptote simply corresponded to the percent correct for all history-conditioned trials with PT < 100 ms, whereas the high asymptote corresponded to the percent correct for all history-conditioned trials with PT > 150 ms. For both quantities, 95% confidence intervals (**CIs**) were evaluated using binomial statistics; specifically, the Agresti-Coull method (Agresti and Coull, 1998). Significance levels were established via resampling methods (Siegel and Castellan, 1988; Hesterberg, 2014). Mean RT together with standard error of the mean (**SEM**) were also calculated for the data set associated with each history sequence. Such history-conditioned mean RTs are based on both correct and incorrect trials at all PTs.

Data pooling

The effects of trial history on performance were qualitatively similar for the two monkey subjects (see Fig. 1 — figure supplement 1and last section of Results), but their PTs were somewhat different. Thus, to pool the data across the two subjects, we first aligned their mean tachometric curves (Fig. 1b). Optimal alignment was determined as follows. First, one of the tachometric curves was rescaled and shifted so that the average absolute difference between the tachometric curves from the two monkeys was minimized. The formula for the transformed curve is

$$z(j) = g * (y_2(j + \delta) + b) \tag{1}$$

where y_2 and z are the original and transformed values (for fraction correct), j indicates a PT bin,

and the transformation parameters are the gain (g), baseline shift (b), and the shift along the PT axis (δ). These parameters were set to minimize

$$\sum_j \left(|z(j) - y_1(j)|\right) \tag{2}$$

which is the total deviation between tachometric curve 1 (y_1) and the transformed version of tachometric curve 2 (z). The variable of interest for the purposes of optimally aligning the two data sets was δ, and it indicated that the tachometric curve for monkey N (Fig. 1b, bottom) should be shifted leftward by 29 ms. Thus, before pooling the data from the two monkeys, we first subtracted 29 ms from all the RTs from monkey N, which in turn shifted all its PTs by the same amount. Although the applied shifts in the tachometric curve (Fig. 1 — figure supplement 1e, gray bar) and the RTs (Fig. 1 — figure supplement 1h, dark bars) of monkey N were in no way critical to the results, the temporal alignment in the pooled data reduced the variance of the measured quantities.

Feature history

Trials were sorted post hoc based on the number of consecutive occurrences in which the target color on immediately preceding trials was either the same or different than that on the current trial. Because target history effects were qualitatively similar for red and green targets, trials were grouped into same (S) versus different (D) categories for subsequent analyses of the effects of target color history (Fig. 1 — figure supplement 2).

Trials classified as S were those preceded by N trials with a target of the same color (e.g., Fig. 2a, blue spectrum). The classification was such that trials preceded by at least 1 trial with a target of the same color were termed 1S trials (sequences red-**red** and green-**green**, where the bolded word indicates the trial of interest). Trials preceded by at least 2 trials with a target of the same color were termed 2S trials (red-red-**red** and green-green-**green**), etc. The S trials reveal the impact of color repetition.

Analogously, trials classified as D (e.g., Fig. 2b, red spectrum) were those preceded by N trials with a target of the opposing color, which was different. As such, trials preceded by at least 1 trial with a target of the opposite color were designated 1D trials (green-**red** and red-**green**), trials preceded by at least 2 trials with a target of the opposite color were designated 2D trials (green-green-**red** and red-red-**green**), and so on. The D trials reveal the impact of a color switch.

Location history

The orientation of the 4-stimulus array and the eccentricity of the stimuli were varied pseudorandomly across experimental sessions. For the purpose of subsequent analysis, each target location was assigned to one of four designated quadrants: right, up, left, and down. Progressing counterclockwise, target location was categorized as follows: right, if corresponding to a positive value along the x axis or anywhere within the first (upper-right) quadrant; up, if corresponding to a positive value along the y axis or within the second (upper-left) quadrant; left, if corresponding to a negative value along the x axis or within the third (lower-left) quadrant; down, if corresponding to a negative value along the y axis or within the fourth (lower-right) quadrant.

Independently of feature history, trials were sorted based upon the number of consecutive occurrences in which the target location (i.e., quadrant) on immediately preceding trials was the same or different than for the current trial. Trials classified as S were those preceded by N trials with a target in the same location (e.g., Fig. 2b, blue spectrum). Trials preceded by at least 1 trial with a target in the same location were 1S trials (sequences right-**right**, up-**up**, left-**left**, down-**down**, where the bolded word again indicates the trial of interest). Trials preceded by at least 2

trials with a target in the same location were 2S trials (right-right-**right**, up-up-**up**, left-left-**left**, down-down-**down**), etc. Trials classified as D (e.g., Fig. 2b, red spectrum) were those preceded by N trials with a target in a different location. Trials preceded by at least 1 trial with a target of a different location were 1D trials (examples: up-**right**, down-**right**, left-**down**, right-**up**; 12 possible combinations), trials preceded by at least 2 trials with a target of (any) different location were 2D trials (examples: left-up-**right**, left-left-**up**, up-down-**left**, etc; 36 possible combinations), and so on. As with color, S and D trials reveal the impact of repeats and switches, but of target location in this case.

Outcome history

Trials were additionally sorted based upon the number of correct or error trials that preceded the trial of interest. Trials preceded by at least 1 correct trial were designated 1C (sequence correct-**x**, where the **x** indicates the outcome of the trial of interest). A trial preceded by at least 2 correct trials was designated 2C (correct-correct-**x**), etc. The complementary classification scheme was applied for error trials such that trials preceded by at least 1 error were labeled 1E (error-**x**), while those preceded by at least 2 error trials were 2E (error-error-**x**), etc.

Feature and location histories combined with outcome history

Utilizing the classification schemes described above, feature and location history effects were further conditioned on prior successes or failures. Here, the S and D sequences were separated into same/correct, same/error, different/correct, and different/error subgroups. Accordingly, trials immediately preceded by at least 1 successful trial with a target of the same color (or location) were designated 1SC trials; trials preceded by at least 2 correct trials with targets of the same color (or location) were considered 2SC trials, and so on (e.g., Fig. 5b, green spectrum). Mirroring this analysis, S trials were conditioned on error history: trials preceded by at least 1 error trial with a target of the same color (or location) were designated 1SE trials, while trials preceded by at least 2 errors with targets of the same color (or location) were designated 2SE trials, etc (e.g., Fig. 5b, purple spectrum). This sorting structure was duplicated for the complementary set of D analyses, which produced sequences 1DC, 2DC,... and 1DE, 2DE,... for correct and error combinations, respectively (e.g., Fig. 5f).

Quantifying functional interactions between history variables

To understand and quantify the functional interactions between the three variables of interest, target color, target location, and trial outcome, we devised an analysis based on the concept of conditional independence (Dawid, 1979), which provides a useful characterization of how three stochastic variables may relate to each other. This widely applicable methodology is described in detail elsewhere (Salinas and Stanford, 2024), but the following is a brief overview.

Two events A and B, are conditionally independent relative to a third event C if their joint probability given C is such that

$$P(A, B|C) = P(A|C)\, P(B|C) \tag{3}$$

where $P(A, B|C)$ is the probability that events A and B occur given that event C has occurred. In other words, if C is known, A and B occur independently of each other. An equivalent formulation of this idea is to say that, once condition C is known, the probability that B occurs is fixed, regardless of whether A occurs or not. That is,

$$P(B|A, C) = P(B|C) \tag{4}$$

is an alternate definition. Note that these expressions are distinct from the standard notion of independence, which for events A and B would correspond to

$$P(A, B) = P(A)\, P(B)\,. \tag{5}$$

This condition is different, and irrelevant to the problem at hand. Conditional independence is essentially a constraint on the three-way relationships between the three variables of interest, as captured by their joint probability, $P(A, B, C)$. Here, it is useful because it provides a quantitative reference for the effect on C that A and B are expected to have when considered individually versus when considered simultaneously.

In the current study, we identify event C with the outcome of a trial (correct or incorrect) and events A and B with selection histories preceding that trial. For this application, we are interested in determining whether different histories (for target color, target location, or success) exert independent effects on trial outcome, or alternatively, whether they interact to reinforce each other (producing a stronger-than-expected effect) or interfere with each other (producing a weaker-than-expected effect). To generate a baseline expectation or benchmark for the no-interaction case, we consider how event C depends on events A and B. The quantity that we seek is

$$P(C|A, B) = P(A, B|C)\, \frac{P(C)}{P(A, B)} \tag{6}$$

where the right side is the standard identity based on Bayes theorem. When A and B are conditionally independent with respect to C, Equation 3 can be used to factorize all the terms where A and B appear jointly on the right-hand side of Equation 6, and everything can be put in terms of $P(C|A)$ and $P(C|B)$, which describe how choice outcome depends on each of two history variables separately. If $C = 1$ and $C = 0$ indicate a correct and an incorrect choice, respectively, then the result of combining Equations 3 and 6 is, after some algebra,

$$P(C{=}1|A, B) = \frac{P(C{=}1|A)\, P(C{=}1|B)}{P(C{=}1|A)\, P(C{=}1|B) + P(C{=}0|A)\, P(C{=}0|B)\, \frac{P(C=1)}{P(C=0)}} \tag{7}$$

(Salinas and Stanford, 2024). This expression represents a prediction based on three quantities: the effect of A alone on choice outcome, the effect of B alone on choice outcome, and the overall probability of a correct choice (i.e., the prior, $P(C = 1)$). It is a model-free benchmark for quantifying how much knowledge is gained when two sources of evidence are considered together rather than separately — if they do not interact. This prediction applies to any pairwise combination of separate history terms that depend on A and B, and is to be contrasted with the empirical result for the probability of a correct choice given that A and B are known *simultaneously*.

For example, suppose that the following numbers are measured from the data: (1) the probability of a correct choice given that the oddball color in the preceding trial was the same as in the current trial

$$P(C = 1|\text{1S color}) = 0.80 \tag{8}$$

(2) the probability of a correct choice given that the preceding outcome was correct

$$P(C = 1|\text{1C outcome}) = 0.75 \tag{9}$$

and (3), the overall probability of a correct choice

$$P(C = 1) = 0.70 \tag{10}$$

Then compute the expected probability of success given the joint condition,

$$P(C = 1|\text{1S color}, \text{1C outcome}) \tag{11}$$

by inserting the above values into Equation 7. The result is 0.84; this is the expectation for the probability a correct choice when both the target color and outcome histories are known. It is larger than either of the probabilities conditioned on a single history variable because they are in agreement (both above the prior) and because together they are more informative of choice outcome. This prediction is to be contrasted with the actual data, that is, with the probability $P(C = 1|\text{1SC})$ obtained by considering all the trials in which the preceding trial was both correct and of the same target color. Any significant discrepancy from the predicton is indicative of an interdependence between color and outcome histories with respect to their impact on the outcome of the current choice.

This analysis was used to uncover functional interactions between target-color, target-location, and outcome histories in various pairwise combinations (Fig. 7). To obtain a complete description of how each pair of history variables influences the current choice, this procedure was applied to all the possible history sequence pairs going from 1 to 4 trials back, with each pair providing a predicted value (based on limited data and the conditional independence assumption) and an empirical value (based on the full data set).

