Adaptive coding for dynamic sensory inference

Wiktor Młynarski[1] & Ann M. Hermundstad[2]

1) Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology
2) Janelia Research Campus, Howard Hughes Medical Institute

mlynar@mit.edu, hermundstada@janelia.hhmi.org

Abstract

Behavior relies on the ability of sensory systems to infer properties of the environment from incoming stimuli. The accuracy of inference depends on the fidelity with which behaviorally-relevant properties of stimuli are encoded in neural responses. High-fidelity encodings can be metabolically costly, but low-fidelity encodings can cause errors in inference. Here, we discuss general principles that underlie the tradeoff between encoding cost and inference error. We then derive adaptive encoding schemes that dynamically navigate this tradeoff. These optimal encodings tend to increase the fidelity of the neural representation following a change in the stimulus distribution, and reduce fidelity for stimuli that originate from a known distribution. We predict dynamical signatures of such encoding schemes and demonstrate how known phenomena, such as burst coding and firing rate adaptation, can be understood as hallmarks of optimal coding for accurate inference.

INTRODUCTION

Biological systems must make inferences about the environment in order to successfully plan and accomplish goals. Inference is the process of estimating behaviorally-relevant properties of the environment from low-level sensory signals registered by neurons in the early sensory periphery [1]. Many perceptual tasks, such as color perception [2], visual speed estimation [3], or sound localization [4, 5], can be understood as probabilistic inference. All of these tasks rely on the estimation of features (such as the speed of an object) that are not explicitly represented by low-level sensory stimuli (such as light signals incident on photoreceptors).

To accurately perform inference, the nervous system can construct an internal model that relates incoming sensory stimuli to behaviorally-relevant properties of the environment [1, 6, 7, 8, 9]. As the environment changes, this internal model must be continually updated with new stimuli [10, 11, 12, 13, 14], and therefore the accuracy of this internal model depends on the fidelity with which incoming stimuli are encoded in neural responses.

The process of encoding sensory stimuli, however, is metabolically expensive [15, 16, 17, 18, 19, 20], and a large body of evidence suggests that sensory systems have evolved to reduce the energetic costs of stimulus coding [15, 21, 22]. These findings provide empirical support for the efficient coding hypothesis [23], which postulates that sensory systems minimize metabolic cost while maximizing the amount of information that is encoded about a stimulus [24, 25, 26].

The goal of maximizing stimulus information does not reflect the fact that different stimuli can have different utility to a system for making inferences about the environment [27, 28, 29, 30]. The relative utility of a stimulus is determined by the potential impact that it can have on the system's belief about the state of the environment; stimuli that sway this belief carry high utility, while stimuli that do not affect this belief are less relevant. Moreover, physically different stimuli can exert the same influence the observer's belief, and can therefore be encoded in the same neural activity pattern without affecting the inference process. Such an encoding strategy decreases the fidelity of the neural representation by using the same activity pattern to represent many stimuli, and consequently reduces the amount of metabolic resources required to perform inference.

When the distribution of stimuli changes in time, as in any natural environment, both the belief about the environment [11] and the relative impact of different stimuli on this belief also change in time. Any system that must perform accurate inference with minimal energy must therefore *dynamically* balance the

cost of encoding stimuli with the error that this encoding can introduce in the inference process. While studies have separately shown that sensory neurons dynamically adapt to changing stimulus distributions in manners that reflect either optimal encoding [31] or inference [10], the interplay between these two objectives is not understood.

In this work, we develop a general framework for relating low-level sensory encoding schemes to the higher-level processing that ultimately supports behavior. We use this framework to explore the dynamic interplay between efficient encoding, which serves to represent the stimulus with minimal metabolic cost, and accurate inference, which serves to estimate behaviorally-relevant properties of the stimulus with minimal error. To illustrate the implications of this framework, we consider three neurally-plausible encoding schemes in a simple model environment. Each encoding scheme reflects a different limitation on the representational capacity of neural responses, and consequently each represents a different strategy for reducing metabolic costs. We then generalize this framework to a visual inference task with natural stimuli.

We find that encoding schemes optimized for inference differ significantly from encoding schemes that are designed to accurately reconstruct all details of the stimulus. The latter produce neural responses that are more metabolically costly, and the resulting inference process exhibits qualitatively different inaccuracies.

Together, these results predict dynamical signatures of encoding strategies that are designed to support accurate inference, and differentiate these strategies from those that are designed to reconstruct the stimulus itself. These dynamical signatures provide a new interpretation of experimentally-observed phenomena such as burst coding and firing-rate adaptation, which we argue could arise as a consequence of a dynamic tradeoff between coding cost and inference error.

RESULTS

A general framework for dynamically balancing coding cost and inference error

Sensory systems use internal representations of external stimuli to build and update models of the environment. As an illustrative example, consider the task of avoiding a predator (Fig. 1A, left column). The predator is signaled by sensory stimuli, such as patterns of light intensity or chemical odorants, that change over time. To avoid a predator, an organism must first determine whether a predator is present, and if so, which direction the predator is moving, and how fast. This inference process requires that incoming stimuli first be encoded in the spiking activity of sensory neurons. This activity must then be transmitted to downstream neurons that infer the position and speed of the predator.

Not all stimuli will be equally useful for this task, and the relative utility of different stimuli could change over time. When first trying to determine whether a predator is present, it might be crucial to encode stimulus details that could discriminate fur from grass. Once a predator has been detected, however, the details of the predator's fur are not relevant for determining its position and speed. If encoding stimuli is metabolically costly, energy should be devoted to encoding those details of the stimulus that are most useful for inferring the quantity at hand.

We formalize this scenario within a general Bayesian framework that consists of three components: (*i*) an environment, which is parameterized by a latent state θ_t that specifies the distribution $p(x_t|\theta_t)$ of incoming sensory stimuli x_t, (*ii*) an adaptive encoder, which maps incoming stimuli x_t onto neural responses y_t, and (*iii*) an observer, which uses these neural responses to update an internal belief about the current and future states of the environment. This belief is summarized by the posterior distribution $p(\theta_t|y_{\tau\leq t})$, and is constructed by first decoding the stimulus from the neural response, and then combining the decoded stimulus with the prior belief $p(\theta_{t-1}|y_{\tau<t})$ and knowledge of environment dynamics. A prediction about the future state of the environment can be computed in an analogous manner by combining the posterior distribution with knowledge of environment dynamics (Methods, Box 1). This prediction is then fed back upstream and used to adapt the encoder.

In order to optimize and assess the dynamics of the system, we use the point values $\hat{\theta}_t$ and $\vec{\theta}_{t+1}$ as an estimate of the current state and prediction of the future state, respectively. The optimal point estimate is computed by averaging the posterior and is guaranteed to minimize the mean squared error between the estimated state $\hat{\theta}_t$ and the true state θ_t, regardless of the form of the posterior distribution [32].

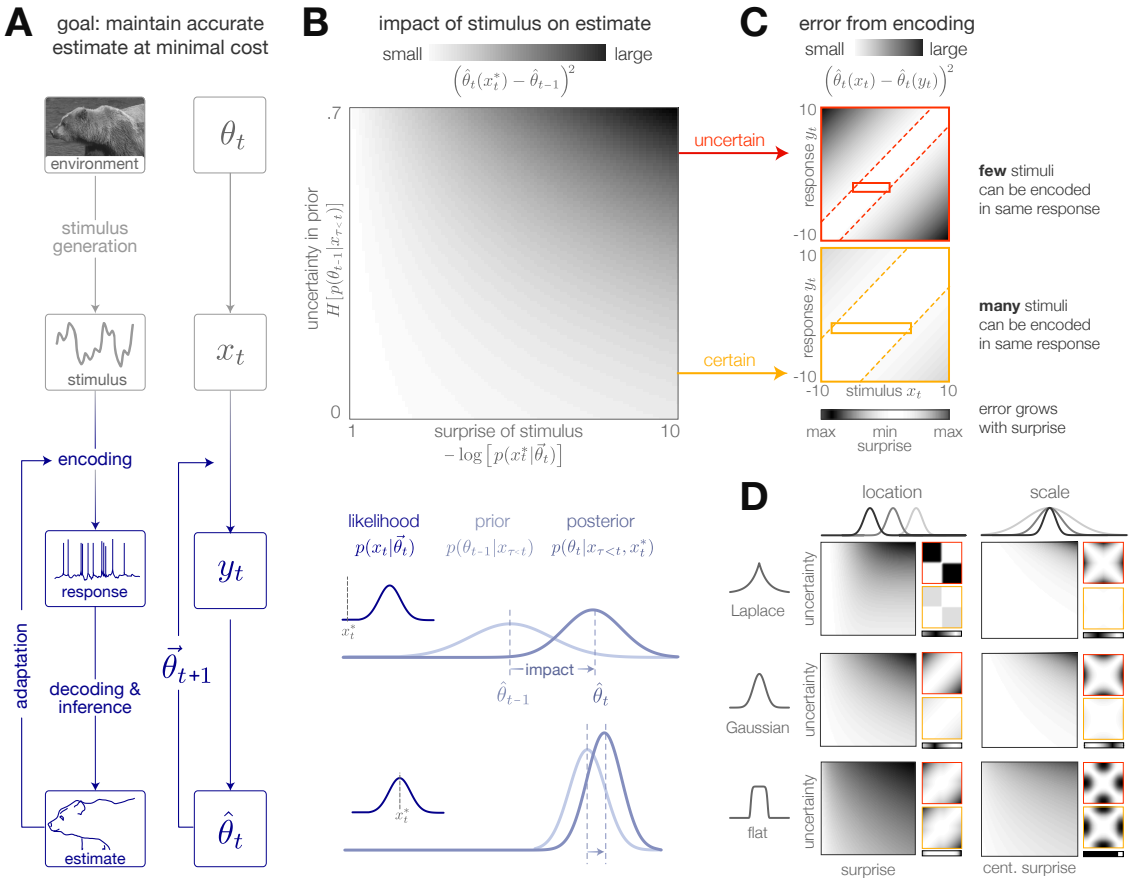

Figure 1: **Surprise and uncertainty determine the impact of incoming stimuli for efficient inference.** **(A)** We consider a framework in which a sensory system infers the state of a dynamic environment at minimal metabolic cost. The state of the environment (θ_t) is signaled by sensory stimuli (x_t) that are encoded in neural responses (y_t). To infer this state, the system must decode stimuli from neural responses and use them to update an internal model of the environment (consisting of an estimate $\hat{\theta}_t$ and a prediction $\vec{\theta}_{t+1}$). This internal model can then be used to adapt the encoding at earlier stages. **(B)** Incoming stimuli can have varying impact on the observer's estimate of the environmental state depending on the relationship between the observer's uncertainty and the surprise of the stimulus (heat map). We use the example of Bayesian estimation of the mean of a stationary Gaussian distribution [33] to demonstrate that when the observer is uncertain (wide prior $p(\theta_{t-1}|x_{\tau<t})$) and the stimulus x_t^* is surprising (x_t^* falls on the edge of the distribution $p(x_t|\vec{\theta}_t)$), the stimulus has high impact and causes a large shift in the posterior $p(\theta_t|x_t^*, x_{\tau<t})$ (schematic, middle row). In contrast, when the observer is certain and the stimulus is expected, the stimulus has a small impact on the observer's estimate (schematic, bottom row). We quantify impact by the squared difference $(\hat{\theta}(x_t^*) - \hat{\theta}_{t-1})^2$ between the estimate before and after incorporating the stimulus (Methods). **(C)** When the observer is certain, a large number of stimuli can be mapped onto the same neural response without inducing error into the observer's estimate (orange panel). When the observer is uncertain, the same mapping from stimulus to response induces higher error (red panel). Error is highest when mapping a surprising stimulus onto an expected neural response, or vice versa. We quantify error by the squared difference $(\hat{\theta}(x_t) - \hat{\theta}_t(y_t))^2$ between the estimate constructed with the stimulus versus the response (Methods). Shown for uncertainty values of 0.1 (orange) and 0.6 (red). **(D)** Qualitatively similar results to those shown in panels B-C are observed for estimating the location and scale of a stationary generalized Gaussian distribution. Stimuli have a larger impact on the observer's estimate when the observer is uncertain and when stimuli are unexpected (quantified by surprise in the case of location estimation, and centered surprise in the case of scale estimation; see main text). The error induced by mapping a stimulus onto a response grows with the surprise of the stimulus. For the case of scale estimation, this error is symmetric to exchanging $+x$ and $-x$, because positive and negative deviations from the mean (taken here to be 0) exert similar influence on the estimation of scale. Shown for the same limits on uncertainty ([0,.7]) and surprise/centered surprise ([1,10]) as in panel B. Stimuli and responses are shown for the intervals [-10,10] (location) and [-20,20] (scale).

Figure 1: **Continued from previous page.** Heatmaps of impact are individually scaled for each stimulus distribution relative to their minimum and maximum values; heatmaps of encoding error are scaled relative to the minimum and maximum error across both uncertainty values for a given stimulus distribution.

In stationary environments with fixed statistics, incoming stimuli can have varying impact on the observer's belief about the state of the environment, depending on the uncertainty in the observer's belief (measured by the entropy of the prior distribution, $H[p(\theta_{t-1}|x_{\tau<t})]$), and on the surprise of a stimulus given this belief (measured by the negative log-likelihood of the stimulus given the current prediction, $-\log p(x_t|\vec{\theta}_t)$. We quantify the impact of a single stimulus x_t^* by measuring the mean-squared difference between the observer's estimate before and after observing the stimulus: $(\hat{\theta}_t(x_t^*) - \hat{\theta}_{t-1})^2$. When the observer is certain about the state of the environment or when a stimulus is consistent with the observer's belief, the stimulus has little impact on the observer's belief (Fig. 1B, illustrated for mean estimation of a stationary Gaussian distribution). Conversely, when the observer is uncertain or when the new observation is surprising, the stimulus has a large impact.

The process of encoding stimuli in neural responses can introduce additional error in the observer's estimate. Some mappings from stimuli onto responses will not alter the observer's estimate, while other mappings can significantly distort this estimate. We measure the error induced by encoding a stimulus x_t in a response y_t using the mean-squared difference between the estimates constructed with each input: $(\hat{\theta}_t(x_t) - \hat{\theta}_t(y_t))^2$. At times when the observer is certain, it is possible to encode many different stimuli in the same neural response without affecting the observer's estimate. However, when the observer is uncertain, some encodings can induce high error, particularly when mapping a surprising stimulus onto an expected neural response, or vice versa. These neural responses can in turn have varying impact on the observer's belief about the state of the environment.

The qualitative features of this relationship between surprise, uncertainty, and the dynamics of inference hold across a range of stimulus distributions and estimation tasks (Fig. 1D). The specific geometry of this relationship depends on the underlying stimulus distribution and the estimated parameter. In some scenarios, surprise alone is not sufficient for determining the utility of a stimulus. For example, when the goal is to infer the spread of a distribution with a fixed mean, a decrease in spread would generate stimuli that are closer to the mean and therefore less surprising than expected. In this case, a simple function of surprise can be used to assess when stimuli are more or less surprising than predicted: $|H[p(x_t|\vec{\theta}_t)] + \log p(x_t|\vec{\theta}_t)|$, where $H[p(x_t|\vec{\theta}_t)]$ is the entropy, or average surprise, of the predicted stimulus distribution. We refer to this as *centered surprise*, which is closely related to the information-theoretic notion of typicality [34].

Together, the relative impact of different stimuli and the error induced by mapping stimuli onto neural responses shape the dynamics of inference. In what follows, we extend this intuition to nonstationary environments, where we show that encoding schemes that are optimized to balance coding cost and inference error exploit these relationships to devote higher coding fidelity at times when the observer is uncertain and stimuli are surprising.

