

Human inference in changing environments

Arthur Prat-Carrabin^{1,2†}, Robert C. Wilson^{3‡}, Jonathan D. Cohen³, and Rava Azeredo da Silveira^{1,2,3,4*}

¹Department of Physics, Ecole Normale Supérieure, PSL Research University, Paris, France

²Laboratoire de Physique Statistique, Centre National de la Recherche Scientifique, Université Pierre et Marie Curie, Université Denis Diderot, Paris, France

³Princeton Neuroscience Institute, Princeton University, Princeton, USA

⁴Department of Neurobiology, Weizmann Institute of Science, Rehovot, Israel

Present address: [†]Department of Economics, Columbia University, USA; [‡]Department of Psychology and Cognitive Science Program, University of Arizona, Tucson, USA

***For correspondence:** rava@ens.fr

Abstract

In past decades, the Bayesian paradigm has gained traction as a principled account of human behavior in inference tasks. Yet this success is tainted by the ubiquity of behavioral suboptimality and variability. We explore these discrepancies using an online inference task, in which we modulate the temporal statistics of hidden change points. We show that humans adapt their inference process to the implicit temporal statistics of stimuli, thereby behaving in an approximate Bayesian fashion. However, they exhibit biases and variability, and these depend on the history of stimuli. A systematic study of a broad family of optimal and sub-optimal models indicates that noise arises ‘internally’—in the inference process itself—rather than at the behavioral output. Specifically, we argue that humans mimic Bayesian inference by approximating the posterior with a modest number of samples. Our results contribute to a growing literature on sample-based cognition and compression by stochastic pruning.

1 Introduction

In a variety of inference tasks, human subjects use sensory cues as well as prior information on the signal in a manner consistent with Bayesian models. In tasks requiring the combination of a visual cue (such as the shape, position, texture, or motion of an object) with a haptic (Ernst and Banks, 2002; Battaglia et al., 2011), auditory (Battaglia et al., 2003), proprioceptive (van Beers et al., 1999), or a secondary visual cue (Jacobs, 1999; Hillis et al., 2004; Knill, 2007), human subjects seem to weigh information coming from each cue according to its uncertainty, in agreement with an optimal, probabilistic approach. Moreover, subjects appear also to optimally integrate prior

knowledge on spatial (Körding and Wolpert, 2004, 2006) and temporal (Miyazaki et al., 2005; Jazayeri and Shadlen, 2010) variables relevant to inference, in line with Bayes' rule.

The Bayesian paradigm hence offers an elegant and mathematically principled account of how humans can carry inference in the presence of uncertainty. In most experimental designs, however, successive trials are unrelated to each other. Yet, in many natural situations, the brain receives streams of new evidence from the environment: inference, then, unfolds in time. Moreover, natural mechanisms introduce sophisticated temporal statistics in the course of events (e.g., rhythmicity in locomotion, day-night cycles, and various structures found in speech). Are these temporal dynamics used by the brain to refine its *online* inference of the state of the environment?

Most studies that support a Bayesian account of human inference discuss average behaviors of subjects, and, thereby, side-step the issue of the variability in human responses. While an optimal Bayesian model provides a unique, deterministic response, human subjects exhibit noisy, and thus suboptimal, responses. Methods commonly used to model response variability, such as 'softmax' and probability-matching response-selection strategies, or, more recently, stochastic inference processes, suggest different forms of departure from Bayesian optimality. Which approach best captures human suboptimal behavior remains unclear. How can the successes and shortcomings of Bayesian theory be reconciled? What deviations from Bayesian models can account for the observed discrepancies from optimality?

To explore these questions, we use an online inference task based on a 'change-point' paradigm, i.e., random stimuli originating from a hidden state that changes occasionally and abruptly. The theoretical and experimental literature on inference problems for this class of signals (Adams and MacKay, 2007; Fearnhead and Liu, 2007; Wilson et al., 2010; Nassar et al., 2010; Wilson et al., 2013; Nassar et al., 2012) focuses primarily on a 'history-independent', Poisson case characterized by the absence of temporal structure: the probability of occurrence of a change is constant and does not depend on the past. Alongside such statistics, we introduce 'history-dependent' statistics that, conversely, include temporal structure. We use these two kinds of statistics in a change-point inference task, which allows us to examine the behavior of human subjects during online inference, and to compare their responses in the presence and in the absence of temporal structure.

Under these two contrasting conditions, the behavior of the Bayesian model and that of human subjects exhibit both similarities and discrepancies. To account for the latter, we investigate in what ways approximations of the Bayesian model alter behavior, in our task. The optimal estimation of a hidden state can be split into two steps: Bayesian posterior inference (computing optimally the belief distribution over the state space) and optimal response selection (using the belief distribution to choose the response that maximizes the expected reward). Suboptimal models introduce systematic errors or stochastic errors in the inference step or in the response-selection step, or in both, thus impacting behavior. Models presented in the change-point literature, along with new models we introduce, provide a wide range of such deviations from optimality, which we compare to experimental data. This allows us to assess how different sources of suboptimality impact behavior, and to what extent they can capture the salient features in human behavior.

The paper is outlined as follows. We first present the main aspects of our task, in which subjects observe a visual stimulus and infer an underlying, changing hidden state. The susceptibility

of subjects to a new stimulus is shown to differ appreciably between the two conditions (with and without temporal structure), and to adapt to the statistics of change points. We then analyze the variability in the subjects' responses, and show that it is not constant, but, rather, subject to modulations during the course of inference. After deriving the optimal, Bayesian solution of the inference problem in the context of our task, we examine its behavior in comparison with experimental data. We then turn to investigating a broad family of suboptimal models. In particular, motivated by the form of the variability present in our human data, we examine stochastic perturbations in both the inference step and in the response-selection step. These models reflect different forms of sampling — model subjects either perform inference with samples of probability distributions or select responses by sampling; the former option includes models with limited memory as well as sequential Monte Carlo (particle-filter) models. Finally, we discuss model fitting, which suggests that humans carry out stochastic approximations of the optimal Bayesian calculations through sampling-based inference (rather than sampling-based response selection).

Our observations confirm and extend the results reported in the change-point literature in the context of Poisson statistics (Wilson et al., 2010; Nassar et al., 2010; Wilson et al., 2013; Nassar et al., 2012), by exploring a more ecological (Nunes Amaral et al., 2004; Nakamura et al., 2007, 2008; Anteneodo and Chialvo, 2009; Hausdorff et al., 1995; Griffin et al., 2000; Ramus et al., 1999; Low et al., 2000; Campione and Véronis, 2002), non-Poisson, temporally structured environment. Likewise, our results broaden those of a number of studies in perception and decision-making that also investigate inference from stimuli with temporal statistics (Janssen and Shadlen, 2005; Ghose and Maunsell, 2002; Li and Dudman, 2013; Miyazaki et al., 2005; Jazayeri and Shadlen, 2010; ten Oever et al., 2014). Our experimental results show that human subjects implicitly learn the temporal statistics of the stimuli. Moreover, we highlight the variability ubiquitous in behavioral data, and show that it is not fixed in the context of our task: instead, it depends on the temporal statistics of the signal, and it is modulated over the course of inference. In fact, the variability in the responses of subjects appears to be in close relation to the Bayesian posterior. We find that a model in which the Bayesian posterior is approximated with a set of samples replicates the modulations of the behavioral variability during inference. This proposal adds to the growing literature on cognitive 'sample-based representations' of probability distributions (Goodman et al., 2008; Moreno-Bote et al., 2011; Gershman et al., 2012; Vul et al., 2014). Our results suggest that the brain carries out complex inference by manipulating a modest number of samples, as a low-dimensional approximation of the optimal, Bayesian posterior.

2 Results

2.1 Behavioral task, and history-independent vs. history-dependent stimuli

In our computer-based task, subjects are asked to infer, at successive trials, t , the location, on a computer screen, of a hidden point, the state, s_t , based on an on-screen visual stimulus, x_t , presented as a white dot on a horizontal line (Fig. 1A,B). Subjects can only observe the white dots, whose positions are generated around the hidden state according to a likelihood probability,

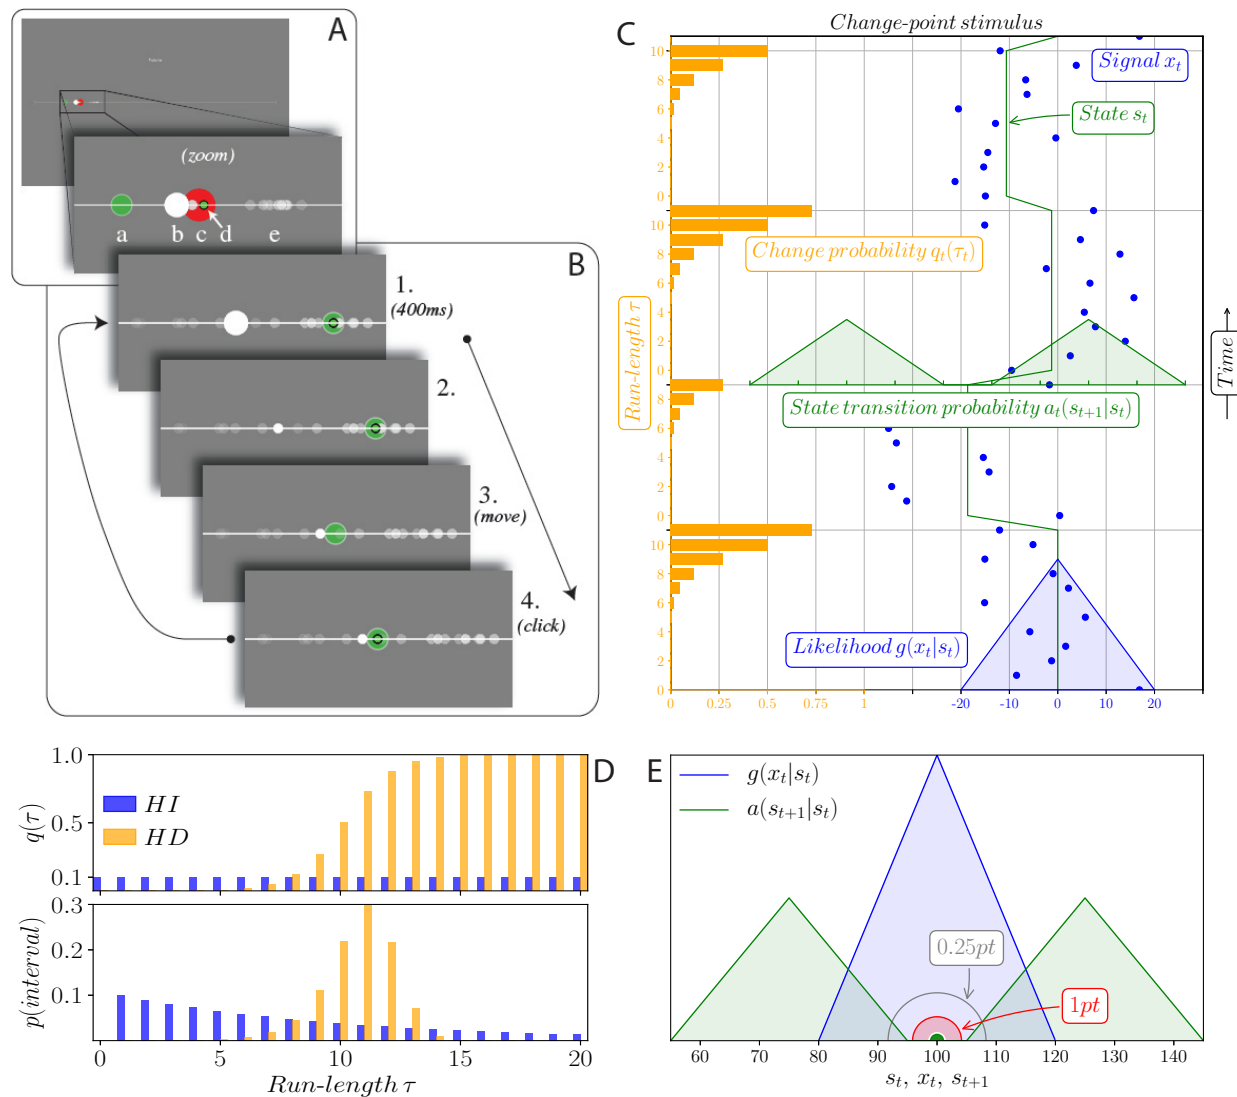


Figure 1. Illustration of the inference task. **A**. The various elements in the task appear on a horizontal white line in the middle of a gray screen. **a**: subject's pointer (green disk). **b**: new stimulus (white disk). **c**: state (red disk, only shown during tutorial runs). **d**: position of subject's previous click (green dot). **e**: for half of subjects, previous stimuli appear as dots decaying with time. **B**. Successive steps of the task: **1**, **2**: a new stimulus is displayed; to attract the subject's attention, it appears as a large, white dot for 400ms, after which it becomes smaller. **3**: the subject moves the pointer. **4**: The subject clicks to provide an estimate of the position of the state. After 100ms, a new stimulus appears, initiating the next trial. **C**. The position of the stimulus on the horizontal axis, x_t , is generated randomly around the current state, s_t , according to the triangle-shaped likelihood function, $g(x_t|s_t)$. The state itself is constant except at change points, at which a new state, s_{t+1} , is generated around s_t from the bimodal-triangle-shaped transition probability, $a(s_{t+1}|s_t)$. The run-length, τ_t , is defined as the number of trials since the last change point. Change points occur with the change probability $q(\tau_t)$ (orange bars), which depends on the run-length in the HD condition (depicted here). **D**. Top panel: change probability, $q(\tau)$, as a function of the run-length, τ . It is constant and equal to 0.1 in the HI condition, while it increases with the run-length in the HD condition. Consequently, the distribution of intervals between two consecutive change points (bottom panel) is geometric in the HI condition whereas it is peaked in the HD condition; in both conditions, the average duration of inter-change-point intervals is 10. **E**. Compared extents of the likelihood, $g(x_t|s_t)$ (green), the state transition probability, $a(s_{t+1}|s_t)$ (blue), the 'shot' resulting from a click (green dot), and the radii of the 1-point (red disk) and 0.25-point (grey circle) reward areas. A shot overlapping the red (gray) area yields 1 (0.25) point.

$g(x_t|s_t)$ (Fig. 1C,E, blue distribution). The state itself, s_t , follows a change-point process, i.e., it is constant except when it ‘jumps’ to a new location, which happens with probability q_t (the ‘hazard rate’ or ‘change probability’). The dynamics of change points are, hence, determined by the change probability, q_t . To examine the behavior of models and human subjects in different ‘environments’, we choose two kinds of signals which differ in their temporal structure. History-independent (**HI**) signals are memoryless, Poisson signals: q_t is constant and equal to 0.1. Consequently, the intervals between two change points last, on average, for 10 trials, and the distribution of these intervals is geometric (Fig. 1D, blue bars). Conversely, history-dependent (**HD**) signals are characterized by temporal correlation. Change points also occur every 10 trials, on average, but the distribution of the duration of inter-change-point intervals is peaked around 10. This corresponds to a change probability, q_t , that is an increasing function of the number of trials since the last change point — a quantity referred to as the ‘run-length’, τ_t . We thus denote it by $q(\tau_t)$. In HD signals, change points occur in a manner similar to a ‘jittered periodic’ process, though the regularity is not readily detected by subjects.

When a change point occurs, the state randomly jumps to a new state, s_{t+1} , according to a state transition probability, $a(s_{t+1}|s_t)$ (Fig. 1C,E, green distribution). The likelihood, g , and the state transition probability, a , overlap, thus allowing for ambiguity when a new stimulus is viewed: is it a random excursion about the current state, or has the state changed? At each trial, subjects click with a mouse to give their estimate, \hat{s}_t , of the state. The reward they receive for each response is a decreasing function, R , of the distance between the state and the estimate, $|\hat{s}_t - s_t|$: one reward point if the estimate falls within a given, short distance from the state, 0.25 point if it falls within twice that distance, and 0 point otherwise (Fig. 1E). The task is presented as a game to subjects: they are told that someone is throwing snowballs at them. They cannot see this hidden person (whose location is the state, s_t), but they observe the snowballs as white dots on the screen (the stimulus, x_t). After several tutorial runs (in some of which the state is shown), they are instructed to use the snowballs to guess the location of the person (i.e., produce an estimate, \hat{s}_t). Additional details on the task are provided in *Methods*.

2.2 Learning rates adapt to the temporal statistics of the stimulus

A typical example of a subject’s responses is displayed in Fig. 2A. To describe the data, we focus, throughout this paper, on three quantities: the learning rate, defined as the ratio of the ‘correction’, $\hat{s}_{t+1} - \hat{s}_t$, to the ‘surprise’, $x_{t+1} - \hat{s}_t$; the repetition propensity, defined as the proportion of trials in which the learning rates vanishes ($\hat{s}_{t+1} = \hat{s}_t$); and the standard deviation of the responses of the subjects. The learning rate represents a normalized measure of the susceptibility of a subject to a new stimulus. If the new estimate, \hat{s}_{t+1} , is viewed as a weighted average of the previous estimate, \hat{s}_t , and the new stimulus, x_{t+1} , the learning rate is the weight given to x_{t+1} . A learning rate of 0 means that the subject has not changed its estimate upon observing the new stimulus; a learning rate of 0.5 means that the new estimate is equidistant from the previous estimate and the new stimulus; and a learning rate of 1 means that the new estimate coincides with the new stimulus, and the past is ignored (Fig. 2A).

Our data show that for human subjects the learning rate is not constant, and can vary from no correction at all (learning rate ≈ 0) to full correction (learning rate ≈ 1). We investigated how the average learning rate behaved in relation to the run-length, in the HI and HD conditions. As the run-length is not directly accessible to subjects, in our analyses we used the empirical run-length, $\tilde{\tau}$, a similar quantity derived from the subjects' data (see Methods). Unless otherwise stated, we focus our analyses on cases in which the surprise, $x_{t+1} - \hat{s}_t$, is in the [8,18] window, in which there is appreciable ambiguity in the signal.

A first observation emerging from our data is that the learning rate changes with the run-length, in a quantitatively different fashion depending on the condition (HI or HD). In the HI condition, learning rates at short run-length ($\tilde{\tau} \in [5, 6]$) are significantly higher than at long run-length ($\tilde{\tau} \in [9, 10]$), i.e., the learning rate decreases with run-length (Fig. 2B, blue bars). In the HD condition, the opposite occurs: learning rates are significantly higher at long run-lengths (Fig. 2B, orange bars), indicating that subjects modify their inference depending on the temporal structure of the signal. In addition, at short run-lengths, learning rates are significantly lower in the HD condition than in the HI condition; this suggests that subjects take into account the fact that a change is less likely at short run-lengths in the HD condition. The opposite holds at long run-lengths: HD learning rates are markedly larger than HI ones (Fig. 2B).

Inspecting the dependence of the learning rate on the run-length (Fig. 2C), we note that the HD learning-rate curve adopts a 'smile shape', unlike the monotonic curve in the HI condition. The HI curve is consistent with a learning rate that simply decreases as additional information is accumulated on the state. In the HD condition, initially the learning rate is suppressed, then boosted at longer run-lengths, reflecting the modulation in the change probability.

These observations demonstrate that subjects adapt their learning rate to the run-length, and that in the HD condition subjects make use of the temporal structure in the signal. These results are readily intuited: shortly after a change point, the learning rate should be high, as little is known about the new state, while at longer run-lengths the learning rate should tend to zero as the state is more and more precisely inferred. This decreasing behavior is observed, but only in the HI condition. The HD condition introduces an opposing effect: as the run-length grows, new stimuli are increasingly likely to divulge the occurrence of a new state, which advocates for adopting a higher learning rate. This tradeoff is reflected in our data in the 'smile shape' of the HD learning-rate curve (Fig. 2C). The increase in learning rate at long run-lengths is reminiscent of the behavior of a driver waiting at a red light: as time passes, the light is increasingly likely to turn green; as a result, the driver is increasingly susceptible to react and start the car.