Timecourse of history effects

The effect of a single past event fades with time, i.e., with the number of subsequent trials, and to estimate the corresponding time constant we designed an analysis that runs as follows. First consider trial sequences in which the color of the target was the same for the current trial and for the Nth trial back, but regardless of what happened in the trials in between. Thus, for $N = 1$ we have the 1S case mentioned earlier, which includes sequences red-**red** and green-**green**, where the bolded word indicates the trial of interest; for $N = 2$ we have the case 1S1x, which includes sequences red-x-**red** and green-x-**green**, where the x indicates a trial with either target color, red or green; for $N = 3$ we have the case 1S2x, which includes sequences red-xx-**red** and green-xx-**green** with two intervening trials of either color; for $N = 4$ we have the case 1S3x, which includes sequences red-xxx-**red** and green-xxx-**green**; and so on. By plotting the probability of a correct choice for each one of these conditions as a function of N, the effect of a single S stimulus can be tracked as a function of time, i.e., trial number (Fig. 4a, b, blue data).

The same procedure can be applied to reveal the timecourse of the modulatory effect caused by a single D stimulus (Fig. 4a, b, brown data). In that case, one would consider the sequences labeled as 1D (green-**red** and red-**green**), 1D1x (green-x-**red** and red-x-**green**), 1D2x (green-xx-**red** and red-xx-**green**), 1D3x (green-xxx-**red** and red-xxx-**green**), and so forth. Together, the timecourses for the S and D cases measure how rapidly the information about target color dissipates as further trials of the task are performed. Entirely analogous procedures were used to determine the timecourses corresponding to same or different target locations, which also involved sequences labeled 1S, 1S1x, 1S2x, 1S3x, and so on (Fig. 4c, d, blue data), and 1D, 1D1x, 1D2x, 1D3x, and so on (Fig. 4c, d, brown data).

Data availability

The trial-by-trial behavioral data that support the findings of this study are publicly available from Zenodo at https://doi.org/10.5281/zenodo.11391884. Matlab scripts for reproducing analysis results and figures are included as part of the shared data package.

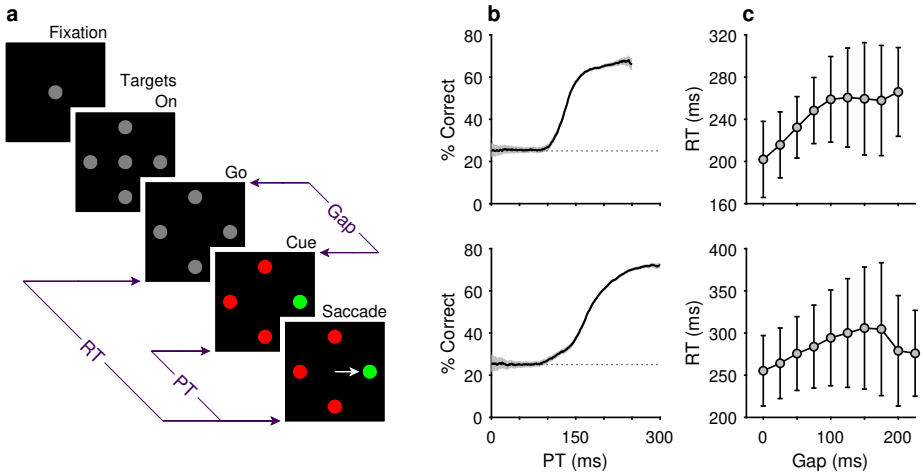

Figure 1. Behavioral task and psychophysical performance from two monkeys. (**a**) The compelled oddball (CO) task. Trials begin when the subject fixates on the gray spot at the center of the screen (Fixation). Four gray placeholders then appear (Targets On) separated by 90° surrounding the fixation spot (at variable eccentricities and orientations across days). Urgency is incorporated by next extinguishing the fixation spot (Go), which instructs the subject to respond. After a variable time delay (Gap, 0–225 ms), the target and distracter identities are revealed (Cue). A correct choice requires a saccade to the stimulus with a unique color, the oddball stimulus, within 450 ms of the go signal. The reaction time (RT) is measured as the time from the go signal to saccade onset. The processing time (PT) interval is the maximum amount of time during which the cue can inform the choice. It is calculated by subtracting the gap duration from the RT in each trial. (**b**) Tachometric curves showing the percentage of correct responses as a function of PT for monkey C (top; 26,038 trials) and monkey N (bottom; 45,173 trials). Gray shades represent 95% confidence intervals (CIs). Chance performance (25%) is indicated by dotted lines. PT bin size is 50 ms. (**c**) Mean RT ($\pm$ 1 SD) as a function of gap for monkey C (top) and monkey N (bottom).

Results

Accuracy of urgent singleton search depends on perceptual processing time

Two monkeys performed the compelled oddball (**CO**) task (Fig. 1a), a singleton search paradigm that, via the imposition of a response time limit, compels subjects to begin preparing a response in advance of making the perceptual judgement that determines the correct choice (Methods). In this task, urgency is incorporated by delivering the signal to make a saccade (Fig. 1a, Go) before presenting the target-distracter stimulus array, or color cue (Fig. 1a, Cue); this cue is withheld for a variable amount of time (Fig. 1a, Gap). By design, the task mandates that saccades be initiated within a normal RT range (RT $\leq$ 450 ms). This, together with the variable Gap interval, yields choices based on a wide range of nominal cue viewing times, or processing times (**PTs**, where PT is calculated by subtracting $\text{RT} - \text{Gap}$ in each trial).

In practice, both monkeys performed the task as designed. Plotted for each Gap, average RTs ranged between 200–250 ms for monkey C (Fig. 1c, top) and 250–300 ms for monkey N (Fig. 1c, bottom), in each case demonstrating only a weak positive dependency on Gap, a trend described and modeled in previous studies (Stanford et al,. 2010; Salinas et al., 2010; Shankar et al., 2011). Importantly, the tachometric curve, which plots performance as a function of PT (Fig. 1b), shows that choice accuracy rises sharply as the amount of cue viewing time increases. Thus, PT readily distinguishes between choices that are likely to be guided by the color cue (informed) and those that are not (guesses). For PT $<$ 100 ms, the performance of both monkey subjects was at chance (probability correct $=$ 0.25), indicating that such choices were not informed by the color

cue. Conversely, for PT $\gtrsim$ 150 ms for monkey C, and PT $\gtrsim$ 200 ms for monkey N, most choices were informed, as indicated by performance asymptotes that approached 70% correct.

Such level of performance may seem unusually low for a visuomotor task in which the response rule is so simple and the stimuli so highly discriminable. Indeed, this was the initial motivation for investigating the impact of past trials on performance. In this case, two observations already point to selection histories as major drivers of monkey behavior. First, although the subjects practiced the CO task extensively (Fig. 1b, caption) throughout several months of data collection, their asymptotic accuracies did not improve (the Pearson correlation between experimental session and asymptotic performance was -0.2, $p = 0.1$ for for monkey C and 0.05, $p = 0.6$ for for monkey N). And second, much higher accuracies ($>$ 95% correct in some cases) were reported in studies in which monkeys performed standard, non-urgent versions of the task with the target colors alternating either between blocks of trials (Bichot and Schall, 1999, 2002) or between days (McPeek and Keller, 2004). Importantly, though, the apparent difficulty is not due to the urgent nature of our task, because similarly low accuracies have also been reported with standard, non-urgent versions of the task in which both color and location were randomized (Song and McPeek, 2015); and furthermore, other monkeys from our laboratory also reached high asymptotic accuracies ($\sim$95% correct) in the urgent CO task when the target color remained fixed or when it switched only between blocks of trials (Fig. 1 — figure supplement 3; Scerra et al., 2019). As will be shown next, when both the target color and location change unpredictably across trials, performance during oddball search is swayed by selection histories to a remarkable degree.

Perceptual and non-perceptual selection history biases are dissociable on the basis of PT

The ability to differentiate informed from uninformed choices is the basis for distinguishing the influences of selection history biases that derive from perceptual versus non-perceptual sources. If present, perceptual priming due to target color repetition could, in principle, affect the timing, slope, and/or asymptote of the tachometric curve, each an attribute that reflects the likelihood that a choice is perceptually informed (Shankar et al., 2011). Perceptual priming cannot, however, influence uninformed choices — saccades made at short PTs — before the processing of color information has had any impact. Importantly, these perceptually uninformed choices are precisely those predicted to be most sensitive to non-perceptual biases. In the context of an urgent feature discrimination task, motor plans developing in advance of cue information (i.e., during the Gap interval) might be expected to incorporate spatial biases that reflect the recent history of target locations. If so, the expression of a history induced location bias would be maximal for guesses and diminish as choices become increasingly guided by perceptual evidence.

Because the effects of selection history on performance were highly consistent for the two monkeys, in this and the following subsections these effects are presented and discussed based on the data pooled across subjects. We return to the differences and similarities between the two individuals in more detail in the last section of the Results, once the main phenomena have been characterized.

The impacts of target color (Fig. 2a-d) and target location (Fig. 2e-h) repetition on performance were assessed by post hoc sorting of trials based on their respective histories (Methods). Separate tachometric curves were computed from trials conditioned on the number of preceding trials (1–4) for which either target color (Fig. 2a) or target location (Fig. 2e) was the same (S) or different (D). For example, the 1S condition includes the color sequences red-**red** and green-**green** (bold type indicates the target color on the trial that contributes to the given tachometric curve), while the 3S condition includes the color sequences red-red-red-**red** and green-green-green-**green**. Likewise,

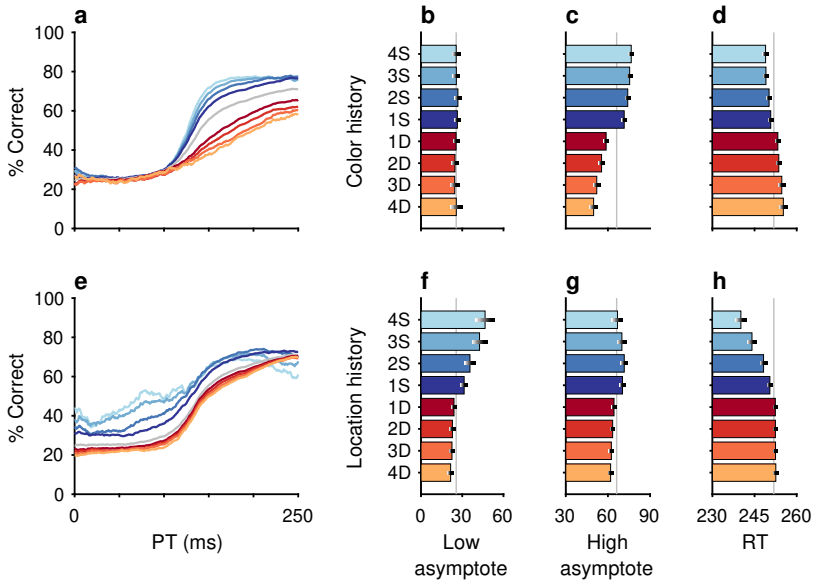

Figure 2. Repetitions of target color and target location cause distinct modulations of performance in the CO task. Each panel shows results for performance on the current trial conditioned on the history of preceding trials when current and past trials are the same (S, blue spectrum; 1–4 preceding trials) or different (D, red spectrum; 1–4 preceding trials) for a given feature. (**a**) Tachometric curves conditioned on target color history. Performance is shown as a function of PT for eight different history sequences. Average performance is shown in gray. Bin width is 50 ms for all curves. For clarity, error bars are omitted. (**b**) Choice accuracy (x axis, percent correct) in short-PT trials (PT $<$ 100 ms, low asymptote) conditioned on color history (y axis). Error bars represent 95% CIs across trials. Gray vertical line denotes value for the average tachometric curve. (**c**) Choice accuracy in long-PT trials (PT $>$ 150 ms, high asymptote) conditioned on color history. Same format as in **b**. (**d**) Average RT (x axis, in ms) conditioned on color history (y axis). Error bars show $\pm$ 1 standard error of the mean (SEM). The gray vertical line denotes average RT across all trials. (**e**) Tachometric curves conditioned on target location history. (**f**) Choice accuracy in short-PT trials (PT $<$ 100 ms, low asymptote) conditioned on location history. (**g**) Choice accuracy in long-PT trials (PT $>$ 150 ms, high asymptote) conditioned on location history. (**h**) Average RT conditioned on location history.