Adaptive coding for inference in nonstationary environments

To make our considerations concrete, we model an optimal Bayesian observer in a two-state environment (Fig. 2A). Despite its simplicity, this model has been used to study the dynamics of inference in neural and perceptual systems, and can generate a range of complex behaviors [11, 35, 12, 36, 37]. Within this model, the state variable θ_t switches randomly between a "low" state ($\theta = \theta^L$) and a "high" state ($\theta = \theta^H$) at a small but fixed hazard rate h (we use $h = 0.01$). We take θ_t to specify either the mean or the standard deviation of a Gaussian stimulus distribution, and we refer to these as "mean-switching" and "variance-switching" environments, respectively. At each point in time, a single stimulus sample x_t is drawn randomly from this distribution. This stimulus is encoded in a neural response and used to update the observer's belief about the environment. For a two-state environment, this belief is fully specified by the posterior probability P_t^L that the environment is in the low state at time t. The predicted posterior is computed based on the likelihood that the environment will switch states in the next timestep: $P_{t+1}^L = P_t^L(1-h) + (1-P_t^L)h$. The posterior can then be used to construct a point estimate of the environmental state at time t: $\hat{\theta}_t = P_t^L\theta^L + (1-P_t^L)\theta^H$ (the point prediction $\vec{\theta}_{t+1}$ can be constructed from the predicted posterior P_{t+1}^L in an analogous manner). For small hazard rates (as considered here), the predicted posterior is very close to the current posterior, and thus the prediction

$\vec{\theta}_{t+1}$ can be approximated by the current estimate $\hat{\theta}_t$. Note that although the environmental states are discrete, the posterior distributions, and the point estimates constructed from them, are continuous (Methods).

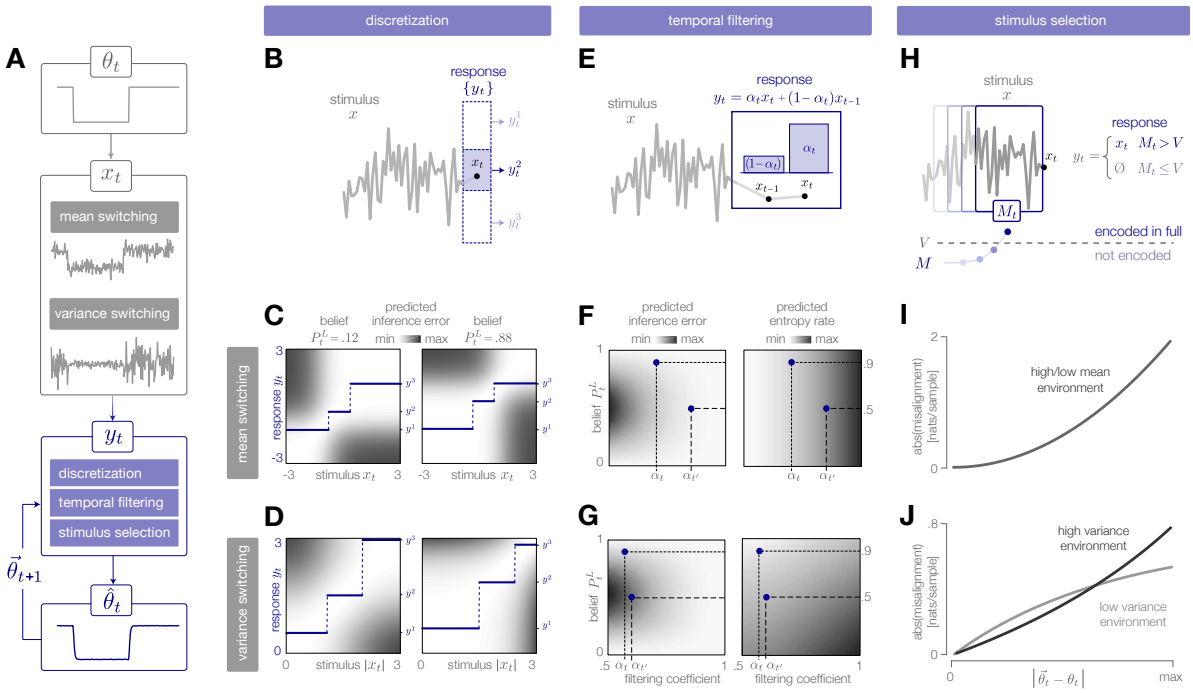

Figure 2: **Adaptive encoding schemes.** **(A)** We consider a specific implementation of our general framework in which an environmental state θ_t switches between two values with fixed probability. This state parameterizes the mean or variance of a Gaussian stimulus distribution. Stimuli x_t are drawn from this distribution and encoded in neural responses y_t. We consider three encoding schemes that perform discretization (panels B-D), temporal filtering (panels E-G), or stimulus selection (panels H-J) on incoming stimuli. **(B)** (Schematic) At each timestep, an incoming stimulus x_t (black dot) is mapped onto a discrete neural response level y_t^i (solid blue rectangle) chosen from a set $\{y_t^i\}$ (dotted rectangles). **(C-D)** The predicted inference error induced by mapping a stimulus x_t onto a neural response y_t varies as a function of the observer's belief P_t^L about the state of the environment (shown for $P_t^L = 0.12$, left column; $P_t^L = 0.88$, right column). At each timestep, the optimal response levels $\{y^1, y^2, y^3\}$ (solid lines) are chosen to minimize this error when averaged over the predicted stimulus distribution. **(E)** (Schematic) At each timestep, incoming stimuli are combined via a linear filter with a coefficient α_t. **(F-G)** The average predicted inference error (left column) depends on the filter coefficient α_t and on the observer's belief P_t^L about the state of the environment. At each timestep, the optimal filter coefficient (blue dot) is found by balancing error and entropy given a prediction of the environmental state (α_t and $\alpha_{t'}$ are shown for $P_t^L = .9$ and $P_{t'}^L = .5$, respectively). **(H)** (Schematic) At each timestep, the encoder computes the misalignment M_t between the predicted and measured surprise of incoming stimuli. If the misalignment exceeds a threshold V, the stimulus is encoded with perfect fidelity; otherwise, the stimulus is not encoded. **(I-J)** The misalignment signal (computed here analytically; see Methods) depends on the relationship between the predicted and true state of the environment. When the mean is changing over time (panel I), the misalignment depends only on the absolute difference between the true and predicted mean. When the variance is changing over time (panel J), the misalignment also depends on the true variance of the environment.

We consider three neurally-plausible encoding schemes that reflect limitations in representational capacity. In one scheme, the encoder is constrained in the total number of distinct responses it can produce at a given time, and uses a discrete set of neural response levels to represent a stimulus ('discretization'; Fig. 2B-D). In second scheme, the encoder is constrained in dynamic range and temporal acuity, and filters incoming stimuli in time ('temporal filtering'; Fig. 2E-G). Finally, we consider an encoder that is constrained in the total amount of activity that can be used to encode a stimulus, and must therefore selectively encode certain stimuli and not others ('stimulus selection'; Fig. 2H-J). For

each scheme, we impose a global constraint that controls the maximum fidelity of the encoding. We then adapt the instantaneous fidelity of the encoding subject to this global constraint. We do so by choosing the parameters of the encoding to minimize the *error in inference*, $(\hat{\theta}_t(x_t) - \hat{\theta}_t(y_t))^2$, when averaged over the predicted distribution of stimuli, $p(x_t|\vec{\theta}_t)$. (In what follows, we will use $\hat{\theta}_t$ and $\vec{\theta}_{t+1}$ to denote the estimates and predictions constructed from the neural response y_t. When differentiating between $\hat{\theta}_t(x_t)$ and $\hat{\theta}_t(y_t)$, we will use the shorthand notation $\hat{\theta}_{x,t}$ and $\hat{\theta}_{y,t}$, respectively). We compare this minimization to one in which the goal is to reconstruct the stimulus itself; in this case, the *error in reconstruction* is given by $(x_t - y_t)^2$. In both cases, the goal of minimizing error (in either inference or reconstruction) is balanced with the goal of minimizing metabolic cost. Because the encoding is optimized based on the internal prediction of the environmental state, the entropy of the neural response will depend on how closely this prediction aligns with the true state of the environment. The entropy specifies the minimal number of bits required to accurately represent the neural response [34], and becomes a lower bound on energy expenditure if each bit requires a fixed metabolic cost [38]. We therefore use the entropy of the response as a general measure of the metabolic cost of encoding.

We expect efficient encoding schemes to operate on uncertainty and surprise. The observer's uncertainty, given by $H[P_t^L] = P_t^L\theta^L + (1 - P_t^L)\theta^H$, is largest when the posterior is near 0.5, and the observer believes that the environment is equally likely to be in either state. The degree to which incoming stimuli are surprising depends on the entropy of the stimulus distribution, and on the alignment between this distribution and the observer's belief. When the mean of the Gaussian distribution is changing in time, the entropy is constant, and surprise depends symmetrically on the squared difference between the true and predicted mean, $(\mu - \vec{\mu})^2$. When the variance is changing, the entropy is also changing in time, and centered surprise depends asymmetrically on the ratio of true and predicted variances, $\sigma^2/\vec{\sigma}^2$. As a result, encoding strategies that rely on stimulus surprise should be symmetric to changes in mean but asymmetric to changes in variance.

To illustrate the dynamic relationship between encoding and inference, we use a "probe" environment that switches between two states at fixed intervals of $1/h$ timesteps. This specific instantiation is not unlikely given the observer's model of the environment [11], and allows us to illustrate average behaviors over many cycles of the environment.

Encoding via discretization

Neurons use precise sequences of spikes [39] or discrete firing rate levels [26, 40] to represent continuous stimuli. This inherent discreteness imposes a fundamental limitation on the number of distinct neural responses that can be used to represent a continuous stimulus space. Many studies have argued that sensory neurons make efficient use of limited response levels by appropriately tuning these levels to match the steady-state distribution of incoming stimuli (e.g. [26, 41]).

Here, we consider an encoder that adaptively maps an incoming stimulus x_t onto a discrete set of neural response levels $\{y_t^i\}$ (Fig. 2B). Because there are many more stimuli than levels, each level must be used to represent multiple stimuli. The number of levels reflects a global constraint on representational capacity; fewer levels indicates a stronger constraint and results in a lower-fidelity encoding.

The encoder can adapt this mapping by expanding, contracting, and shifting the response levels to devote higher fidelity to different regions of the stimulus space. We consider an optimal strategy in which the response levels are chosen at each timestep to minimize the predicted inference error, subject to a constraint on the number of levels:

$$\underbrace{\left\langle \left(\hat{\theta}_{x,t} - \hat{\theta}_{y,t}\right)^2 \right\rangle_{p(x_t|\vec{\theta}_{y,t})}}_{\substack{\text{predicted} \\ \text{inference error}}} \tag{1}$$

When the mean of the stimulus distribution is changing over time, we define these levels with respect to the raw stimulus value x_t. When the variance is changing, we define these levels with respect to the absolute deviation from the mean, $|x_t - \mu|$ (where we take $\mu = 0$). The predicted inference error induced by encoding a stimulus x_t in a response y_t changes over time as a function of the observer's prediction of the environmental state (Fig. 2C-D). Because some stimuli have very little effect on the estimate at a given time, they can be mapped onto the same neural response level without inducing error in the estimate (white regions in Fig. 2C-D). The optimal response levels are chosen to minimize this error when averaged over the predicted distribution of stimuli.

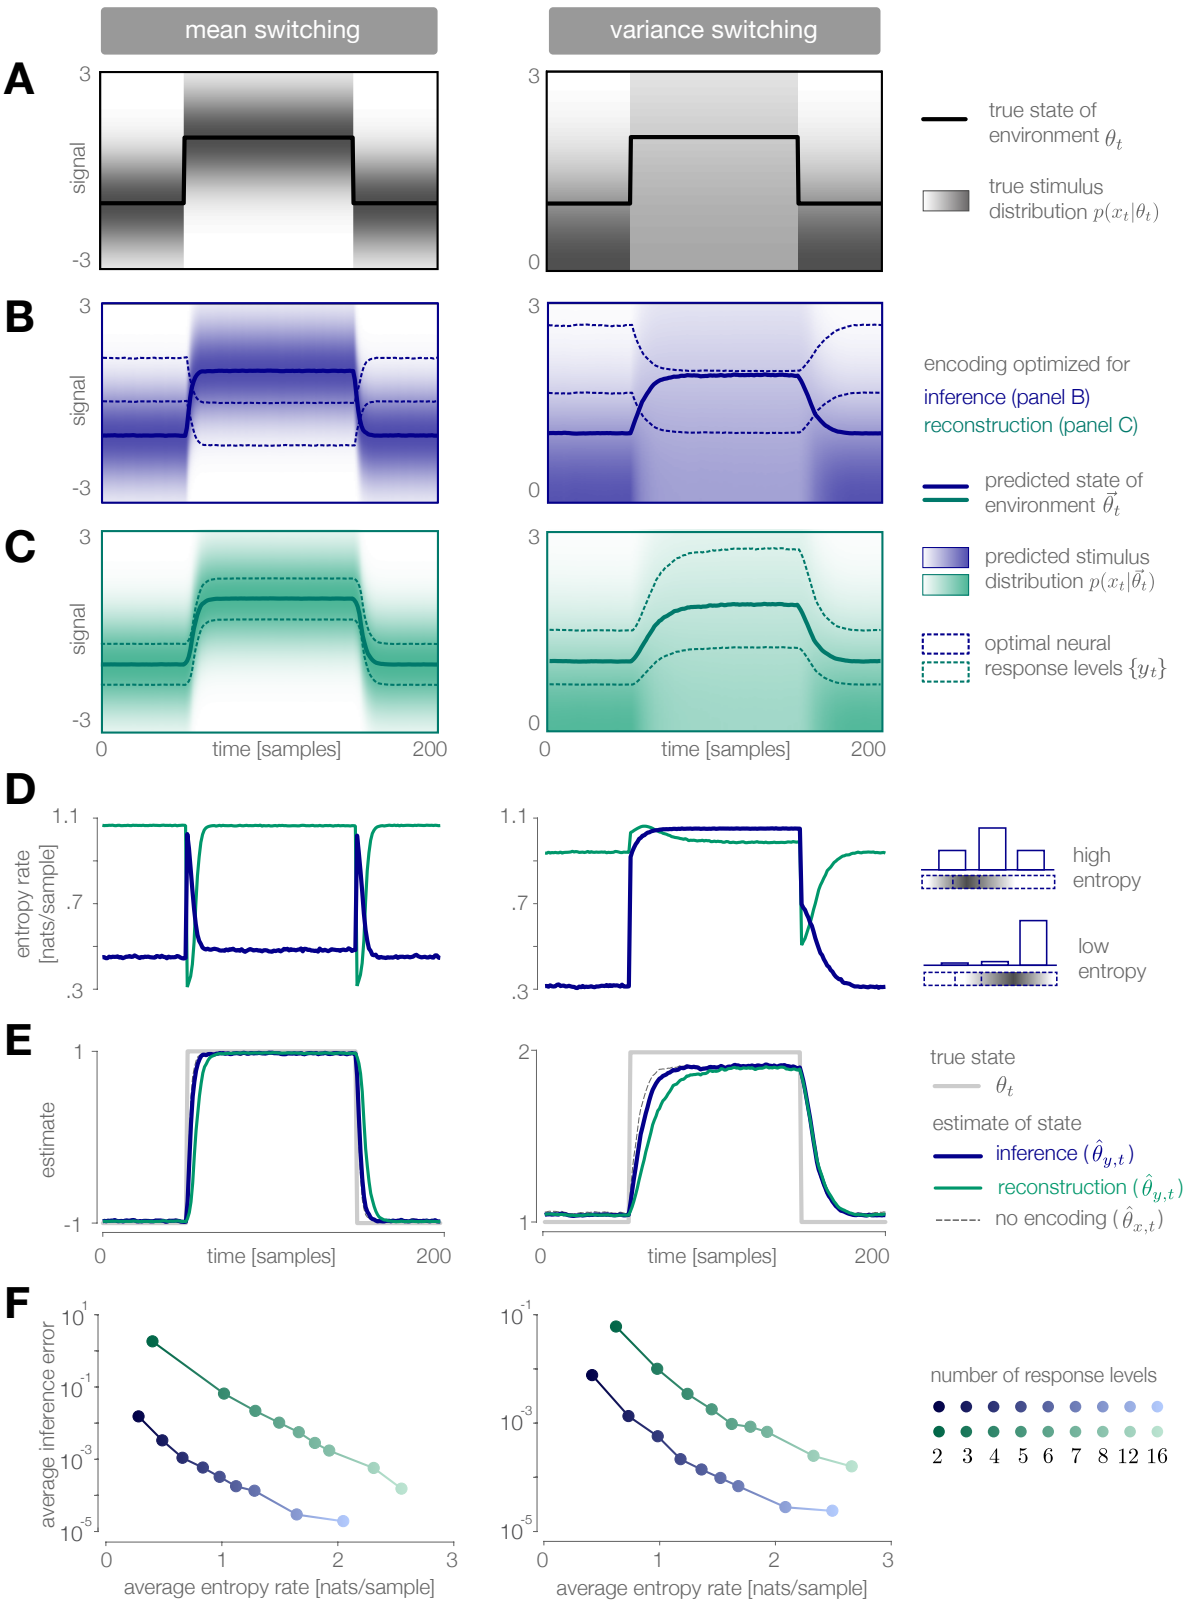

Figure 3: **Dynamic inference with optimally-adapted response levels. (A)** We consider a probe environment in which a state θ_t (solid line) switches between two values at fixed time intervals. This state parametrizes the mean (left) or the variance (right) of a Gaussian stimulus distribution (heatmap). **(B, C)** Optimal response levels (dotted lines) are chosen to minimize error in inference (blue) or stimulus reconstruction (green) based on the predicted stimulus distribution $p(x_t|\vec{\theta}_t)$ (heatmap). Results are shown for three response levels. All probability distributions in panels A-C are scaled to the same range, $[0, .4]$. **(B)** Response levels optimized for inference devote higher resolution (narrower levels) to stimuli that are surprising given the current prediction of the environment. **(C)** Response levels optimized for stimulus reconstruction devote higher resolution to stimuli that are likely.