2.3 Human repetition propensity

A closer look at the data presented in the previous section reveals that in a number of trials the learning rate vanishes, i.e., $\hat{s}_{t+1} = \hat{s}_t$. The distribution of the subjects' corrections, $\hat{s}_{t+1} - \hat{s}_t$, exhibits a distinct peak at zero (Fig. 3A). In other words, in a fraction of trials, subjects click twice consecutively on the same pixel. We call such a response a 'repetition', and the fraction of repetition trials the 'repetition propensity'. The latter varies with the run-length: it increases with

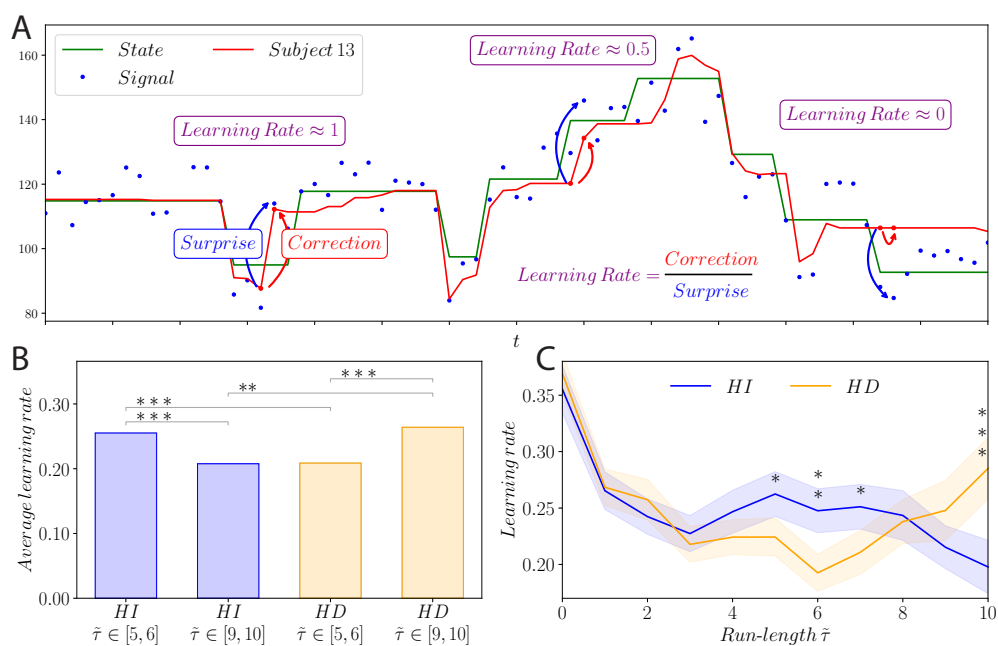


Figure 2. Human learning rates for different temporal statistics. **A.** Illustration of the learning rate using a sample of a subject’s responses (red line). The ‘surprise’ (blue arrow) is the difference, $x_{t+1} - \hat{s}_t$, between the estimate at trial t , \hat{s}_t (red), and the new stimulus at trial $t + 1$, x_{t+1} (blue). The ‘correction’ (red arrow) is the difference between the estimate at trial t and the estimate at trial $t + 1$, $\hat{s}_{t+1} - \hat{s}_t$. The ‘learning rate’ is the ratio of correction to surprise. **B.** Average learning rates in HI (blue) and HD (orange) conditions, at short run-lengths ($\tilde{\tau} \in [5, 6]$) and long run-lengths ($\tilde{\tau} \in [9, 10]$). **C.** Average learning rates in HI (blue) and HD (orange) conditions, vs. run-length $\tilde{\tau}$. Shaded bands indicate the standard error of the mean. **B,C.** Stars indicate p-values of one-sided Welch’s t-tests, which do not assume equal population variance. Three stars: $p < 0.01$; two stars: $p < 0.05$; one star: $p < 0.1$.

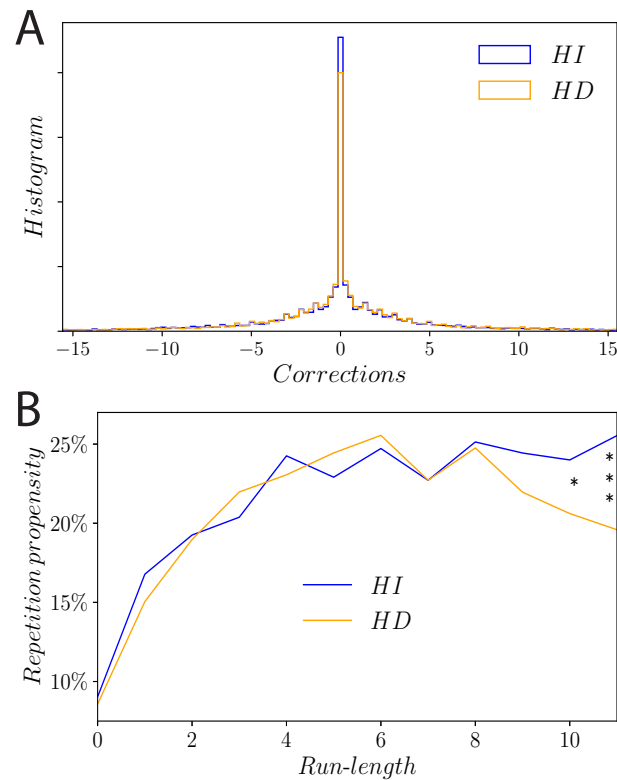


Figure 3. Human repetition propensity. **A.** Histogram of subject corrections (difference between two successive estimates, $\hat{s}_{t+1} - \hat{s}_t$), in the HI (blue) and HD (orange) conditions. The width of bins corresponds to one pixel on the screen, thus the peak at zero represents the repetition events ($\hat{s}_{t+1} = \hat{s}_t$). **B.** Repetition propensity, i.e., proportion of occurrences of repetitions in the responses of subjects, as a function of run-length, in the HI (blue) and HD (orange) conditions. Stars indicate p-values of Fisher's exact test of equality of the repetition propensities between the two conditions, at each run-length.

τ in both HI and HD conditions, before decreasing in the HD condition for long run-lengths (Fig. 3B).

What may cause the subjects' repetition behavior? The simplest explanation is that, after observing a new stimulus, a subject may consider that the updated best estimate of the state lands on the same pixel as in the previous trial. The width of one pixel in arbitrary units of our state space is 0.28. As a comparison, the triangular likelihood, g , has a standard deviation, σ_g , of 8.165. An optimal observer estimating the center of a Gaussian density of standard deviation σ_g , using 10 samples from this density, comes up with a posterior density with standard deviation $\sigma_g/\sqrt{10} \approx 2.6$. Therefore, after observing even 10 successive stimuli, the subjects' resolution is not as fine as a pixel (it is, in fact, 10 times coarser). This indicates that the subjects' repetition propensity is higher than the optimal one (the behavior of the optimal model, presented below,

indeed exhibits a lower average repetition propensity than that of the subjects). Another possible explanation is that even though the new estimate falls on a nearby location, a motor cost prohibits a move if it is not sufficiently extended to be ‘worth it’ (Morasso, 1981; Wolpert, 1997; Shadmehr, 2009; Rigoux and Guigon, 2012). A third, heuristic explanation is that humans are subject to a ‘repetition bias’ according to which they repeat their response irrespective of their estimate of the state.

Regardless of its origin, the high repetition propensity in data raises the question of whether it dominates the behavior of the average learning rate. As a control, we excluded all occurrences of repetitions in subjects’ data and carried out the same analyses on the truncated dataset. We reached identical conclusions, namely, significant discrepancies between the HI and HD learning rates at short and long run-lengths, albeit with, naturally, higher average rates overall (see Supplementary Fig. 1).

2.4 The variability in subject responses evolves over the course of inference

In the previous two sections, we have examined two aspects of the distribution of responses: the average learning rate and the probability of response repetition. We now turn to the variability in subject responses. Although all subjects were presented with identical series of stimuli, x_t , their responses at each trial were not the same (Fig. 4A). This variability appears in both HI and HD conditions. The distribution of responses around their averages at each trial has a width comparable to that of the likelihood distribution, $g(x_t|s_t)$ (Fig. 4B). More importantly, the variability in the responses (as measured by the standard deviation) is not constant, but decreases with successive trials following a change point, at short run-lengths (Fig. 4C). Comparing the HI and HD conditions, we observe that for run-lengths shorter than 7, the standard deviation in the HD condition is significantly lower than that in the HI condition. At longer run-lengths, the two curves cross and the variability in the HD condition becomes significantly higher than in the HI condition. The HD curve adopts, again, a ‘smile shape’ (Fig. 4C). What is the origin of the response variability? Because it changes with the run-length and the HI vs. HD condition, it cannot be explained merely by the presence of noise independent from the inference process, such as pure motor noise. In order to encompass human behavior in a theoretical framework and to investigate potential sources of this inference-dependent variability, we start by comparing the recorded behavior with that of an optimal observer.

2.5 Optimal estimation: Bayesian update and maximization of expected reward

We derive the optimal solution for the task of estimating the hidden state, s_t , given random stimuli, x_t . The first step (the ‘inference step’) is to derive the optimal posterior distribution over the state, s_t , using Bayes’ rule. Because the state is a random variable coupled with the run-length, τ_t , another random variable, it is convenient to derive the Bayesian update equation for the (s_t, τ_t) pair (more precisely, the (s_t, τ_t) pair verifies the Markov property, whereas s_t alone does not, in

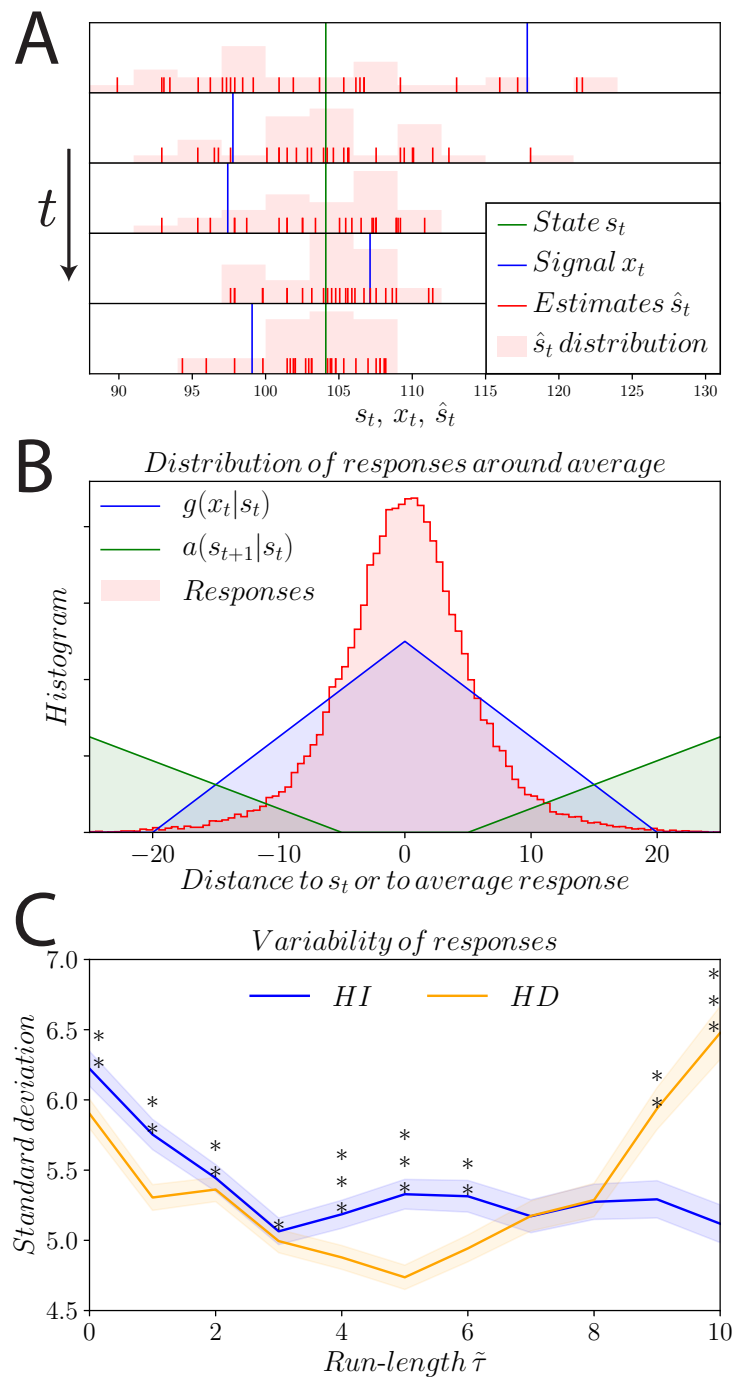


Figure 4. Variability in human inference. **A.** Responses of subjects in an example of 5 consecutive stimuli. In this example, there is no change point and the state (green) is constant. At each trial (from top to bottom), subjects observe the stimuli (blue) and provide their responses (red bars). A histogram of the locations of the responses is obtained by counting the number of responses in bins of width 3 (light red). **B.** Distribution of the responses of subjects around their average (red), compared to the likelihood, g (blue), and the state transition probability, a (green). **C.** Standard deviation of the responses of subjects vs. run-length, $\tilde{\tau}$, in the HI (blue) and HD (orange) conditions. Stars indicate p-value of Levene’s test of equality of variance between the two conditions, at each $\tilde{\tau}$. Shaded bands indicate the standard error of the standard deviation (Ahn and Fessler, 2003).

the HD condition). We denote by $x_{1:t}$ the string of stimuli received between trial 1 and trial t , and by $p_t(s, \tau | x_{1:t})$ the probability distribution over (s, τ) , at trial t , after having observed the stimuli $x_{1:t}$. At trial $t + 1$, Bayes' rule yields $p_{t+1}(s, \tau | x_{1:t+1}) \propto g(x_{t+1}|s)p_{t+1}(s, \tau | x_{1:t})$. Furthermore, we have the general transition equation,

$$p_{t+1}(s, \tau | x_{1:t}) = \sum_{\tau_t} \int_{s_t} p_{t+1}(s, \tau | s_t, \tau_t) p_t(s_t, \tau_t | x_{1:t}) ds_t, \quad (1)$$

given by the change-point statistics. As the transition probability, $p_{t+1}(s, \tau | s_t, \tau_t)$, can be expressed using $q(\tau_t)$ and $a(s|s_t)$ (see Methods for details), we can reformulate the update equation as

$$p_{t+1}(s, \tau | x_{1:t+1}) = \frac{1}{Z_{t+1}} g(x_{t+1}|s) \left[\mathbb{1}_{\tau=0} \sum_{\tau_t} q(\tau_t) \int_{s_t} a(s|s_t) p_t(s_t, \tau_t | x_{1:t}) ds_t + \mathbb{1}_{\tau>0} (1 - q(\tau - 1)) p_t(s, \tau - 1 | x_{1:t}) \right], \quad (2)$$

where $\mathbb{1}_C = 1$ if condition C is true, 0 otherwise; and Z_{t+1} is a normalization constant. This equation includes two components: a 'change-point' one ($\tau = 0$) and a 'no change-point' one ($\tau > 0$). We call the model that performs this Bayesian update of the posterior the *OptimalInference* model.

Finally, following the inference step just presented (i.e., the computation of the posterior), a 'response-selection step' determines the behavioral response. At trial t and for a response \hat{s}_t , the expected reward is $\mathbb{E}_s R = \int R(|\hat{s}_t - s|) p_t(s | x_{1:t}) ds$. The optimal strategy selects the response, \hat{s}_t , that maximizes this quantity. Before exploring the impact of relaxing the optimality in the inference step, in the response-selection step, or both, we examine, first, the behavior of the optimal model.

2.6 The optimal model captures qualitative trends in learning rate and repetition propensity

Equipped with the optimal model for our inference task, we compare its output to experimental data. For short run-lengths ($\tau < 8$), the learning rates in both HI and HD conditions decrease as a function of the run-length, and the HD learning rates are lower than their HI counterparts. They increase, however, at longer run-lengths ($\tau \geq 8$) and ultimately exceed the HI learning rates; these, by contrast, decrease monotonically (Fig. 5A, solid line). These trends are similar to those observed in behavioral data (Fig. 5A, dashed line). Hence, the modulation of the subjects' learning rates with the temporal statistics of the stimuli, and over the course of inference, is consistent, at least qualitatively, with that of a Bayesian observer.

Although a Bayesian observer can, in principle, hold a continuous posterior distribution, we discretize, instead, the posterior, in order to reproduce the experimental condition of a pixelated screen. This discretization allows for repetitions. The repetition propensity of the optimal model varies with the run-length: it increases with τ in both HI and HD conditions, and decreases in the HD condition for long run-lengths, a pattern also found in experimental data (Fig. 5B).

Hence, the optimal model captures the qualitative trends in learning rate and repetition propensity present in the responses of the subjects. Quantitative differences, however, remain. The learning rates of the subjects, averaged over both HI and HD conditions, are 43% higher than the average learning rate in the optimal model, and the average repetition propensity of the subjects is 9 percentage points higher than that in the optimal model.

Moreover, subject responses are variable, while the optimal model is deterministic: at each trial, the optimal estimate is a deterministic function of the observed stimuli. As such, the optimal model does not capture the variability inherent in the behavior of subjects. The modulations of the behavioral variability as a function of the run-length and of the temporal structure of the signal, exhibited above, suggest that this variability depends on the inference process. A look at the standard deviation of the posterior distribution of the optimal model reveals that it decreases with the run-length, in the HI condition: following a change point, the posterior becomes narrower as new stimuli are observed. In the HD condition, the standard deviation of the posterior exhibits a ‘smile shape’ as a function of the run-length: it decreases until the run-length reaches 5, then increases for larger run-lengths (Fig. 5C). This behavior is similar to that of the standard deviation of the responses of the subjects, and it supports the hypothesis that the behavioral variability in the data is, at least in part, related to the underlying inference and decision processes.

In what follows, we introduce an array of suboptimal models, with the aim of resolving the quantitative and qualitative discrepancies between the behavior of the optimal model and that of the subjects. In particular, we formulate several stochastic models which include possible sources of behavioral variability. Two scenarios are consistent with the scaling of behavioral variability with the width of the posterior: stochasticity in the inference step (i.e., in the computation of the posterior) and stochasticity in the response-selection step (i.e., in the computation of an estimate from the posterior). The models we examine below cover both these scenarios.

2.7 Suboptimal models reflecting cognitive limitations

In the previous sections, we have examined the learning rate of the subjects, their repetition propensity, and the variability in their responses; comparison of the behaviors of these quantities to that of the Bayesian, optimal model, revealed similarities (namely, the qualitative behaviors of the learning rate and of the repetition propensity) and discrepancies (namely, quantitative differences in these two quantities, and lack of variability in the optimal model). Although the latter call for a non-Bayesian account of human behavior, the former suggest not to abandon the Bayesian approach altogether (in favor, for instance, of *ad hoc* heuristics). Thus, we choose to examine a family of sub-optimal models obtained from a sequence of deviations away from the Bayesian model, each of which captures potential cognitive limitations hampering the optimal performance.

In the Bayesian model, three ingredients enter the generation of a response upon receiving a stimulus: first, the belief on the structure of the task and on its parameters; second, the inference algorithm which produces a posterior on the basis of stimuli; third, the selection strategy which maps the posterior into a given response. The results presented above, exhibiting the similarity between the standard deviation of the Bayesian posterior and the standard deviation of the




Figure 5. Behavior of the learning rate, the repetition propensity, and the standard deviation of responses, as functions of the run-length, for the optimal model and the human subjects. **A.** Average learning rate as a function of the run-length. In the HI condition, the learning rate decreases with the run-length, for both the optimal model and the subjects. In the HD condition, learning rates in the optimal model are lower than in the HI condition, for short run-lengths, and higher for long run-lengths. The learning rate of subjects exhibits a similar smile shape, in the HD condition. **B.** Repetition propensity, i.e., proportion of repetition trials, as a function of the run-length. **C.** Standard deviation of the responses of the subjects (dashed lines) and of the optimal model (solid lines), and standard deviation of the optimal, Bayesian posterior distribution (long dashes), as a function of the run-length. The optimal model is deterministic and, thus, exhibits no variability in its responses. The optimal posterior distribution, however, has a positive standard deviation which decreases with the run-length, in the HI condition, and exhibits a smile shape, in the HD condition.

responses of the subjects (Fig. 5C), points to a potential departure from the optimal selection strategy, in which the posterior is sampled rather than maximized. This sampling model, which we implement (see below), captures qualitatively the modulated variability of responses; sizable discrepancies in the three quantities we examine, however, remain (see *Methods*). Hence, we turn to the other ingredients of the estimation process, and we undertake a systematic analysis of the effects on behavior of an array of deviations away from optimality.

Below, we provide a conceptual presentation of the resulting models; we fit them to experimental data, and comment on what the best-fitting models suggest about human inference and estimation processes. For the detailed mathematical descriptions of the models, and an analysis of the ways in which their predictions depart from the optimal behavior, we refer the reader to the *Methods* section.

Models with erroneous beliefs on the statistics of the signal Our first model challenges the assumption, made in the optimal model, of a perfectly faithful representation of the set of parameters governing the statistics of the signal. Although subjects were exposed in training phases to blocs of stimuli in which the state, s_t , was made visible, they may have learned the parameters of the generative model incorrectly. We explore this possibility, and, here, we focus on the change probability, $q(\tau)$, which governs the dynamics of the state. (We found that altering the value of this parameter had a stronger impact on behavior than altering the values of any of the other parameters.) In the HD condition, $q(\tau)$ is a sigmoid function shaped by two parameters: its slope, $\lambda = 1$, which characterizes ‘how suddenly’ change points become likely, as a function of τ ; and the average duration of inter-change-points intervals, $T = 10$. In the HI condition, $q = 0.1$ is constant; it can also be interpreted as an extreme case of a sigmoid in which $\lambda = 0$ and $T = 1/q = 10$. We implement a suboptimal model, referred to as *IncorrectQ*, in which these two quantities, λ and T , are treated as free parameters, thus allowing for a broad array of different beliefs in the temporal structure of the signal (Fig. 6A).