D trials are preceded by N trials with a target of the opposing color. The 1D condition, then, includes the color sequences green-**red** and red-**green**, while the 3D condition includes the color sequences green-green-green-**red** and red-red-red-**green**. The same nomenclature was applied to target location repetitions (Methods). In this case, S trials are preceded by N (1–4) trials having the same target location (e.g., right-right-**right** or down-down-**down**) and D trials are those preceded by N trials in which the singleton target appeared at any one of the three alternate locations (e.g., left-left-**right** or up-up-**left**).

After combining the data from the two monkey subjects (Methods; Fig. 1 — figure supplement 1), tachometric curves were generated conditioned on the specific histories just described (Fig. 2a, e). Strong selection history effects became evident, as performance depended systematically on the number of preceding trials in which the target color (Fig. 2a-d) or the target location (Fig. 2e-i) was either the same or different from that in the current trial. But critically, the respective influences of these two history variables on choice behavior were readily distinguishable by their PT dependencies, as anticipated.

The influence of target color repetition only became apparent at PT $>$ 100 ms, an inflection point that marks the transition from uninformed choices to those increasingly steered by color information. From there, each additional preceding trial of the same target color (1S–4S) was as-

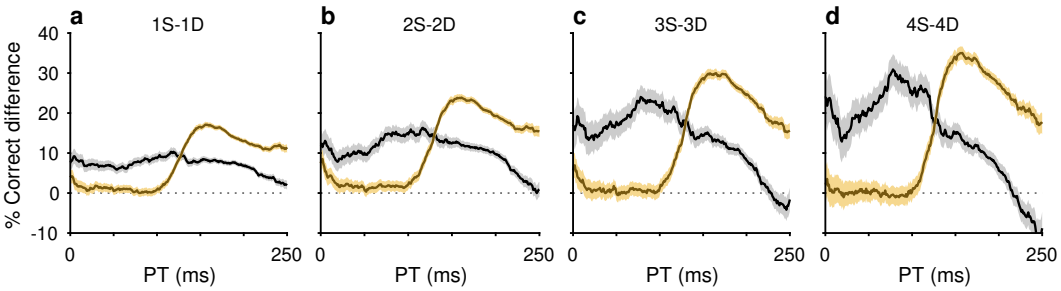

Figure 3. History effects due to target color and target location have distinct temporal manifestations within the timecourse of a trial. Each curve plots the difference between two tachometric curves, one conditioned on same (S) trials and another conditioned on different (D) trials. Each panel shows two difference curves, one for color history (gold traces) and another for location history (black traces), with shaded error bands indicating $\pm$ 1 SE across trials. Dotted lines at zero mark no difference in performance given S and D histories. (**a**) Difference curves for histories going back 1 trial (1S $-$ 1D). (**b**) Difference curves for histories going back 2 trials (2S $-$ 2D). (**c**) Difference curves for histories going back 3 trials (3S $-$ 3D). (**d**) Difference curves for histories going back 4 trials (4S $-$ 4D). Note that the location bias is strongest at short PTs and then declines, whereas the color bias rises sharply after $\sim$100 ms of PT.

sociated with a performance increment, as evident in both the slopes and upper asymptotes of the corresponding tachometric curves (Fig. 2a, blue curves; Fig. 2c, blue bars). Correspondingly, decreasing slopes and upper asymptotes depending on the number of preceding trials of a different target color (Fig. 2a, red curves; Fig. 2c, red bars) were indicative of history-induced performance decrements. Thus, the trial-by-trial history of target color decreased (same color) or increased (different color) the amount of cue viewing time needed to achieve a given level of accuracy. Importantly, this perceptual bias altered the efficacy with which color information guided the choice, but, as expected, had no measurable effect on uninformed choices, as indicated by uniformity of the lower asymptote across the range of history conditions (Fig. 2a, b).

In contrast, the principal effect of target location repetition was on uninformed choices (Fig. 2e, f), those made at PT $<$ 100 ms, before cue information began to guide behavior. For these perceptually uninformed choices, or guesses, the likelihood of looking toward a given location increased as a function of the number of preceding trials in which the target appeared at that location (Fig. 2e, blue curves; Fig. 2f, blue bars), and decreased slightly when preceding targets appeared elsewhere (Fig. 2e, red curves; Fig. 2f, red bars).

The PT-based dichotomy between color and location history biases and their respective selection history dependencies (i.e., effect of number of S or D trials) are well summarized by a contrast computed as the difference between two tachometric curves, one conditioned on a given number of S repetitions minus another conditioned on the same number of D repetitions (Fig. 3). Thus, for color (gold traces), the resulting difference curve reveals a PT-dependent bias that tracks closely with the increasing likelihood of making a correct color discrimination; the bias is near zero for uninformed choices but its magnitude rises sharply after the inflection point (at PT $\gtrsim$ 100 ms) that demarcates the transition to increasingly informed choices. In contrast, for location (black traces), the corresponding bias demonstrates the inverse trend: it is strongest for uninformed choices and decreases progressively after the inflection point. For both color and location bias, the size of the computed differential trend increases monotonically as the number of S and D repetitions increases from one (Fig. 3a, 1S $-$ 1D) to four (Fig. 3d, 4S $-$ 4D).

In summary, the results so far validate the key qualitative prediction — that selection history biases modulate performance in the CO task substantially, with perceptual and non-perceptual

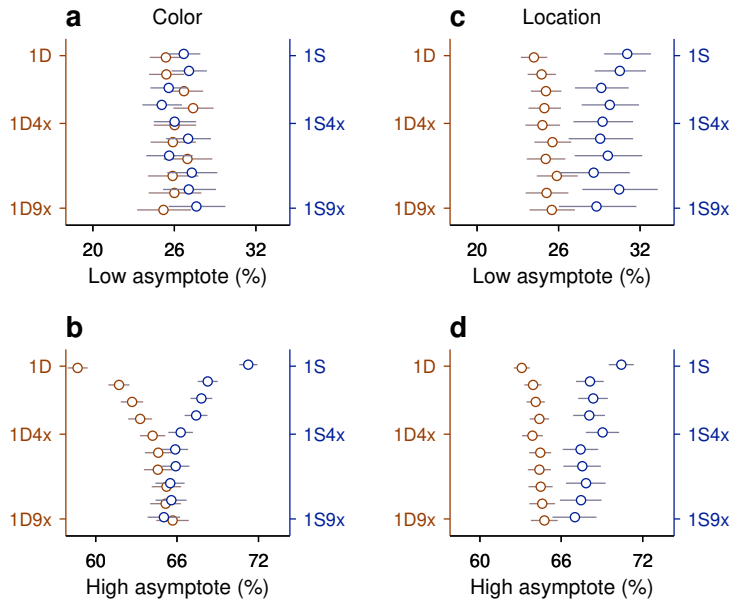

Figure 4. The timescale of history effects induced by a single repeat or a single switch varies across history types. (**a**) Task performance during uninformed choices (low asymptote) following a single color repetition (1S; blue axis) or color switch (1D; brown axis) that occurred 1–10 trials prior to the trial being evaluated. The notation Nx indicates N intervening trials of any type between the first event and the repeat or switch. (**b**) As in **a**, but for performance during informed choices (high asymptote). (**c**) Task performance during uninformed choices (low asymptote) following a target location repetition (1S, blue axis) or target location switch (1D, red axis) that occurred 1–10 trials prior to the trial being evaluated. (**d**) As in **c**, but for performance during informed choices (high asymptote). In all plots, error bars represent 95% CIs.

components being dissociable based on the characteristic ways in which they impact the temporal evolution of the choice process. These two components were characterized further.

Color- and location-driven biases have different inter-trial timescales

Biases accumulate across trials (Fig 3), indicating that the influence of a unitary color or location event persists over time (or trials). Notably, the duration of this influence, or analogously, the time course over which it wanes, is another potential criterion for distinguishing the biases induced by color and location selection histories. We quantified the impact on performance of a single repeat (in S sequences) or switch (in D sequences) for both color (Fig. 4a, b) and location (Fig. 4c, d) by computing tachometric curves from trials that occurred $N = 1$ to $N = 10$ trials after a reference event (Methods). For example, for color, the condition 1S4x denotes a color repeat separated by 5 trials, and includes two possible sequences, red-xxxx-**red** and green-xxxx-**green**, where the color of the first trial is the same as that of the trial being evaluated (bolded) and each of the four intervening trials can be of either color (x). Similarly, for location, the condition 1D2x denotes a location switch separated by 3 trials of any type, and would include sequences such as up-xx-**down** or left-xx-**right**, for instance.

As anticipated, a single color event, whether corresponding to a repeat or a switch, had no impact on a future uninformed choice, as indicated by the low asymptotes of the tachometric curves obtained for different trial separations (Fig. 4a). However, single color events did impact future informed choices, as indicated by the high asymptotes of the same curves (Fig. 4b). The impact of a color repeat (Fig. 4b, S conditions, blue circles) versus a switch (Fig. 4b, D conditions, brown cir-

cles) was fairly symmetric about a mean performance baseline, and accounted for a difference of approximately 13 percentage points in the likelihood of making a correct informed choice on the trial immediately following the reference event (1S versus 1D). This color bias decayed within 5 trials or so; it had a demonstrable influence over informed choices for the next 4 trials (Fig. 4b, 1S4x versus 1D4x; note separation between 95% CIs) but largely faded thereafter. Such time course is highly consistent with prior measurements based on non-urgent versions of the oddball task (Bichot and Schall, 2002).

Notably, though, the inter-trial temporal dynamic for location bias (Fig. 4c, d) was markedly different from that of color. In this case, the biasing effect of a location event, whether a repeat (S conditions) or a switch (D conditions), was smaller overall, producing a maximum difference of 7 percentage points in high asymptote (Fig. 4d, 1S versus 1D). More importantly, the location bias was evident for both uninformed (low asymptote; Fig. 4c) and informed choices (high asymptote; Fig. 4d), and showed greater persistence than the color bias. In this case, a single event had a significant influence up to 9 trials into the future (Fig. 4c, d, 1S8x versus 1D8x; note separation between 95% CIs).