Figure 3: **Continued from previous page. (D)** The entropy rate of the encoding is found by partitioning the true stimulus distribution (heatmap in panel A) based on the optimal response levels (dotted lines in panels B-C). Abrupt changes in the environment induce large changes in entropy rate that are symmetric for mean estimation (left) but asymmetric for variance estimation (right). Apparent differences in the baseline entropy rate for low- versus high-mean states arise from numerical instabilities. **(E)** Encoding induces error in the estimate $\hat{\theta}_t$. Errors are larger if the encoding is optimized for stimulus reconstruction than for inference. The error induced by upward and downward switches is symmetric for mean estimation (left) but asymmetric for variance estimation (right). In the latter case, errors are larger when inferring upward switches in variance. **(F)** Increasing the number of response levels decreases the average inference error but increases the cost of encoding. Across all numbers of response levels, an encoding optimized for inference (blue) achieves lower error at lower cost than an encoding optimized for stimulus reconstruction (green). All results in panels A-C and E are averaged over 500 cycles of the probe environment. Results in panel D were computed using the average response levels shown in panels B-C. Results in panel F were determined by computing time-averages of the results in panels D-E.

The relative width of each level is a measure of the resolution devoted to different regions of the stimulus space; narrower levels devote higher resolution (and thus higher fidelity) to the corresponding regions of the stimulus space. The output of these response levels is determined by their alignment with the true stimulus distribution. An encoding that devotes higher resolution to stimuli that are likely to occur in the environment will produce a higher entropy rate (and thus higher cost), because many different response levels will be used with relatively high frequency. In contrast, if an encoding scheme devotes high resolution to surprising stimuli, very few response levels will be used, and the resulting entropy rates will be low.

When designed for accurate inference, we find that the optimal encoder devotes its resolution to stimuli that are surprising given the current prediction of the environment (Fig. 3B). In a mean-switching environment (left column of Fig. 3), stimuli that have high surprise fall within the tails of the predicted stimulus distribution. As a result, when the observer's prediction is accurate, the bulk of the stimulus distribution is mapped onto the same response level (Fig. 3B, left), and entropy rates are low (blue curve in Fig. 3D, left). When the environment changes abruptly, the bulk of the new stimulus distribution is mapped onto different response levels. This results in a large spike in entropy rate, which enables the observer to quickly adapt its estimate to the change (blue curve in Fig. 3E, left).

In a variance-switching environment (right column of Fig. 3), stimuli that have high centered surprise fall either within the tails of the predicted stimulus distribution (when variance is low), or within the bulk (when variance is high). As a result, entropy rates are low in the low-variance state, but remain high during the high-variance state (blue curve in Fig. 3D, right).

When designed for accurate reconstruction of the stimulus, we find that the optimal encoder devotes its resolution to stimuli that are likely given the current prediction of the environmental state (Fig. 3C). As a result, entropy rates are high when the observer's prediction is accurate, regardless of the environment (green curves in Fig. 3D). Entropy rates drop when the environment changes, because likely stimuli become mapped onto the same response level. This drop slows the observer's detection of changes in the environment (green curve in Fig. 3E, left). An exception occurs when the variance abruptly increases, because likely stimuli are still given high resolution by the encoder following the change in the environment.

Whether optimizing for inference or stimulus reconstruction, the entropy rate, and thus the coding cost, changes dynamically over time in a manner that is tightly coupled with the inference error. The average inference error can be reduced by increasing the number of response levels, but this induces a higher average coding cost (Fig. 3F). As expected, a strategy optimized for inference achieves lower inference error than a strategy optimized for stimulus reconstruction (across all numbers of response levels), but it also does so at significantly lower coding cost.

Encoding via temporal filtering

Neural responses have limited gain and temporal acuity, a feature that is often captured by linear filters. For example, neural receptive fields are often characterized as linear temporal filters, sometimes followed by a nonlinearity [42, 39]. The properties of these filters are known to dynamically adapt to changing stimulus statistics (e.g. [43, 44]), and numerous theoretical studies have suggested that such filters are adapted to maximize the amount of information that is encoded about the stimulus [24, 45].

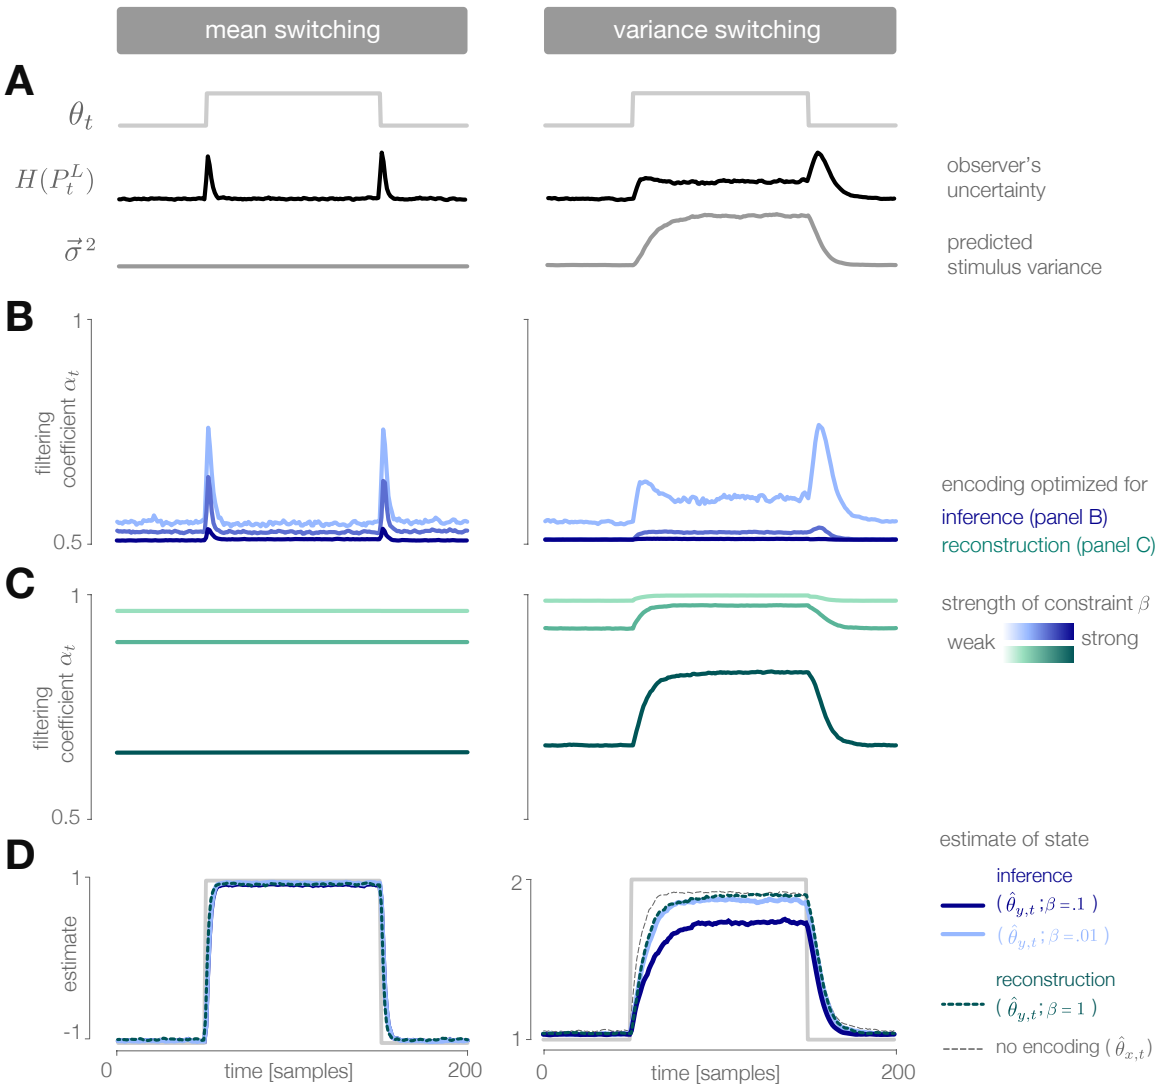

Figure 4: **Dynamic inference with optimally-adapted temporal filters.** **(A)** The observer's uncertainty ($H[P_t^L]$) is largest when the environment is changing. The predicted stimulus variance (a proxy for both the predicted magnitude of the stimulus distribution, and the predicted surprise of incoming stimuli) is constant in a mean-switching environment (left) but variable in a variance-switching environment (right) (computed using a filter coefficient optimized for inference with a weak entropy constraint, corresponding to the lightest blue curves in panel B). **(B, C)** Optimal values of the filter coefficient α_t are chosen at each timestep to minimize error in inference (blue) or stimulus reconstruction (green), subject to a constraint on predicted entropy. Darker colors indicate stronger constraints. **(B)** Filters optimized for inference devote high fidelity at times when the observer is uncertain and stimuli are predicted to be surprising. Shown for $\beta = 0.02, 0.1, 1$ (left) and $\beta = 0.01, 0.1, 1$ (right). **(C)** Filters optimized for reconstruction devote fidelity at times when the magnitude of the stimulus is predicted to be high. Shown for $\beta = 0.01, 0.1, 1$. **(D)** Filtering induces error into the estimate $\hat{\theta}_t$. Strong filtering has minimal impact on mean estimation (left), but induces large errors in the estimation of high variances (right). All results in panels A-D are averaged over 800 cycles of the probe environment.

Here, we consider an encoder that implements a very simple temporal filter (Fig. 2E):

$$y_t = \alpha_t x_t + (1 - \alpha_t) x_{t-1} \tag{2}$$

where $\alpha_t \in [0.5, 1]$ is a coefficient that specifies the shape of the filter and controls the instantaneous fidelity of the encoding. When $\alpha_t = 0.5$, the encoder computes the average of current and previous

stimuli by combining them with equal weighting, and the fidelity is minimal. When $\alpha_t = 1$, the encoder transmits the current stimulus with perfect fidelity (i.e., $y_t = x_t$). In addition to introducing temporal correlations, the filtering coefficient changes the gain of the response y_t by rescaling the inputs $\{x_t, x_{t-1}\}$.

The encoder can adapt α_t in order to manipulate the instantaneous fidelity of the encoding (Fig. 2E). We again consider an optimal strategy in which the value of α_t is chosen at each timestep to minimize the predicted inference error, subject to a constraint on the predicted entropy rate of the encoding:

$$\underbrace{\left\langle \left(\hat{\theta}_{x,t} - \hat{\theta}_{y,t}\right)^2 \right\rangle_{p(x_t|\vec{\theta}_{y_t})}}_{\substack{\text{predicted}\\ \text{inference error}}} + \underbrace{\beta H(y_t, y_{t+1})}_{\substack{\text{predicted}\\ \text{entropy rate}}} \tag{3}$$

Both terms depend on the strength of averaging α_t and on the observer's belief P_t^L about the state of the environment (Fig. 2F-G). The inference error depends on belief through the observer's uncertainty; when the observer is uncertain, strong averaging yields low fidelity representation. When the observer is certain, however, incoming stimuli can be strongly averaged without impacting the observer's estimate. The entropy rate depends on belief through the predicted entropy rate (variance) of the stimulus distribution; when the predicted entropy rate is high, incoming stimuli are more surprising on average. The multiplier β reflects a global constraint on representational capacity; larger values of β correspond to stronger constraints and reduce the maximum fidelity of the encoding. This, in turn, results in a reduction in coding fidelity through a decrease in gain and an increase in temporal correlation.

When designed for accurate inference, we find that the optimal encoder devotes higher fidelity at times when the observer is uncertain and the predicted stimulus variance is high. In a mean-switching environment, the stimulus variance is fixed (Fig. 4A, left), and thus the fidelity depends only on the observer's uncertainty. This uncertainty grows rapidly following a change in the environment, which results in a transient increase in coding fidelity (Fig. 4B, left) and a rapid adaptation of the observer's estimate (Fig. 4D, left). This estimate is highly robust to the strength of the entropy constraint; even when incoming stimuli are strongly averaged ($\alpha_t = 0.5$), the encoder transmits the mean of two consecutive samples, which is precisely the statistic that the observer is trying to estimate.

In a variance-switching environment, the predicted stimulus variance also changes in time (Fig. 4A, right). This results in an additional increase in fidelity when the environment is in the high- versus low-variance state, and an asymmetry between the filter responses for downward versus upward switches in variance (Fig. 4B, right). Both the encoder and the observer are slower to respond to changes in variance than to changes in mean, and the accuracy of the inference is more sensitive to the strength of the entropy constraint (Fig. 4D, right).

When designed to accurately reconstruct the stimulus, the fidelity of the optimal encoder depends only on the predicted stimulus variance. In a mean-switching environment, the variance is fixed (Fig. 4A), and thus the fidelity is flat across time. In a variance-switching environment, the fidelity increases with the predicted variance of incoming stimuli, not because variable stimuli are more surprising, but rather because they are larger in magnitude and can lead to higher errors in reconstruction (Fig. 4C). As the strength of the entropy constraint increases, the encoder devotes proportionally higher fidelity to high-variance stimuli because they have a greater impact on reconstruction error.

Encoding via stimulus selection

Sensory neurons show sparse activation during natural stimulation [46, 47, 48], an observation that is often interpreted as a signature of coding cost minimization [49, 38]. In particular, early and intermediate sensory neurons may act as gating filters, selectively encoding only highly-informative features of the stimulus [50, 51]. Such a selection strategy reduces the number of spikes transmitted downstream.

Here, we consider an encoder that selectively transmits only those stimuli that are surprising and are therefore likely to change the observer's belief about the state of the environment. When the observer's prediction is inaccurate, the predicted average surprise $H[p(x_t|\vec{\theta}_t)]$ will differ from the true average surprise $H[p(x_t|\theta_t)]$ by an amount equal to the KL-divergence of the predicted from the true stimulus distributions (Methods). In principle, this difference could be used to selectively encode stimuli at times when the observer's estimate is inaccurate.