Models with limited memory Aside from operating with an inexact representation of the generative model, human subjects may use a suboptimal form of inference. In the HD condition, the optimal model maintains ‘in memory’ a probability distribution over the entire (s, τ) -space (see Eq. (2)), thus keeping track of a rapidly increasing number of possible histories, each characterized by a sequence of run-lengths. Such a process entails a large computational and memory load. We explore suboptimal models that alleviate this load by truncating the number of possible scenarios stored in memory; this is achieved through various schemes of approximations of the posterior distribution. More specifically, in the following three suboptimal models, the true (marginal) probability of the run-lengths, $p_t(\tau|x_{1:t})$, is replaced by an approximate probability distribution.

A first, simple way of approximating the marginal distribution of the run-lengths is to consider only its mean, i.e., to replace it by a Kronecker delta which takes the value 1 at an estimate of the expected value of the run-lengths. Nassar et al. (2010) introduce a suboptimal model based on this idea, some details of which depend on the specifics of the task; we implement a generalization

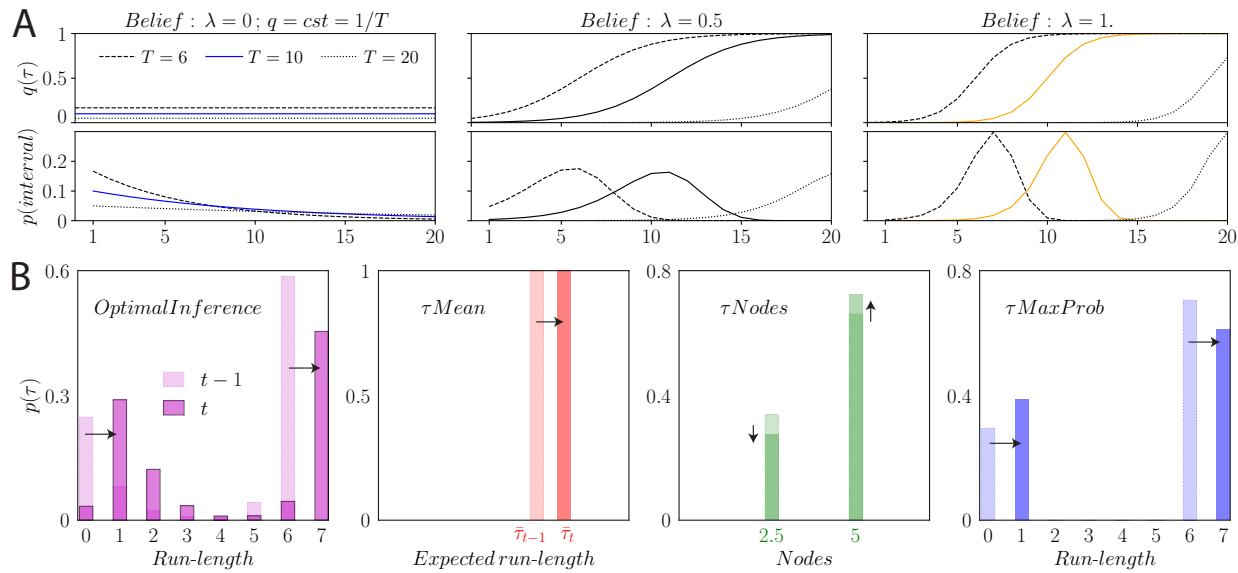


Figure 6. Illustration of the erroneous belief in the *IncorrectQ* model and of the approximations made in the τMean , τNodes , and $\tau\text{MaxProb}$ models. **A. Change probability, $q(\tau)$, as a function of the run-length (first row), and distribution of intervals between two consecutive change-points (second row), for various beliefs on the parameters of the change probability: the slope, λ , and the average duration of intervals, T . For a vanishing slope ($\lambda = 0$), the change probability is constant and equal to $1/T$ (first panel). With $T = 10$ this corresponds to the HI condition (blue lines). For a positive slope ($\lambda > 0$), the change probability increases with the run-length (i.e., a change-point becomes more probable as the time since the last change-point increases), and the distribution of intervals between two successive change-points is peaked. The HD condition (orange lines) corresponds to $\lambda = 1$ and $T = 10$. **B.** Schematic illustration of the marginal distribution of the run-length, $p(\tau)$, in each model considered. The *OptimalInference* model assigns a probability to each possible value of the run-length, τ , and optimally updates this distribution upon receiving stimuli (first panel). The τMean model uses a single run-length which tracks the inferred expected value, $\bar{\tau}_t$ (second panel). The τNodes model holds in memory a limited number, N_τ , of fixed hypotheses on τ (“nodes”), and updates a probability distribution over these nodes; $N_\tau = 2$ in this example (third panel). The $\tau\text{MaxProb}$ model reduces the marginal distribution by discarding less likely run-lengths; in this example, 2 run-lengths are stored in memory at any given time (fourth panel).**

of this model, adapted to the parameters of our task. We call it the τ *Mean* model. While the optimal marginal distribution of the run-lengths, $p_t(\tau|x_{1:t})$, spans the whole range of possible values of the run-length, it is approximated, in the τ *Mean* model, by a delta function parameterized by a single value, which we call the ‘approximate expected run-length’ and which we denote by $\bar{\tau}_t$. Upon the observation of a new stimulus, x_{t+1} , the updated approximate expected run-length, $\bar{\tau}_{t+1}$, is computed as a weighted average between two values of the run-length, 0 and $\bar{\tau}_t + 1$, which correspond to the two possible scenarios: with and without a change point at trial $t + 1$. Each scenario is weighted according to the probability of a change point at trial $t + 1$, given the stimulus, x_{t+1} . This model has no free parameter (Fig. 6B, second panel).

In a second limited-memory model, contrary to the τ *Mean* model just presented, the support of the distribution of the run-lengths is not confined to a single value. This model generalizes the one introduced by Wilson et al. (2013). In this model, the marginal distribution of the run-lengths, $p_t(\tau|x_{1:t})$, is approximated by another discrete distribution defined over a limited set of constant values, called ‘nodes’ (Fig. 6B, third panel). We call this model τ *Nodes*. A difference with the previous model (τ *Mean*) is that the support of the distribution is fixed, i.e., the set of nodes remains constant as time unfolds, whereas in the τ *Mean* model the single point of support, $\bar{\tau}_t$, depends on the stimuli received. The details of the implementation of this algorithm, and, in particular, of how the approximate marginal distribution of the run-lengths is updated upon receiving a new stimulus, are provided in *Methods*. The model is parameterized by the number of nodes, N_τ , and the values of the nodes. We implement it with up to five nodes.

The two models just presented are drawn from the literature. We propose a third suboptimal model that relieves the memory load in the inference process. We also approximate, in this model, the marginal distribution of the run-lengths, $p_t(\tau|x_{1:t})$, by another, discrete distribution. We call N_τ the size of the support of our approximate distribution, i.e., the number of values of the run-length at which the approximate distribution does not vanish. A simple way to approximate $p_t(\tau|x_{1:t})$ is to identify the N_τ most likely run-lengths, and set the probabilities of the other run-lengths to zero. More precisely, if, at trial t , the run-length takes a given value, τ_t , then, upon the observation of a new stimulus, at trial $t + 1$ it can only take one of two values: 0 (if there is a change point) or $\tau_t + 1$ (if there is no change point). Hence, if the approximate marginal distribution of the run-lengths at trial t is non-vanishing for N_τ values, then the updated distribution is non-vanishing for $N_\tau + 1$ values. We approximate this latter distribution by identifying the most unlikely run-length, $\arg \min p_{t+1}(\tau|x_{1:t+1})$, and setting its probability to zero. In other words, at each step, the N_τ most likely run-lengths are retained while the least likely run-length is eliminated. We call this algorithm τ *MaxProb* (Fig. 6B, fourth panel). It is parameterized by the size of the support of the marginal distribution, N_τ , which can be understood as the number of ‘memory slots’ in the model.

A model with limited run-length memory through sampling-based inference The five models considered hitherto (*OptimalInference*, *IncorrectQ*, τ *Mean*, τ *Nodes*, and τ *MaxProb*) are deterministic: a given sequence of stimuli implies a given sequence of responses, in marked contrast with the variability exhibited in the responses of subjects. To account for this experimental

observation, we suggest several models in which stochasticity is introduced in the generation of a response. Response stochasticity can stem from the inference step, the response-selection step, or both. We present, first, a model with stochasticity in the inference step.

This model, which we call τ *Sample*, is a stochastic version of the τ *MaxProb* model: instead of retaining deterministically the N_τ most likely run-lengths at each trial, the τ *Sample* model samples N_τ run-lengths using the marginal distribution of the run-lengths, $p_t(\tau|x_{1:t})$. More precisely, if at trial $t + 1$ the marginal distribution of the run-lengths, $p_{t+1}(\tau|x_{1:t+1})$, is non-vanishing for $N_\tau + 1$ values, then a run-length is sampled from the distribution $[1 - p_{t+1}(\tau|x_{1:t+1})] / z_{t+1}$, where z_{t+1} is a normalization factor, and the probability of this run-length is set to zero (Fig. 7). In other words, while the τ *MaxProb* model eliminates the least likely run-length deterministically, the τ *Sample* model eliminates one run-length stochastically, in such a fashion that less probable run-lengths are more likely to be eliminated. The τ *Sample* model has one parameter, N_τ , the size of the support of the marginal distribution of the run-lengths.

Stochastic inference model with sampling in time *and* in state space: the particle filter

Although the τ *Mean*, τ *Nodes*, τ *MaxProb*, and τ *Sample* models introduced above relieve the memory load by prescribing a form of truncation on the set of run-lengths, inference in these models is still executed on a continuous state space labeled by s (or, more precisely, on a discrete space with resolution as fine as a pixel). Much as subjects may retain only a compressed representation of probabilities along the τ axis, it is conceivable that they may not maintain a full probability function over the 1089-pixel-wide s axis, as they carry out the behavioral task. Instead, they may infer using a coarser spatial representation, in order to reduce their memory and computational loads. Monte Carlo algorithms perform such approximations by way of randomly sampling the spatial distribution; sequential Monte Carlo methods, or ‘particle filters’, were developed in the 1990s to address Hidden Markov Models, a class of hidden-state problems within which falls our inference task (Gordon et al., 1993; Arulampalam et al., 2002; Doucet and Johansen, 2008). Particle filters approximate a distribution by a weighted sum of delta functions. In our case, a *particle* i at trial t is a triplet, (s_t^i, τ_t^i, w_t^i) , composed of a state, a run-length, and a weight; a particle filter with N_P particles approximates the posterior, $p_t(s, \tau|x_{1:t})$, by the distribution

$$\tilde{p}_t(s, \tau|x_{1:t}) = \sum_{i=1}^{N_P} w_t^i \delta(s - s_t^i) \delta_{\tau, \tau_t^i}, \quad (3)$$

where $\delta(s - s_t^i)$ is a Dirac delta function, and δ_{τ, τ_t^i} a Kronecker delta. In other words, a distribution over the (s, τ) space is replaced by a (possibly small) number of points, or samples, in that space, along with their probability weights.

To obtain the approximate posterior at trial $t + 1$ upon the observation of a new stimulus, x_{t+1} , we note, first, that the Bayesian update (Eq. (2)) of the approximate posterior, $\tilde{p}_t(s, \tau|x_{1:t})$, is a mixture (a weighted sum) of the N_P Bayesian updates of each single particle (i.e., Eq. (2) with the prior, $p_t(s, \tau|x_{1:t})$, replaced, for each particle, by $\delta(s - s_t^i) \delta_{\tau, \tau_t^i}$). Then, we sample independently each component of the mixture (i.e., each Bayesian update of a particle), to obtain stochastically

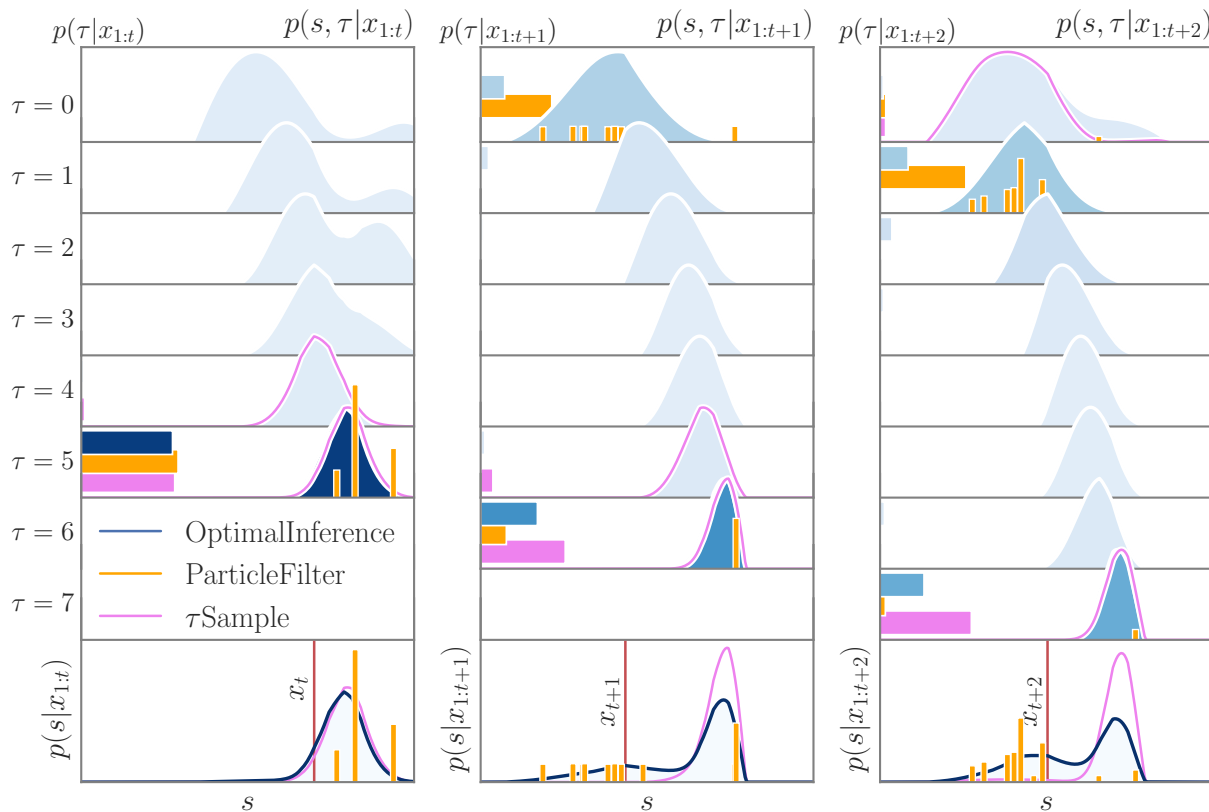


Figure 7. Illustration of the posterior density over three successive trials for the *OptimalInference* model, the τ *Sample* model with $N_\tau = 2$, and the *ParticleFilter* model with ten particles. The three panels correspond to the three successive trials. Each row except the last one corresponds to a different run-length, τ . In these rows, the *horizontal* bars show the marginal probability of the run-length, $p(\tau|x_{1:t})$. The posterior (i.e., the joint distribution of the run-length and the state, $p(s, \tau|x_{1:t})$) is shown as a function of the state, s , for the *OptimalInference* model (blue shaded curve), the τ *Sample* model (pink line), and the *ParticleFilter* model (orange vertical bars). The marginal probability of the run-length, $p(\tau|x_{1:t})$, for the *OptimalInference* model, is additionally reflected in the hue of the curve (darker means higher probability). For the *ParticleFilter* model, the heights of the bars are proportional to the weights of the particles. When the state, s , of two or more particles coincide, a single bar is shown with a height proportional to the sum of the weights. The last row shows the marginal distributions of the states, $p(s|x_{1:t}) = \sum_\tau p(s, \tau|x_{1:t})$, along with the location of the stimulus at each trial (red vertical line). At trial t (left panel), the probability of the run-length $\tau = 5$ dominates in the three models. In the τ *Sample* model, it vanishes at run-lengths from 0 to 3, and it is very small for $\tau = 4$. In the *ParticleFilter* model, the run-lengths of the ten particles are all 5, and thus the probability of all other run-lengths is zero. At trial $t + 1$ (middle panel), upon observation of the new stimulus, x_{t+1} , the marginal probability of the vanishing run-length ($\tau = 0$), which corresponds to a ‘change-point’ scenario, becomes appreciable in the *OptimalInference* model (top row). The probability of the run-length $\tau = 6$ (a ‘no change-point scenario’) is however higher. As a result, a ‘bump’ appears in the marginal distribution of the state, around the new stimulus (bottom row). In the τ *Sample* model, the optimal update of the posterior results in a non-vanishing probability for three run-lengths ($\tau = 0, 5$, and 6), more than the number of ‘memory slots’ available ($N_\tau = 2$). One run-length is thus randomly chosen, and its marginal probability is set to zero; in the particular instantiation of the model presented here, the run-length $\tau = 0$ is chosen, and thus the resulting marginal probability of run-length is non-vanishing for $\tau = 5$ and 6 only. In the *ParticleFilter* model, the stochastic update of the particles results in seven particles adopting a vanishing run-length, and the probability of a ‘change-point’ scenario ($\tau = 0$) becomes higher than that of the ‘no change-point’ scenario ($\tau = 6$) supported by the remaining three particles. The various marginal distributions of the states obtained in these three models (bottom row) illustrate how the τ *Sample* model and the *ParticleFilter* model approximate the optimal posterior: the τ *Sample* model assigns a negligible probability to a set of states whose probability is substantial under the *OptimalInference* model, while the *ParticleFilter* yields a coarse approximation reduced to a support of ten states (as opposed to a continuous distribution).

the updated particles, $(s_{t+1}^i, \tau_{t+1}^i)$, and to each particle is assigned the weight of the corresponding component in the mixture. The details of the procedure just sketched, in particular the derivation of the mixture and of its weights, and how we handle the difficulties arising in practical applications of the particle filter algorithm, can be found in **Methods**. This model, which we call **ParticleFilter**, has a single free parameter: the number of particles, N_P (Fig. 7).

Models with variability originating in the response-selection step The τ *Sample* and *ParticleFilter* models presented above reduce the dimensionality of the inference problem by pruning stochastically the posterior, in the inference step. But, as we pointed out, the behavior of the standard deviation of the responses of the subjects, as compared to that of the width of the Bayesian posterior (Fig. 5C), hints at a more straightforward mechanism at the origin of response variability. The model we now introduce features stochasticity not in the inference step, but rather in the response-selection step of an otherwise optimal model. In this model, the response is sampled from the marginal posterior on the states, $p_t(s|x_{1:t})$, i.e., the response, \hat{s}_t , is a random variable whose density is the posterior. This contrasts with the optimal response-selection strategy, which maximizes the expected score based on the Bayesian posterior, and which was implemented in all the models presented above. Henceforth, we denote the optimal, deterministic response-selection strategy by **Max**, and the suboptimal, stochastic strategy just introduced by **Sampling**. It has no free parameter.

Another source of variability in the response-selection step might originate in a limited motor precision, in the execution of the task. To model this motor variability, in some implementations of our models we include a fixed, additive, Gaussian noise, parameterized by its standard deviation, σ_m , to obtain the final estimate. Both this motor noise and the *Sampling* strategy entail stochasticity in response selection. The former, however, has a fixed variance, σ_m^2 , while the variance of the latter depends on the posterior which varies over the course of inference (Fig. 5C). When we include motor noise in the *Max* or in the *Sampling* strategies, we refer to these as **NoisyMax** and **NoisySampling**, respectively.

In sum, we have described four response-selection strategies (*Max*, *Sampling*, *NoisyMax*, and *NoisySampling*), and seven inference strategies, of which five are deterministic (*OptimalInference*, *IncorrectQ*, τ *Mean*, τ *Nodes*, and τ *MaxProb*) and two are stochastic (τ *Sample* and *ParticleFilter*). We can combine any inference strategy with any response-selection strategy: thus, we have at hand $4 \times 7 = 28$ different models, 27 of which are suboptimal, obtained from pairings of the inference and selection strategies. We label each of the 28 models by the combination of the two names referring to the two steps in the process, e.g., *ParticleFilter+Sampling*.