Color- and location-driven biases are strongly modulated by outcome

Goal-directed behavior is, by definition, sensitive to outcome (Wolfensteller and Ruge, 2012; Dolan and Dayan, 2013). When an action aims to achieve a particular objective (e.g., looking toward the color singleton to obtain reward), the result influences subsequent behavioral strategy. But in addition, as we show here, it also has powerful effects on the development of non-strategic perceptual- and motor-based selection history biases.

The impact of outcome (reward versus no reward) on perception was quite apparent in the way that it modulated the expression of history-induced color bias (Fig. 5). Tachometric curves conditioned on previous trial repetitions of either the same (Fig. 5a) or different (Fig. 5e) target color were jointly conditioned on outcome; that is, different curves were generated depending on whether the color repetitions were associated with correct (Fig. 5, green color spectrum) or incorrect (Fig. 5, magenta color spectrum) choices. The biasing impact of outcome was large. For example, for tachometric curves in which each choice was preceded by four trials of the same target color, we observed a difference of nearly 50 percentage points in upper asymptotic performance depending on whether the previous same-color trials were correct or incorrect (Fig. 5a, light green versus light magenta traces; Fig. 5c, 4SC versus 4SE). Thus, for repeats of the same target color, a history of successful outcomes increased the likelihood of correctly choosing the target, whereas a history of unsuccessful outcomes decreased this likelihood. In other words, the same stimulus sequences led to vastly different performance depending on the their association with reward. For repeats of the alternative target color, the effect of outcome on asymptotic performance was more subtle but qualitatively consistent (Fig. 5e, g). In this case, when the different target color was associated with success on previous trials, choosing the current target color was less likely (Fig. 5e, g, green); and vice versa, when the different target color was previously associated with errors, choosing the current target color was comparatively more likely (Fig. 5e, g, magenta).

The impact of outcome on location bias was similarly strong (Fig. 6). When the same target location was repeated, outcome either reinforced or diminished the likelihood of choosing the target, depending on whether the repeats were associated with correct (Fig. 6a-c, green) or incorrect (Fig. 6a-c, magenta) choices. Conversely, repeated occurrences of different target locations yielded the complementary result; for these, correct choices diminished the likelihood of choosing the current target location (Fig. 6e-g, green) and errors increased this likelihood (Fig. 6e-g, magenta). As expected, these outcome-dependent variations in motor bias were most evident for uninformed

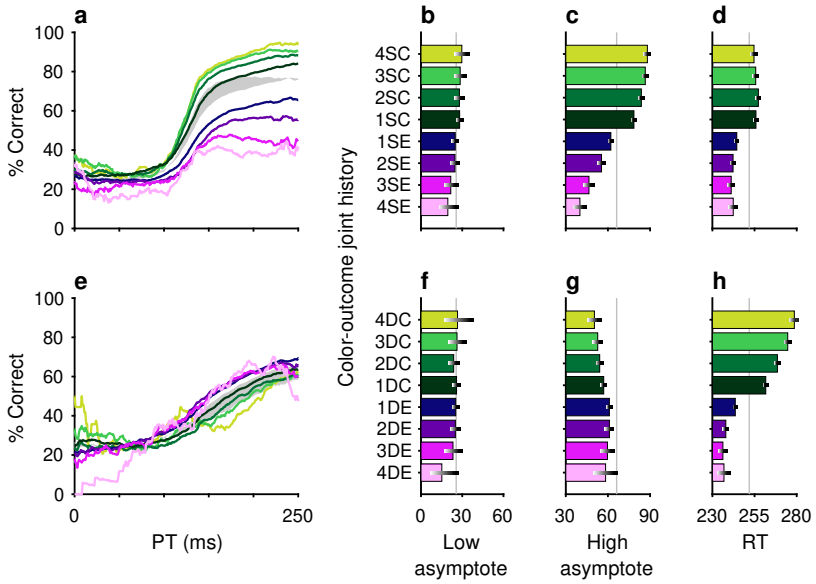

Figure 5. Perceptual accuracy is strongly modulated by the joint history of target color and trial outcome. (**a**–**d**) Task performance conditioned on the histories of preceding trials having the same target color (S) and resulting in either correct outcomes (C, green spectrum; 1–4 preceding correct trials) or error outcomes (E, purple spectrum; 1–4 preceding error trials). (**e**–**h**) Task performance conditioned on the histories of preceding trials having a different (D) target color and resulting in either correct outcomes (C, green spectrum; 1–4 preceding correct trials) or error outcomes (E, purple spectrum; 1–4 preceding error trials). Formatting for the tachometric curves (**a**, **e**), choice accuracies in short-PT trials (**b**, **f**), choice accuracies in long-PT trials (**c**, **g**), and mean RT (**d**, **h**) is the same as in Fig. 2. The gray shaded area in **a** represents the range of the history effects for color repeats obtained without taking outcome history into account (from 1S–4S in Fig. 2a, blue curves), and the gray shaded area in **e** represents the range of the history effects for color switches obtained without taking outcome history into account (from 1D–4D in Fig. 2a, red curves).

choices, as indicated by modulation of the low asymptotes (Fig. 6b, f) of the conditioned tachometric curves (Fig. 6a, e).

The effects of unrewarded (error) outcomes on accuracy are consistent with the idea that these expressions of bias were not the product of a deliberate strategy to optimize behavior. This is true for repetitions of both the same target color and the same target location. For example, if the color singleton was red on 4 consecutive trials and the subject made an error (i.e., looked to a green distracter) on each of those trials, the results suggest a progressive perceptual devaluation of red stimuli (Fig. 5a, c, magenta). This, in turn, suggests that the attribution of outcome on each trial was to the color of the singleton (e.g., red), not the color of the distracter (e.g., green) that was erroneously targeted by the saccade — an attribution that ignores the role of the motor act in causing the outcome. A similar logic applies to repeated target location. If the target repeatedly appeared in the same location and was missed (i.e., the subject looked elsewhere), the subject was less likely to select the target at that location on the current trial (Fig. 6a-c, magenta). Again, to the extent that error history biased behavior, the attribution of the unrewarding outcome appears to be tied to the singleton location, not the erroneous motor act.

A common trend for the influence of outcome on RT was observed for both color and location bias, and for same and different target repetitions within each of these feature dimensions. Specifically, RTs tended to be shorter when preceded by one or more error trials (Figs. 5d, h, 6d, h, magenta) and longer when preceded by correct trials (Figs. 5d, h, 6d, h, green). This effect on RT is consistent with a general association between time and accuracy that, by design, is particularly

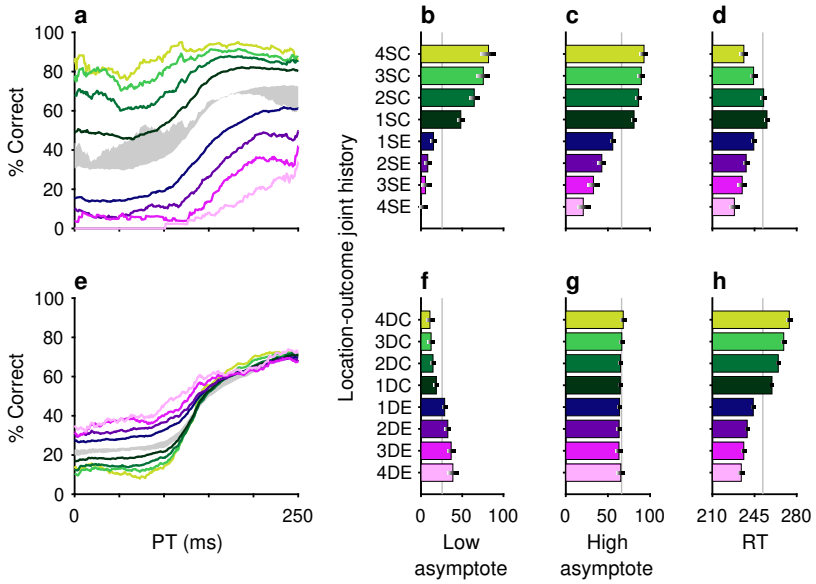

Figure 6. Motor bias is strongly modulated by the joint history of target location and trial outcome. The format of this figure is exactly the same as that of Fig. 5, except that it considers the joint history of target location and trial outcome (rather than target color and trial outcome).

acute in urgent choice tasks; namely, long RTs typically lead to long PTs and thereby to high accuracy, whereas short RTs typically lead to short PTs and thereby to near-chance performance. In general, however, relatively modest differences in mean RT, like those between SC and DC conditions (green bars in Figs. 5d, h, 6d, h), are difficult to attribute because variations in motor urgency (and thus RT) can be very large even under relatively simple conditions, such as during saccades to lone visual targets (Hikosaka et al., 2006; Hauser et al., 2018). Moreover, in an urgent choice paradigm like the CO task, numerous mechanisms are likely to determine the RT distribution (Shankar et al., 2011; Goldstein et al., 2022; Oor et al., 2023). As discussed below, while differences in choice accuracy were slight across the two monkey subjects, their differences in mean RT were larger, consistent with the notion that, in general, RTs have more complex variability.

Little evidence of interaction between color and location histories

Up to now, the results indicate that both the history of target colors and the history of target locations produce strong biases on task performance that are heavily gated by the outcomes (success/failure) of the trials within those histories. Although the resulting color- and location-driven effects have distinct temporal signatures, it is unclear whether the two sources of bias interact or, alternatively, act independently to determine performance at long processing times ($>$ 150 ms), when motor and perceptual mechanisms must coordinate and overlap. To assess the functional interaction between color- and location-driven biases, as well as its dependence on outcome, we devised an analysis in which the individual, separate effects of two trial histories on performance (during informed trials) are compared to their joint, combined effect (Methods; Salinas and Stanford, 2024).

In this approach, predictability serves as a criterion for interaction. Take event C to represent the outcome of an informed trial, which can be correct ($C=1$) or incorrect ($C=0$), and take A and B to represent two history sequences for different variables, say one for target color and another for target location. Then $P(C|A,B)$ represents the probability of an outcome given that histories A and B were observed jointly. The main idea is to estimate this probability in two ways: (1) by

calculating it directly from the experimental data, and (2) by predicting it based on the separate probabilities associated with the two histories, $P(C|A)$ and $P(C|B)$, assuming no statistical interaction between the history variables. If the prediction agrees with the measurement, the natural interpretation is that the two history variables exert independent influences on performance. Conversely, a significant discrepancy is interpreted as evidence of a functional interaction between the two history variables with respect to task performance.