In practice, however, the encoder does not have access to the entropy of the true stimulus distribution. Instead, it must measure surprise directly from incoming stimulus samples. The measured surprise of each incoming stimulus sample is given by its negative log-likelihood, $-\log[p(x_t|\vec{\theta}_t)]$. We consider an

encoder that compares the predicted surprise to a running average of the measured surprise. In this way, the encoder can heuristically assess whether a change in the stimulus distribution had occurred by computing the "misalignment" M_t between the predicted and measured stimulus distributions:

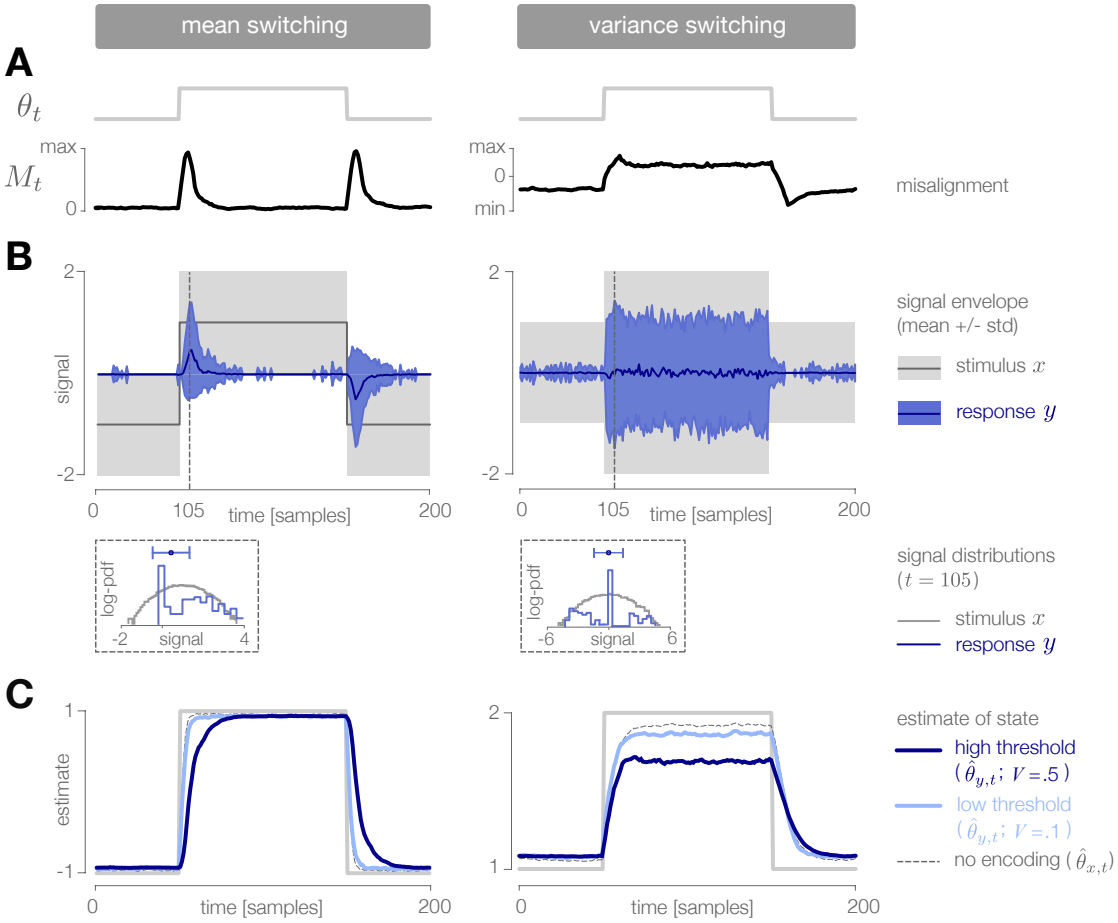

Figure 5: **Dynamic inference with stimulus selection. (A)** When the environment is changing, the observer's prediction is misaligned with the state of the environment. When this misalignment M_t is large, stimuli are transmitted in full ($y_t = x_t$). When this misalignment falls below a threshold V, stimuli are not transmitted at all ($y_t = 0$). **(B)** The distribution of encoded stimuli changes over time, as can be seen by comparing the envelope of the stimulus distribution (gray) with the envelope of the neural responses (blue). Left: When the mean of the stimulus distribution changes abruptly, a large proportion of stimuli are encoded, and the mean of the neural response (blue line) approaches the mean of the stimulus distribution (black line). At times when the mean of the stimulus distribution is stable, very few stimuli are encoded, and the mean of the neural response drops to zero. Right: When the variance is low, very few stimuli are encoded. When the variance increases, the average surprise of incoming stimuli increases, and a large proportion of stimuli are encoded. The envelope of the neural response expands and approaches the envelope of the stimulus distribution. Insets: At times when the environment is changing (shown for $t = 105$), the distribution of responses (blue) is sparser than the distribution of stimuli (gray), due to the large proportion of stimuli that are not encoded (indicated by the large peak in probability mass at 0). Shown for $V = 0.5$. **(C)** Higher thresholds slow the observer's detection of changes in the mean (left), and cause the observer to underestimate high variances (right). Threshold values are scaled relative to the maximum analytical value of the misalignment signal in the mean- and variance-switching environment (shown in Fig. 2I and J, respectively). Results in panels B and C are averaged over 800 cycles of the probe environment.

$$M_t = H[p(x_t|\vec{\theta}_t)] + \frac{1}{T}\sum_{\tau=0}^{T} \log p(x_{t-\tau}|\vec{\theta}_t) \tag{4}$$

The misalignment is computed over a time window T, which ensures that the observer's prediction does not gradually drift from the true value in cases where surprising stimuli are not indicative of a change in the underlying stimulus distribution (we use $T = 10$). Because the misalignment signal is directly related to the surprise of incoming stimuli, it is symmetric to upward and downward switches in the mean of the stimulus distribution, but it is asymmetric to switches in variance and has a larger magnitude in the high-variance state (shown analytically in Fig. 2I-J).

The misalignment signal is both non-stationary and non-Gaussian. Optimizing an encoding scheme based on this signal would require deriving the corresponding optimal observer model, which is difficult to compute in the general case. We instead propose a heuristic (albeit sub-optimal) solution, in which the encoder selectively encodes the current stimulus with perfect fidelity ($y_t = x_t$) when recent stimuli are sufficiently surprising and the magnitude of the misalignment signal exceeds a threshold V (Fig. 2H). When the magnitude of the misalignment signal falls below the threshold, stimuli are not encoded ($y_t = \emptyset$). At these times, the observer does not receive any information about incoming stimuli, and instead marginalizes over its internal prediction to update its estimate (Methods). The value of the threshold reflects a constraint on overall activity; higher thresholds result in stronger criteria for stimulus selection, which decreases the maximum fidelity of the encoding.

When the mean of the stimulus distribution changes in time, very few stimuli are required to maintain an accurate estimate of the environmental state (Fig. 5A-B, left). When the environment changes abruptly, the observer's prediction is no longer aligned with the environment, and the misalignment signal increases until incoming stimuli are encoded and used to adapt the observer's prediction. Because it requires several stimulus samples for the misalignment to exceed threshold, there is a delay between the switch in the environment and the burst of encoded stimuli. This delay, which is proportional to the size of the threshold, slows the observer's detection of the change (Fig. 5C, left).

When the variance changes in time, the average surprise of incoming stimuli also changes in time. When the variance abruptly increases, the misalignment signal grows both because the observer's prediction is no longer accurate, and because the average surprise of the incoming stimulus distribution increases. A large proportion of stimuli are transmitted, and the observer quickly adapts to the change. If the threshold is sufficiently high, however, the observer's prediction never fully aligns with the true state. When the variance abruptly decreases, the incoming stimulus distribution is less surprising on average, and therefore a greater number of stimulus samples is needed before the misalignment signal exceeds threshold. As a result, observer is slower to detect decreases in variance than increases (Fig. 5C, right).

Dynamical signatures of adaptive coding

The preceding sections examined the dynamics of optimal encoding strategies as seen through the internal parameters of the encoder itself. The alignment between these internal parameters and the external dynamics of the environment determine the output response properties of each encoder. It is these output response properties that would give experimental access to the underlying encoding scheme, and that could potentially be used to distinguish an encoding scheme optimized for inference from one optimized for stimulus reconstruction.

To illustrate this, we simulate output responses of each encoder to repeated presentations of the probe environment. In the case of discretization, we use a simple entropy coding procedure to map each of four response levels to four spike patterns ([00], [01], [10], [11]) based on the predicted probability that each response level will be used, and we report properties of the estimated spike rate (see spike rasters in Fig. 6A; Methods). In the cases of filtering and stimulus selection, we report properties of the response y_t.

We find that encodings optimized for inference typically show transient changes in neural response properties after a switch in the environment, followed by a return to baseline. This is manifested in a burst in firing rates in the case of discretization, and a burst in response variability in the cases of filtering and stimulus selection. Filtering is additionally marked by a transient decrease in the temporal correlation of the response. The magnitude of these transient changes relative to baseline is most apparent in the case of mean estimation, where the variability in the environment remains fixed over time. Because periods of higher variability in the environment are intrinsically more surprising, baseline response properties change during variance estimation, and bursts relative to baseline are less pronounced. Nevertheless,

we see a transient decrease in temporal correlation in the case of filtering, and a transient increase in response variability in the case of stimulus selection, following switches in variance.

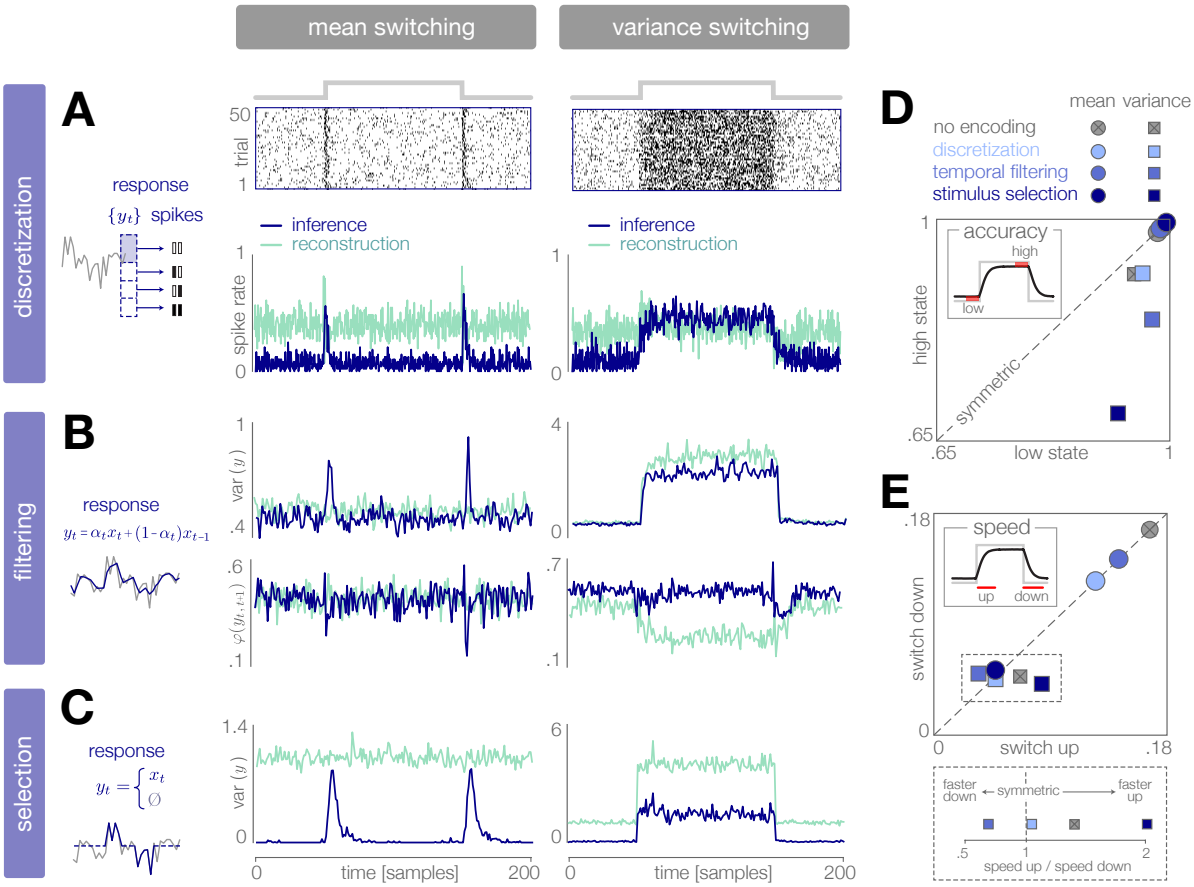

Figure 6: **Dynamical signatures of adaptive coding schemes. (A-C)** We simulate the output of each encoder to repeated cycles of the probe environment. In the case of discretization (panel A), we use a simple entropy coding procedure to map optimal response levels $\{y_t\}$ onto spike patterns, as shown by the spike rasters. In the case of temporal filtering (panel B) and stimulus selection (panel C), we measure properties of the response y_t. When encodings are optimized for inference (dark blue traces), abrupt changes in the mean of the stimulus distribution (panels A-C, left) are followed by transient increases in spike rate (discretization, panel A) and response variability (filtering, panel B; stimulus selection, panel C). In the case of temporal filtering, these changes are additionally marked by decreases in the temporal correlation $\varphi(y_{t,t+1})$ of the response. In contrast, the response properties of encoders optimized for stimulus reconstruction (light green traces) remain more constant over time. Abrupt changes in variance (panels A-C, right) are marked by changes in baseline response properties. Responses show transient deviations away from baseline when encodings are optimized for inference, but remain fixed at baseline when encodings are optimized for reconstruction. In all cases, encodings optimized for inference maintain lower baseline firing rates, lower baseline variability, and higher baseline correlation than encodings optimized for stimulus reconstruction. Spike rates (panel A) are averaged over 500 cycles of the probe environment. Response variability (panels B-C) is computed at each timepoint across 800 cycles of the probe environment. Temporal correlation $\varphi(y_{t,t+1})$ (panel B) is computed between consecutive timepoints across 800 cycles of the probe environment. **(D-E)** Encoding schemes impact both the accuracy (panel D) and speed (panel E) of inference. In all cases, the dynamics of inference are symmetric for changes in mean (points lie along the diagonal) but asymmetric for changes in variance (points lie off the diagonal). Encodings decrease the accuracy of estimating high-variance states (panel D), and they alter the speed of responding to changes in both mean and variance. The response to upward versus downward switches (dotted box) separates encoding schemes based on whether they are faster (right of dotted vertical line) or slower (left of dotted vertical line) to respond to increases versus decreases in variance. Speed and accuracy are measured from the trial-averaged trajectories of $\hat{\theta}_{y,t}$ shown in Fig. 3E, 4D ($\beta = 0.01$), and 5C ($V = 0.5$) (Methods).

The same dynamical features are not observed in encoders optimized for stimulus reconstruction. For mean estimation, firing rates and response variability remain nearly constant over time, despite abrupt changes in the mean of the stimulus distribution. Discretization shows a brief rise and dip in firing rate following a switch, which has been observed experimentally [31]. For variance estimation, response properties show sustained (rather than transient) changes following a switch.

Differences in response properties are tightly coupled to the speed and accuracy of inference, as mediated by the feedforward and feedback interactions between the encoder and the observer. Note that these measures of speed and accuracy (as well as the comparisons made in Fig. 3E, 4D, and 5C) intrinsically favor encodings optimized for inference; we therefore restrict our comparison to this set of encodings. We find that both the speed and accuracy of inference are symmetric to changes in the mean of the stimulus distribution, but asymmetric to changes in variance. This is qualitatively consistent with the optimal Bayesian observer in the absence of encoding [11]. We find that encoding schemes optimized for inference have a more significant impact on the speed and accuracy of variance estimation than of mean estimation. Interestingly, the speed of variance adaptation deviates from optimality in a manner that could potentially be used to distinguish between encoding strategies. In the absence of encoding, the ideal observer is faster to respond to increases than to decreases in variance. We find that encoding via stimulus selection increases this asymmetry, encoding via discretization nearly removes this asymmetry, and encoding via stimulus selection reverses this asymmetry.

Together, these observations suggest that both the dynamics of the neural response and the patterns of deviation from optimal inference could be used to infer features of the underlying sensory coding scheme. Moreover, these results suggest that an efficient system could prioritize some encoding schemes over others, depending on whether the goal is to reconstruct the stimulus or infer its underlying properties, and if the latter, whether this goal hinges on speed, accuracy, or both.

Adaptive coding for inference under natural conditions

The simplified task used in previous sections allowed us to explore the dynamic interplay between encoding and inference. To illustrate how this behavior might generalize to more naturalistic settings, we consider a visual inference task with natural stimuli (Fig. 7A, Methods). As before, the goal of the system is to infer a change in the statistics of the environment from incoming sensory stimuli. We consider a sequence of image patch stimuli drawn randomly from a local region of a natural image; this sequence could be determined by, e.g., saccadic fixations. Each image patch is encoded in the responses of a population of sensory neurons using a well-known sparse-coding model [25]. After adapting to natural stimulus statistics, the basis functions of each model neuron resemble receptive fields of simple cells in V1. A downstream observer decodes the stimulus from this population response and normalizes its contrast. The contrast-normalized patch is then projected onto a set of curvature filters. The variance in the output of these filters is used as an estimate of the underlying statistics of the image region. Both the computation of local image statistics and visual sensitivity to curvature are known to occur in V2 [52, 53, 54].

The encoder reconstructs each stimulus subject to a sparsity constraint λ; large values of λ decrease the population activity at the cost of reconstruction accuracy (see Supplementary Fig. 10). In contrast to the encoding models discussed previously, this encoder is explicitly optimized to reconstruct each stimulus, rather than to support accurate inference. Even in this scenario, however, the observer can manipulate the sparsity of the population response to decrease resource use while maintaining an accurate estimate of the environmental state. It has been proposed that early sensory areas, such as V1, could manipulate the use of metabolic resources depending on top-down task demands (e.g. [55]).