2.8 Fitting models to experimental data favors sample-based inference

The 27 suboptimal models introduced in the previous section yield a range of discrepancies from the optimal behavior. The ways in which each deviation from the optimal model impacts behavior is examined in **Methods**; here, we ask how well these models account for the behavior of human subjects. Whereas the optimal model, *OptimalInference+Max*, computes the Bayesian posterior

(*OptimalInference*) and selects the maximizing response (*Max*), the suboptimal models mimic cognitive limitations that may prevent the brain from reaching optimality: incorrect belief in the temporal structure of the signal (*IncorrectQ*), compressed representation of the Bayesian posterior, either deterministically (τ *Mean*, τ *Nodes*, and τ *MaxProb*) or stochastically (τ *Sample* and *ParticleFilter*), and noise introduced in the response-selection step, with a width either scaling with that of the posterior (*Sampling*), or constant (*NoisyMax*), or a combination of the two (*NoisySampling*).

To evaluate the ability of each of these models to account for human behavior, we compare quantitatively their respective outputs with the responses of human subjects. For the three quantities we examine (the learning rate, the repetition propensity, and the standard deviation of the responses), we compute the normalized mean squared error (NMSE) (sometimes referred to as the ‘Fraction of Variance Unexplained’ in the context of linear regressions). It is defined, for a given model and for each quantity, as the ratio of the mean squared error in the model output as compared to data, and the variance of the quantity under scrutiny in the behavioral data. We fit each of our models to human data, using the average of the three NMSEs as our error measure.

We find that the five best-fitting models make use of stochastic compression in the inference step, in either the τ *Sample* approximation or the *ParticleFilter* approximation (Table 1). These models all reproduce the qualitative trends in the behavior of subjects with respect to our three measures: for the learning rate and the standard deviation, the ‘smile shape’ of the HD curve, which crosses a decreasing HI curve; for the repetition propensity, conversely, an inverted U shape of the HD curve which crosses an increasing HI curve (Fig. 8, results from the τ *Sample+Max* and *ParticleFilter+Sampling* models are not shown, but the corresponding curves are similar).

The τ *Sample* and *ParticleFilter* strategies have one or two parameters, depending on whether they include motor noise or not. Other models, including all models with a deterministic inference step, have an error at least 30% higher than the best five models (and 2.45 times higher than the best model), despite the fact that other strategies come with up to five parameters (Table 1). The best-fitting model is *ParticleFilter+NoisyMax* with $N_P = 9$ particles. The fitted standard deviation, σ_m , of the Gaussian motor noise is approximately equal to 0.77 pixels; as a consequence, in about half of the trials, the noise component is within the width of a pixel, and thus has no impact. The second best model also follows a *ParticleFilter* inference strategy, with $N_P = 14$, combined with a *NoisySampling* response selection (with $\sigma_m = 0.70$ pixels).

The third and fourth best-fitting models use the τ *Sample* inference strategy, with $N_\tau = 1$, and the *NoisyMax* (with $\sigma_m = 0.45$ pixels, for the third one) and the *Max* (for the fourth one) selection strategies. At any given trial, these two models retain only a single assumption, τ_t , on the run-length. Upon receiving a new stimulus, x_{t+1} , a model subject computes $p_{change} = p_{t+1}(\tau = 0|x_{1:t+1})$ and $1 - p_{change} = p_{t+1}(\tau = \tau_t + 1|x_{1:t+1})$, and decides whether there was a change-point by sampling this simple, Bernoulli distribution. This sampling process, over a marginalization of the posterior, is similar to that in the particle filter model, which samples over the full (s, τ) -dependent posterior. As a consequence of sampling, the τ *Sample* strategy also exhibits variability, which behaves in a fashion similar to the variability in the *ParticleFilter* strategy (Fig. 8, bottom right). As for response selection, we note that with the *Sampling* and *NoisySampling* selection strategies (instead of the *Max* and *NoisyMax* strategies), these models do not perform as well, and

| | | Deterministic | | | | | Stochastic | |
|----------------------|-----------------|---------------|------------|-------------|--------------|----------------|---------------|----------------|
| Inference strategy → | | OptimalInf. | IncorrectQ | τ Mean | τ Nodes | τ MaxProb | τ Sample | ParticleFilter |
| Selection strategy ↓ | | | | | | | | |
| Det. | Max | 22.2 | 21.8 | 22 | 22.1 | 21.9 | 1.58 | 5.76 |
| | Noisy Max | 3.25 | 2.45 | 3.33 | 2.9 | 2.79 | 1.5 | 1 |
| Stochastic | Fixed width | | | | | | | |
| | Posterior-based | | | | | | | |
| | Sampling | 4.4 | 3.14 | 4.28 | 3.82 | 3.78 | 2.96 | 1.89 |
| | Noisy Sampling | 4.32 | 2.51 | 4.32 | 3.79 | 2.77 | 2.95 | 1.46 |

Table 1. Results of model fitting. Ratios of the normalized mean squared error (NMSE) in each model to that of the best-fitting model, *ParticleFilter+NoisyMax*. Each model is a combination of an inference strategy (columns) with a response-selection strategy (rows). The second best model, also a *ParticleFilter* but with a *NoisySampling* response-selection strategy, yields an NMSE 46% higher than the best-fitting model.

result in errors larger by 86% (*Sampling* vs. *Max*) and 96% (*NoisySampling* vs. *NoisyMax*). In fact, for all the seven inference models, the *NoisyMax* response-selection strategy results in errors lower or equal (but more often, lower) than the other three selection strategies (*Max*, *Sampling* and *NoisySampling*) (Table 1). This suggests that the variability in human responses does not originate from a posterior-sampling strategy in the response-selection step, but, rather, from an intrinsically stochastic inference process. In order to seek further validation of this finding, we explore, below, a generalization of the *Sampling* strategy.

Robustness of the results To substantiate the picture that emerges from the results summarized above, we perform two supplementary analyses. First, we investigate whether a generalized *Sampling* strategy yields smaller errors than the *NoisyMax* strategy. Second, we consider our choice of fitting-performance measure (the average of the NMSEs in the three quantities we examine), and we check for the robustness of model fitting to changes in the relative weights of each quantity in the fitting performance measure.

Sampling from the posterior function is only one of many possible sampling strategies for re-

sponse selection. Furthermore, in practice sampling may be difficult to tease apart from maximizing a perturbed posterior function. Acerbi et al. (2014) argue that, for some forms of random perturbations of a posterior probability density, maximizing the randomly perturbed function yields similar results to sampling from a modified posterior density function obtained as a power of the correct posterior: $p_\kappa(s) \propto p(s|x_{1:t})^\kappa$. To establish the equivalence, the exponent, κ , is chosen as inversely related to the magnitude of the perturbing noise. Sampling from the modified posterior yields a behavior that interpolates between posterior sampling (for $\kappa = 1$) and maximizing (for $\kappa \rightarrow \infty$); it yields a family of softmax operations over the posterior (Vul, 2011; Yu and Huang, 2014). Another interpretation of this sampling strategy is proposed by Battaglia et al. (2011): in the case of an integer κ and a Gaussian posterior, the mean of κ samples drawn from the posterior is a Gaussian random variable, with a standard deviation equal to that of the posterior scaled by $1/\sqrt{\kappa}$; i.e., a distribution equal to the posterior raised to the power κ , and normalized. Hence, sampling from the exponentiated posterior can be interpreted as drawing κ samples from the unexponentiated posterior, and taking the mean.

We implement this strategy of response selection by sampling a modified posterior, which we denote κ *Sampling*. We find that it performs better than the *Sampling* strategy, as expected since the *Sampling* strategy is a special case of the κ *Sampling* (with the parameter, κ , set to unity). However, in the case of all seven inference models, the κ *Sampling* strategy, which has one free parameter, performs worse than the *NoisyMax* strategy, which has, also, a single parameter (Fig. 9B). Hence, a random, additive perturbation of the maximization strategy remains a better account of human behavior than a posterior-sampling strategy.

Our results, which suggest that the variability in the responses of subjects originate in the inference step rather than in the response-selection step, rely upon the fitting performance measure used for model comparison. We chose a measure that weighted equally the three NMSEs (on the learning rate, repetition propensity, and standard deviation), so as to obtain a model performing well on all fronts, but that choice was arbitrary. Hence, one may be concerned, for instance, that the goodness-of-fit of the *ParticleFilter* model be due to our choice of weighing errors. As a control, we computed the three ‘two-measure errors’, each excluding one of the three measures and averaging the errors in the two remaining ones. We found that, regardless of the choice of the combination, the relative order of the models in terms of performance stays identical, with only a few exceptions. Most importantly, the *ParticleFilter* remains, in all three cases of error combinations, the best-fitting model (Fig. 9B).

Taken together, our results suggest that variability in human behavior, at least in the context of our task, is dictated primarily by stochasticity in the inference step — i.e., in the manipulation and update of probabilities — rather than by ‘output noise’ such as stochasticity in the response-selection step or motor noise. This view agrees with the conclusion of a recent study of a cue combination task (Drugowitsch et al., 2016); its authors argue that a “dominant fraction” of human choice suboptimality arises from random fluctuations in the inference step.

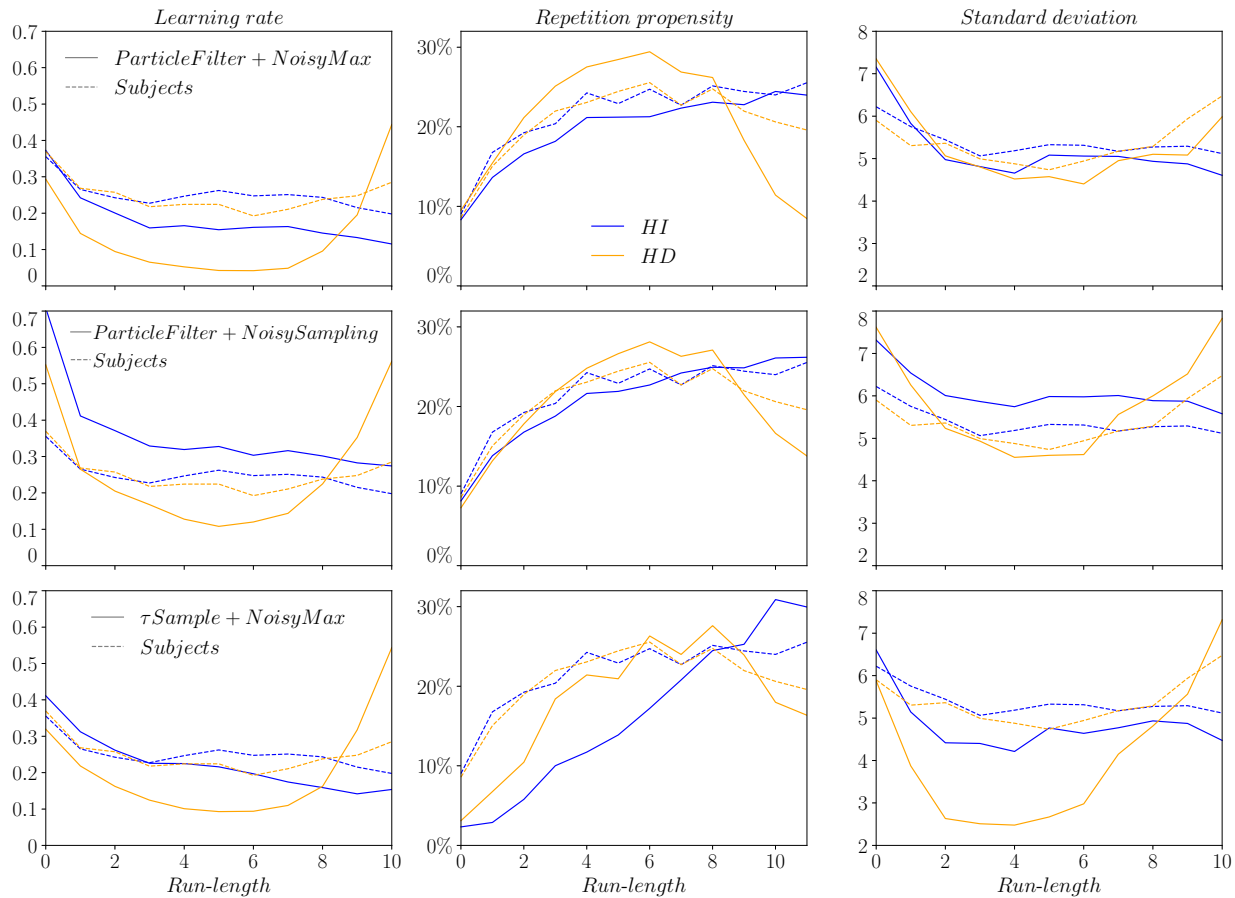


Figure 8. Behaviors of the three best-fitting models. In HI (blue curves) and HD (yellow curves) conditions, average learning rate (first column), repetition propensity (second column), and standard deviation of responses (third column), as a function of run-length, for the subjects (solid lines) and the three best-fitting models (dashed lines): *ParticleFilter+NoisyMax* (first row), *ParticleFilter+NoisySampling* (second row), and *τ Sample+NoisyMax* (third row).

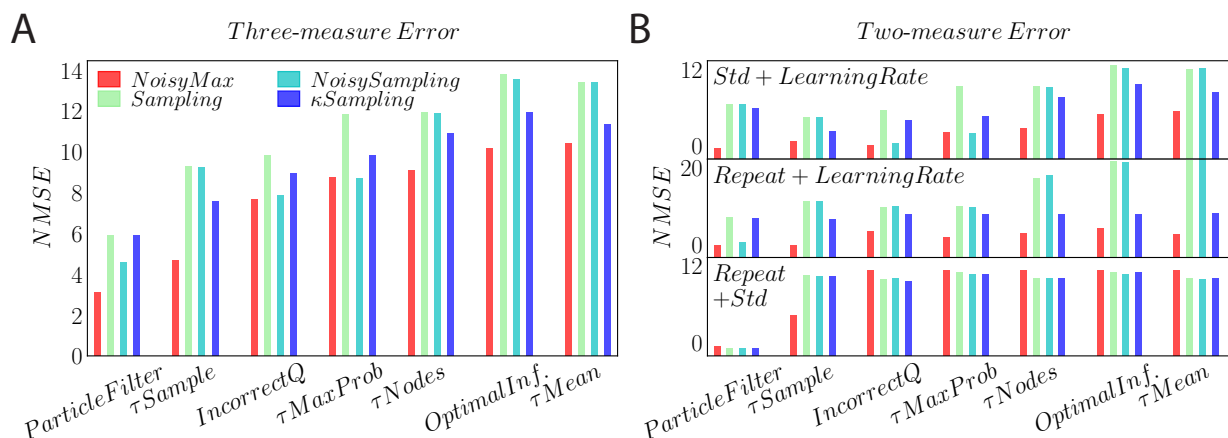


Figure 9. Normalized Mean Squared Error of fitted models. **A.** NMSE of models fitted to subjects data, averaged over the three measures (learning rate, repetition propensity, and standard deviation of responses), grouped by inference models. **B.** NMSE between fitted models and subjects data, averaged over two out of the three measures.

3 Discussion

3.1 Summary

This study investigates the behavior of human subjects in an online inference task, and examines mechanisms that can account for behavioral trends found in experimental data. An important aspect of this task is that it makes use of both a history-independent (HI) condition with no temporal structure, and a history-dependent (HD) condition in which a hidden state is almost periodical and, hence, highly structured in time (Fig. 1D). We find that subjects display different behaviors in the two conditions, adapting their learning rate to the temporal structure of the hidden state. We also note a propensity in subjects to repeat their response in consecutive trials; this repetition propensity increases with the run-length, and in the HD condition drops again for larger run-lengths. Moreover, we observe that subjects exhibit a greater variability in their responses shortly after a change point, in both conditions, and at long run-lengths in the HD condition, i.e., the *variability* in behavior also depends on the temporal statistics of the stimuli.

The distinctive behaviors of the learning rate and the repetition propensity in the HI and HD conditions are reproduced qualitatively by a Bayesian model of inference which yields optimal updates of the probability density of the hidden state. As for the variability in subjects responses, we find that its behavior is similar to that of the standard deviation of the Bayesian posterior. We therefore use the Bayesian model as a starting point to elaborate variant models which can account for the trends exhibited in human responses. We find that the variability in human behavior, and its dynamics, can be reproduced by supoptimal models in which inference is executed in a stochastic manner. Specifically, the τ Sample and the ParticleFilter models alter the optimal inference step

by maintaining an approximate version of the posterior, by means of random sampling. This alteration of the optimal model at once introduces variability in the behavior and relieves the memory capacity, through sampling either in the ‘time dimension’ (in the τ Sample model) or in the ‘time and space dimensions’ (in the *ParticleFilter* model).

3.2 Temporal structures in nature and their behavioral and neural counterparts

In order to make appropriate decisions in relation to their environment, humans and animals must infer the state of the surrounding world on the basis of the sensory signals they receive. If these signals are noisy and if the environment is changing, their inference task is complicated by the fact that a new stimulus may reflect either noise or a change in the underlying state. However, if events in the world present some kind of temporal structure, such as in our HD signal, it is possible to use this structure to refine one’s inference. Conversely, if events follow a Poisson process, as in the HI signal, their occurrences present no particular temporal structure, and what just happened conveys no information on what is likely to happen next. Hence, there is a fundamental difference between the HI and HD conditions, which impacts the inference of an optimal observer.

Many natural events are not Poisson-distributed in time, and exhibit strong regularities. Nunes Amaral et al. (2004), Nakamura et al. (2007, 2008), and Anteneodo and Chialvo (2009) have recorded the motor activity of both rodents and human subjects over the course of several days. In both species, they found that the time intervals between motion events were distributed as a power law, a distribution characterized by a long tail, leading to bursts, or clusters, of events followed by long waiting epochs. The durations of motion episodes also exhibited heavy tails. These kinds of distribution are incompatible with Poisson processes, which yield exponentially distributed inter-event epochs. Moreover, both rodent and human activity exhibited long-range autocorrelations, another feature that cannot be explained by a Poisson process. A particular form of autocorrelation is periodicity, which occurs in a wide range of phenomena. In the context of human motor behavior, walking is a highly rhythmical natural activity (Hausdorff et al., 1995; Griffin et al., 2000). More complex patterns exist (neither clustered nor periodic), such as in human speech which presents a variety of temporal structures, whether at the level of syllables, stresses, or pauses (Ramus et al., 1999; Low et al., 2000; Campione and Véronis, 2002). In all these examples, natural mechanisms produce series of temporally structured events. The ubiquity of history-dependent statistics of events in nature begs for explorations of inference mechanisms in their presence.

In the case of studies of perception and decision-making, in both humans and animals, history-dependent signals have been used widely. In a number of experiments (Janssen and Shadlen, 2005; Ghose and Maunsell, 2002; Li and Dudman, 2013; Miyazaki et al., 2005; Jazayeri and Shadlen, 2010), a first event (a sensory cue, or a motor action such as a lever press) is followed by a second event, such as the delivery of a reward, or a ‘go’ signal triggering the next behavior. The time elapsed between these two events – the ‘reward delay’ or the ‘waiting time’ – is randomized and sampled from distributions that, depending on the studies, vary in mean, variance, or shape. For instance, both Janssen and Shadlen (2005) and Ghose and Maunsell (2002) use unimodal and

bimodal temporal distributions. Because of the stochasticity of the waiting time, the probability of occurrence of the second event varies with time, similarly to the probability of a change point in our HD condition; these studies explore whether variations of this probability are used by human and animal subjects. Janssen and Shadlen (2005)'s recordings from the V4 cortical area in rhesus monkey indicate that, for both unimodal and bimodal waiting times distributions, the attentional modulation of sensory neurons varies consistently with the event probability. Ghose and Maunsell (2002) note that the reaction times of macaques are inversely related to the event probability, for both unimodal and bimodal distributions, and that the activity of neurons in the lateral intraparietal (LIP) area is correlated with the evolution of this probability over time. Li and Dudman (2013) manipulate another aspect of the distribution of reward delays: between blocks of trials, the standard deviation of this distribution is varied, while the mean is left unchanged. Mice, in this situation, are shown to adapt their waiting times to this variability of reward delays, consistently with a probabilistic inference model of reward timing.

Akin to the tasks just outlined are 'ready-set-go time-reproduction tasks', in which subjects are asked to estimate the random time interval between 'ready' and 'set' cues, and to reproduce it immediately afterwards. Miyazaki et al. (2005) and Jazayeri and Shadlen (2010) show that human subjects combine optimally the cue (consisting in the perceived ready-set interval) with their prior on the interval length. Different priors are learned in training runs: they differ by the variances of the interval distributions (Miyazaki et al. (2005)) or by their means (Jazayeri and Shadlen (2010)). In both cases, subjects integrate the prior in a fashion consistent with Bayesian inference. Adopting a different approach, ten Oever et al. (2014) show that attentional resources can be dynamically allocated to points in time at which input is expected: when asked to detect auditory stimuli (beeps) of low intensity embedded in a continuous white noise, human subjects perform better when detecting periodic beeps rather than random beeps, suggesting that they are able to identify the temporal regularity and use it in their detection process.