This is best understood with an example. First, from the data, we calculate the probability of a correct outcome given a combined sequence; say, that the target in the prior trial had a different color *and* the same location. The measured value is $P(C=1|1DS)=0.68$ in this case. Then we calculate the outcome probabilities given the two separate histories, namely $P(C=1|1D\ \mathrm{color})$, which is the probability correct given that the prior target color was different, and $P(C=1|1S\ \mathrm{location})$, which is the probability correct given that the prior target location was the same. The values from the experimental data are 0.59 and 0.71, respectively. Then, by assuming a form of functional independence (conditional independence; Methods; Salinas and Stanford, 2024) between the two history variables, color and location, it is possible to use these two values measured separately to predict the probability correct when the two history conditions occur jointly. The predicted probability correct is 0.64 in this case (using Equation 7 and the fact that the overall probability correct was 0.66). The discrepancy between the two resulting numbers (0.68 versus 0.64) is small, so in a plot of predicted versus measured accuracy they correspond to a point near the diagonal (Fig. 7a, top row, point marked 1DS).

When the same analysis is performed for all the possible sequences that combine color and location going one trial back (sequences 1DD, 1SD, 1DS, 1SS), the four data points lie fairly close to the diagonal (Fig. 7a, top row), suggesting that the interaction between the two history variables is, at most, weak. To obtain a more complete characterization of the interdependence between color and location histories, the analysis is extended to sequences that go back further in time, up to 4 trials before the trial being evaluated (Fig. 7a, $H=1$ to $H=4$, from top to bottom row). As the history length increases, more sequence combinations arise and the contrast between predicted and measured accuracies includes more data points (with the number of trials represented by each point decreasing accordingly). For instance, going back 2 trials into the past, the 4 possible color sequences (SS, SD, DS, DD) can be combined with 4 possible location sequences to yield 16 unique combinations of color and location, so 16 data points (Fig. 7a, $H=2$). But critically, the pattern that emerges is similar for all the history lengths: the data generally lie along the diagonal, with regression slopes close to 1 (Fig. 7a, blue lines and β_2 coefficients). This indicates that, to a first approximation, color and location histories influence asymptotic performance independently of each other.

To validate this analysis method, we applied it to a different pair of history variables, target color and trial outcome. A strong interaction was expected for this combination of variables based on the earlier results, which showed that the same color history can either increase or decrease asymptotic accuracy depending on the outcome of the past trials considered (Fig. 5, green versus magenta). In this case, the effect of a given color history was paired with the effect of a given success history. For instance, $P(C=1|1D\ \mathrm{color})$, which is the same as above, was now paired with $P(C=1|1E\ \mathrm{outcome})$, the probability that the outcome of an informed trial is correct given that the prior outcome was an error, to estimate the joint effect of the two histories; that is, the probability $P(C=1|1DE)$, which is the probability correct when in the preceding trial the target color was different *and* the choice was erroneous (Fig. 7b, top row, point marked 1DE). For this example, the discrepancy between predicted (0.54) and empirical (0.62) probabilities is larger, and the four data points for the 1-trial histories deviate further from the diagonal (Fig. 7b, $H=1$, top row). Now, as longer histories are considered, the pattern that emerges demonstrates a consistent deviation

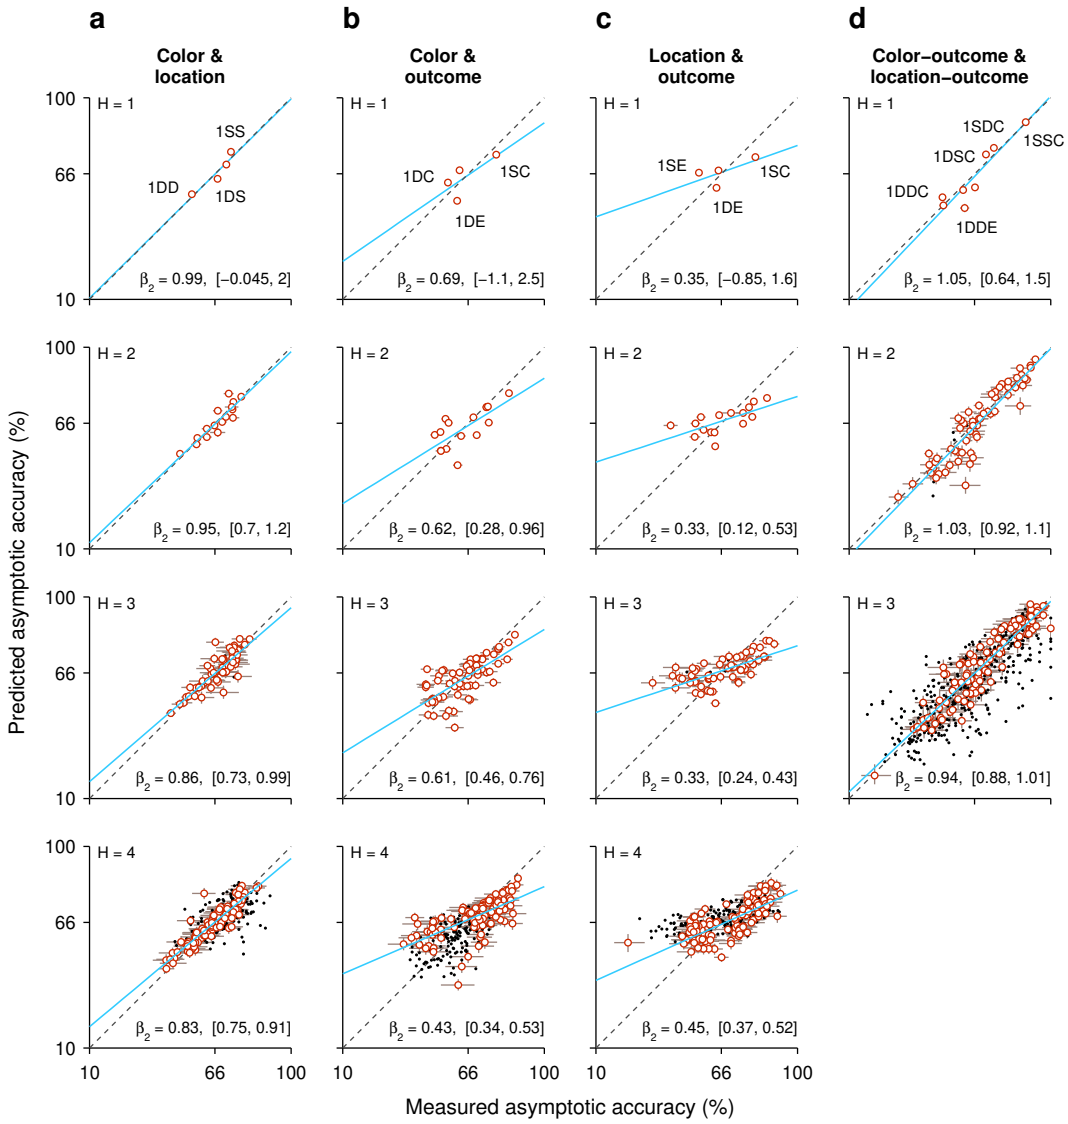

Figure 7. Tests of interaction between individual histories. Each panel plots the predicted task performance at long PTs (high asymptote) based upon independent contributions from two history effects (y axis), contrasted with the actual, measured task performance (x axis). Rows correspond to comparisons based on histories going 1, 2, 3, or 4 trials back (H, from top to bottom). Each point represents one specific history sequence. Labels in top panels identify a few combined history sequences in each case. Error bars represent 95% CIs. Red and black identify individual points with CI spans smaller (red) or larger (black) than 15 percentage points, which divides the data into high- and low-reliability points. Regression lines based on the red points are plotted in blue, with the slope (β_2) and its 95% CI indicated in each case. Dotted diagonal lines indicate equality. Note that the overall probability of success across all trials is 66% correct. (**a**) Interaction between target color and target location histories. (**b**) Interaction between target color and outcome histories. (**c**) Interaction between target location and outcome histories. (**d**) Interaction between target color and target location histories, each combined with outcome. Note that the joint effect of color and location histories is highly predictable based on their individual effects on performance.

from the diagonal and a shallower slope (Fig. 7b, blue lines and β_2 coefficients), indicating that the predicted asymptotic accuracy typically underestimates the magnitude of the joint history effect: when the empirical accuracy is above the mean ($>$ 66% correct along the x axis), the prediction is typically too low, and when the empirical accuracy is below the mean ($<$ 66% correct along the x axis), the prediction is typically too high. Thus, as expected, for color and outcome histories the results of this independence analysis are consistent with an underlying functional interaction.

Very similar results were obtained when we considered the combined effect of target location and outcome histories, except that the discrepancies between predicted and measured probabilities were slightly larger than for color and outcome (Fig. 7c, blue lines and β_2 coefficients). Again, this is consistent with earlier results (Fig. 6) and denotes a strong functional interaction, in this case between outcome and target-location histories.

Finally, we considered again the interaction between target-color and target-location histories but this time pairing each of them with success history. In this case the goal was to predict the probability of a correct outcome given the recent history of the three variables. For example, we measured $P(C=1|1SSC)$, which is the probability that a choice is correct when, in the preceding trial, the target color was the same, the target location was the same, and the outcome was correct. Then we compared the measured probability with an estimate based on two quantities, $P(C=1|1SC\ \text{color})$, which is the probability correct given that the prior trial had the same color and was correct, and $P(C=1|1SC\ \text{location})$, which is the probability correct given that the prior trial had the same location and was correct. For this example, the measured and predicted probabilities were nearly identical (0.892 versus 0.897; Fig. 7d, top row, point marked 1SSC). There were 8 unique sequences going back 1 trial, and most of the corresponding data points fell close to the equality diagonal (Fig. 7d, $H=1$, top row). For longer sequences, the resulting data points spanned a wide range of accuracies (for $H=3$, going from $<$ 20% correct to almost 100% correct), but the alignment was consistent with a slope of 1 in all cases. The data show that color- and location-driven biases exert largely independent effects on asymptotic performance, especially once outcome history is taken into consideration.

In summary, the results from our analyses based on predictability are consistent with the notion that although both color and location histories interact strongly with outcome history, the resulting color- and location-driven biases make approximately independent contributions to performance during informed choices.

Selection history effects are consistent across monkeys

Considering the diverse sources of history-driven biases that impact performance in the CO task, it would not be surprising if selection history effects were to manifest to widely varying degrees across individual subjects. Interestingly, though, we found that while RTs indeed tended to be somewhat idiosyncratic, variations in choice accuracy across history conditions were highly consistent in both sign and magnitude across the two monkey subjects.