We model a change in the stimulus distribution by a gaze shift from one region of the image to another (Fig. 7B). This shift induces an increase in the variance of curvature filters. Following this change, the observer must update its estimate of local curvature using image patches drawn from the new image region. We empirically estimated the impact of stimulus surprise and observer uncertainty on this estimation and found it to be consistent with results based on model environments (Fig. 7D; compare with Fig. 1B). Surprising stimuli that project strongly on curvature filters exert a large impact on inference, while expected stimuli (characterized by low centered-surprise) exert little impact (Fig. 7C-E). Similarly, individual stimuli exert a larger impact on the estimate when the observer is uncertain than when the observer is certain (Fig. 7C-D,F).

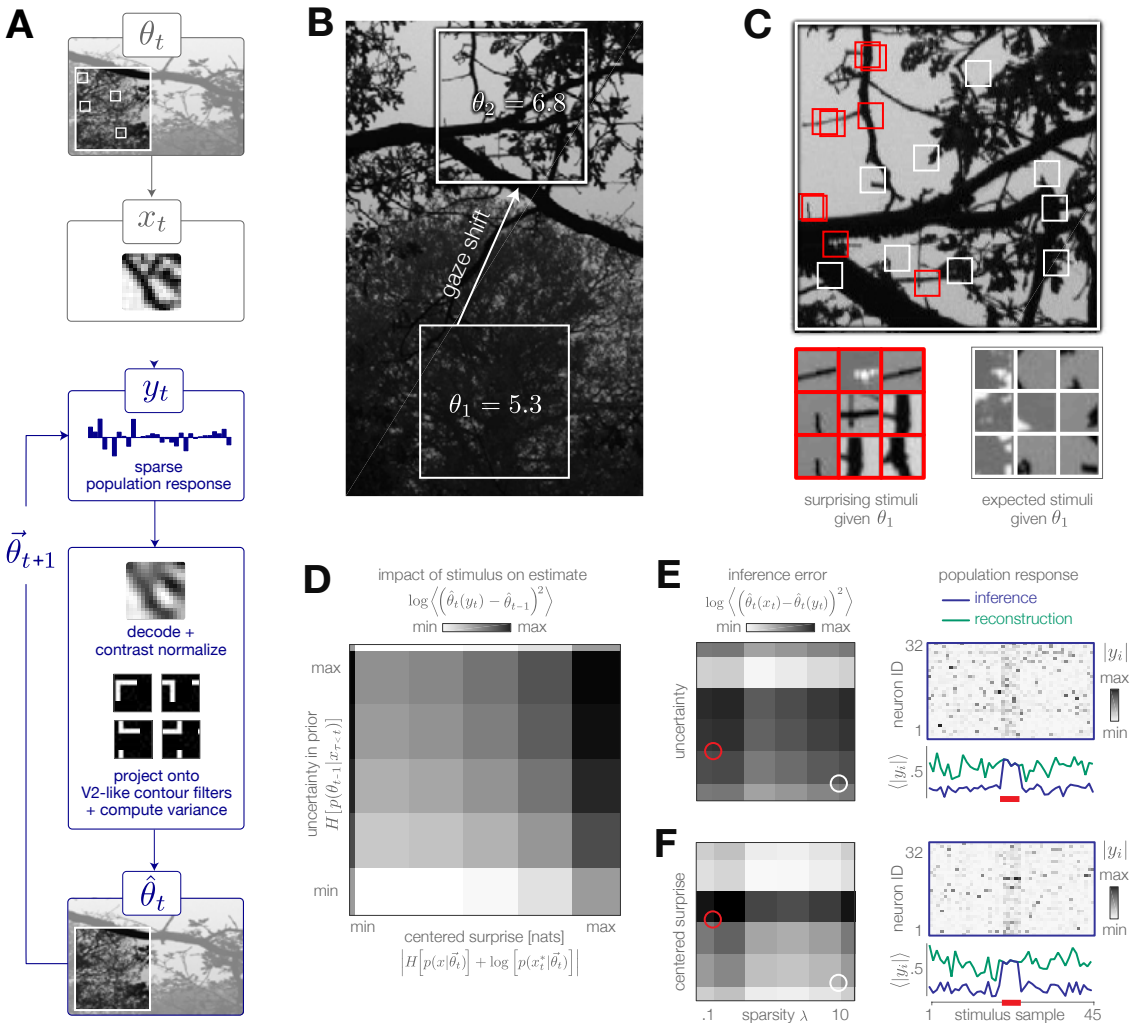

Figure 7: **Model inference task with natural stimuli.** **(A)** We model a simple task of inferring the variance of local curvature in a region of an image. The system encodes randomly-drawn image patches that model saccadic fixations. Individual image patches are encoded in sparse population activity via V1-like receptive fields (see Supplementary Fig. 10). Image patches are then decoded from the population activity, contrast-normalized, and projected onto V2-like curvature filters. The observer computes the variance of these filter outputs. **(B)** After a gaze shift from an area of low curvature (bottom square, $\theta = \theta_1$) to an area of high curvature (top square, $\theta = \theta_2$), the observer must update its estimate of local curvature. **(C)** Image patches that are surprising given the observer's estimate (red) have larger variance in curvature, while expected patches (white) have low variance in curvature. Frames of highly overlapping patches were slightly shifted for display purposes. **(D)** Individual image patches have a large impact on the observer's estimate when the observer is uncertain and when image patches have high centered surprise, analogous to the behavior observed in simple model environments (see Fig. 1B). Shown for $\lambda = 0.1$. **(E)** The observer can exploit its uncertainty to adapt the sparsity of the sensory encoding (heatmap; blue trace). When the observer is certain (white marker), population activity can be significantly reduced without changing the inference error. Increases in uncertainty (red marker) result in bursts of activity (red bar). An encoder optimized for constant reconstruction error produces activity that remains constant over time (green trace). **(F)** The observer can similarly exploit the predicted surprise of incoming stimuli to reduce population activity when stimuli are expected.

The system can modulate the sparsity of the population response based on surprise and uncertainty. To illustrate this, we simulated neural population activity in response to a change in each of these quantities (Fig. 7E and F, respectively). To do this, we selected a sequence of 45 image patches, 5 of which were chosen to have high centered-surprise (Fig. 7E; red marker) or to correspond to an observer with high uncertainty (Fig. 7F; red marker). An increase in either surprise or uncertainty requires a higher fidelity response to maintain an approximately constant level of inference error. This results in a

burst of population activity (blue traces in Fig. 7E-F). Similar population bursts were recently observed in V1 in response to violations of statistical regularities in stimulus sequences [56]. When optimized for constant reconstruction error, the sparsity of the population response remains fixed in time. The resulting population response does not adapt, and instead fluctuates around a constant value determined by λ (green traces in Fig. 7E-F).

DISCUSSION

Organisms rely on incoming sensory stimuli to infer behaviorally-relevant properties of their environment, and hierarchical inference is postulated to be a computational function of a broad range of neural circuits [57, 7]. Representing and transmitting these stimuli, however, is energetically costly, and such costs are known to constrain the design and function of the nervous system [38]. Here, we explored the interplay between efficient encoding and accurate inference, and we identified two general principles that can be used to balance these objectives. First, when the environment is changing over time, the relative utility of incoming stimuli for inference can also change. Second, physically different signals can exert similar influence on the observer's model of the environment and can therefore be encoded in the same neural representation without negatively affecting the inference process.

We introduced a general theoretical framework that could exploit these two principles in order to dynamically reduce metabolic costs while maintaining accurate inferences about the environment. This framework employs a well-known computational motif consisting of a feedback loop between an observer and an encoder. We demonstrated that when the goal is accurate inference, the encoder can optimally adapt depending on the uncertainty in the observer's belief about the state of the environment, and on the surprise of incoming stimuli given this belief. This optimal adaptation enables the system to efficiently infer high-level features from low-level inputs, which we argue is a broad goal of neural circuits across the brain. We therefore expect this framework to bear relevance for many different stages of sensory processing, from the periphery through the midbrain to central brain areas.

Transient increases in fidelity signal salient changes in the environment

To maintain low metabolic costs, we found that encoders optimized for inference adapt their encoding strategies in response to the changing utility of incoming stimuli. This adaptation was signaled by elevated periods of response variability, temporal decorrelation, or total activity. Transient, burst-like changes in each of these properties served to increase the fidelity of the neural response, and enabled the system to quickly respond to informative changes in the stimulus distribution. In the nervous system, bursts of high-frequency activity are thought to convey salient changes in an organism's surroundings [58]. For example, in the lateral line lobe of the weakly electric fish, neurons burst in response to electric field distortions similar to those elicited by prey [59], and these bursts are modulated by predictive feedback from downstream neurons [58]. Similarly, in the auditory system of the cricket, bursts signal changes in frequency that are indicative of predators, and the amplitude of these bursts is closely linked to the amplitude of behavioral responses [60, 61]. In the visual system, retinal ganglion cells fire synchronously in response to surprising changes in the motion trajectory of a stimulus [62], and layer 2/3 neurons in primary visual cortex show transient elevated activity in response to stimuli that violate statistical regularities in the environment [56]. Neurons in IT cortex show strong transient activity in response to visual stimuli that violate predicted transition rules [63], and recent evidence suggests that single neurons in IT encode latent probabilities of stimulus likelihood during behavioral tasks [64]. In thalamus, burst firing is modulated by feedback from cortex [65] and is thought in signal the presence of informative stimuli [66, 51, 50]. In the auditory forebrain of the zebra finch, neural activity is better predicted by the surprise of a stimulus than by its spectrotemporal content [67], and brief synchronous activity is thought to encode a form of statistical deviance of auditory stimuli [68]. We propose that this broad range of phenonema could be indicative of an active data selection process controlled by a top-down prediction of an incoming stimulus distribution, and could thus serve as an efficient strategy for encoding changes in the underlying statistics of the environment. While some of these phenomena appear tuned to specific stimulus modulations (such as those elicited by specific types of predators or prey), we argue that transient periods of elevated activity and variability more generally reflect an optimal strategy for efficiently inferring changes in high-level features from low-level input signals.

In some cases, it might be more important to reconstruct details of the stimulus itself, rather than to infer its underlying cause. In such cases, we found that the optimal encoder maintained consistently higher firing rates and more heterogeneous response patterns. In both the cricket [69] and the weakly

electric fish [58], heterogeneous neural responses were shown to encode stimulus details relevant for evaluating the quality of courtship signals (in contrast to the bursts of activity that signal the presence of aggressors). While separate circuits have been proposed to implement these two different coding schemes (inferring the presence of an aggressor versus evaluating the quality of a courtship signal), these two strategies could in principle be balanced within the same encoder.

Importantly, whether the goal is inference or stimulus reconstruction, the encoders considered here were optimized based on predictive feedback from a downstream unit and thus both bear similarity to hierarchical predictive coding as formulated by Rao and Ballard [55]. The goal, however, crucially determines the difference between these strategies: sustained heterogeneous activity enables reconstruction of stimulus details, while transient bursts of activity enable rapid detection of changes in their underlying statistics.

Periods of stationarity give rise to ambiguous stimulus representations

A central idea of this work is that stimuli that are not useful for a statistical estimation task need not be encoded. This was most notably observed during periods in which an observer maintained an accurate prediction of a stationary stimulus distribution. Here, different stimuli could be encoded by the same neural response without impacting the accuracy of the observer's prediction. This process ultimately renders stimuli ambiguous, and it predicts that the discriminability of individual stimuli should decrease over time as the system's internal model becomes aligned with the environment (Methods, Supplementary Fig. 9). Ambiguous stimulus representation have been observed in electrosensory pyramidal neurons of the weakly electric fish, where adaptation to the envelope of the animal's own electric field (a second-order statistic analogous to the variance step considered here) reduces the discriminability of specific amplitude modulations [70]. Similarly, in the olfactory system of the locust, responses of projection neurons to chemically-similar odors are highly distinguishable following an abrupt change in the odor environment, but become less distinguishable over time [71]. The emergence of ambiguous stimulus representations has recently been observed in human perception of auditory textures that are generated from stationary sound sources such as flowing water, humming wind, or large groups of animals [72]. Human listeners are readily capable of distinguishing short excerpts of sounds generated by such sources. Surprisingly, however, when asked to tell apart long excerpts of auditory textures, performance sharply decreases. We propose that this steady decrease in performance with excerpt duration reflects adaptive encoding for accurate inference, where details of the stimulus are lost over time in favor of their underlying statistical summary.

Efficient use of metabolic resources yields diverse signatures of suboptimal inference

We used an ideal Bayesian observer to illustrate the dynamic relationship between encoding and inference. Ideal observer models have been widely used to establish fundamental limits of performance on different sensory tasks [29, 73, 3]. The Bayesian framework in particular has been used to identify signatures of optimal performance on statistical estimation tasks [74, 32], and a growing body of work suggests that neural systems explicitly perform Bayesian computations [14, 7, 75, 8]. In line with recent studies [76, 77], we examined the impact of limited metabolic resources on such probabilistic neural computations.

While numerous studies have identified signatures of near-optimal performance in both neural coding (e.g. [10]) and perception (e.g. [30, 3]), the ideal observer framework can also be used to identify deviations from optimality. Such deviations have been ascribed to noise [73] and suboptimal neural decoding [78]. Here, we propose that statistical inference can deviate from optimality as a consequence of efficient, resource-constrained stimulus coding. We observed deviations from optimality in both the speed and accuracy of inference, and we found that some of these deviations (namely asymmetries in the speed of variance adaptation) could potentially be used to differentiate the underlying scheme that was used to encode incoming stimuli. It might therefore be possible to infer underlying adaptation strategies by analyzing patterns of suboptimal inference.

Limitations and future work

We discussed general principles that determine optimal encoding strategies for accurate inference, and we demonstrated the applicability of these principles in simple model systems. Understanding the applicability in more complex settings and for specific neural systems requires further investigation.

We considered a simple nonstationary environment whose dynamics varied on a single timescale. These dynamics were parameterized by a single latent variable that specified either the mean or the variance of a Gaussian stimulus distribution. These first- and second-order moments are basic properties of an input distribution and often correspond to interpretable, physical properties such as luminance or local contrast. Similar stimulus distribution have been used to study a range of neural and perceptual dynamics, including adaptation of fly visual neurons to changes in luminance and contrast [31], neural representations of electric field modulations in the weakly electric fish [70], and human perceptual decision making [12]. In natural settings, signals vary over a range of temporal scales, and neurons are capable of adapting to multiple timescales in their inputs [79, 10]. Here, we find differences in the timescale of responses to stationary versus transient periods of the environment, and differences in the timescale of encoders optimized for reconstruction versus inference. By combining both goals (reconstruction and inference) in an environment with a richer set of temporal dynamics, it might be possible to determine efficient strategies for inferring a hierarchy of timescales. In such cases, we expect the dimensionality of the latent variable space to determine the lower bound on coding costs for inference. Even in the limit of highly complex models, however, we expect accurate inference and reconstruction to impose qualitatively different constraints on neural response properties.

We considered three encoding schemes that approximate known features of neural responses, and as such could be implemented broadly across the brain. Discretization is a non-linear encoding scheme that specifies a finite set of instantaneous response levels (such as spiking patterns or discriminable firing rates) and provides a good model of retinal ganglion cells responses (e.g. [80]). Temporal filtering, on the other hand, is a linear encoding scheme that forms the basis of a broad class of linear-nonlinear (LN) models. These models have been used to describe neural responses in a range of systems [81], and can capture temporal dependencies in the neural response. To more closely approximate spiking nonlinearities observed in real neurons, the linear output of this encoder could be followed by a nonlinearity whose parameters are also adapted over time, thereby enabling the system to more strongly suppress irrelevant stimuli. Finally, our model of stimulus selection implements an extreme form of gating, whereby unsurprising stimuli are not encoded. This nonlinear encoding scheme produces bimodal responses (either strongly active or completely silent), and we would therefore expect such a mechanism to be useful when transmitting signals over long distances.