In all these studies, the event of interest has a probability of occurrence that varies with time. The resulting temporal structure in the signal appears to be captured by human and animal subjects, and reflected in behavior and in its neural correlate. Various probability distributions used in the reported tasks can be compared directly to our HD sigmoid-shaped change probability, with adjusted parameters. In line with these studies, our results confirm that human subjects adapt their behavior depending on the temporal structure of stimuli. Additionally, we provide a comparison between two different conditions, a HD condition akin to a 'jittered periodic' process, and the Poisson, HI condition; the latter produces a memoryless process. Importantly, it plays the role of a benchmark from the point of view of probability theory: in discrete time it yields a geometric distribution, and in continuous time it yields an exponential distribution; both distributions maximize the entropy, subject to the constraint of a fixed event rate. In this study, we compared a specific, temporally structured HD condition to this benchmark, HI condition.

3.3 Online Bayesian inference

Our first observation is that the average learning rate of subjects and their repetition propensity are captured by a Bayesian model. The Bayesian paradigm has been viewed as an extension of logic that enables reasoning with propositions whose truth or falsity is uncertain (Cox, 1946; Jaynes, 2003). In cognitive science, it has successfully accounted for a wide range of observations, including cue combination in humans (Jacobs, 1999; van Beers et al., 1999; Ernst and Banks, 2002; Knill, 2003; Battaglia et al., 2003; Hillis et al., 2004; Knill, 2007; Battaglia et al., 2011), sensorimotor control (Körding and Wolpert, 2004, 2006; Berniker and Kording, 2011), integration of temporal statistics (Ghose and Maunsell, 2002; Janssen and Shadlen, 2005; Miyazaki et al., 2005; Jazayeri and Shadlen, 2010; Li and Dudman, 2013), perceptual multistability (Sundareswara and Schrater, 2008; Moreno-Bote et al., 2011; Gershman et al., 2012), and various aspects of cognition (Griffiths and Tenenbaum, 2006; Blaisdell et al., 2006; Stocker and Simoncelli, 2008; Goodman et al., 2008; Griffiths and Tenenbaum, 2011).

The literature on Bayesian *online* inference in cognition, where belief is updated iteratively as a function of incoming information, is scarcer, but examples can be found in word segmentation (Pearl et al., 2011), sentence processing (Levy et al., 2008), conditioning (Daw and Courville, 2008), as well as in the change-point literature (Nassar et al., 2010; Wilson et al., 2013, 2010). In change-point tasks, subjects are presented with a long sequence of consecutive inference problems (1000 of them, in our case). Each trial is a slightly different task, in which one has to handle the uncertainty resulting from the belief distribution, from the signal likelihood, and from the possibility of a change point. The latter, in the HD condition, bears the added complexity of a change-point probability, $q(\tau)$, that depends on the time of the last change point. How these uncertainties are handled determines the behavior, and in particular to what extent an observer reacts to a new stimulus: either shift the estimate towards it, or not move at all. We quantify this response through the learning rate and the repetition propensity, and we find that the ideal Bayesian observer and the subjects obey similar trends (Fig. 5A-B).

The success of the Bayesian paradigm, however, is limited, and comes with three shortcomings. First, subjects do not behave quantitatively like the ideal Bayesian observer, and hence there remains unexplained suboptimality. Second, we find variability in the responses of subjects (Fig. 5C), an observation incompatible with optimal Bayesian inference. Third, inference problems in the real world are complex and high-dimensional, rendering Bayesian reasoning computationally heavy and memory-intensive. This suggests that humans use approximations when carrying out inference and estimation tasks (Gershman et al., 2015; Gershman and Beck, 2016; Sanborn and Chater, 2016). These three observations call for the investigation of alternatives to the optimal Bayesian paradigm. Our study explores several scenarios.

3.4 Sampling versus noisy maximization

Although the behavior of the subjects and of the Bayesian model differ in that the former exhibits variability while the latter is deterministic, the temporal modulation of the human variability

follows a similar course to that of the standard deviation of the Bayesian posterior (Fig. 5C). Therefore, it is natural to propose that response selection operates by sampling the Bayesian posterior instead of maximizing the expected score. Decision by posterior sampling, or ‘probability matching’, has been suggested by other decision-making experiments (Herrnstein, 1961; Koehler and James, 2009; Denison et al., 2013) and, more recently, perceptual experiments (Wozny et al., 2010; Battaglia et al., 2011; Moreno-Bote et al., 2011). Although close to optimal in some specific paradigms (Kaufmann et al., 2012), sampling is suboptimal in the context of our behavioral task. When fitting models to human data, we observe that the *ParticleFilter+Sampling* and the *ParticleFilter+NoisySampling* models yield larger NMSE than the *ParticleFilter+NoisyMax* model (the best-fitting model), by 90% and 48%, respectively. More generally, we observe that for each of our seven inference strategies, the *NoisyMax* response-selection strategy results in better fits (lower NMSE) than the *Sampling* and the *NoisySampling* strategies (Fig. 9A). Relaxing the *Sampling* model by allowing the posterior to be exponentiated before sampling, as in the κ *Sampling* strategy, does not yield better fits than the *NoisyMax* strategy either. We conclude that posterior sampling accounts less successfully for our data than a simple perturbation of the optimal maximization strategy by an additive, fixed-width, Gaussian noise.

3.5 Stochastic pruning of the posterior during inference

After rejecting the sampling hypothesis for response selection, we are left with an unexplained modulation of the variability – specifically, the relation between the magnitude of behavioral variability and the width of the Bayesian posterior at successive times (Fig. 4). Noisy maximization, which makes use of an additive random perturbation with fixed variance, leads to behavioral variability with constant variance; it is thus insufficient to explain the experimental observations. If modulated variability does not originate in the response-selection step, it must derive from the inference step. Out of the 28 models we consider, the five best-fitting models implement either the *ParticleFilter* or the τ *Sample* inference strategy, both of which rely on sampling *during inference*. These strategies capture the trends in the variability of human responses, in both HI and HD conditions (Fig. 8).

Both these inference strategies reduce the computational load of the inference problem by stochastically trimming the posterior, in a fashion akin to the ‘pruning’ model proposed by Huys et al. (2012) in the context of a decision-tree task. In their decision model, the evaluation of a possible sequence of decisions is more likely to be curtailed (thus alleviating the dimensionality of the problem) if it appears to have a low value. Similarly, the *ParticleFilter* and the τ *Sample* inference strategies ignore with a higher probability possible run-lengths and states that are less likely to be correct. Furthermore, we note that the τ *MaxProb* inference strategy also relies on pruning unlikely run-lengths, but it deterministically eliminates the most unlikely, in contrast to its stochastic counterpart, τ *Sample* — which yields a better fit of the data. In explore-exploit problems, ‘Thompson sampling’ (Thompson, 1933) refers to a strategy in which one “explores” by randomly choosing an action with the probability that this action maximizes the reward (instead of deterministically choosing the action most likely to maximize the reward). In this perspective, the

stochastic pruning of the posterior in our best-fitting models appears as an exploration strategy, deployed during inference.

3.6 Inference through sample-based representations of probability

The *ParticleFilter* strategy is noteworthy in a number of respects. First, it is our best-fitting model. Second, it constitutes a generic approach to inference; it was reported to account successfully for other inference and learning behaviors, such as category learning (Sanborn et al., 2006, 2010), conditioning in pigeons (Daw and Courville, 2008), sentence processing (Levy et al., 2008), hidden state inference (Brown and Steyvers, 2009), and visual tracking of multiple objects (Vul et al., 2009). Third, out of all the models we consider, it is by far the less demanding on memory: with nine particles, one needs to store 27 numbers (for s , τ , and the weight of each particle) in memory. As a comparison, the optimal model stores a discretized probability distribution over the (s, τ) space, which amounts to about 16000 numbers (the optimal posterior could be well approximated with less memory-intensive methods, but this would require further hypotheses.) Previous uses of particle filter methods for cognitive tasks report best-fitting numbers of particles of the same order of magnitude: a mean of 56 in Brown and Steyvers (2009), around 20 in Levy et al. (2008), and as few as one particle in Daw and Courville (2008) and Sanborn et al. (2010).

A fourth aspect of particle filters is that they provide a natural interpretation of the high repetition propensity observed in subjects (Fig. 5B). As the support of the probability distribution is reduced to $N_P = 9$ points on the (s, τ) plane, there is a fair chance that the posterior-maximizing particle at trial t , (s_t, τ_t) , remains the posterior-maximizing particle at trial $t + 1$. Hence, the response s_t is likely to be repeated. In a similar spirit, particle filters have been shown to account for order effects in category learning (Sanborn et al., 2010) and observations about online sentence comprehension (such as the processing of ‘garden-path sentences’ (Levy et al., 2008)).

The success of particle filters, also known as Sequential Monte Carlo method, in accounting for human behavior in an online inference task adds to a growing literature on sample-based representations in cognitive processes (Goodman et al., 2008; Moreno-Bote et al., 2011; Gershman et al., 2012; Vul et al., 2014). Monte Carlo methods, which approximate probability distributions with sets of samples, constitute a major element of a family of techniques used in machine learning to address a wide range of problems (inference, optimization, numerical integration, etc); they have also been put forth as candidate cognitive algorithms (Sanborn, 2015; Gershman and Beck, 2016). Moreover, they account for a range of cognitive biases in the laboratory, such as base-rate neglect, conjunction fallacy, and the unpacking effect, as well as for human performance in complex, real-world tasks, and specific observations such as response variability and autocorrelation in perception and reasoning tasks (Gershman et al., 2012; Sanborn and Chater, 2016). At the implementation level, sample-based representations are well suited to learning in neural networks (Fiser et al., 2010). Here, the variability in neural activity can be interpreted in terms of sampling-based representations of probability (Hoyer and Hyvärinen, 2003; Fiser et al., 2010; Buesing et al., 2011; Gershman et al., 2012), and a number of neural network models performing probability sampling have been proposed (Shi and Griffiths, 2009; Moreno-Bote et al., 2011; Hennequin et al., 2014; Savin and Denève, 2014;

Aitchison and Lengyel, 2016).

4 Methods

4.1 Details of the behavioral task

The computer-based task was programmed and run with Psychopy (Peirce, 2009). In this task, white dots appeared on a horizontal line in the middle of a grey screen. Subjects were told that these white dots were snowballs thrown by a hidden person, the ‘enemy’ (also located on the horizontal line). The horizontal location of a snowball was the stimulus, x_t , and the position of the hidden person was the state, s_t . The state space was arbitrarily chosen to be $[0, 300]$; this scale did not appear on the screen. By clicking with a mouse (whose pointer moved on the horizontal axis only), subjects could indicate where they thought the hidden person was (i.e., give their estimate, \hat{s}_t , of the state). The time of response was not constrained. A green dot provided a visual feedback of the location of the click. After 100ms, a new white dot appeared, starting the next trial (Fig. 1A-B). If a subject’s ‘shot’ was within a fixed distance around the state (the radius of the enemy), the subject was rewarded with 1 point. If the shot was ‘outside the enemy’ but within a distance equal to twice the enemy radius, the reward was 0.25 point (Fig. 1E). Otherwise, the reward was zero. Subjects were not informed of the reward immediately after each shot, as this would have provided additional information on the location of the state. The total score was given every 100 trials, to allow for an assessment of average self-performance and to foster motivation.

4.2 Subjects

We ran the computer-based task on 30 paid subjects; all gave informed consent. The study was approved by Princeton University’s Institutional Review Board for Human Subjects. The sample size was determined so as to be comparable to that used in similar experiments (Nassar et al., 2010, 2012). Four subjects performed significantly worse than the other ones: their average error, defined as the absolute difference between their estimate and the state, $|\hat{s}_t - s_t|$, was 10.4 (standard deviation (s.d.): 0.93), while the average error of the other 26 subjects was 6.5 (s.d.: 0.62). Because of this difference of more than 5 standard deviations, these four subjects were excluded from the analyses. Hence, a total of 26 subjects were included in the analyses.

4.3 Details of the signal

The stimulus, x_t , was generated around the state, s_t , according to the likelihood probability, $g(x_t|s_t)$, which was chosen to be triangular, centered at s_t , and of half-width 20. The state, s_t , was piecewise-constant with respect to time, i.e., constant in the absence of a change point. In the HI condition, the probability of a change point, q_t , was constant and equal to 10%. In the HD condition, q_t depended on the run-length, τ_t , defined as the number of trials since the last change point, and had a sigmoid shape: $q_t(\tau_t) = 1/(1 + e^{-(\tau_t-10)})$. At $\tau_t = 0$ (i.e., immediately

after a change point), the probability of another change point was very small. Six trials after a change point it was still small, less than 2%, before growing appreciably (50% at $\tau_t = 10$, 95% at $\tau_t = 13$). This led to more regular intervals between change points than in the HI condition, with a change point roughly every 10 trials (Fig. 1C-D). The average number of trials between two change points in both conditions was 10. At a change point, the state randomly jumped to a new state, s_{t+1} , according to the state transition probability, $a(s_{t+1}|s_t)$. This distribution was chosen to be bimodal, symmetric, and centered at s_t (two triangles of half-width 20 each, centered at $s_t \pm d$, where $d = 25$, Fig. 1E). This prevented new states to be too close or too far from the previous state, which would have made change-point detection too difficult or too obvious.

4.4 Training runs

All subjects did the task in both HI and HD conditions. 14 started with the HD task and 16 started with the HI task (no significant difference were found in results between these two groups). Subjects were not told the specificity of each situation. An explanatory text indicated that there were “differences” between each condition but no further indications were given. Each condition started with a series of explanations and tutorial runs. In a first tutorial run, the enemy (i.e., the state) was visible and moved according to the current (HI or HD) condition, and successive snowballs appeared without any action from the user (as in passive video viewing). In a second run, the enemy was still visible and subjects had to click at each trial, after which the next snowball would appear. This run was a very simple version of the actual task, because subjects were seeing the state. In a third run, the half-width of the triangular likelihood, $g(x_t|s_t)$, was 10, i.e., half the value it took in the actual task. In this run, the state was not visible, except after a change point: in the occurrence of a change point, the position of the state before the change point was shown, along with the shots of the subject since the previous change point. This run had two goals: first, to emphasize the timing of change points, and, second, to allow for self-performance assessment and to illustrate that a strategy consisting in ‘following the white dots’, i.e., clicking on the stimuli, was inefficient. A fourth tutorial run was an ‘easy’ version of the actual task: the state was always hidden, but the likelihood, g , had a half-width of 10. A fifth and last tutorial run reproduced the third run, but with the likelihood, g , with half-width 20. During the task, 15 subjects (7 amongst the HD-first group and 8 amongst the HI-first group) were also shown the positions of past stimuli, as white dots with decreasing contrast, gradually merging with the grey background (Fig. 1). The other 15 subjects were not shown past stimuli. No significant differences were found in data between the two groups. The number of stimuli presented in the tutorial runs totaled 297 for each condition. During the actual task there were 1000 trials in each condition, leading to a total of 2000 data points per subjects.

4.5 Empirical run-length

As subjects did not know the true run-length, τ , we computed an empirical run-length, $\tilde{\tau}$, based on the responses of subjects. Whereas the true run-length is defined as the number of trials since

the last change point, the empirical run-length is defined as the number of trials since the last ‘large correction’; a large correction is defined as a correction with absolute value larger than the 90th percentile of corrections. This percentile level is chosen in relation to the average frequency of change points, 1 for every 10 trials, in both HI and HD conditions. In some occasions, a subject “misses” a change point: the run-length and the empirical run-length, consequently, differ. For instance, $\tilde{\tau} = 10$, while $\tau = 0$ or 1. In such a case, because the change point did occur, the subject experiences a large surprise and is thus likely to subsequently opt for a large correction, i.e., to increase the learning rate. In the HD condition, because of the temporal statistics of change points, this situation is more likely to occur at empirical run-lengths around 10. Hence, this effect could bias the learning rates to higher values at these empirical run-lengths, in this condition. This effect, however, does not originate in the inference process, but rather in the temporal statistics of the HD signal. In other words, even an observer whose inference algorithm is not adapted to the HD condition would have higher learning rates at empirical run-lengths around 10. In the results presented, we removed all trials with a (true) run-length of 0 or 1, in order to avoid this artifact.

4.6 Bayesian update equation

We derive the Bayesian update equation for a general case with $q = q(s_t, \tau_t)$, $p(x_t|s_t, \tau_t) = g(x_t|s_t, \tau_t)$, and $p(s_{t+1}|\tau_{t+1} = 0, s_t, \tau_t) = a(s_{t+1}|s_t, \tau_t)$, which includes the case used in our task with $q = q(\tau_t)$, $g = g(x_t|s_t)$, and $a = a(s_{t+1}|s_t)$.

Our goal is to obtain an update rule for the posterior, $p_t(s, \tau|x_{1:t})$ upon the observation of a new stimulus, x_{t+1} . Bayes’ rule yields

$$p_{t+1}(s, \tau|x_{1:t+1}) = \frac{1}{Z_{t+1}} g(x_{t+1}|s, \tau) p_{t+1}(s, \tau|x_{1:t}), \quad (4)$$

where $Z_{t+1} = p_{t+1}(x_{t+1}|x_{1:t})$ is a normalization constant. The third term in this product can be written as

$$p_{t+1}(s, \tau|x_{1:t}) = \sum_{\tau_t} \int_{s_t} p_{t+1}(s, \tau|s_t, \tau_t) p_t(s_t, \tau_t|x_{1:t}) ds_t. \quad (5)$$

The transition probability, $p_{t+1}(s, \tau|s_t, \tau_t)$, is determined by q and a . An absence of change point occurs with probability $1 - q(s_t, \tau_t)$, and in such a case a state (s_t, τ_t) evolves into the state $(s_{t+1} = s_t, \tau_{t+1} = \tau_t + 1)$. In the case of a change point, an event which occurs with probability $q(s_t, \tau_t)$, possible states at $t + 1$ have the form $(s_{t+1}, \tau_{t+1} = 0)$. Hence the transition probability from (s_t, τ_t) to (s, τ) at $t + 1$ is given by

$$p_{t+1}(s, \tau|s_t, \tau_t) = \mathbb{1}_{\tau=0} q(s_t, \tau_t) a(s_t, \tau_t, s) + \mathbb{1}_{\tau=\tau_t+1, s=s_t} (1 - q(s_t, \tau_t)). \quad (6)$$

Combining Equations (4), (5), and (6), we obtain the Bayesian update equation, as

$$p_{t+1}(s, \tau | x_{1:t+1}) = \frac{1}{Z_{t+1}} g(x_{t+1} | s, \tau) \left[\mathbb{1}_{\tau=0} \sum_{\tau_t} \int_{s_t} q(s_t, \tau_t) a(s_t, \tau_t, s) p(s_t, \tau_t | x_{1:t}) ds_t + \mathbb{1}_{\tau>0} (1 - q(s, \tau - 1)) p_t(s, \tau - 1 | x_{1:t}) \right]. \quad (7)$$

In the special case with $q = q(\tau_t)$, $a = a(s_{t+1} | s_t)$, and $g = g(x_t | s_t)$, we obtain the slightly simpler Eq. (2). In addition, we note that, in the HI condition, the change probability is constant, $q(\tau) = q$; in this condition, we can marginalize over the variable τ to obtain a closed recursion over the state posterior, as

$$p_{t+1}(s | x_{1:t+1}) = \frac{1}{Z_{t+1}} g(x_{t+1} | s) \left[q \int a(s | s_t) p_t(s_t | x_{1:t}) ds_t + (1 - q) p_t(s | x_{1:t}) \right]. \quad (8)$$

4.7 Derivation of the suboptimal models and analysis of their behaviors

4.7.1 *IncorrectQ* model

In the *IncorrectQ* model, the two quantities governing the shape of $q(\tau)$, λ and T , are treated as free parameters, and we explore how varying these parameters impacts behavior, as compared to the *OptimalInference* model.

Keeping T constant at 10, and varying λ from 0 (HI condition) to 1 (HD condition), we find that the learning rate as a function of the run-length gradually morphs from the HI, monotonically decreasing curve, to the HD, non-monotonic and ‘smile-shaped’ curve (Fig. 10B). A similar behavior obtains at any fixed value of T , with the difference that the minimum of the HD curve is shifted to smaller τ for smaller T , and to larger τ for larger T . In other words, for fixed T , a higher value of λ , i.e., a sharper slope of the change probability, leads to a higher learning rate at run-lengths comparable to T . (Note that, for $T = 20$, the minimum of the learning rate occurs at run-lengths larger than 10, hence the non-monotonicity is not apparent in Fig. 10B.) Conversely, for a fixed $\lambda > 0$, the minimum of the learning rate occurs at a run-length comparable to T which, precisely, determines when change points become likely. Finally, for $\lambda = 0$, the change probability is constant and there is no increase in the learning rates; these are, however, slightly higher for smaller T , because in that case the change probability, $q = 1/T$, is larger, so a new stimulus is more likely to be interpreted as stemming from a change point.