Data from monkeys C and N were compared side-by-side for the history conditions already discussed in the preceding sections. For choice accuracy, the results recapitulated the trends identified earlier. During uninformed choices (Fig. 8a), color history had essentially no impact, whereas saccades were strongly biased toward prior target locations, especially when they were associated with success (or reward). Across subjects, the corresponding modulations of the low asymptote of the tachometric curve were nearly identical. During informed choices (Fig. 8b), the effect of location history alone was more subtle than that of color alone, but both feature dimensions had a remarkably strong impact on performance when prior outcomes were taken into account. Across subjects, the corresponding modulations of the high asymptote of the tachometric curve

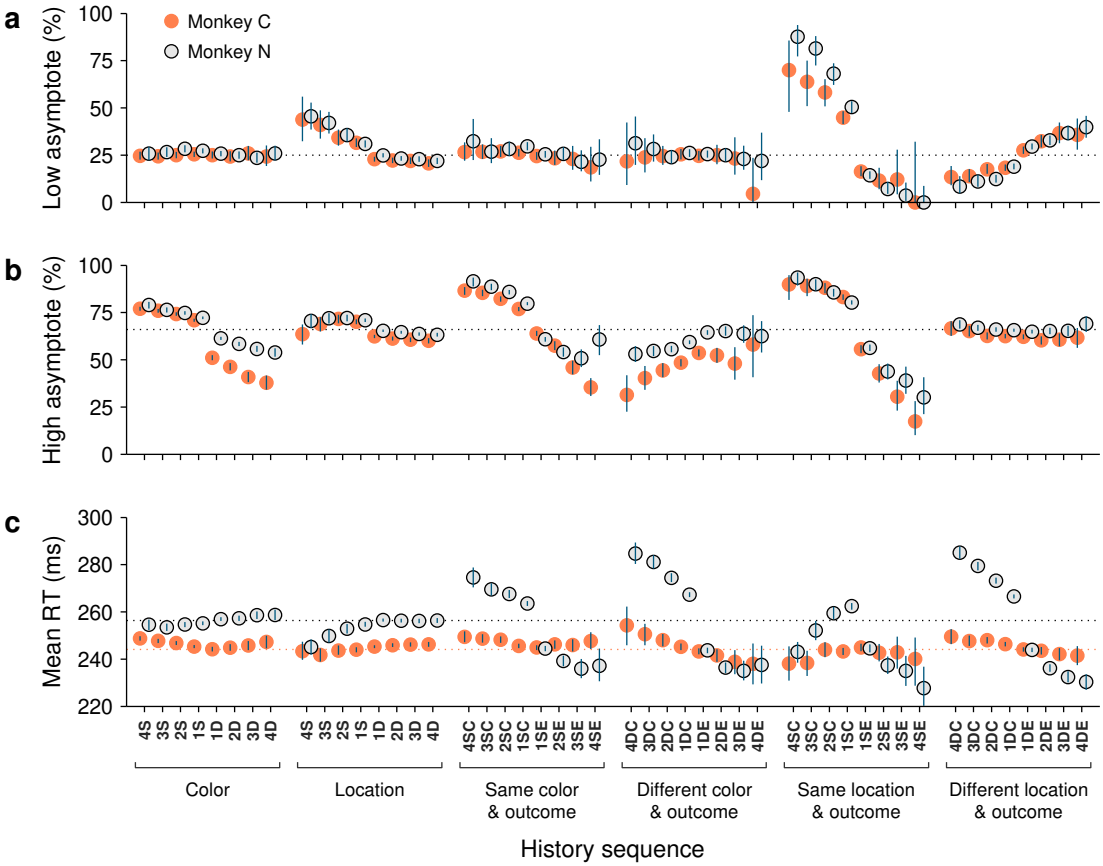

Figure 8. Consistency of history effects across individual subjects. The data in this figure are analogous to the bar plots shown in Figs. 2, 5, and 6, but for each of the two monkeys. In each panel, the x axis indicates specific history sequences (see labels in **c**). On the y axes, orange filled circles are results from monkey C and open circles are from monkey N. For clarity, the two sets of data points are slightly offset from each other along the x axis. Blue error bars indicate 95% CIs across trials. (**a**) Choice accuracy (percent correct) in short-PT trials (PT $<$ 100 ms, low asymptote) conditioned on each history sequence. The dotted line marks chance performance (25% correct). (**b**) Choice accuracy (percent correct) in long-PT trials (PT $>$ 150 ms, high asymptote) conditioned on each history sequence. The dotted line marks average performance for informed choices (66% correct). (**c**) Mean RT conditioned on each history sequence. Data points include all trials in each condition (correct and incorrect, all PTs). The dotted lines mark the overall mean RTs for the two monkeys (with the RTs of monkey N shifted by -29 ms; Methods).

were again similar in this case, except that monkey C was somewhat more sensitive to switches in target color (Fig. 8b; note slightly larger deviations from average in orange points for D color conditions).

In contrast to choice accuracy, the mean RTs were different across monkeys in several respects. First, as noted earlier, monkey N generally responded more slowly (Fig. 1c). This was visible in the data (Fig. 8c, note difference between dotted lines) even after having down-shifted all the RTs of this subject (by 29 ms; Methods). And second, the magnitudes of the history-driven modulations in RT were considerably larger for monkey N than for monkey C when trial outcome was taken into account (Fig. 8c, histories conditioned on outcome). Interestingly, however, for each type of history condition, the trends in these modulations were not dissimilar: for the D sequences, for instance, the changes across repetitions and between correct and error conditions were qualitatively

the same for the two animals (Fig. 8c, histories conditioned on different color and outcome, and different location and outcome).

These findings are in line with the notion that RTs are the product of numerous interacting factors, including sensory, cognitive, and motor mechanisms. They also illustrate a key property of the tachometric curve, which is that it is minimally sensitive to variations in urgency or motor preparation (Stanford et al., 2010; Shankar et al., 2011; Seideman et al., 2018; Stanford and Salinas, 2021). This is because, when trials are parsed according to PT, the variability in performance due to urgency is drastically curtailed. As a consequence, derived measures of choice accuracy, including the low and high asymptotes, are largely indicative of sensory and cognitive processing but not motor processing. Therefore, it is not surprising to find that the two monkey subjects showed relatively large differences in their sensitivities to history when measured by mean RT, but showed only slight differences when measured by choice accuracy.

Discussion

When searching for a color oddball, the performance of trained monkeys is excellent as long as the color of the target remains fixed (Bichot and Schall, 1999, 2002; McPeek and Keller, 2004; Scerra et al., 2019; Fig. 1 — figure supplement 3c, d). Otherwise, when target and distracter colors are randomly swapped across trials, accuracy drops precipitously (Song and McPeek, 2015; Fig. 1b). Here we investigated why. As in prior studies (Busse et al., 2011; Gupta et al., 2024), a large lapse rate ($\sim$35% in our case) that could be naively interpreted as stochastic behavior due to internal noise sources turned out to correspond, instead, to the manifestation of short-term biases driven by the recent history of events preceding each choice. In contrast to prior studies, we examined this problem under high urgency, which means that (uninformed) motor plans must be initiated early on. We exploited the temporal resolution afforded by the urgent-choice design to dissociate bias mechanisms that may start acting before the relevant sensory information becomes available from those that must start acting afterward. That is, mechanisms that bias the initial motor selection process versus mechanisms that bias the later perceptual evaluation of sensory features (color) and the subsequent choice. With this approach, we found that the histories associated with target location and target color made large contributions to choice accuracy that (1) were associated with motor and perceptual mechanisms, respectively, (2) had distinct time constants, (3) in both cases depended strongly on the success of the past trials considered, and (4) to a large degree acted independently of each other during perceptually informed trials. Having teased apart these selection history biases, we found that the magnitude of their effect on visuomotor performance could be remarkably large.

History-driven biases are ubiquitous, but in the context of visuomotor choice tasks, or sensory-guided tasks more generally, they are typically observed when the relevant perceptual information is weak and the choice difficult. When associated with threshold performance in perceptual tasks, such effects are often referred to as "serial dependencies" (Kiyonaga et al., 2017; Manassi et al., 2023; Cicchini et al., 2024). Many of these effects are specifically related to sensory representations kept in working memory, such as when a saccade is made to the location of a remembered visual cue (Papadimitriou et al., 2015, 2017), or during discrimination tasks in which two stimuli are presented sequentially, one after the other (Akrami et al., 2018; Boboeva et al., 2023). In other cases, the categorization of stimuli requires comparison with respect to an implicit boundary that must be represented internally, and the dependence on past trials is consistent with a boundary representation that is automatically updated as new stimuli are experienced (Ashourian and Loewenstein, 2011; Raviv et al., 2012; Hachen et al., 2018; Mendonça et al., 2020; Sheehan and Serences, 2022; Boboeva et al., 2023). In these types of study, the data can be generally understood as

perception resulting from current sensory evidence being combined with or compared to prior information stored internally, and the observed biases are most obvious when the available sensory evidence is weak, unreliable, so prior information takes precedence. In this framework, history-dependent perceptual biases are the manifestations of involuntary mechanisms for updating the relevant prior information (Lak et al., 2020; Mendonça et al., 2020). Such mechanisms are most useful under natural conditions, when statistical regularities are present (Kiyonaga et al., 2017; Fritsche et al., 2020; Hermoso-Mendizabal et al., 2020), but they may be counterproductive during typical laboratory tasks in which trials are uncorrelated (Akrami et al., 2018).

As in the studies just discussed, the task used here is a simple visuomotor paradigm — and yet, the differences are stark. In those studies, variations in performance were primarily driven by variations in stimulus strength or discriminability, and the observed history-driven biases typically led to relatively small deviations in accuracy of a few percentage points (but see Mendonça et al., 2020). Here, in contrast, the variability in performance was mainly driven by the amount of time available for processing the visual cue, and the magnitude of the history effects was extremely large (e.g., Figs. 5, 6), spanning nearly the full range from 0% to 100% correct for certain conditions. This is even more remarkable given that the task had no short-term memory requirements other than the task rule itself, and that the discriminability of the stimuli was always high — a condition under which the effects discussed in the previous paragraph typically vanish. Clearly, the framework developed for the abovementioned studies, based on the update of task-relevant information held in short-term memory, does not apply in this case.

Rather, our results are better understood as the confluence of two phenomena that have been described before. On one hand, the motor bias we observed is highly consistent with experiments showing that monkeys make much faster and much more accurate saccades to spatial locations that have been recently associated with a large reward than to locations that have not (Lauwereyns et al., 2002; Ikeda and Hikosaka, 2003; Hikosaka et al., 2006; Isoda and Hikosaka, 2008; Hauser et al., 2018). Such effects are caused by the recent history of reward-driven reinforcement going back a few (1–15) trials. And importantly, they are entirely compatible, both behaviorally and neurophysiologically, with the endogenous deployment of attention to a specific location tagged by reward (Maunsell, 2004; Hauser et al., 2018). Thus, although such spatial effects have been obtained under more minimalistic conditions (i.e., during saccades to lone visual targets), all indications are that they correspond closely to the motor biases in the CO task (i.e., during perceptually guided saccadic choices). Interestingly, however, expression of the strong bias that we observed in favor of previously rewarded target locations is likely to depend on the urgent nature of the CO task, because under non-urgent but otherwise comparable conditions, target location repetitions had no discernible effect on accuracy, and only a small influence on RT (Bichot and Schall, 2002). This contrast suggests that the suspicion mentioned in the Introduction was correct, namely, that the impact of motor biases on *informed* choices is greatly amplified and much more evident when motor activity is already ongoing before the target selection occurs.