Noise is known to have varying impact on efficient coding strategies, depending on the stage of processing [82, 24, 83]. If designed appropriately, lossy encoding schemes can be beneficial when incoming signals are polluted by noise, as they can guarantee that metabolic resources are not wasted in the process of representing noise fluctuations. We expect that noise in either the encoding or decoding stages would introduce additional error in the inference process and additional ambiguity in the representation of individual stimuli. The degree of noise can also impact the goal of the system. In noisy environments where the separation between latent environmental states is low, an optimal system might need to remain stable to noise, rather than flexible to environmental changes. We expect the balance between stability and flexibility to be modulated by the spread of the stimulus distribution relative to the separation between environmental states.

In order to achieve optimal bounds on performance, encoding parameters were computed and updated on each timestep. While it is known that neural systems can adapt on timescales approaching physical limits [31], it is possible that more complex neural circuits might implement a heuristic version of this adaptation that operates on slower timescales. To measure the optimal bound on inference error, we used the mean-squared difference between point estimates derived in the presence and absence of an encoding step. This metric is general and makes no assumptions about the form of the posterior distribution [84, 32]. Other measures, such as KL-divergence, could be used to capture not only changes in point estimates, but also changes in uncertainty underlying these estimates.

Connections to existing theoretical frameworks

Efficient coding of task-relevant information has been studied before, primarily within the framework of the Information Bottleneck (IB) method [27, 85, 86]. The IB framework provides a general theoretical approach for extracting task-relevant information from sensory stimuli, and it has been successfully applied to the study of neural coding in the retina [28] and in the auditory cortex [87]. In parallel, Bayesian Efficient Coding (BEC) has recently been proposed as a framework through which a metabolically-constrained sensory system could minimize an arbitrary error function that could, as in IB, be chosen to reflect task-relevant information [88]. However, neither framework (IB nor BEC) explicitly addresses the issue of adaptive sensory coding in non-stationary environments, where the relevance of different stimuli

can change in time. Here, we frame general principles that constrain the dynamic balance between coding cost and task relevance, and we pose neurally-plausible implementations.

Our approach bears conceptual similarities to the predictive coding framework proposed by Rao and Ballard [55], in which low-level sensory neurons support accurate stimulus reconstruction by encoding the residual error between an incoming stimulus and a top-down prediction of the stimulus. Our encoding schemes similarly use top-down predictions to encode useful deviations in the stimulus distribution. Importantly, however, the goal here was not to reconstruct the stimulus itself, but rather to infer the underlying properties of a changing stimulus distribution. To this end, we considered encoding schemes that could use top-down predictions to adaptively adjust their strategies over time based on the predicted utility of different stimuli for supporting inference.

This work synthesizes different theoretical frameworks in an effort to clarify their mutual relationship. In this broad sense, our approach aligns with recent studies that aim to unify frameworks such as efficient coding and Bayesian inference [88], as well as concepts such as efficient, sparse, and predictive coding [89].

Outlook

Efficient coding and probabilistic inference are two prominent frameworks in theoretical neuroscience that address the separate questions of how stimuli can be encoded at minimal cost, and how stimuli can be used to support accurate inferences. In this work, we bridged these two frameworks within a dynamic setting. We examined optimal strategies for encoding sensory stimuli while minimizing the error that such encoding induces in the inference process, and we contrasted these with strategies designed to optimally reconstruct the stimulus itself. These two goals could correspond to different regimes of the same sensory system [17], and future work will explore strategies for balancing these regimes depending on task requirements. In order to test the implications of this work for physiology and behavior, it will be important to generalize this framework to more naturalistic stimuli, noisy encodings, and richer inference tasks. At present, our results identify broad signatures of a dynamical balance between metabolic costs and task demands that could potentially explain a wide range of phenomena in both neural and perceptual systems.

Acknowledgements

We thank John Briguglio, Vivek Jayaraman, Yarden Katz, Emily Mackevicius, and Josh McDermott for useful discussions and feedback on the manuscript. WM was supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216. AMH was supported by the Howard Hughes Medical Institute.

Competing Interests

The authors declare no competing interests.

METHODS

A Optimal Bayesian inference with adaptively-encoded stimuli

We describe a class of discrete-time environmental stimuli x_t whose statistics are completely characterized by a single time-varying environmental state variable θ_t.

We then consider the scenario in which these stimuli are encoded in neural responses, and it is these neural responses that must be used to construct the posterior probability over environmental states. In what follows, we derive the optimal Bayesian observer for computing this posterior given the history of neural responses. The steps of this estimation process are summarized in Box 1.

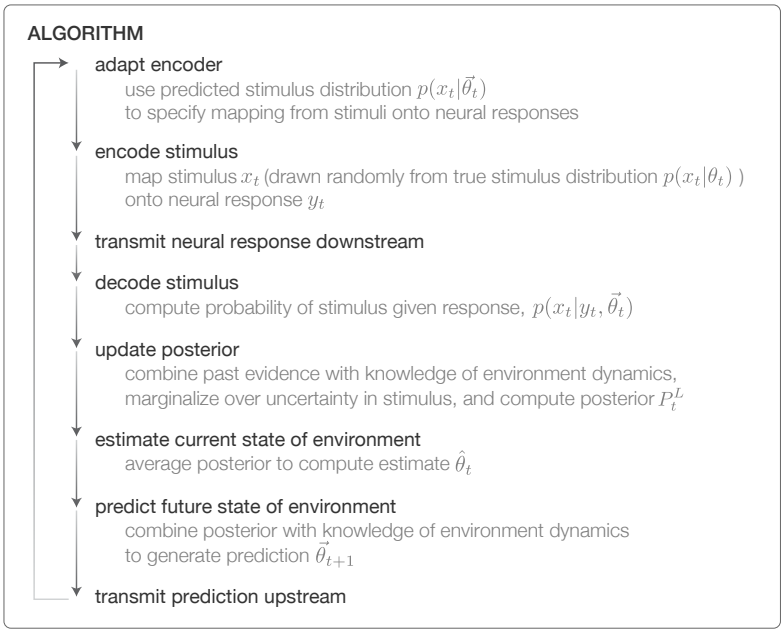

Box 1: **Algorithm for performing Bayesian inference with adaptively-encoded stimuli.**

In a full Bayesian setting, the observer should construct an estimate of the stimulus distribution, $p(x_t)$, by marginalizing over its uncertainty in the estimate of the environmental state θ_t (i.e., by computing $p(x_t) = \int d\theta_t p(x_t|\theta_t)p(\theta_t)$). For simplicity, we avoid this marginalization by assuming that the observer's belief is well-summarized by the average of the posterior, which is captured by the point value $\hat{\theta}_t = \int d\theta_t\, \theta_t p(\theta_t)$ for estimation, and $\vec{\theta}_{t+1} = \int d\theta_{t+1}\, \theta_{t+1} p(\theta_{t+1})$ for prediction. The average of the posterior is an optimal scalar estimate that minimizes the mean-squared-error between the estimated and true states of the environment, and is known to provide a good description of both neural [11] and perceptual [12] dynamics. The observer then uses these point values to condition its prediction of the stimulus distribution, $p(x_t|\vec{\theta}_t)$. Conditioning on a point estimate guarantees that the observer's prediction of the environment belongs to the same family of distributions as the true environment. This is not guaranteed to be the case when marginalizing over uncertainty in θ_t. For example, if the posterior assigns non-zero probability mass to two different mean values of a unimodal stimulus distribution, the predicted stimulus distribution could be bimodal, even if the true stimulus distribution is always unimodal. We verified numerically that the key results of this work are not affected by approximating the full marginalization with point estimates.

When the timescale of the environment dynamics is sufficiently slow, the point prediction $\vec{\theta}_{t+1}$ can be approximated by the point estimate $\hat{\theta}_t$. In the two-state environments considered here, the probability that the environment remains in the low state from time t to time $t+1$ is equal to $P_{t+1}^L = P_t^L(1-h) + (1-P_t^L)h$, where h is the hazard rate [11]. For the small hazard rate used here ($h = 0.01$),

$P^L_{t+1} = 0.99P^L_t + 0.01(1 - P^L_t)$, and the estimate $\hat{\theta}_t$ is therefore a very close approximation of the prediction $\vec{\theta}_{t+1}$. All results presented in the main text were computed using this approximation (i.e., $\vec{\theta}_{t+1} \approx \hat{\theta}_t$). With this approximation, the optimal Bayesian observer computes the approximate posterior distribution $p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t})$, conditioned on the history of neural responses $y_{\tau\leq t}$ and the history of point estimates $\hat{\theta}_{\tau<t}$. In the remainder of the Methods, we will formulate all derivations and computations in terms of the history of past estimates (up to and including time $t-1$), with the understanding that these estimates can be used as approximate predictions of the current state at time t.

With these simplifications, the general steps of the inference process can be broken down as follows:

1. **Encoder:** maps incoming stimuli $x_{\tau\leq t}$ onto a neural response y_t by sampling from the "encoding distribution" $p(y_t|x_{\tau\leq t}, \hat{\theta}_{\tau<t})$
2. **Decoder:** uses Bayes' rule to estimate the likelihood of a stimulus x_t given the neural response y_t, as given by the "decoding distribution" $p(x_t|y_t, \hat{\theta}_{\tau<t})$
3. **Observer:** uses the neural response y_t to update the posterior $p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t})$. This can be broken down into the following steps, in which the observer:
   (a) Combines the previous posterior $p(\theta_{t-1}|y_{\tau<t}, \hat{\theta}_{\tau<t-1})$ with knowledge of environment dynamics $p(\theta_t|\theta_{t-1})$ to compute the probability of θ_t given all past data, $p(\theta_t|y_{\tau<t}, \hat{\theta}_{\tau<t-1})$
   (b) Uses Bayes' rule to incorporate a new stimulus x_t and form $p(\theta_t|x_t, y_{\tau<t}, \hat{\theta}_{\tau<t-1})$
   (c) Marginalizes over the uncertainty in x_t using the decoding distribution $p(x_t|y_t, \hat{\theta}_{\tau<t})$, thereby obtaining the updated posterior $p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t})$ (which can be averaged to compute the point estimate $\hat{\theta}_t$)
   (d) Combines the updated posterior with knowledge of environment dynamics $p(\theta_{t+1}|\theta_t)$ to generate a predicted posterior $p(\theta_{t+1}|y_{\tau\leq t}, \hat{\theta}_{\tau<t})$ (which can be averaged to compute the point prediction $\vec{\theta}_{t+1}$)
4. **Feedback loop:** sends the prediction back upstream to update the encoder.

In what remains of Section A, we derive the general equations for the full inference process in the presence of both encoding and decoding. In Section B, we derive the specific forms of the inference equations in a simplified, two-state environment. We first focus on the general equations of the observer model (Section B.2). We then describe the forms of the encoding and decoding distributions implemented by the three different encoding schemes considered in this paper, and detail how the parameters of each encoder can be optimized based on the observer's prediction of the environmental state (Sections B.3-B.6). In Section C, we describe the numerical approximations used to simulate the results presented in the main paper.

A.1 Environment dynamics

We consider a non-stationary environment with Markovian dynamics. The dynamics of the environmental state variable θ_t are then specified by the distribution $p(\theta_t|\theta_{t-1})$. At each time t, the value of θ_t specifies the distribution of stimuli $p(x_t|\theta_t)$.

A.2 Encoder

We consider an encoder that maps incoming stimuli $x_{\tau\leq t}$ onto a neural response y_t. We assume that the encoder has access to the history of estimates $\hat{\theta}_{\tau<t}$ (fed back from a downstream observer) to optimally encode incoming stimuli via the "encoding distribution", $p(y_t|x_{\tau\leq t}, \hat{\theta}_{\tau<t})$.

A.3 Decoder

Because the observer does not have direct access to the stimulus, it must first decode the stimulus from the neural response. We assume that the decoder has access to the instantaneous neural response y_t and this history of past estimates $\hat{\theta}_{\tau<t}$. The decoder must use these signals to marginalize over past stimuli $x_{\tau<t}$ and compute the likelihood of the response y_t conditioned on the current stimulus x_t (this likelihood will later be used to update the observer's posterior):

$$p(y_t|x_t, \hat{\theta}_{\tau<t}) = \int dx_{\tau<t} \underbrace{p(y_t|x_t, x_{\tau<t}, \hat{\theta}_{\tau<t})}_{\substack{\text{encoding}\\ \text{distribution}}} p(x_{\tau<t}|\hat{\theta}_{\tau<t}) \tag{5}$$

The decoder must then invert this distribution (using Bayes' rule) to estimate the probability of the stimulus x_t given the response y_t and past estimates $\hat{\theta}_{\tau<t}$:

$$\begin{aligned} p(x_t|y_t, \hat{\theta}_{\tau<t}) &= \frac{p(y_t|x_t, \hat{\theta}_{\tau<t})p(x_t|\hat{\theta}_{\tau<t})}{p(y_t|\hat{\theta}_{\tau<t})} \\ &= \frac{p(y_t|x_t, \hat{\theta}_{\tau<t})p(x_t|\hat{\theta}_{t-1})}{Z(y_t, \hat{\theta}_{\tau<t})} \end{aligned} \tag{6}$$

where we have written the distribution in the denominator as a normalization constant obtained by integrating the numerator:

$$Z(y_t, \hat{\theta}_{\tau<t}) = \int dx_t \, p(y_t|x_t, \hat{\theta}_{\tau<t})p(x_t|\hat{\theta}_{t-1}) \tag{7}$$

In what follows, we refer to $p(x_t|y_t, \hat{\theta}_{\tau<t})$ (defined in Eq. 6) as the "decoding distribution".

A.4 Observer

The optimal observer should use the decoding distribution to marginalize over its uncertainty about the true value of the stimulus x_t and thereby obtain the posterior probability of θ_t given past stimuli $y_{\tau\leq t}$ and past estimates $\hat{\theta}_{\tau<t}$. To do this, we first write an expression for the probability of θ_t given all data up to (but not including) the current timestep:

$$p(\theta_t|y_{\tau<t}, \hat{\theta}_{\tau<t-1}) = \int d\theta_{t-1} \, p(\theta_t|\theta_{t-1}) \underbrace{p(\theta_{t-1}|y_{\tau<t}, \hat{\theta}_{\tau<t-1})}_{\substack{\text{posterior from}\\ \text{previous timestep}}} \tag{8}$$

where the prior is taken to be the posterior from the last timestep, and the distribution $p(\theta_t|\theta_{t-1})$ governs the dynamics of the environment.

This distribution can then be combined with a new stimulus x_t:

$$\begin{aligned} p(\theta_t|x_t, y_{\tau<t}, \hat{\theta}_{\tau<t-1}) &= \frac{p(x_t|\theta_t, y_{\tau<t}, \hat{\theta}_{\tau<t-1})p(\theta_t|y_{\tau<t}, \hat{\theta}_{\tau<t-1})}{p(x_t|y_{\tau<t}, \hat{\theta}_{\tau<t-1})} \\ &= \frac{p(x_t|\theta_t)p(\theta_t|y_{\tau<t}, \hat{\theta}_{\tau<t-1})}{\Omega(x_t, y_{\tau<t}, \hat{\theta}_{\tau<t-1})}. \end{aligned} \tag{9}$$

As before, we have written the distribution in the denominator as a normalization constant obtained by integrating the numerator:

$$\Omega(x_t, y_{\tau<t}, \hat{\theta}_{\tau<t-1}) = \int d\theta_t p(x_t|\theta_t) p(\theta_t|y_{\tau<t}, \hat{\theta}_{\tau<t-1}) \tag{10}$$

Finally, we marginalize over the unknown value of the signal x_t using the decoding distribution $p(x_t|y_t, \hat{\theta}_{\tau<t})$ to obtain the updated posterior distribution:

$$p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t}) = \int dx_t \, p(\theta_t|x_t, y_{\tau<t}, \hat{\theta}_{\tau<t-1}) p(x_t|y_t, \hat{\theta}_{\tau<t}) \tag{11}$$

To form a prediction about the future state of the environment, the observer should combine its belief $p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t})$ about the current state of the environment with the knowledge $p(\theta_{t+1}|\theta_t)$ of the environment dynamics in a manner analogous to Eq. 8.

A.5 Computing point estimates

The posterior can be used to compute a point estimate $\hat{\theta}_t$ and prediction $\vec{\theta}_{t+1}$ of the environmental state:

$$\hat{\theta}_t = \int d\theta_t \theta_t p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t}) \equiv \hat{\theta}_{y,t} \tag{12}$$

$$\begin{aligned} \vec{\theta}_{t+1} &= \int d\theta_{t+1} \, \theta_{t+1} p(\theta_{t+1}|y_{\tau\leq t}, \hat{\theta}_{\tau<t}) \\ &= \int d\theta_{t+1} \theta_{t+1} \int d\theta_t \, p(\theta_{t+1}|\theta_t) p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t}) \end{aligned} \tag{13}$$

The point estimate given in Eq. 12 is referred to in the main text as "$\hat{\theta}_{y,t}$". We distinguish this from the point estimate "$\hat{\theta}_{x,t}$", which was derived in [11] in the absence of encoding/decoding.