A subtlety that any analysis has to grapple with is that the statistics of responses depend not only on the inference process, but also, of course, on the statistics of the stimuli. To tease the two effects apart, for each *IncorrectQ* model subjects (with differing values of λ and of T) we computed the response behavior in presence of either signal: the HI signal, characterized by

a constant change probability, $q = 0.1$, and the HD signal, characterized by a change probability, $q(\tau)$, varying with the run-length as a sigmoid with parameters $\lambda = 1$ and $T = 10$. We note that the impact on behavior of changing the signal is modest, as compared to the impact of changing the model subject's beliefs (Fig. 10B). This indicates that the discrepancy in human behavior in the HI and HD conditions does not originate primarily from the statistics of the signals, but rather from the different beliefs on the temporal statistics of the signals, held by the subjects.

Paralleling the behavior of learning rates, the repetition propensity in the HD condition peaks earlier or later depending on the value of T , and its shallowness depend on the value of λ (Fig. 10C). A belief in a shorter average inter-change-point interval, T , leads to a smaller repetition propensity: assuming frequent change points enhances the frequency of changes in one's estimate.

Human subjects correctly believe that q is not constant in the HD condition, and they use this belief in their inference process, but they may hold an inexact representation of the shape of $q(\tau)$ (Fig. 10). This, however, is not sufficient to capture data quantitatively: subjects exhibit both higher learning rates and more frequent repetitions than in the optimal model (Fig. 5), an observation that cannot be explained by manipulating λ and T in the *IncorrectQ* model; in the latter, high learning rates are accompanied by lower repetition propensity, and vice versa. Thus, and letting alone the issue of variability, an erroneous belief on the change probability, $q(\tau)$, is insufficient to model experimental data.

4.7.2 τ Mean model

Derivation This model is a generalization of the model introduced in Ref. Nassar et al. (2010). The approximate joint probability of the state and the run-length, which we denote by $\tilde{p}_t(s, \tau|x_{1:t})$, is assumed, in this model, to vanish at all values of the run-length, except for one, which we call the 'approximate expected run-length' and which we denote by $\bar{\tau}_t$. Hence,

$$\tilde{p}_t(s, \tau|x_{1:t}) = \delta_{\tau, \bar{\tau}_t} \tilde{p}_t(s, \bar{\tau}_t|x_{1:t}), \quad (9)$$

where $\delta_{\tau, \bar{\tau}_t}$ is the Kronecker delta. As in the optimal model (see Eq. (2)), we use Baye's rule and the parameters of the task to derive the update equation, as

$$p_{t+1}(s, \tau|x_{1:t+1}) = \frac{1}{Z_{t+1}} g(x_{t+1}|s) \left[\mathbb{1}_{\tau=0} q(\bar{\tau}_t) \int_{s_t} a(s|s_t) p_t(s_t, \bar{\tau}_t|x_{1:t}) ds_t + \mathbb{1}_{\tau=\bar{\tau}_t+1} (1 - q(\bar{\tau}_t)) p_t(s, \bar{\tau}_t|x_{1:t}) \right]. \quad (10)$$

This distribution is non-vanishing for two values of the run-length, 0 and $\bar{\tau}_t + 1$, which correspond to the two possible scenarios: with and without a change point at trial t . We use this distribution to compute the approximate expected run-length at trial $t + 1$, $\bar{\tau}_{t+1}$, and the approximate posterior at trial $t + 1$, $\tilde{p}_{t+1}(s, \tau|x_{1:t+1})$. First, we obtain the probability of a change point at trial $t + 1$,

$$\Omega_{t+1} \equiv p_{t+1}(\tau = 0|x_{1:t+1}) = \frac{1}{Z_{t+1}} q(\bar{\tau}_t) \int g(x_{t+1}|s) \int a(s|s_t) p_t(s_t, \bar{\tau}_t|x_{1:t}) ds_t ds, \quad (11)$$

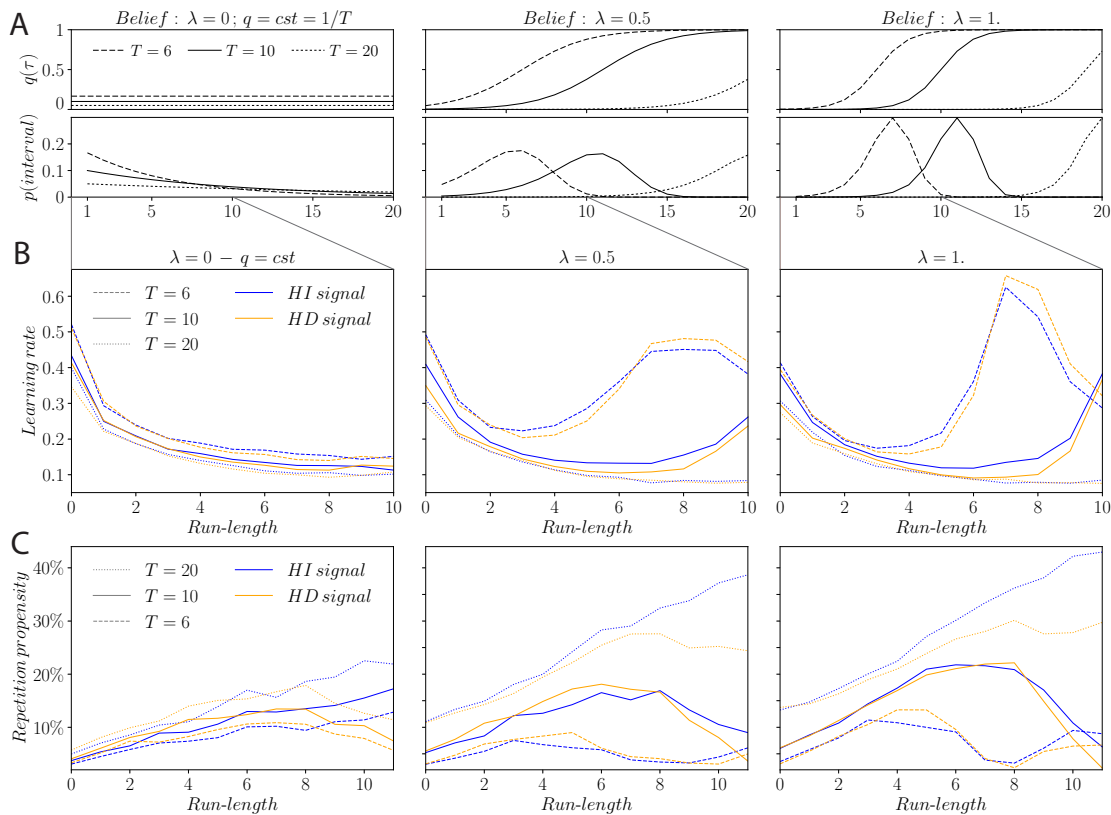


Figure 10. Illustration of *IncorrectQ* models with various beliefs on the shape of the change probability. **A.** Examples of beliefs in *IncorrectQ* models. Change probability, $q(\tau)$ (top line), and resulting inter-change-point interval distribution (bottom line), for constant q (left column), sigmoid-shaped $q(\tau)$ with slope parameter $\lambda = 0.5$ (middle) and $\lambda = 1$ (right); and for average interval length, T , of 6 (dashed line), 10 (solid), and 20 (dotted). The ‘true’ HI signals used in the task correspond to $\lambda = 0, T = 10$, while the HD signals correspond to $\lambda = 1, T = 10$. **B, C.** Average learning rate (B) and repetition propensity (C) as a function of the run-length, in *IncorrectQ* models performing optimal inference with various beliefs on the change probability, $q(\tau)$, and presented with HI signals (blue) and HD signals (orange).

and we use it to compute the approximate expected run-length at trial $t + 1$:

$$\bar{\tau}_{t+1} = \Omega_{t+1} \cdot 0 + (1 - \Omega_{t+1})(\bar{\tau}_t + 1). \quad (12)$$

Second, we approximate the posterior (Eq. (10)) by marginalizing it over the run-lengths, and multiplying the result by a Kronecker delta which takes the value 1 at $\bar{\tau}_{t+1}$:

$$\begin{aligned} p_{t+1}(s|x_{1:t+1}) &= \sum_{\tau} p_{t+1}(s, \tau|x_{1:t+1}) = p_{t+1}(s, \tau = 0|x_{1:t+1}) + p_{t+1}(s, \tau = \bar{\tau}_t + 1|x_{1:t+1}), \\ \tilde{p}_{t+1}(s, \tau|x_{1:t+1}) &= \delta_{\tau, \bar{\tau}_{t+1}} p_{t+1}(s|x_{1:t+1}). \end{aligned} \quad (13)$$

This model has no parameter.

Behavior While the optimal marginal distribution of the run-lengths, $p_t(\tau|x_{1:t})$, spans the whole range of possible values of the run-length, it is approximated, in the τ Mean model, by a delta function on a single value, $\bar{\tau}_t$. In the HI condition, the change probability, q , does not depend on the run-length, τ ; the approximate joint distribution evaluated at $\tau = \bar{\tau}_t$, $\tilde{p}_t(s, \bar{\tau}_t|x_{1:t})$, is equal to the optimal posterior distribution on the state, $p_t(s|x_{1:t})$. (Compare Eq. (8) to the combination of Eqs. (10) and (13).) As a result, the τ Mean model computes the optimal posterior on the state, and, thus, the responses in this model are the same as those in the optimal model, i.e., the τ Mean model is optimal in the HI condition. In the HD condition, by contrast, the change probability, $q(\tau)$, depends on the run-length. The τ Mean model evaluates this function at only one run-length, $\bar{\tau}_t$, an approximation of the mean run-length; as compared to the optimal model, it fails to capture fully the consequences of the dependence of the change probability on the run-length (Fig. 1D). The learning rates in this model are higher than the optimal ones for short run-lengths, and lower than the optimal ones for long run-lengths (Fig. 11B); and the repetition propensities are lower than the optimal ones for short run-lengths, and higher than the optimal ones for long run-lengths (Fig. 11C).

4.7.3 τ Nodes model

Derivation This model generalizes the one introduced in Ref. Wilson et al. (2013). In this paper, the authors interpret a change-point setting similar to ours as a ‘message-passing’ graph where run-lengths are nodes, weighted by their marginal probability $p_t(\tau|x_{1:t})$, edges are characterized by the change probability, q , and ‘messages’ are passed along edges from one node to another. More precisely, we compute the marginal probability of the run-length, using Eqs. (5) and (6):

$$p_{t+1}(\tau|x_{1:t}) = \begin{cases} \sum_{\tau_t} q(\tau_t) p_t(\tau_t|x_{1:t}) & \text{if } \tau = 0 \\ (1 - q(\tau - 1)) p_t(\tau - 1|x_{1:t}) & \text{otherwise} \end{cases}. \quad (14)$$

Hence, at trial $t + 1$, the weight of a node τ (i.e., the marginal probability of the corresponding run-length) is equal, if $\tau = 0$, to a sum of the marginal probabilities of all nodes at trial t , τ_t ,

weighted by their corresponding change probabilities, $q(\tau_t)$; and, if $\tau > 0$, it is the probability of the node $\tau - 1$, at trial t , weighted by the probability that there was no change, $1 - q(\tau - 1)$. Taking a different view, one can reformulate these weighted sums of probabilities as transfers of probability masses, as follows. Each node τ sends two ‘messages’: a ‘no-change-point’ message is sent to node $\tau + 1$ so as to set its weight to $(1 - q(\tau))p_t(\tau|x_{1:t})$, and a ‘change-point’ message is sent to node $\tau = 0$ to increase its probability by $q(\tau)p_t(\tau|x_{1:t})$. This is the message-passing algorithm. In Ref. Wilson et al. (2013), the authors assume a constant change probability, q , a likelihood, g , which belongs to the exponential family, and a state transition probability, $a(s_{t+1})$, that does not depend on the previous state, s_t , and which is a conjugate prior of the likelihood. They show that in this case each node can be seen as implementing a delta-rule, and the optimal Bayesian model amounts to the weighted sum of these delta-rules.

The authors then ‘reduce’ this model by removing nodes and accordingly revise the message-passing algorithm and each node’s update rule. We only focus on the aspects of the model that will be used in our τ Nodes implementation. The set of new nodes comprises ‘virtual’ run-lengths $l \in \{l^0, l^1, \dots, l^N\}$. A node l^i , with $i \neq 0$, now sends three messages: one to l^0 , one to the next node l^{i+1} , and one to itself. The ‘change-point’ message remains the same as in the previous algorithm, i.e., the quantity $q(l^i)p_t(l^i|x_{1:t})$ is sent to l^0 (i.e., this quantity is added to the probability of this node). The ‘no-change-point’ message is now split in two, one message being sent to the next node, l^{i+1} , and the other one being a self-passing message (i.e., sent to itself, l^i). The authors seek the relative weight $w(l^i)$ assigned to the self-passing message which gives an average run-length increase of 1 (i.e. $\mathbb{E}(l_{t+1}|l_t^i, \text{no change}) = l^i + 1$). They find $w(l^i) = \frac{l^{i+1} - l^i - 1}{l^{i+1} - l^i}$ for $i \neq N$ and $w(l^N) = 1$. The next node, l^{i+1} , hence receives the message $(1 - w(l^i))(1 - q(l^i))p_t(l^i|x_{1:t})$. With the assumptions mentioned above on q , g , and a , the model can again be understood as ‘a mixture of delta-rules’.

We implement these ideas in our τ Nodes model. Instead of $p_t(s, \tau|x_{1:t})$, we consider the probability distribution $p_t(s, l|x_{1:t})$ and apply the same Bayesian and marginalization equations used in Eqs. (4) and (5). The main difference of the new model resides in the transition probability, $p_{t+1}(s, l|s_t, l_t)$, which becomes

$$\begin{aligned} p_{t+1}(s, l|s_t, l_t) &= \mathbb{1}_{l=l^0} q(s_t, l_t) a(s_t, \tau_t, s) \\ &\quad + \mathbb{1}_{l=l_t, s=s_t} (1 - q(s_t, l_t)) w(l_t) \\ &\quad + \mathbb{1}_{l=l_t+1, s=s_t} (1 - q(s_t, l_t)) (1 - w(l_t)). \end{aligned} \tag{15}$$

Combining Eq. (15) to Eqs. (4) and (5) adapted with l , we obtain an update equation similar to the full Bayesian update equation (Eq. (7)), with an additional term corresponding to the ability of nodes for self-passing messages. The model is parameterized by the number of nodes, N_τ , but also by the values of the nodes. We chose the possible values of the nodes to be in the set $\{0, 2.5, 5, 7.5, 10, 12.5\}$. When fitting the model for a given N_τ , all the models corresponding to every possible choice of N_τ nodes within these values, were computed, and the best-fitting one was chosen.

Behavior In the HI condition, the model computes the optimal posterior on the state, $p_t(s|x_{1:t})$. Thus, as for the τ Mean model above, the responses in the τ Nodes model are the same as those in the optimal model, in the HI condition. In the HD condition, the greater the number of nodes, the more faithfully the model approximates optimal behavior (Fig. 11B,C). The learning rates are higher than the optimal ones for short run-lengths, and lower than the optimal ones for long run-lengths (Fig. 11B). The repetition propensities are higher than the optimal ones for long run-lengths; for short run-lengths, they are appreciably closer to the optimal ones, in the model with five nodes, than in the model with one node (Fig. 11C).

4.7.4 τ MaxProb model

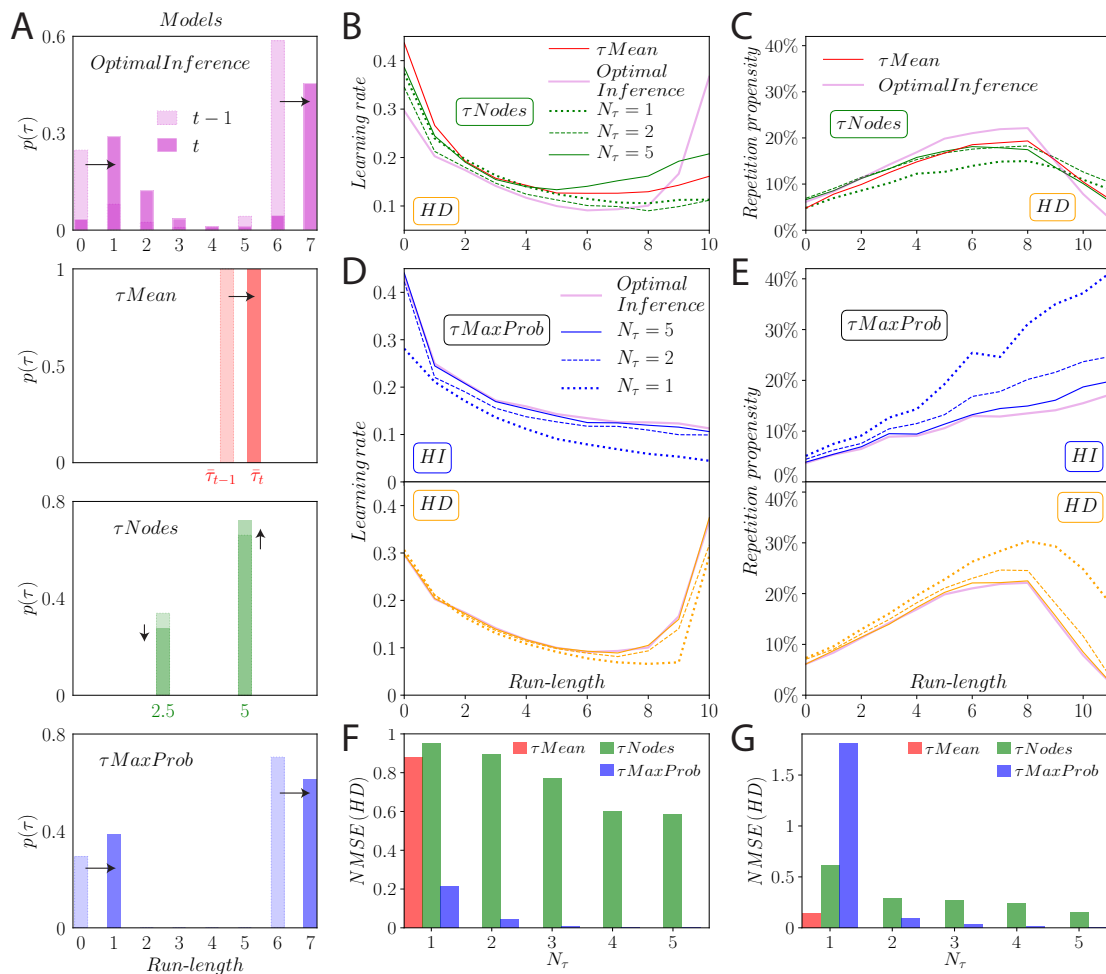
Derivation In the τ MaxProb model, we assume that we have, at trial t , an approximation of the joint distribution of the state and the run-length, which we denote by $\tilde{p}_t(s, \tau|x_{1:t})$, and we assume that the approximate marginal distribution of the run-lengths, $\tilde{p}_t(\tau|x_{1:t})$, is non-vanishing for no more than N_τ values of the run-length. Upon receiving a new stimulus, we perform a Bayesian update of the approximate joint distribution, $\tilde{p}_t(s, \tau|x_{1:t})$, as in Eq. (2), and obtain the posterior, $p_{t+1}(s, \tau|x_{1:t+1})$, from which we derive the marginal distribution of the run-lengths, $p_{t+1}(\tau|x_{1:t+1})$. If, at trial t , the run-length takes a given value, τ_t , then, at trial $t+1$, it can only take one of two values: 0 (if there is a change point) or $\tau_t + 1$ (if there is no change point). Hence, if the marginal distribution at trial t , $\tilde{p}_t(\tau|x_{1:t})$, is non-vanishing for at most N_τ values, as we assume, then the updated distribution, $p_{t+1}(\tau|x_{1:t+1})$, is non-vanishing for at most $N_\tau + 1$ values. In the case that this distribution is non-vanishing for *less* than $N_\tau + 1$ values, we do not perform further approximations, at this stage. In the other, more generic case, i.e., if $N_\tau + 1$ values of the run-length have a non-zero probability, then we identify the most unlikely run-length, $\tau^* = \arg \min p_{t+1}(\tau|x_{1:t+1})$, and we approximate the posterior as

$$\tilde{p}_{t+1}(s, \tau|x_{1:t+1}) = \begin{cases} 0 & \text{if } \tau = \tau^* \\ \frac{1}{Z} p_{t+1}(s, \tau|x_{1:t+1}) & \text{if } \tau \neq \tau^* \end{cases}, \quad (16)$$

where Z is a normalization constant equal to $1 - p_{t+1}(\tau^*|x_{1:t+1})$.