On the other hand, the perceptual bias we observed is highly consistent with color priming, a well-known effect whereby the perceptual processing of a given color in one trial facilitates the processing of the same color in the next trial. Although priming effects are typically reported as differences in RT, the color bias in the CO task is most likely the same as the "priming of pop-out" reported in human participants (Maljkovic and Nakayama, 1994; McPeek et al., 1999), which is characterized as an attention-focusing mechanism that builds up cumulatively (as in Figs. 2a-d, 3), decays within 5–8 trials (as in Fig. 4b), and is largely involuntary (consistent with permanence in spite of months of practice). As with the motor bias, neurophysiological studies again indicate that the phenomenology is entirely compatible with known attention mechanisms; that is, the priming results can be explained by short-lived, adaptive changes in the responses of identified

neural circuits in visual and oculomotor areas that convey information about object salience and thereby guide the selection of saccade targets (Westerberg and Schall, 2021). In contrast to the motor bias, the color priming effect seems to be minimally affected by urgency, because its impact on choice accuracy under non-urgent conditions (Bichot and Schall, 2002) was comparable to that found here.

History-based influences on performance are ubiquitous and take many forms; they probably involve a wide diversity of neural mechanisms and circuits (Yashar et al., 2017; Failing and Theeuwes, 2018; Jiang and Sisk, 2019; Theeuwes 2019; Anderson et al., 2021). Consistent with this notion, we found little evidence of a functional interaction between the location- and feature-based biases (Fig. 7a, d), in agreement with a prior psychophysical study (Di Pace et al., 1997). Also, although both serial dependencies and color priming are history-driven and perceptual in nature, they likely represent mechanistically distinct phenomena. Indeed, a study that directly compared them found that their effects were uncorrelated, and concluded that they are mediated by separate mechanisms (Galluzzi et al., 2022). History effects are relevant not only to attention and visuomotor behaviors, as in the present study, but also to other cognitive functions such as working memory (Papadimitriou et al., 2015, 2017; Kuo, 2016; Akrami et al., 2018; Boboeva et al., 2023), value assessment (Failing and Theeuwes, 2018; Constantinople et al., 2019), or task switching (Wylie and Allport, 2000; Monsell, 2003; Stoet and Snyder, 2007). Given this, and the remarkable degree to which the performance of our monkey subjects was swayed by history-driven biases, we conclude that the contribution of recent history to behavioral variability is likely much larger than is generally appreciated, particularly under more naturalistic, less constrained conditions during which multiple such effects can be expressed. This is an important topic for further investigation.

Acknowledgments

We thank Denise Anderson for logistical and administrative support.

Funding

Research was supported by the National Institutes of Health through grant R01EY025172 (to ES and TRS) from the National Eye Institute.

References

Abrahamyan A, Silva LL, Dakin SC, Carandini M, Gardner JL (2016) Adaptable history biases in human perceptual decisions Proc Natl Acad Sci USA 113:E3548-57.

Akrami A, Kopec CD, Diamond ME, Brody CD (2018) Posterior parietal cortex represents sensory history and mediates its effects on behaviour. Nature 554:368–372.

Ashourian P, Loewenstein Y (2011) Bayesian inference underlies the contraction bias in delayed comparison tasks. PLoS ONE 6:e19551.

Anderson BA (2019) Neurobiology of value-driven attention. Curr Opin Psychol 29:27–33.

Anderson BA, Kim H, Kim AJ, Liao MR, Mrkonja L, Clement A, Grégoire L (2021) The past, present, and future of selection history. Neurosci Biobehav Rev 130:326–350.

Anderson BA, Laurent PA, Yantis S (2011) Value-driven attentional capture. Proc Natl Acad Sci. USA 108: 10367–10371.

Awh E, Belopolsky AV, Theeuwes J (2012) Top-down versus bottom-up attentional control: a failed

theoretical dichotomy. Trends Cogn Sci 16:437–443.

Barraclough DJ, Conroy ML, Lee D (2004) Prefrontal cortex and decision making in a mixed-strategy game. Nat Neurosci 7: 404–410.

Bichot NP, Schall JD (1999) Saccade target selection in macaque during feature and conjunction visual search. Vis Neurosci 16:81–89.

Bichot NP, Schall JD (2002) Priming in macaque frontal cortex during popout visual search: feature-based facilitation and location-based inhibition of return. J Neurosci 22:4675–4685.

Boboeva V, Pezzotta A, Clopath C, Akrami A (2023) From recency to central tendency biases in working memory: a unifying network model. eLife 12:RP86725 (preprint).

Brascamp JW, Blake R, Kristjánsson Á (2011) Deciding where to attend: priming of pop-out drives target selection. J Exp Psychol Hum Percept Perform 37:1700–1707.

Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Schölvinck ML, Zaharia AD, Carandini M (2011) The detection of visual contrast in the behaving mouse. J Neurosci 31:11351–11361.

Busse L, Katzner S, Treue S (2008) Temporal dynamics of neuronal modulation during exogenous and endogenous shifts of visual attention in macaque area MT. Proc Natl Acad Sci USA 105:16380–16385.

Carrasco M (2011) Visual attention: the past 25 years. Vision Res 51:1484–1525.

Chelazzi L, Estocinova J, Calletti R, Lo Gerfo E, Sani I, Della Libera C, Santandrea E (2014) Altering spatial priority maps via reward-based learning. J Neurosci 34:8594–8604.

Chen X, Mihalas S, Niebur E, Stuphorn V (2013) Mechanisms underlying the influence of saliency on value-based decisions. J Vis 13:18.

Chen X, Zirnsak M, Vega GM, Govil E, Lomber SG, Moore T (2020) Parietal cortex regulates visual salience and salience-driven behavior. Neuron 106:177–187.

Cicchini GM, Mikellidou K, Burr DC (2024) Serial Dependence in Perception. Annu Rev Psychol 75:129–154.

Constantinople CM, Piet AT, Bibawi P, Akrami A, Kopec C, Brody CD (2019) Lateral orbitofrontal cortex promotes trial-by-trial learning of risky, but not spatial, biases. Elife 8:e49744.

Dawid AP (1979) Conditional independence in statistical theory. J R Statist Soc B 41:1–31.

Di Pace E, Marangolo P, Pizzamiglio L (1997) Response bias in color priming. Acta Psychol (Amst) 95:3–14

Dolan RJ, Dayan P (2013) Goals and habits in the brain. Neuron 80:312–325.

Failing M, Theeuwes J (2018) Selection history: How reward modulates selectivity of visual attention. Psychon Bull Rev 25:514–538.

Fecteau JH, Munoz DP (2003) Exploring the consequences of the previous trial. Nat Rev Neurosci 4:435–443.

Fritsche M, Spaak E, de Lange FP (2020) A Bayesian and efficient observer model explains concurrent attractive and repulsive history biases in visual perception. Elife 9:e55389.

Galluzzi F, Benedetto A, Cicchini GM, Burr DC (2022) Visual priming and serial dependence are mediated by separate mechanisms J Vis 22:1.

Goldstein AT, Stanford TR, Salinas E (2022) Exogenous capture accounts for fundamental differ-

ences between pro- and antisaccade performance. Elife 11:e76964.

Gupta D, DePasquale B, Kopec CD, Brody CD (2024) Trial-history biases in evidence accumulation can give rise to apparent lapses in decision-making. Nat Commun 15:662.

Hachen I, Reinartz S, Brasselet R, Stroligo A, Diamond ME (2021) Dynamics of history-dependent perceptual judgment. Nat Commun 12:6036.

Hauser CK, Zhu D, Stanford TR, Salinas E (2018) Motor selection dynamics in FEF explain the reaction time variance of saccades to single targets. eLife pii:e33456.

Hermoso-Mendizabal A, Hyafil A, Rueda-Orozco PE, Jaramillo S, Robbe D, de la Rocha J (2020) Response outcomes gate the impact of expectations on perceptual decisions. Nat Commun 11:1057.

Hesterberg T (2014) What teachers should know about the bootstrap: resampling in the undergraduate statistics curriculum. arXiv: 1411.5279.

Hickey C, Chelazzi L, Theeuwes J (2014) Reward-priming of location in visual search PLoS One 9:e103372.

Hikosaka O, Nakamura K, Nakahara H (2006) Basal ganglia orient eyes to reward. J Neurophysiol 95:567–584.

Ikeda T, Hikosaka O (2003) Reward-dependent gain and bias of visual responses in primate superior colliculus. Neuron 39:693–700.

Isoda M, Hikosaka O (2008). A neural correlate of motivational conflict in the superior colliculus of the macaque. J Neurophysiol 100:1332–1342.

Itti L, Koch C (2001) Computational modelling of visual attention Nat Rev Neurosci 2:194–203.

Jiang YV, Sisk CA (2019) Habit-like attention. Curr Opin Psychol 29:65–70.

Kiyonaga A, Scimeca JM, Bliss DP, Whitney D (2017) Serial dependence across perception, attention, and memory. Trends Cogn Sci 21:493–497.

Kristjánsson Á, Ásgeirsson ÁG (2019) Attentional priming: recent insights and current controversies. Curr Opin Psychol 29:71–75.

Kristjánsson Á, Campana G (2010) Where perception meets memory: a review of repetition priming in visual search tasks. Atten Percept Psychophys 72:5–18.

Kuo BC (2016) Selection history modulates working memory capacity. Front Psychol 7:1564.

Lak A, Hueske E, Hirokawa J, Masset P, Ott T, Urai AE, Donner TH, Carandini M, Tonegawa S, Uchida N, Kepecs A (2020) Reinforcement biases subsequent perceptual decisions when confidence is low, a widespread behavioral phenomenon. Elife 9:e49834.

Lauwereyns J, Watanabe K, Coe B, Hikosaka O (2002) A neural correlate of response bias in monkey caudate nucleus. Nature 418:413–417.

Maljkovic V, Nakayama K (1994) Priming of pop-out: I. Role of features. Mem Cognit 22:657–672.

Maljkovic V, Nakayama K (1996) Priming of pop-out: II. The role of position. Percept Psychophys 58:977–991.

Manassi M, Murai Y, Whitney D (2023) Serial dependence in visual perception: A meta-analysis and review. J Vis 23:18.

Markowitz DA, Shewcraft RA, Wong YT, Pesaran B (2011) Competition for visual selection in the

oculomotor system. J Neurosci 31:9298–9306.

Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? Trends Cogn Sci 8:261–265.

McPeek RM, Keller EL (2004) Deficits in saccade target selection after inactivation of superior colliculus. Nat Neurosci 7:757–763.

McPeek RM, Maljkovic V, Nakayama K (1999) Saccades require focal attention and are facilitated by a short-term memory system. Vision Res 39:1555–1566.

Mendonça AG, Drugowitsch J, Vicente MI, DeWitt EEJ, Pouget A, Mainen ZF (2020) The impact of learning on perceptual decisions and its implication for speed-accuracy tradeoffs. Nat Commun 11:2757.

Monsell S (2003) Task switching. Trends Cogn Sci 7(3):134–140.

Oor EE, Stanford TR, Salinas E (2023) Stimulus salience conflicts and colludes with endogenous goals during urgent choices. iScience 26:106253.

Papadimitriou C, Ferdoash A, Snyder LH (2015) Ghosts in the machine: memory interference from the previous trial. J Neurophysiol 113:567–577.

Papadimitriou C, White RL 3rd, Snyder LH (2017) Ghosts in the machine II: neural correlates of memory interference from the previous trial. Cereb Cortex 27:2513–2527.