B Model environments

B.1 Environment dynamics

We consider a two-state environment in which the state θ_t can take one of two values, θ^L and θ^H. At each timestep, the environment can switch states with a constant probability h, referred to as the "hazard rate". The hazard rate fully specifies the dynamics of the environment:

$$\theta_t = z_t \theta_{t-1} + (1 - z_t)(\theta^L + \theta^H - \theta_{t-1}) \tag{14}$$

where z_t is a binary random variable equal to 1 with probability h and 0 with probability $1 - h$.

We take θ_t to parametrize either the mean μ or the standard deviation σ of a Gaussian stimulus distribution:

$$p(x_t|\theta_t) = \begin{cases} \mathcal{N}(x_t; \theta_t, \sigma^2), & \text{mean-switching environment } (\theta_t = \mu) \\ \mathcal{N}(x_t; \mu, \theta_t^2), & \text{variance-switching environment } (\theta_t = \sigma) \end{cases} \tag{15}$$

B.2 Observer

In a two-state environment, the posterior distribution $p(\theta_t|y_{\tau\leq t}, \hat{\theta}_{\tau<t})$ can be summarized by a single value $P_t^L = p(\theta_t = \theta^L|y_{\tau\leq t}, \hat{\theta}_{\tau<t})$, which is the probability that the environment is in the low state at time t.

Given the posterior P_{t-1}^L at the previous timestep, the distribution for θ_t given all past data is given by:

$$p(\theta_t = \theta^L|y_{\tau<t}, \hat{\theta}_{\tau<t}) = (1-h)P_{t-1}^L + h(1-P_{t-1}^L) \tag{16}$$

where h is the a priori probability that a switch occurred at the current timestep. This distribution can then be combined with a new stimulus x_t:

$$\begin{aligned} p(\theta_t = \theta^L|x_t, y_{\tau<t}, \hat{\theta}_{\tau<t}) &= \frac{p(x_t|\theta_t = \theta^L)p(\theta_t = \theta^L|y_{\tau<t}, \hat{\theta}_{\tau<t})}{\Omega(x_t, y_{\tau<t}, \hat{\theta}_{\tau<t})} \\ &= \frac{\mathcal{N}(x_t; \mu_L, \sigma_L^2)\left[(1-h)P_{t-1}^L + h(1-P_{t-1}^L)\right]}{\Omega(x_t, y_{\tau<t}, \hat{\theta}_{\tau<t})} \end{aligned} \tag{17}$$

The variables (μ_L, σ_L) and (μ_H, σ_H) correspond to mean and standard deviation of the stimulus distribution in the low and high states, respectively, and their values vary depending on the type of the environment (mean-switching versus variance-switching).

To obtain the updated posterior P_t^L, we marginalize over the decoding distribution $p(x_t|y_t, \hat{\theta}_{\tau<t})$:

$$\begin{aligned} P_t^L = p(\theta_t = \theta^L|y_{\tau\leq t}, \hat{\theta}_{\tau<t}) &= \int dx_t\, p(\theta_t = \theta^L|x_t, y_{\tau<t}, \hat{\theta}_{\tau<t})p(x_t|y_t, \hat{\theta}_{\tau<t}) \\ &= \int dx_t\, \frac{\mathcal{N}(x_t; \mu_L, \sigma_L^2)\left[(1-h)P_{t-1}^L + h(1-P_{t-1}^L)\right]}{\Omega(x_t, y_{\tau<t}, \hat{\theta}_{\tau<t})} p(x_t|y_t, \hat{\theta}_{\tau<t}) \end{aligned} \tag{18}$$

The posterior can be used to construct a new point-estimate $\hat{\theta}_t$ of the environmental state:

$$\hat{\theta}_t = P_t^L\theta^L + (1-P_t^L)\theta^H, \tag{19}$$

where $1 - P_t^L = P_t^H$ is the probability that the environment is in the high state at time t. Note that although the environmental states are discrete, the optimal Bayesian observer maintains a continuous estimate $\hat{\theta}_t$.

To form a prediction about the future state of the environment, the observer first combines the posterior P_t^L with knowledge of environment dynamics (in a manner analogous to Eq. 16), and then computes a point prediction (in a manner analogous to Eq. 19):

$$P_{t+1}^L = P_t^L(1-h) + (1-P_t^L)h \tag{20}$$

$$\vec{\theta}_{t+1} = P_{t+1}^L\theta^L + (1-P_{t+1}^L)\theta^H \tag{21}$$

For small hazard rates (as considered here), the predicted value $\vec{\theta}_{t+1}$ is very close to the current estimate $\hat{\theta}_t$. For simplicity, we approximate the prediction $\vec{\theta}_{t+1}$ by the estimate $\hat{\theta}_t$. This estimate is then fed back upstream and used to update the encoder. In the general case, however, one should compute the full predicted posterior via Eq. 20, and use this full prediction to optimize the encoder.

B.3 Encoder/Decoder

The posterior (given in Eq. 18) is a function of the decoding distribution $p(x_t|y_t, \hat{\theta}_{\tau<t})$, which depends on the encoding distribution $p(y_t|x_{\tau\leq t}, \hat{\theta}_{\tau<t})$ through Eqs. 5-6. In what follows, we derive the encoding and decoding distributions for the three encoding schemes considered in this paper. All three encoding schemes are noiseless; as a result, the encoding distribution $p(y_t|x_{\tau\leq t}, \hat{\theta}_{\tau<t})$ reduces to a delta function in each case. This encoding distribution can then be used to derive the decoding distribution, from which it is straightforward to compute the posterior P_t^L via Eq. 18 (and similarly any estimates and predictions derived from the posterior).

Each encoding scheme considered here was parametrized by one or more encoding parameters. In two of the three encoding schemes, these parameters were chosen to minimize an error function $E(x, y)$, subject to a constraint on the fidelity of the encoding. We defined this error function with respect to inference or stimulus reconstruction:

$$E(x,y) = \begin{cases} \left(\hat{\theta}_x - \hat{\theta}_y\right)^2, & \text{error in inference} \\ (x-y)^2, & \text{error in reconstruction} \end{cases} \tag{22}$$

where $\hat{\theta}_y$ was defined in Eq. 12, and $\hat{\theta}_x$ was derived in [11].

B.4 Limited neural response levels: encoding via discretization

B.4.1 Encoder

Here, we consider a quantization (instantaneous discretization) of the stimulus space that maps the current stimulus x_t onto one of a discrete set of values $\{y_t^i\}$, where $i = 1, 2, ...N$ labels distinct response levels. This mapping is performed deterministically by choosing the response level that minimizes the instantaneous error $E(x_t, \{y_t^i\})$:

$$\begin{aligned} y_t &= \underset{y_t^i}{\text{argmin}}\ E(x_t, \{y_t^i\}) \\ &= y_t^i \end{aligned} \tag{23}$$

We can therefore write the encoding distribution as a delta function:

$$p(y_t|x_{\tau\leq t}, \hat{\theta}_{\tau<t}) = \delta(x_t - y_t^i), \tag{24}$$

where the set of response levels $\{y_t^i\}$ implicitly contains the dependence on $\hat{\theta}_{t-1}$.

B.4.2 Decoder

The decoder must estimate the probability of a stimulus x_t, given that the observed response was y_t^i. In principle, the response y_t^i could have been generated by any stimulus in the range $[y_t^{i,L}, y_t^{i,H}]$, where $y_t^{i,L}$ and $y_t^{i,H}$ are the lower and upper bounds of the bin represented by level y_t^i, respectively.

The decoding distribution can then be written as a truncated Gaussian distribution:

$$p(x_t|y_t, \hat{\theta}_{\tau<t}) = \begin{cases} \frac{1}{Z(y_t^{i,L}, y_t^{i,H}, \hat{\theta}_{t-1})}\mathcal{N}(x_t; \hat{\mu}_{t-1}, \hat{\sigma}_{t-1}^2), & y_t^{i,L} < x_t < y_t^{i,H} \\ 0, & \text{otherwise} \end{cases} \tag{25}$$

where $Z(y_t^{i,L}, y_t^{i,H}, \hat{\theta}_{t-1})$ is a normalization constant. For simplicity, we approximated this truncated Gaussian distribution with a delta function:

$$p(x_t|y_t, \hat{\theta}_{\tau<t}) \approx \delta(x_t - y_t^i) \tag{26}$$

We verified numerically that this approximation did not impact our results.

B.4.3 Determining the optimal response levels

At each point in time, the optimal set of response levels $\{y_t^i\}^*$ was found by minimizing the following equation:

$$\{y_t^i\}^* = \underset{\{y_t^i\}}{\operatorname{argmin}} \left\langle E(x_t, \{y_t^i\}) \right\rangle_{p(x_t|\hat{\theta}_{t-1})} \tag{27}$$

subject to a hard constraint on the number of response levels. When optimizing for mean-switching environments, we defined the error function with respect to the raw stimulus and neural response (i.e., $E = E(x, y)$). When optimizing for variance-switching environment, we defined the error function with respect to the absolute value of the stimulus and neural response (i.e., $E = E(|x|, |y|)$). We computed $\langle E(x_t, \{y_t^i\})\rangle$ numerically; see Section C.3.1.

B.5 Limited gain and temporal acuity: encoding via temporal filtering

B.5.1 Encoder

In this encoding scheme, we consider a simple temporal filter parameterized by the coefficient α_t. This filter linearly combines current (x_t) and past (x_{t-1}) stimuli:

$$y_t = \alpha_t x_t + (1 - \alpha_t) x_{t-1} \tag{28}$$

The encoding distribution is then given by:

$$p(y_t|x_{\tau\leq t}, \hat{\theta}_{\tau<t}) = \delta\left(y_t - (\alpha_t x_t + (1 - \alpha_t) x_{t-1})\right), \tag{29}$$

where the filtering coefficient α_t implicitly contains the dependence on $\hat{\theta}_{t-1}$.

B.5.2 Decoder

The encoding is a function of both current and past stimuli. The decoder, however, only has access to the current response y_t. In order to estimate the probability that this response was generated by the stimulus x_t, the decoder must first use the internal estimates $\hat{\theta}_{\tau<t}$ to marginalize over uncertainty in past stimuli $x_{\tau<t}$. This was first outlined in Eq. 5, which reduces here to:

$$
\begin{aligned}
p(y_t|x_t, \hat{\theta}_{\tau<t}) &= \int dx_{\tau<t}\, p(y_t|x_t, x_{\tau<t}, \hat{\theta}_{\tau<t}) p(x_{\tau<t}|\hat{\theta}_{\tau<t}) \\
&= \int dx_{t-1}\, p(y_t|x_t, x_{t-1}, \vec{\theta}_t) p(x_{t-1}|\hat{\theta}_{t-1}) \\
&= \int dx_{t-1}\, \delta\left(y_t - (\alpha_t x_t + (1-\alpha_t)x_{t-1})\right) \mathcal{N}(x_{t-1}; \hat{\mu}_{t-1}, \hat{\sigma}^2_{t-1}) \\
&= \frac{1}{(1-\alpha_t)} \mathcal{N}\left(\frac{y_t - \alpha_t x_t}{(1-\alpha_t)}; \hat{\mu}_{t-1}, \hat{\sigma}^2_{t-1}\right) \\
&= \mathcal{N}\left(y_t; \alpha_t x_t + (1-\alpha_t)\hat{\mu}_{t-1}, (1-\alpha_t)^2 \hat{\sigma}^2_{t-1}\right)
\end{aligned} \tag{30}
$$

The decoder can then use Bayes' rule to invert this distribution and determine the probability of the stimulus x_t given the response y_t:

$$
\begin{aligned}
p(x_t|y_t, \hat{\theta}_{\tau<t}) &= \frac{p(y_t|x_t, \hat{\theta}_{\tau<t}) p(x_t|\hat{\theta}_{\tau<t})}{Z(y_t, \hat{\theta}_{\tau<t})} \\
&= \frac{1}{Z(y_t, \hat{\theta}_{\tau<t})} \mathcal{N}\left(y_t; \alpha_t x_t + (1-\alpha_t)\hat{\mu}_{t-1}, (1-\alpha_t)^2 \hat{\sigma}^2_{t-1}\right) \mathcal{N}(x_t; \hat{\mu}_{t-1}, \hat{\sigma}^2_{t-1})
\end{aligned} \tag{31}
$$

In its current form, this decoding distribution is written as a Gaussian over the variable y_t. Ultimately, the observer must use this decoding distribution to marginalize over uncertainty in x_t. In Appendix I, we walk through the algebra needed to rewrite this distribution as Gaussian over x_t. The final form of this distribution in given by:

$$
p(x_t|y_t, \hat{\theta}_{\tau<t}) = \mathcal{N}\left(x_t; \frac{\alpha_t y_t - (1-\alpha_t)(2\alpha_t - 1)\hat{\mu}_{t-1}}{1 - 2\alpha_t + 2\alpha_t^2}, \left(\frac{(1-\alpha_t)^2}{1 - 2\alpha_t + 2\alpha_t^2}\right) \hat{\sigma}^2_{t-1}\right) \tag{32}
$$

B.5.3 Determining the optimal filter coefficient

The optimal filtering coefficient α_t^* was found by minimizing the following equation:

$$
\alpha_t^* = \underset{\alpha_t}{\operatorname{argmin}} \left\langle E(x_t, y_t) \right\rangle_{p(x_t|\hat{\theta}_{t-1})} + \beta H(y_t, y_{t+1}|\hat{\theta}_{\tau<t}) \tag{33}
$$

The error term, $\langle E(x_t, y_t) \rangle$, was computed numerically; see Section C.3.2. The entropy term, $H(y_t, y_{t+1}|\hat{\theta}_{\tau<t})$, can be computed analytically (see Appendix II for details):

$$
\begin{aligned}
H(y_t, y_{t+1}|\hat{\theta}_{\tau<t}) &= H(y_{t+1}|y_t, \hat{\theta}_{\tau<t}) + H(y_t|\hat{\theta}_{\tau<t}) \\
&= \frac{1}{2} \log_e \left(4\pi^2 e^2 \left(\alpha_t^2 \hat{\sigma}^2_{t-1} + \frac{(1-\alpha_t)^4}{1 - 2\alpha_t + 2\alpha_t^2} \hat{\sigma}^2_{t-1}\right) \left(\alpha_t^2 \hat{\sigma}^2_{t-1} + (1-\alpha_t)^2 \hat{\sigma}^2_{t-1}\right)\right)
\end{aligned} \tag{34}
$$

B.6 Limited neural activity: encoding via dynamic stimulus selection

B.6.1 Encoder

In this encoding scheme, the encoder uses the misalignment signal M_t to determine whether or not to encode and transmit the stimulus x_t. If the magnitude of the misalignment signal exceeds the threshold

V, the stimulus is encoded and transmitted. Otherwise, the stimulus is not encoded, and a "null symbol" is transmitted to the observer. For the purposes of computing the encoding and decoding distributions, we use $y_t = 0$ to denote the null symbol (in the main text, we denoted the null symbol by $y_t = \emptyset$).

This encoding is a deterministic mapping of the stimulus x_t onto the response y_t, dependent upon the misalignment signal M_t. The encoding distribution can thus be written in a probabilistic form as a mixture of two delta functions:

$$p(y_t|x_t, \hat{\theta}_{\tau<t}) = \begin{cases} \delta(y_t - x_t), & |M_t| > V \\ \delta(y_t), & |M_t| \leq V \end{cases} \tag{35}$$

where M_t implicitly contains the dependence on $\hat{\theta}_{t-1}$.

B.6.2 Decoder

In this scheme, the form of the decoding distribution depends on whether or not the encoder transmits the stimulus x_t. If the stimulus was encoded and transmitted, there is no uncertainty in its value, and the decoding distribution is a delta function about y_t. If the stimulus was not encoded and the null symbol was instead transmitted, the decoder can only assume that the stimulus came from the estimated stimulus distribution $p(x_t|\hat{\theta}_{t-1})$.