Behavior Even with $N_\tau = 1$ (one memory slot), the τ MaxProb model captures qualitatively the optimal, high learning rates for large run-lengths, in the HD condition (Fig. 11D, second panel, dotted line). However, in other situations (HD condition for shorter run-lengths, and HI condition for all run-lengths), change points are not likely ($q < 0.5$). Hence, in most cases, a vanishing run-length, i.e., the hypothesis of a change point, minimizes the marginal distribution, $p_{t+1}(\tau|x_{1:t+1})$, and its probability vanishes in our approximation: $\tilde{p}_{t+1}(\tau = 0|x_{1:t+1}) = 0$. In other words, change points tend to go by undetected. Consequently, suppressed learning rates and enhanced repetition propensity obtain in a model with a single memory slot (Fig. 11D, E).

To compare our suboptimal models to the *OptimalInference* model, we compute their normalized mean squared errors (NMSE) with regard to the responses of the optimal model (as opposed



to the responses of the human subjects, as we do in the main text). With a NMSE for learning rates at 0.21, the $\tau MaxProb$ model with $N_\tau = 1$ is closer to optimality than the $\tau Mean$ model (NMSE of 0.88) and the $\tau Nodes$ model with one node (NMSE of 0.95), in the HD condition (Fig. 11F). The high repetition propensity of the $\tau MaxProb$ model, however, leads to a larger error for this measure (NMSE of 1.81), as compared to the $\tau Mean$ (0.14) and $\tau Nodes$ (0.61) models (Fig. 11G). Adding a second memory slot allows for a better approximation of the marginal distribution, $p_t(\tau|x_{1:t})$, in the $\tau MaxProb$ model, as demonstrated by its close-to-optimal behavior with $N_\tau = 2$, both in terms of learning rates (NMSE of 0.043; compare to the $\tau Nodes$ model: 0.89) and repetition propensity (0.097; compare to the $\tau Nodes$ model: 0.29) (Fig. 11F, G).

4.7.5 $\tau Sample$ model

The $\tau Sample$ model is identical to the $\tau MaxProb$ model, except that the run-length τ^* is chosen randomly, i.e., sampled from the distribution $[1 - p_{t+1}(\tau|x_{1:t+1})]/z_{t+1}$, where z_{t+1} is a normalization factor. The stochastic nature of the update rule on the probable run-lengths influences the learning rate and the repetition propensity. In the case $N_\tau = 1$, there is, at trial t , a single run-length, τ_t , with non-vanishing probability. At trial $t + 1$, the model subject chooses randomly between the no-change-point scenario, with $\tau_{t+1} = \tau_t + 1$, and the change-point scenario, with $\tau_{t+1} = 0$. Hence, the model can incur ‘false positives’ (a change-point scenario is opted for in the absence of a true change point) and ‘false negatives’ (a true change point goes undetected by the model subject), and these occur stochastically. In most trials, the change-point scenario is less likely than the no-change-point scenario; in the $\tau MaxProb$ model, the former would be eliminated, but it occurs with some probability in the $\tau Sample$ model, leading to false positives, which induce higher learning rates. Similarly, average learning rates of the $\tau Sample$ model are higher than the optimal ones (Fig. 12A). The false negatives, in which change points go undetected, result, as in the $\tau MaxProb$ model, in higher repetition propensities (Fig. 12C). With increasing memory capacity, N_τ , the behavior of the model approaches optimality, as reflected in the decrease of the NMSEs for the learning rates (Fig. 12B) and the repetition propensities (Fig. 12D).

A qualitatively new aspect brought in by the $\tau Sample$ model is the stochasticity in the inference step, which is reflected in behavioral variability and measured by the standard deviation of the responses of a model subject. Quantitatively, false negatives have a large impact on the behavioral variability. A model subject can however correct for a false negative during the few trials that follow a true change point, i.e., at short run-lengths. This occurs randomly, in the $\tau Sample$ model, resulting in variability in responses. At longer run-lengths, the posterior probability of a change point, $p_t(\tau = 0|x_{1:t})$, is dominated by the shape of the change probability, $q(\tau)$, rather than by the observed evidence. In the HI condition, q is constant; hence, the variability reaches a plateau for run-lengths larger than about 2 (Fig. 12E, top panel). In the HD condition, as $q(\tau)$ is an increasing function of the run-length, the variability increases for run-lengths larger than 5, resulting in the ‘smile shape’ of the curve (Fig. 12E, bottom panel). As the parameter N_τ is increased, the behavior of the model approaches optimality, and, correspondingly, the standard deviation of the responses of the model subject decreases (Fig. 12F).

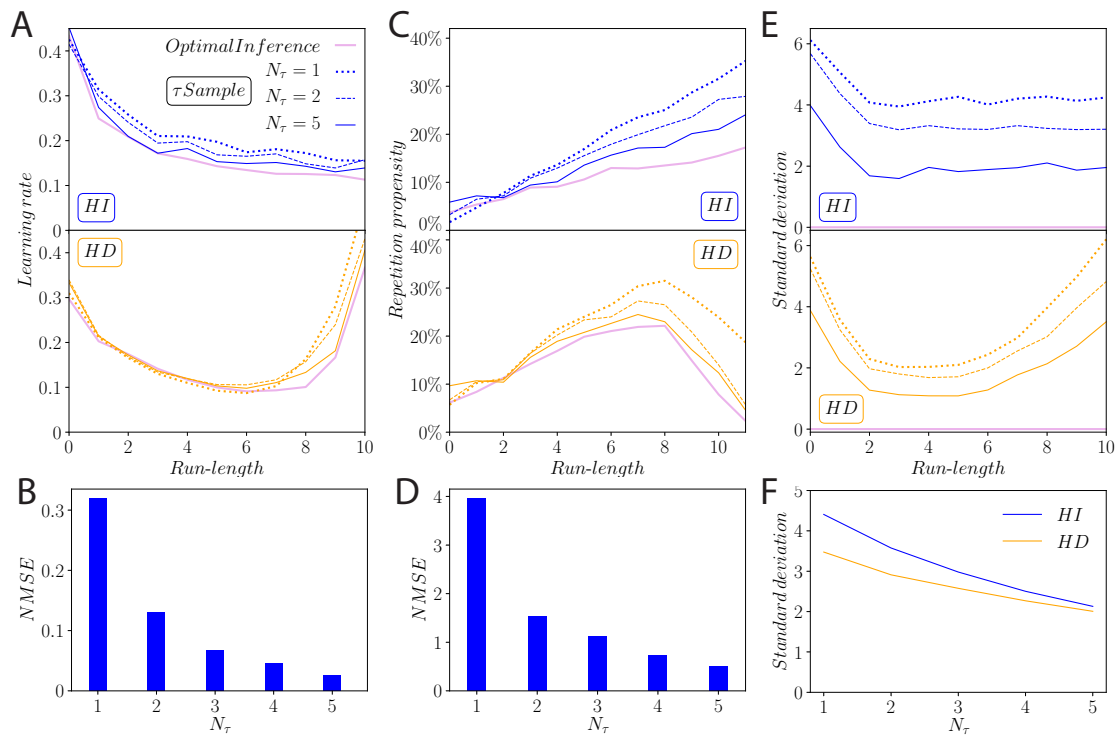


Figure 12. Behavior of the τ Sample model, as compared to the *OptimalInference* model. A, C, E. Average learning rate (A), repetition propensity (C), and standard deviations of responses (E), as a function of run-length, in the *OptimalInference* model and the τ Sample model with $N_\tau = 1, 2,$ and $5,$ in the HI condition (top panels) and in the HD condition (bottom panels). **B, D.** Normalized Mean Squared Error on learning rates (B), and on repetition propensity (D), as compared to the *OptimalInference* model, for the τ Sample model, with $N_\tau = 1$ to $5.$ **F.** Standard deviation of the responses of the τ Sample model, as a function of $N_\tau.$

4.7.6 ParticleFilter

Derivation The *ParticleFilter* approximates the posterior by a weighted sum of delta functions (Eq. (3)). To obtain the approximate posterior at trial $t + 1$, upon receiving a new observation, x_{t+1} , we start by writing the Bayesian update (Eq. (2)) of the approximate posterior at trial t , $\tilde{p}_t(s, \tau|x_{1:t})$, as

$$p_{t+1}(s, \tau|x_{1:t+1}) = \frac{1}{Z_{t+1}} g(x_{t+1}|s) \sum_{\tau_t} \int p_{t+1}(s, \tau|s_t, \tau_t) \tilde{p}_t(s, \tau|x_{1:t}) ds_t, \quad (17)$$

with the transition probability, $p_{t+1}(s, \tau|s_t, \tau_t)$, defined in Eq. (6). Injecting the expression of the approximate posterior at trial t (Eq. (3)), we can rewrite the Bayesian update as a sum of N_P functions:

$$p_{t+1}(s, \tau|x_{1:t+1}) = \frac{1}{Z_{t+1}} \sum_{i=1}^{N_P} w_t^i g(x_{t+1}|s) p_{t+1}(s, \tau|s_t^i, \tau_t^i), \quad (18)$$

where

$$p_{t+1}(s, \tau|s_t^i, \tau_t^i) = \sum_{\tau_t} \int p_{t+1}(s, \tau|s_t, \tau_t) \delta(s - s_t^i) \delta_{\tau, \tau_t^i} ds_t. \quad (19)$$

The interpretation of this form becomes apparent if we introduce, for each particle, a probability distribution over (s, τ) , defined as

$$\pi_{t+1}(s, \tau|s_t^i, \tau_t^i, x_{t+1}) \equiv \frac{g(x_{t+1}|s) p_{t+1}(s, \tau|s_t^i, \tau_t^i)}{p(x_{t+1}|s_t^i, \tau_t^i)}, \quad (20)$$

where the denominator is obtained by normalization,

$$p(x_{t+1}|s_t^i, \tau_t^i) = \sum_{\tau} \int g(x_{t+1}|s) p_{t+1}(s, \tau|s_t^i, \tau_t^i) ds. \quad (21)$$

The distribution $\pi_{t+1}(s, \tau|s_t^i, \tau_t^i, x_{t+1})$ is none other than the Bayesian update of a single particle (i.e., Eq. (17) with the approximate prior, $\tilde{p}_t(s, \tau|x_{1:t})$, replaced by $\delta(s - s_t^i) \delta_{\tau, \tau_t^i}$), and the full Bayesian update is a weighted sum of these N_P functions:

$$p_{t+1}(s, \tau|x_{1:t+1}) = \sum_{i=1}^{N_P} \frac{w_t^i p(x_{t+1}|s_t^i, \tau_t^i)}{Z_{t+1}} \pi_{t+1}(s, \tau|s_t^i, \tau_t^i, x_{t+1}). \quad (22)$$

To complete the definition of the particle filter, we have to formulate a prescription for selecting the N_P particles at trial $t + 1$. Following the literature, instead of sampling the full Bayesian update, $p_{t+1}(s, \tau|x_{1:t+1})$, we sample independently each component of the mixture, $\pi_{t+1}(s, \tau|s_t^i, \tau_t^i, x_{t+1})$, to obtain the updated particles, $(s_{t+1}^i, \tau_{t+1}^i)$. To each sample, i.e., to each particle, is assigned the weight of the corresponding component in the mixture, $w_{t+1}^i = w_t^i p(x_{t+1}|s_t^i, \tau_t^i) / Z_{t+1}$. In the rare

cases in which $p(x_{t+1}|s_t^i, \tau_t^i) = 0$, i.e., if new data invalidate particle i , and thus, $\pi_{t+1}(s, \tau|s_t^i, \tau_t^i, x_{t+1}) = 0$, we resample a new particle i from the other particles.

In practical applications of particle filters, there exists a ‘weight degeneracy’ risk, whereby the weight of one particle may overwhelm the combined weight of the others. A common method to mitigate this shortcoming is called ‘resampling’. It is a stochastic method in which the particles with high weights are likely to be duplicated, while the particles with low weights are likely to be eliminated. To achieve this, we use the N_P -dimensional categorical distribution parameterized by the N_P weights of the particles, i.e., $p(j) = w_t^j$. We sample this distribution N_P times, and obtain, thus, a set of N_P indexes, $\{j_i\}_{i=1}^{N_P}$. We use those to define the new N_P particles: for each particle i , we replace (s_t^i, τ_t^i) by $(s_t^{j_i}, \tau_t^{j_i})$, and we set all the weights to $1/N_P$. In other words, the set of particles is randomly sampled with replacement, N_P times. Particles with low weights are unlikely to survive this scheme, as compared to particles with high weights. For the sake of simplicity, we resample at each trial. A possible consequence of resampling is the ‘sample impoverishment’ problem, i.e., the loss of particles diversity (all particles bear the same state). This issue is mitigated naturally by the structure of our problem, as a new state is sampled from the distribution $a(s|s_t)$ every time a new particle carries a change-point run-length ($\tau = 0$), thus renewing the set of particles.

Behavior With a single particle ($N_P = 1$), the posterior is reduced to a unique sample, (s_t^1, τ_t^1) , and, thus, the model subject has access to a single ‘hypothesis’ on the change probability, $q(\tau_t^1)$. The particle can then evolve in one of two ways: either it opts, with this probability, for a change-point scenario, in which the new stimulus, x_{t+1} , is the only information available on the new state, along with the prior transition probability, and thus the learning rate is close to 1, or, with probability $1 - q(\tau_t^1)$, a no-change-point scenario is opted for, and the particle stays put ($s_{t+1}^1 = s_t^1$), i.e., the learning rate vanishes. As a result, when averaged over several instantiations of the particle filter, the behavior of the learning rate as a function of run-length resembles that of the change probability, $q(\tau)$, i.e., constant in the HI condition, and increasing in the HD condition (Fig. 13B), a behavior qualitatively different from either that of the *OptimalInference* model or the human responses. But it is sufficient to add no more than a second particle for the model to capture the main trends in the learning rate (a decreasing learning rate in the HI condition and smile shape in the HD condition). The NMSE drops sharply from 2.2 for $N_P = 1$ to less than 0.7 for $N_P = 2$. As additional particles are included in the model, the latter approaches optimality (the NMSE becomes less than 0.1 for $N_P \geq 70$) (Fig. 13C).

As mentioned, sampling in the particle filter induces variability in behavior: two particle filters receiving the same sequence of observations do not respond with the same sequence of estimates. Since the stochasticity stems from the sampling of an (approximate) posterior, the resulting variability scales with the width of the posterior. As measured by the standard deviation of responses, it decreases with the run-length, in the HI condition. In the HD condition, it decreases at short run-lengths before increasing at longer run-lengths (Fig. 13F). This behavior reproduces, at least qualitatively, that of the subjects (compare to Fig. 5C). The greater the number of particles

in a particle filter, the closer the latter approximates the *OptimalInference* model; the standard deviation of the responses is a decreasing function of the number of particles (Fig. 13G).

Since it operates on a low-dimensional spatial representation, the particle filter naturally predicts a higher repetition propensity than the *OptimalInference* model does. More specifically, the posterior is non-vanishing for only a finite (possibly small) set of values at each trial, and it is more likely than in the optimal case that the subject model’s estimate remains unchanged following stimulus presentation. This effect is quantitatively appreciable, and leads to repetition propensities which are multiples of those in the *OptimalInference* model. Again, the repetition propensities decrease toward their optimal values as the number of particles, N_P , increases. The corresponding NMSE drops from 202 for $N_P = 1$ to 24 for $N_P = 150$ (Fig. 13D, E).

4.7.7 Sampling model

In the *Sampling* model, instead of using the Bayesian posterior to maximize its expected score, a model subject samples its response from the marginal posterior on the states, $p_t(s|x_{1:t})$. In spite of this suboptimal selection rule, the average learning rate as a function of the run-length has a behavior similar to the optimal one (decreasing in the HI condition, smile-shaped in the HD condition), albeit with higher average values (Fig. 14A). The repetition propensity also behaves similarly to the optimal one, but is suppressed in magnitude due to sampling (Fig. 14B). Finally, as expected by construction, the *Sampling* model leads to behavioral variability, and the amplitude of the latter scales with the width of the posterior distribution (Fig. 14C).

4.7.8 Normalized Mean Squared Error

This section provides some details on the Normalized Mean Squared Error we use to compare the results of the various models to the *OptimalInference* model and to human data. Let $y_i(\tau)$ be the value at run-length τ of the quantity of interest i (learning rate, repetition propensity, or standard deviation of responses), as observed in data or as resulting from the optimal model, and $\hat{y}_i(\tau)$ the value resulting from a suboptimal model. The mean squared error is $MSE(\hat{y}_i) = \frac{1}{n} \sum_{\tau} (\hat{y}_i(\tau) - y_i(\tau))^2$, where n is the number of run-lengths. We want to be able to compare the errors for different quantities of interest. By dividing the MSE by the variance of y_i , we obtain the Normalized Mean Squared Error, which is translation-invariant and scale-invariant:

$$NMSE_i = \frac{MSE(\hat{y}_i)}{\text{Var}[y_i]} = \frac{\sum_{\tau} (\hat{y}_i(\tau) - y_i(\tau))^2}{\sum_{\tau} (\bar{y}_i - y_i(\tau))^2}, \quad (23)$$

where \bar{y}_i is the average $y_i(\tau)$. For model fitting, we then use the average of this quantity over the three quantities of interest (“three-error measure”) or over two of them (“two-error measure”).

5 Acknowledgements

This work was supported by the CNRS through UMR8550, the Global Scholar Program at Princeton University, and the Visiting Faculty Program at the Weizmann Institute of Science. A.P.C. was

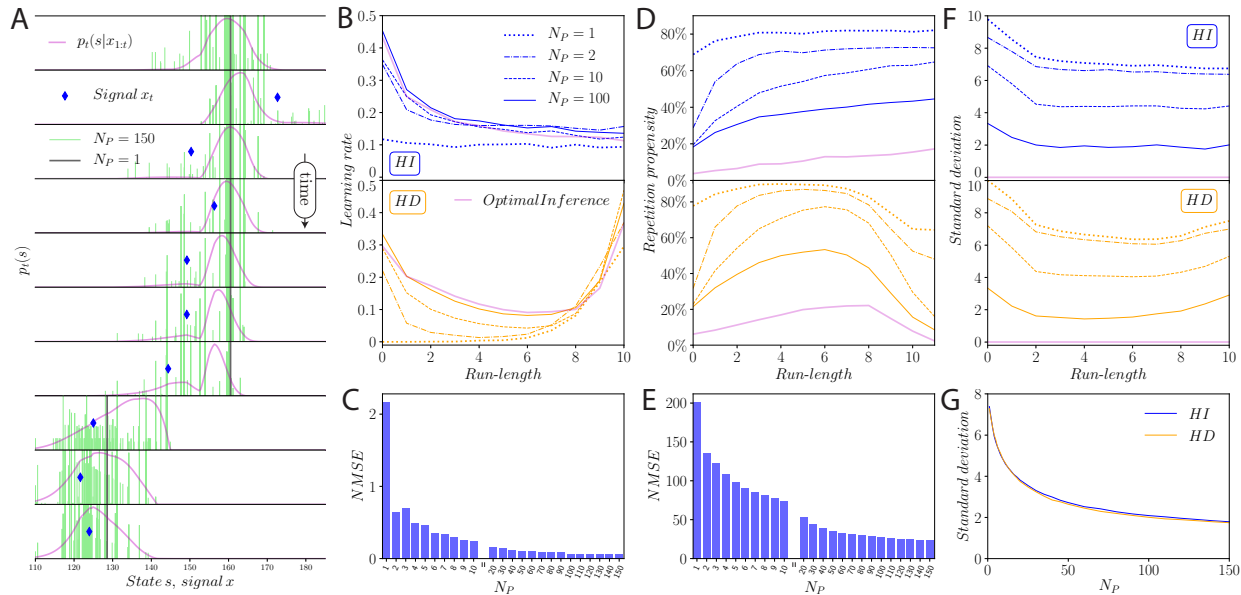


Figure 13. Illustration of the *ParticleFilter* model and its behavior. **A.** Distribution of particles during inference, compared with the optimal posterior distribution, for particle filters with $N_P = 1$ and 150. Each particle is a point on the (s, τ) plane, equipped with a weight. Only the spatial components s are represented here, as vertical bars (grey for $N_P = 1$, green for $N_P = 150$). Bars heights are proportional to the corresponding weights, but some are truncated due to the choice of scale, which emphasizes weight diversity. Upon receiving a new stimulus, x_{t+1} (blue), a particle i is updated by sampling $p_{t+1}(s, \tau | s_t^i, \tau_t^i, x_{t+1})$. This may or may not involve a change point, in which case $s_{t+1}^i \neq s_t^i$. **B, D, F.** Average learning rate (B), repetition propensity (D), and standard deviations of responses (F), as a function of run-length, in the *OptimalInference* model and the *ParticleFilter* model with $N_P = 1, 2, 10,$ and 100 , in the HI condition (top panels) and in the HD condition (bottom panels). **C, E.** Normalized Mean Squared Error on learning rates (C), and on repetition propensity (E), as compared to the *OptimalInference* model, for the *ParticleFilter* model, with $N_P = 1$ to 150. **G.** Standard deviation of the responses of the *ParticleFilter* model, as a function of number of particles, N_P .