Raviv O, Ahissar M, Loewenstein Y (2012) How recent history affects perception: the normative approach and its heuristic approximation. PLoS Comput Biol 8:e1002731.

Salinas E, Shankar S, Costello MG, Zhu D, Stanford TR (2010) Waiting is the hardest part: comparison of two computational strategies for performing a compelled-response task. Front Comput Neurosci 4:153, 1–17.

Salinas E, Stanford TR (2021) Under time pressure, the exogenous modulation of saccade plans is ubiquitous, intricate, and lawful. Curr Opin Neurobiol 70:154–162.

Salinas E, Stanford TR (2024) Conditional independence as an assessment of evidence integration processes. PLoS One 19:e0297792.

Scerra VE, Costello MG, Salinas E, Stanford TR (2019) All-or-none context dependence delineates limits of FEF visual target selection. Curr Biol 29:294–305.e3.

Seideman JA, Stanford TR, Salinas E (2018) Saccade metrics reflect decision-making dynamics during urgent choices. Nat Commun 9:2907.

Seideman JA, Stanford TR, Salinas E (2022) A conflict between spatial selection and evidence accumulation in area LIP. Nat Commun 13:4463.

Shankar S, Massoglia DP, Zhu D, Costello MG, Stanford TR, Salinas E (2011) Tracking the temporal evolution of a perceptual judgment using a compelled-response task. J Neurosci 31:8406–8421.

Sheehan TC, Serences JT (2022) Attractive serial dependence overcomes repulsive neuronal adaptation. PLoS Biol 20:e3001711.

Siegel S, Castellan NJ (1988) Nonparametric statistics for the behavioral sciences. Boston, MA: McGraw-Hill.

Song JH, McPeek RM (2015) Neural correlates of target selection for reaching movements in superior colliculus. J Neurophysiol 113:1414–1422.

Stanford TR, Salinas E (2021) Urgent decision making: resolving visuomotor interactions at high temporal resolution. Annu Rev Vis Sci 7:323–348.

Stanford TR, Shankar S, Massoglia DP, Costello MG, Salinas E (2010) Perceptual decision making in less than 30 milliseconds. Nat Neurosci 13(3):379–385.

Stoet G, Snyder LH (2007) Extensive practice does not eliminate human switch costs. Cogn Affect Behav Neurosci 7:192–197.

Theeuwes J (2010) Top-down and bottom-up control of visual selection. Acta Psychol (Amst) 135:77–99.

Theeuwes J (2018) Visual selection: usually fast and automatic; seldom slow and volitional. J Cogn. 1:29.

Theeuwes J (2019) Goal-driven, stimulus-driven and history-driven selection. Curr Opin Psychol 29:97–101.

Westerberg JA, Schall JD (2021) Neural mechanism of priming in visual search. Atten Percept Psychophys 83:587–602.

Wolfe JM, Horowitz TS (2017) Five factors that guide attention in visual search. Nature Hum Behav 1:0058.

Wolfensteller U, Ruge H (2012) Frontostriatal mechanisms in instruction-based learning as a hallmark of flexible goal-directed behavior. Front Psychol 3:192.

Wyder MT, Massoglia DP, Stanford TR (2003) Quantitative assessment of the timing and tuning of visual-related, saccade-related, and delay period activity in primate central thalamus. J Neurophysiol 90:2029–2052.

Wylie G, Allport A (2000) Task switching and the measurement of "switch costs". Psychol Res 63:212–233.

Yashar A, White AL, Fang W, Carrasco M (2017) Feature singletons attract spatial attention independently of feature priming. J Vis 17:7.

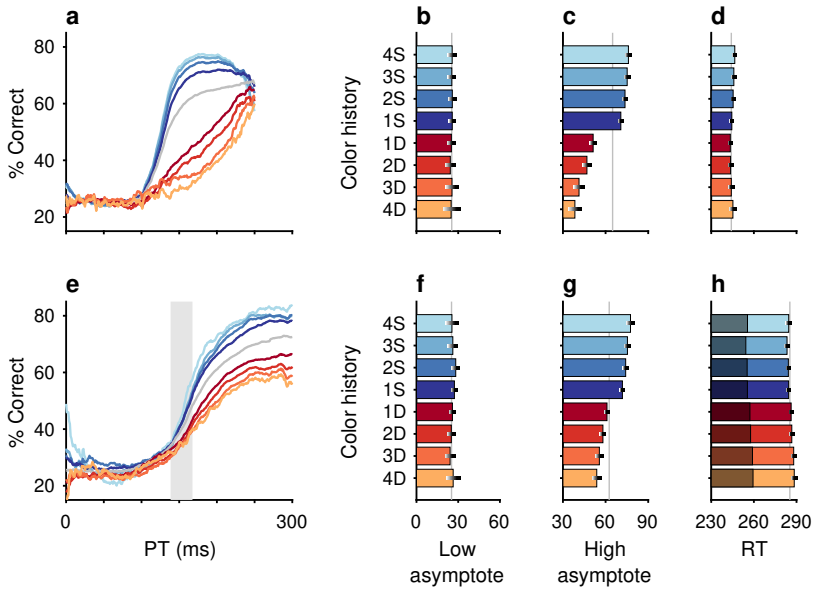

Figure 1 — figure supplement 1. Target color repetition affects perceptual accuracy in a consistent manner across subjects. (**a**–**d**) Task performance for monkey C. (**e**–**h**) Task performance for monkey N. (**a**, **e**) Individual subject tachometric curves denoting performance on the current trial conditioned on histories of preceding trials having the same target color (S, blue spectrum; 1–4 preceding trials) or different target color (D, red spectrum; 1–4 preceding trials). Average performance is shown in gray. The vertical gray shade represents the −29 ms shift used to align the data from monkey N with those from monkey C; it represents the centerpoints of the average tachometric curve of monkey N before (right edge of the shade) and after the shift (left edge of the shade). (**b**, **f**) Choice accuracies in short-PT trials (PT $<$ 100 ms, low asymptote). Error bars represent 95% CIs. Gray vertical lines denote values for the average tachometric curves. (**c**, **g**) Choice accuracies in long-PT trials (PT $>$ 150 ms, high asymptote). Error bars represent 95% CIs. Gray vertical lines denote values for the average tachometric curves. (**d**, **h**) Average RT $\pm$ 1 SEM for each history condition. The gray vertical lines denote the average RTs across all trials included in the analysis. Dark bars in **h** depict the aligned RT of monkey N shifted by −29 ms.

947

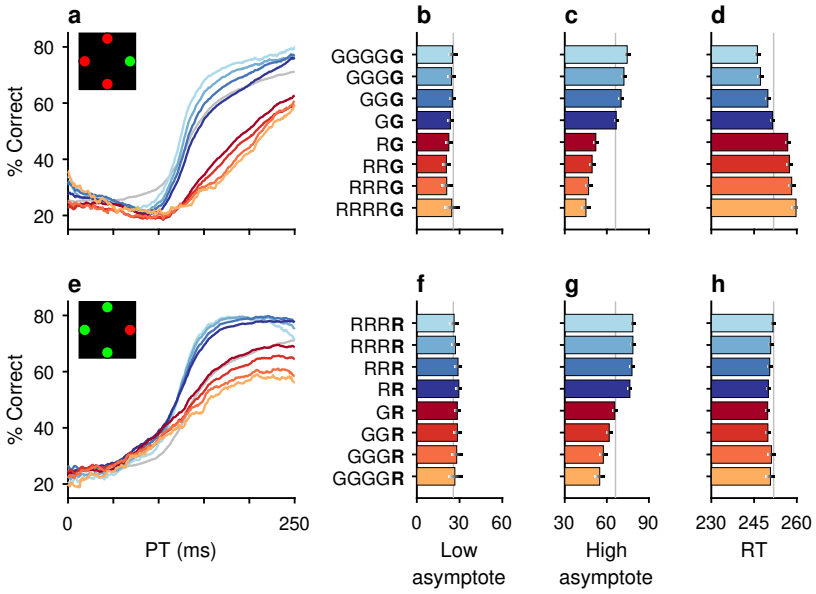

Figure 1 — figure supplement 2. Target color repetition affects decision accuracy in a consistent manner across colors. (**a**) Tachometric curves denoting performance on green-target trials (see target array inset) conditioned on histories of preceding trials having the same target color (blue spectrum; 1–4 preceding green trials) or a different target color (red spectrum; 1–4 preceding red trials). Average performance is shown in gray. (**b**) Choice accuracy in short-PT trials (PT $<$ 100 ms, low asymptote) conditioned on color history. Error bars represent 95% CIs across trials. Gray vertical line denotes value for the average tachometric curve. (**c**) Choice accuracy in long-PT trials (PT $>$ 150 ms, high asymptote) conditioned on color history. Error bars represent 95% CIs across trials. Gray vertical line denotes value for the average tachometric curve. (**d**) Average RT $\pm$ 1 SEM conditioned on color history. Gray vertical line denotes average RT across all trials included in the analysis. (**e–h**) As in **a–d** but for red targets among green distracters (see target array inset).

948

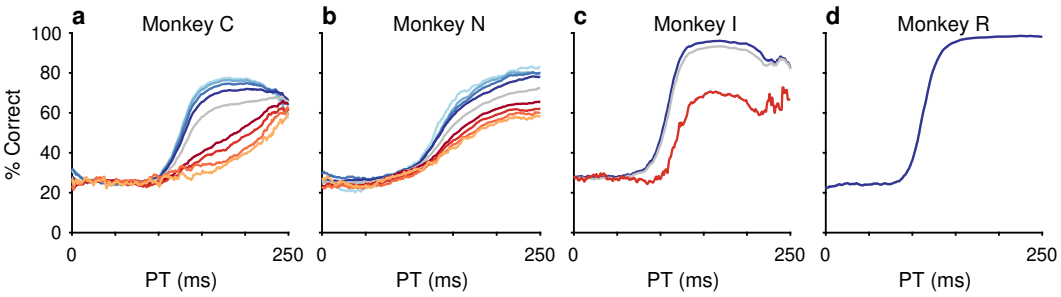

Figure 1 — figure supplement 3. Variability in performance across animals and rule sets. (**a**, **b**) Task performance for monkey C (**a**) and monkey N (**b**) in the current study; during the CO task, target colors and locations were randomly sampled across trials. Tachometric curves show performance on the current trial conditioned on the histories of preceding trials having the same target color (S, blue spectrum; 1S–4S sequences; dark to light) or a different target color (D, red spectrum; 1D–4D sequences; dark to light). (**c**) Task performance for monkey I from a previous study (Scerra et al., 2019). This subject ran a version of the CO task in which target locations were randomly sampled but target color remained constant for each block of (10–200) trials. The red curve represents average performance from the first 5 trials following a color switch. The blue curve denotes average performance on a block of trials but excluding the first 5 trials after a color switch. (**d**) Task performance for monkey R from a previous study (Scerra et al., 2019). This subject ran a version of the CO task in which target locations were randomly sampled but target color did not change. Bin width is 50 ms for all curves. The gray trace in each panel shows the average tachometric curve obtained from all the trials considered.

949