Figure 14. Behavior of the *Sampling* model, as compared to the *OptimalInference* model. **A,B.** Average learning rate (A), and repetition propensity (B), in the *Sampling* model (dashed lines) and the *OptimalInference* model (solid lines), as a function of run-length, in the HI and HD conditions. **C.** Standard deviations of responses, as a function of run-length, of the *Sampling* model in the HI and HD conditions. The *OptimalInference* model exhibits no variability.

supported by a PhD fellowship of the Fondation Pierre-Gilles de Gennes pour la Recherche. This work was granted access to the HPC resources of MesoPSL financed by the Region Ile de France and the project Equip@Meso (reference ANR-10-EQPX-29-01) of the program Investissements d’Avenir supervised by the Agence Nationale pour la Recherche.

References

- Acerbi, L., Vijayakumar, S., and Wolpert, D. M. (2014). On the Origins of Suboptimality in Human Probabilistic Inference - S2 - Noisy Probabilistic Inference. *PLoS Computational Biology*, 10(6):1–5.
- Adams, R. P. and MacKay, D. J. C. (2007). Bayesian Online Changepoint Detection. *ArXiv e-prints*, pages 1–7.
- Ahn, S. and Fessler, J. (2003). Standard Errors of Mean, Variance, and Standard Deviation Estimators. *EECS Department, University of Michigan*, pages 1–2.
- Aitchison, L. and Lengyel, M. (2016). The Hamiltonian Brain: Efficient Probabilistic Inference with Excitatory-Inhibitory Neural Circuit Dynamics. *PLoS Computational Biology*, 12(12):1–24.
- Anteneodo, C. and Chialvo, D. R. (2009). Unraveling the fluctuations of animal motor activity. *Chaos*, 19:1–7.
- Arulampalam, M., Maskell, S., Gordon, N., and Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188.
- Battaglia, P. W., Jacobs, R. a., and Aslin, R. N. (2003). Bayesian integration of visual and auditory signals for spatial localization. *Journal of the Optical Society of America A*, 20(7):1391.

- Battaglia, P. W., Kersten, D., and Schrater, P. R. (2011). How haptic size sensations improve distance perception. *PLoS Computational Biology*, 7(6):e1002080.
- Berniker, M. and Kording, K. (2011). Bayesian approaches to sensory integration for motor control. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(4):419–428.
- Blaisdell, A. P., Sawa, K., Leising, K. J., and Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, 311(5763):1020–1022.
- Brown, S. D. and Steyvers, M. (2009). Detecting and predicting changes. *Cognitive Psychology*, 58(1):49–67.
- Buesing, L., Bill, J., Nessler, B., and Maass, W. (2011). Neural dynamics as sampling: A model for stochastic computation in recurrent networks of spiking neurons. *PLoS Computational Biology*, 7(11):e1002211.
- Campione, E. and Véronis, J. (2002). A Large-Scale Multilingual Study of Silent Pause Duration. *Speech Prosody 2002*, pages 199–202.
- Cox, R. T. (1946). Probability, Frequency and Reasonable Expectation. *American Journal of Physics*, 14(1):1.
- Daw, N. D. and Courville, A. C. (2008). The pigeon as particle filter. *Advances in neural information processing systems*, pages 369–376.
- Denison, S., Bonawitz, E., Gopnik, A., and Griffiths, T. L. (2013). Rational variability in children’s causal inferences: The Sampling Hypothesis. *Cognition*, 126(2):280–300.
- Doucet, A. and Johansen, A. M. (2008). A Tutorial on Particle Filtering and Smoothing: Fifteen years later. In Crisan, D. and Rozovsky, B., editors, *The Oxford Handbook of Nonlinear Filtering*, pages 1–39. Oxford University Press.
- Drugowitsch, J., Wyart, V., Devauchelle, A.-d., Koechlin, E., Drugowitsch, J., Wyart, V., Devauchelle, A.-d., and Koechlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron*, 92(6):1398–1411.
- Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433.
- Fearnhead, P. and Liu, Z. (2007). On-line inference for multiple changepoint problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(4):589–605.
- Fiser, J., Berkes, P., Orbán, G., and Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–130.
- Gershman, S. J. and Beck, J. M. (2016). Complex Probabilistic Inference : From Cognition to Neural Computation. In *Computational Models of Brain and Behavior*, pages 1–17.
- Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278.
- Gershman, S. J., Vul, E., and Tenenbaum, J. B. (2012). Multistability and Perceptual Inference. *Neural Computation*, 24(1):1–24.

- Ghose, G. M. and Maunsell, J. H. R. (2002). Attentional modulation in visual cortex depends on task timing. *Nature*, 419(6907):616–620.
- Goodman, N., Tenenbaum, J., Feldman, J., and Griffiths, T. (2008). A Rational Analysis of Rule-Based Concept Learning. *Cognitive Science: A Multidisciplinary Journal*, 32(1):108–154.
- Gordon, N., Salmond, D., and Smith, A. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F Radar and Signal Processing*, 140(2):107.
- Griffin, L., West, D. J., and West, B. J. (2000). Random stride intervals with memory. *Journal of Biological Physics*, 26(3):185–202.
- Griffiths, T. L. and Tenenbaum, J. B. (2006). Optimal Predictions in Everyday Cognition. *Psychological Science*, 17(9):767–773.
- Griffiths, T. L. and Tenenbaum, J. B. (2011). Predicting the future as Bayesian inference: People combine prior knowledge with observations when estimating duration and extent. *Journal of Experimental Psychology: General*, 140(4):725–743.
- Hausdorff, J. M., Peng, C. K., Ladin, Z., Wei, J. Y., and Goldberger, a. L. (1995). Is walking a random walk? Evidence for long-range correlations in stride interval of human gait. *Journal of applied physiology (Bethesda, Md. : 1985)*, 78(1):349–358.
- Hennequin, G., Aitchison, L., and Lengyel, M. (2014). Fast Sampling-Based Inference in Balanced Neuronal Networks. *Advances in Neural Information Processing Systems*, pages 1–9.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4(3):267–272.
- Hillis, J. M., Watt, S. J., Landy, M. S., and Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, 4(12):1.
- Hoyer, P. O. and Hyvärinen, A. (2003). Interpreting neural response variability as monte carlo sampling of the posterior. *Advances in neural information processing systems*, (1):293–300.
- Huys, Q. J. M., Eshel, N., Nions, E. O., Sheridan, L., Dayan, P., and Jonathan, P. (2012). Bonsai Trees in Your Head : How the Pavlovian System Sculpts Goal-Directed Choices by Pruning Decision Trees. *PLOS Computational Biology*, 8(3).
- Jacobs, R. a. (1999). Optimal integration of texture and motion cues to depth. *Vision Research*, 39(21):3621–3629.
- Janssen, P. and Shadlen, M. N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nature Neuroscience*, 8(2):234–241.
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press.
- Jazayeri, M. and Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, 13(8):1020–1026.

- Kaufmann, E., Korda, N., and Munos, R. (2012). Thompson Sampling: An Asymptotically Optimal Finite Time Analysis. *Algorithmic Learning Theory*, (1):15.
- Knill, D. C. (2003). Mixture models and the probabilistic structure of depth cues. *Vision Research*, 43(7):831–854.
- Knill, D. C. (2007). Robust cue integration: A Bayesian model and evidence from cue-conflict studies with stereoscopic and figure cues to slant. *Journal of Vision*, 7(7):5.
- Koehler, D. J. and James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, 113(1):123–127.
- Körding, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247.
- Körding, K. P. and Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10(7):319–326.
- Levy, R., Reali, F., and Griffiths, T. L. (2008). Modeling the effects of memory on human online sentence processing with particle filters. *Proceedings of NIPS 2008*, pages 1–8.
- Li, Y. and Dudman, J. T. (2013). Mice infer probabilistic models for timing. *Proceedings of the National Academy of Sciences*, 110(42):17154–17159.
- Low, E., Grabe, E., and Nolan, F. (2000). Quantitative characterisations of speech rhythm. *Language and Speech*, 46:377–401.
- Miyazaki, M., Nozaki, D., and Nakajima, Y. (2005). Testing Bayesian models of human coincidence timing. *Journal of neurophysiology*, 94(1):395–399.
- Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, 42(2):223–7.
- Moreno-Bote, R., Knill, D. C., and Pouget, A. (2011). Bayesian sampling in visual perception. *Proceedings of the National Academy of Sciences*, 108(30):12491–12496.
- Nakamura, T., Kiyono, K., Yoshiuchi, K., Nakahara, R., Struzik, Z. R., and Yamamoto, Y. (2007). Universal Scaling Law in Human Behavioral Organization. *Physical Review Letters*, 138103(September):1–4.
- Nakamura, T., Takumi, T., Takano, A., Aoyagi, N., Yoshiuchi, K., Zbigniew, R., and Yamamoto, Y. (2008). Of Mice and Men — Universality and Breakdown of Behavioral Organization. *PLoS ONE*, 3(4):1–8.
- Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., and Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, 15(7):1040–1046.
- Nassar, M. R., Wilson, R. C., Heasly, B., and Gold, J. I. (2010). An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *Journal of Neuroscience*, 30(37):12366–12378.

- Nunes Amaral, L. A., Bezerra Soares, D. J., da Silva, L. R., Lucena, L. S., Saito, M., Kumano, H., Aoyagi, N., and Yamamoto, Y. (2004). Power law temporal auto-correlations in day-long records of human physical activity. *Europhys. Lett.*, 66(3):448–454.
- Pearl, L., Goldwater, S., and Steyvers, M. (2011). Online Learning Mechanisms for Bayesian Models of Word Segmentation. *Research on Language and Computation*, 8(2-3):1–26.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using psychopy. *Frontiers in Neuroinformatics*, 2(10).
- Ramus, F., Nespore, M., and Mehler, J. (1999). Correlates of linguistic rhythm in the speech signal. *Cognition*, 73(3):265–292.
- Rigoux, L. and Guigon, E. (2012). A Model of Reward- and Effort-Based Optimal Decision Making and Motor Control. *PLoS Computational Biology*, 8(10).
- Sanborn, A. N. (2015). Types of approximation for probabilistic cognition: Sampling and variational. *Brain and Cognition*, pages 8–11.
- Sanborn, A. N. and Chater, N. (2016). Bayesian Brains without Probabilities. *Trends in Cognitive Sciences*, xx:1–11.
- Sanborn, A. N., Griffiths, T. L., and Navarro, D. J. (2006). A More Rational Model of Categorization. *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, pages 1–6.
- Sanborn, A. N., Griffiths, T. L., and Navarro, D. J. (2010). Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review*, 117(4):1144–1167.
- Savin, C. and Denève, S. (2014). Spatio-temporal representations of uncertainty in spiking neural networks. *Advances in Neural Information Processing Systems*, pages 1–9.
- Shadmehr, R. (2009). Computational Approaches to Motor Control. In *Encyclopedia of Neuroscience*, volume 3, pages 9–17.
- Shi, L. and Griffiths, T. (2009). Neural implementation of hierarchical Bayesian inference by importance sampling. *Advances in Neural Information Processing Systems 22*, pages 1669–1677.
- Stocker, A. A. and Simoncelli, E. P. (2008). A Bayesian Model of Conditioned Perception. *Advances in neural information processing systems*, 20(May):1409–1416.
- Sundareswara, R. and Schrater, P. R. (2008). Perceptual multistability predicted by search model for Bayesian decisions. *Journal of vision*, 8(5):12.1–19.
- ten Oever, S., Schroeder, C. E., Poeppel, D., van Atteveldt, N., and Zion-Golumbic, E. (2014). Rhythmicity and cross-modal temporal cues facilitate detection. *Neuropsychologia*, 63:43–50.
- Thompson, W. R. (1933). On the Likelihood That One Unknown Probability Exceeds Another In View Of The Evidence Of Two Samples. *Biometrika*, 25(3):285–294.

- van Beers, R. J., Sittig, A. C., and Denier, J. A. N. J. (1999). Integration of Proprioceptive and Visual Position-Information: An Experimentally Supported Model. *Integration The Vlsi Journal*, 81(3):1355–1364.
- Vul, E. (2011). Sampling in human cognition - Chapter 1. *Dissertation Abstracts International: Section B: The Sciences and Engineering*, 72(4-B):2471.
- Vul, E., Frank, M. C., Tenenbaum, J. B., and Alvarez, G. (2009). Explaining human multiple object tracking as resource-constrained approximate inference in a dynamic probabilistic model. *Advances in Neural Information Processing Systems*, 22:1–9.
- Vul, E., Goodman, N., Griffiths, T. L., and Tenenbaum, J. B. (2014). One and Done? Optimal Decisions From Very Few Samples. *Cognitive Science*, 38(4):599–637.
- Wilson, R. C., Nassar, M. R., and Gold, J. I. (2010). Bayesian Online Learning of the Hazard Rate In Change Point Problems. *Neural Computation*, pages 1–25.
- Wilson, R. C., Nassar, M. R., and Gold, J. I. (2013). A Mixture of Delta-Rules Approximation to Bayesian Inference in Change-Point Problems. *PLoS Computational Biology*, 9(7):e1003150.
- Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, 1(6):1–5.
- Wozny, D. R., Beierholm, U. R., and Shams, L. (2010). Probability Matching as a Computational Strategy Used in Perception. *PLoS Computational Biology*, 6(8):e1000871.
- Yu, A. J. and Huang, H. (2014). Maximizing Masquerading as Matching in Human Visual Search Choice Behavior. *Decision*, 1(4):275–287.

Appendix 1

A Supplementary tables: statistical tests

| $\tilde{\tau}$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|-----------------------|------|------|------|------|-------|--------|---------|--------|------|------|---------|
| W | .307 | .448 | .265 | .335 | .180 | .072* | .016** | .073* | .426 | .160 | .008*** |
| MWU | .269 | .391 | .281 | .212 | .207 | .066* | .034** | .022** | .491 | .568 | .006*** |
| S | .307 | .448 | .268 | .337 | .178 | .070* | .015** | .074* | .426 | .158 | .008*** |
| N(HI) | 454 | 521 | 613 | 508 | 461 | 436 | 411 | 360 | 318 | 316 | 240 |
| N(HD) | 603 | 616 | 707 | 632 | 637 | 521 | 499 | 421 | 387 | 300 | 231 |
| Excluding repetitions | | | | | | | | | | | |
| W | .353 | .421 | .329 | .352 | .142 | .042** | .008*** | .151 | .391 | .202 | .034** |
| MWU | .357 | .431 | .484 | .170 | .073* | .025** | .007*** | .058* | .343 | .590 | .050** |
| S | .352 | .421 | .332 | .353 | .140 | .041** | .007*** | .153 | .391 | .201 | .035** |
| N(HI) | 412 | 435 | 495 | 402 | 353 | 333 | 307 | 288 | 237 | 231 | 174 |
| N(HD) | 553 | 529 | 582 | 497 | 493 | 405 | 379 | 318 | 291 | 226 | 186 |

Table 1. p-values for the one-sided statistical tests of equality of the means of the learning rates in the HI and HD conditions, for each run-length (Fig. 2C). W: Welch’s test. MWU: Mann-Whitney’s U test. S: Student’s test. The N(HI) and N(HD) lines report the number of observations. The second half of the table reports the same quantities when excluding all occurrences of repetitions (Suppl. Fig. 1B).

| Condition 1 | Condition 2 | Avg_1 | Avg_2 | W | MWU | S | N_1 | N_2 |
|--------------------------------|--------------------------------|---------|---------|----------|----------|----------|-------|-------|
| HI, $\tilde{\tau} \in [5, 6]$ | HI, $\tilde{\tau} \in [9, 10]$ | 0.255 | 0.208 | .0099*** | .0476** | .0119** | 847 | 556 |
| HI, $\tilde{\tau} \in [5, 6]$ | HD, $\tilde{\tau} \in [5, 6]$ | 0.255 | 0.209 | .0055*** | .0089*** | .0052*** | 847 | 1020 |
| HI, $\tilde{\tau} \in [9, 10]$ | HD, $\tilde{\tau} \in [9, 10]$ | 0.208 | 0.264 | .0103** | .0602* | .0100*** | 556 | 531 |
| HD, $\tilde{\tau} \in [5, 6]$ | HD, $\tilde{\tau} \in [9, 10]$ | 0.209 | 0.264 | .0072*** | .0173** | .0050*** | 1020 | 531 |
| Excluding repetitions | | | | | | | | |
| HI, $\tilde{\tau} \in [5, 6]$ | HI, $\tilde{\tau} \in [9, 10]$ | 0.338 | 0.285 | .0203** | .0780* | .0233** | 640 | 405 |
| HI, $\tilde{\tau} \in [5, 6]$ | HD, $\tilde{\tau} \in [5, 6]$ | 0.338 | 0.272 | .0017*** | .0009*** | .0016*** | 640 | 784 |
| HI, $\tilde{\tau} \in [9, 10]$ | HD, $\tilde{\tau} \in [9, 10]$ | 0.285 | 0.340 | .0342** | .1721 | .0345** | 405 | 412 |
| HD, $\tilde{\tau} \in [5, 6]$ | HD, $\tilde{\tau} \in [9, 10]$ | 0.272 | 0.340 | .0066*** | .0095*** | .0047*** | 784 | 412 |

Table 2. Learning rates averages under various HI/HD and short/long run-length conditions, and p-values for one-sided statistical tests of equality of the means (Fig. 2B). The first two columns indicate the HI/HD and short/long run-length conditions. Learning rates at trials verifying these two conditions have their averages reported in the Avg_1 and Avg_2 columns. Columns W, MWU, and S provide the p-values for the tests of equality of the means between the two conditions. W: Welch’s test. MWU: Mann-Whitney’s U test. S: Student’s test. N_1 and N_2 report the number of observations for each condition. The second half of the table reports the same quantities when excluding all occurrences of repetitions (Suppl. Fig. 1A).

| $\tilde{\tau}$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|-----------------------|---------|---------|---------|--------|----------|----------|---------|------|------|---------|----------|
| p | .0266** | .0331** | .0123** | .0550* | .0048*** | .0000*** | .0387** | .213 | .445 | .0331** | .0000*** |
| N(HI) | 1255 | 1425 | 1611 | 1440 | 1347 | 1268 | 1111 | 975 | 883 | 795 | 722 |
| N(HD) | 1591 | 1674 | 1943 | 1804 | 1650 | 1493 | 1333 | 1168 | 983 | 792 | 584 |
| Excluding repetitions | | | | | | | | | | | |
| p | .0243** | .0293** | .0556* | .4847 | .0094*** | .0000*** | .4853 | .272 | .953 | .0560* | .0004*** |
| N(HI) | 1123 | 1147 | 1291 | 1127 | 1003 | 963 | 818 | 748 | 643 | 583 | 533 |
| N(HD) | 1455 | 1405 | 1559 | 1390 | 1255 | 1115 | 978 | 890 | 723 | 614 | 458 |

Table 3. p-values for Levene’s statistical test of equality of the variances of the corrections in the HI and HD conditions, for each run-length (Fig. 4C). The N(HI) and N(HD) lines report the number of observations. The second part of the table reports the same quantities when excluding all occurrences of repetitions (Suppl. Fig. 1).

| $\tilde{\tau}$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
|----------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|--------|----------|
| p | 0.561 | 0.114 | 0.843 | 0.245 | 0.420 | 0.333 | 0.636 | 1.000 | 0.848 | 0.194 | 0.091* | 0.005*** |
| N(HI) | 2635 | 2211 | 1984 | 1786 | 1620 | 1471 | 1351 | 1215 | 1098 | 978 | 875 | 791 |
| N(HD) | 2647 | 2339 | 2135 | 1943 | 1804 | 1661 | 1534 | 1404 | 1240 | 1111 | 946 | 781 |

Table 4. p-values for Fisher’s exact test of independence of the repetition propensities in the HI and HD conditions, at each run-length (Fig. 3B). The N(HI) and N(HD) lines report the number of observations.

Appendix 2

B Supplementary figure: data analysis excluding all occurrences of repetitions




Figure 1. Human learning rates and standard deviations of responses, excluding all occurrences of repetitions. (A,B) as in Figs. 2B,C; (C) as in Fig. 4C.