The decoding distribution therefore takes the following form:

$$p(x_t|y_t, \hat{\theta}_{\tau<t}) = \begin{cases} \delta(x_t - y_t), & y_t \neq 0 \\ p(x_t|\hat{\theta}_{t-1}), & y_t = 0 \end{cases} \tag{36}$$

B.6.3 Determining the misalignment signal

In defining this encoding scheme, our aim was to construct a heuristic "misalignment" signal that would alert the encoder to a change in the stimulus distribution. One candidate is a signal that tracks the average surprise of incoming stimuli, given the internal estimate of the environmental state.

The surprise associated with a single stimulus x_t is equal to the negative log-likelihood of the stimulus given the estimate $\hat{\theta}_{t-1}$:

$$S(x_t) = -\log p(x_t|\hat{\theta}_{t-1}) \tag{37}$$

The average surprise of incoming stimuli, obtained by averaging over the true stimulus distribution $p(x_t|\theta_t)$, is equal to cross-entropy between the true and estimated stimulus distributions:

$$H(x_t; \theta_t, \hat{\theta}_{t-1}) = \int dx_t \, S(x_t) p(x_t|\theta_t) \tag{38}$$

$$= H(x_t; \theta_t) + D_{KL}[p(x_t|\theta_t) \| p(x_t|\hat{\theta}_{t-1})], \tag{39}$$

where the second term in Eq. 39 is the Kullback-Leibler divergence of the estimated stimulus distribution from the true stimulus distribution.

The cross-entropy is equal to the entropy of the true stimulus distribution when the observer's estimate is accurate (i.e., when $\hat{\theta}_{t-1} = \theta_t$), and increases with the inaccuracy of the observer's estimate. To construct a signal that deviates from zero (rather than from the entropy of the stimulus distribution) whenever

observer's estimate is inaccurate, we subtract the estimated entropy $H(x_t; \hat{\theta}_{t-1})$ from the cross-entropy to define the "misalignment signal":

$$M_t = H(x_t; \theta_t, \hat{\theta}_{t-1}) - H(x_t; \hat{\theta}_{t-1}) \tag{40}$$

$$= H(x_t; \theta_t) + D_{KL}[p(x_t|\theta_t) \| p(x_t|\hat{\theta}_{t-1})] - H(x_t; \hat{\theta}_{t-1}) \tag{41}$$

The magnitude of this signal is large whenever the average surprise of incoming stimuli differs from the estimated surprise, and monotonically increases as a function of the difference between the true and estimated states of the environment. In the case of a Gaussian distribution, the misalignment signal reduces to:

$$M_t = \frac{1}{2}\log(2\pi e \sigma_t^2) + \left(\log\frac{\hat{\sigma}_{t-1}}{\sigma_t} + \frac{\sigma_t^2 + (\mu_t - \hat{\mu}_{t-1})^2}{2\hat{\sigma}_t^2} - \frac{1}{2}\right) - \frac{1}{2}\log(2\pi e \hat{\sigma}_{t-1}^2) \tag{42}$$

where μ_t and σ_t are the mean and standard deviation of the true stimulus distribution, respectively, and $\hat{\mu}_{t-1}$ and $\hat{\sigma}_{t-1}$ are the estimated values of the same parameters. The analytical values of this misalignment signal are plotted in Fig. 2 I-J.

In practice, the encoder does not have access to the parameters of the true stimulus distribution, and must therefore estimate the misalignment signal directly from incoming stimulus samples. This is discussed in more detail in Section C.3.3.

C Numerical Simulations

C.1 Environment parameters

All results were generated using a probe environment in which the state θ_t switched between two fixed values, θ^L and θ^H, every 100 time samples (corresponding to a hazard rate of $h = 0.01$). A single cycle of this probe environment consists of 200 time samples, for which the environment is in the low state ($\theta_t = \theta^L$) for the first 100 time samples and in the high state ($\theta_t = \theta^H$) for the second 100 time samples. In the main text, we averaged results over multiple cycles of the probe environment.

For the mean-switching environment, the state θ_t parametrized the mean of the stimulus distribution and switched between $\mu = \theta^L = -1$ and $\mu = \theta^H = 1$. The standard deviation was fixed to $\sigma = 1$. For the variance-switching environment, θ_t parametrized the standard deviation of the stimulus distribution and switched between $\sigma = \theta^L = 1$ and $\sigma = \theta^H = 2$. The mean was fixed to $\mu = 0$.

C.2 Updating the posterior

On each timestep, a single stimulus x_t was drawn randomly from $p(x_t|\theta_t)$. The stimulus was encoded, decoded, and used to update the posterior P_t^L. Updating the posterior requires marginalizing over the decoding distribution $p(x_t|y_t, \hat{\theta}_{\tau<t})$ (as given by Eq. 11). We approximated this marginalization numerically via Monte-Carlo simulation. At each time step, we generated 200 samples from the decoding distribution specified by each encoding scheme (for reference, the decoding distributions are given in Eqs. 26, 32, and 36). Individual samples were then used to compute separate estimates of the posterior, and the resulting set of estimates was averaged over samples. Results were robust to the number of samples used, provided that this number exceeded 50. In the case of encoding via discretization, we found that results were not sensitive to the inclusion of this marginalization step. We therefore computed all results for the discretization encoding scheme in the absence of marginalization by using the neural response y_t to directly update the posterior. This posterior forms the basis of all estimates $\hat{\theta}_t$ and predictions $\vec{\theta}_{t+1}$.

C.3 Optimizing the encoding

For two of the three encoding schemes (discretization and temporal filtering), the estimate $\hat{\theta}_{t-1}$ was used to optimize a set of encoding parameters (the set of neural response levels $\{y_t^i\}$ in the case of discretization, and the filtering coefficient α_t in the case of temporal filtering). To perform these optimizations, we discretized the posterior P_t^L into 100 values equally spaced between 0 and 1. This resulted in a set of 100 discretized values of the estimate state $\hat{\theta}_{bin}$. We found the optimal encoding parameters for each value of $\hat{\theta}_{bin}$ (described in detail in the following sections); this resulted in 100 sets of optimal response levels (given a fixed number of levels), and 100 values of the filtering coefficient α (given a fixed constraint strength β). On each timestep of the simulation, the true estimate $\hat{\theta}_t$ was mapped onto the closest discretized value $\hat{\theta}_{bin}$. The corresponding encoding parameters were then used to encode the incoming stimulus x_t. Additional details of each optimization procedure are described in the following sections.

C.3.1 Limited neural response levels: encoding via discretization

Response levels were chosen to optimize the following objective function:

$$\{y_t^i\}^* = \underset{\{y_t^i\}}{\operatorname{argmin}} \left\langle E(x_t, \{y_t^i\}) \right\rangle_{p(x_t|\vec{\theta}_t)} \tag{43}$$

The optimal set of response levels $\{y_t^i\}^*$ was found numerically using Lloyd's algorithm [34] (often referred to as K-means clustering). The algorithm takes the following as inputs: a set of points to be clustered $\{x\}$ (corresponding to stimulus samples), a number of quantization levels N (corresponding to the number of neural response levels), and a distortion measure $d(x, y)$ (corresponding to the error function $E(x, y)$). The goal of the algorithm is to find a quantization (what we referred to as a discretization of the stimulus space) that minimizes the average value of the distortion.

The values of the quantization levels, $y^1, \ldots, y^N$, are first randomly initialized. The algorithm then proceeds in two steps:

1. Each point x is assigned to a quantization level y^i that yields the smallest distortion $d(x, y^i)$.
2. Each quantization level is replaced by the average value of the points assigned to it.

The two steps are iterated until convergence.

We computed a set of optimal quantization levels (optimal response levels) for each of the 100 discretized values of $\hat{\theta}_{bin}$ (described in Section C.2). For each value of $\hat{\theta}_{bin}$, we generated a training dataset $\{x\}$ consisting of 50000 values drawn from the estimated stimulus distribution $p(x_t|\hat{\theta}_{bin})$. We determined the boundaries of each quantization level (i.e., the values $y^{i,L}$ and $y^{i,H}$ that bounded the set of stimuli that were mapped to the same quantization level) by assigning points in the training dataset to the quantization level y^i that minimized $d(x, y^i)$.

To compute optimal quantization levels for stimulus reconstruction, we used the standard distortion measure $d(x, y) = (x - y)^2$; in this case, the algorithm is guaranteed to converge to the global optimum. To compute optimal quantization levels for inference, we defined the distortion measure to be $d(x, y) = (\hat{\theta}_x - \hat{\theta}_y)^2$. The algorithm is not guaranteed to converge to a global optimum in this case, but we found empirically that the algorithm converged to a local optimum (Fig. 8). Moreover, the two distortion measures did not produce equivalent results.

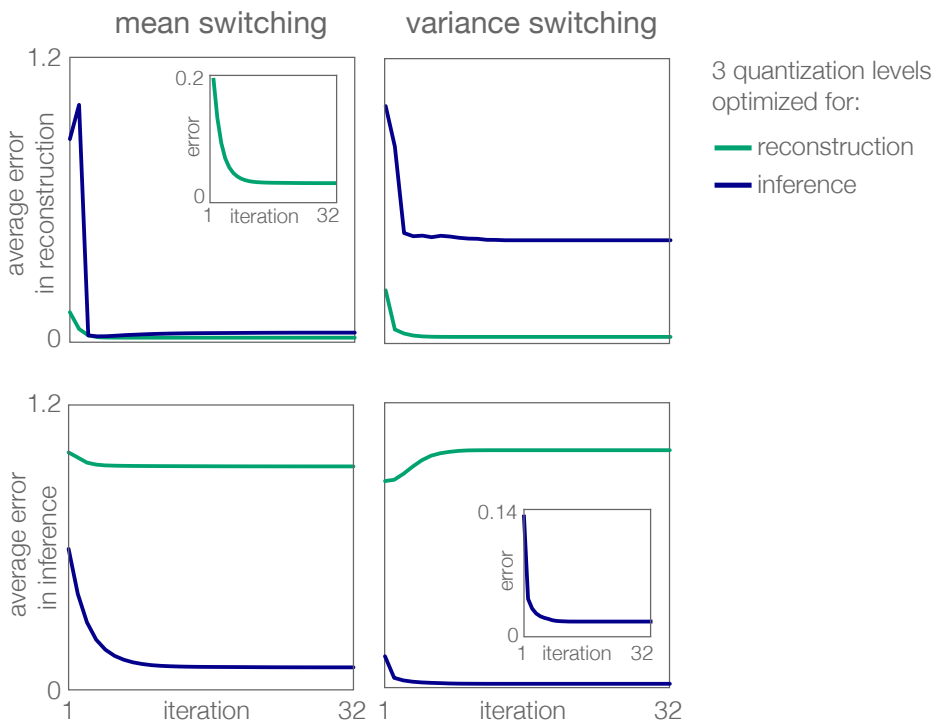

Figure 8: **Learning of optimal quantizers with Lloyd's algorithm.** Optimization of 3 quantization levels for reconstruction (green) differs from optimization for inference (blue). Regardless of the distortion measure, the algorithm converges to an optimum.

C.3.2 Limited gain and temporal acuity: encoding via temporal filtering

The optimal filtering coefficient was chosen to minimize the following objective function:

$$\alpha_t^* = \underset{\alpha_t}{\operatorname{argmin}} \langle E(x_t, y_t) \rangle_{p(x_t|\hat{\theta}_{t-1})} + \beta H(y_t, y_{t+1}|\hat{\theta}_{\tau<t}), \tag{44}$$

where as before, we choose $E(x, y) = (\hat{\theta}_x - \hat{\theta}_y)^2$ when optimizing for inference, and $E(x, y) = (x - y)^2$ when optimizing for reconstruction.

The joint entropy $H(y_t, y_{t+1}|\hat{\theta}_{\tau<t})$ can be determined analytically, as derived in Section B.5.3. We approximated the error term, $\langle E(x_t, y_t) \rangle_{p(x_t|\hat{\theta}_{t-1})}$, numerically. To do so, we first discretized α into 50 values evenly spaced between 0 and 1 (corresponding to 50 discrete values of α_{bin}). As described in Section B.5.3, we also discretized the posterior P_t^L into 100 values (corresponding to 100 discrete values of $\hat{\theta}_{bin}$). For each combination of α_{bin} and $\hat{\theta}_{bin}$, we generated 50000 pairs of stimulus samples (x_{t-1}, x_t) from the distribution $p(x_t|\hat{\theta}_{t-1})$. Each sample was used to compute values of the estimates $\hat{\theta}_x$ and $\hat{\theta}_y$. The errors $(\hat{\theta}_x - \hat{\theta}_y)^2$ and $(x_t - y_t)^2$ were then averaged over all 50000 stimulus pairs.

The optimal value α_t^* was then chosen as the value of α_{bin} that minimized the objective in Eq. 44 for a given choice of the error function $E(x, y)$ and constraint strength β.

C.3.3 Limited neural activity: encoding via dynamic stimulus selection

The misalignment signal, derived in Section B.6.3, was defined in terms of the relative alignment between the true stimulus distribution, $p(x_t|\theta_t)$, and the estimated stimulus distribution, $p(x_t|\hat{\theta}_{t-1})$. When

the parameters of the true stimulus distribution are known, the value of this signal can be computed analytically via Eq. 40. However, when the system does not have access to the stimulus distribution (as is the case here), this signal must be estimated directly from incoming stimulus samples. We consider a scenario in which the encoder can approximate Eq. 40 by computing a running-average of the stimulus surprise:

$$M_t = -\frac{1}{T}\sum_{\tau=0}^{T} \log p(x_{t-\tau}|\hat{\theta}_{t-1}) - H(x_t|\hat{\theta}_{t-1}), \tag{45}$$

where T specifies the number of timebins used to estimate the average surprise. All results in the main text were generated using $T = 10$ timebins.

C.4 Illustration of observer's sensitivity to surprise and uncertainty

Fig. 1B-D illustrated the relative impact of different stimuli on the observer's estimate of an environmental state θ, which is modulated by the observer's uncertainty and the surprise of incoming stimuli.

To illustrate this, we considered optimal Bayesian estimation of the location μ and scale $\alpha^2/2$ of a generalized Gaussian stimulus distribution:

$$p(x;\mu,\alpha,\beta) = \frac{\beta}{2\alpha\Gamma(1/\beta)} \exp\left[-\left(\frac{|x-\mu|}{\alpha}\right)^{\beta}\right] \tag{46}$$

Our derivation is analogous to that outlined in [33] for estimating the mean of a Gaussian stimulus distribution.

We consider a snapshot of the inference process, when the observer's prior is centered around a fixed estimate of the location ($\hat{\mu} = 0$) or scale ($\hat{\alpha}^2/2 = 1$). When estimating location, we fix the scale parameter to be $\alpha = \sqrt{2}$ (corresponding to a Gaussian distribution with variance $\sigma^2 = \alpha^2/2 = 1$ when $\beta = 2$). When estimating scale, we fix the location parameter to be $\mu = 0$. In both cases, we consider three different values of the shape parameter: $\beta = 1, 2, 10$.

The surprise of a single stimulus observation is quantified by the negative log-likelihood of the stimulus value given the observer's estimate. We consider 100 evenly-spaced values of surprise between 1 and 10. For each value of surprise, we compute the value of the stimulus x_t^* that elicits a given surprise.

The observer's uncertainty is captured by the entropy of the prior distribution. When estimating the location parameter, the natural conjugate prior is the Gaussian distribution $\mathcal{N}(\mu;\mu_0,\sigma_0^2)$ with mean $\mu_0 = \hat{\mu}$ (we take this mean to be the observer's point estimate of the environmental state before observing a stimulus sample x_t^*, i.e. $\hat{\theta}_{t-1} = \hat{\mu}$). The entropy of the prior distribution depends only on its variance: $H = \frac{1}{2}\log 2\pi e\sigma_0^2$. We consider 100 evenly-spaced values of the entropy between 0 and 0.7. For each value of entropy, we compute the value $\sigma_0^2 = \exp(2H)/2\pi e$ that elicits a given entropy.

When estimating the scale parameter, the natural conjugate prior is the inverse gamma function $p(\alpha;\alpha_0,\beta_0)$ with mean $\hat{\alpha} = \beta_0/(\alpha_0 - 1)$ (we take $\hat{\theta}_{t-1} = \hat{\alpha}^2/2$ to be the observer's estimate of the environmental state before observing x_t^*). The entropy of the prior depends on both α_0 and β_0: $H = \alpha_0 + \log(\beta_0\Gamma(\alpha_0)) - (1+\alpha_0)\Psi(\alpha_0)$. We fix $\beta_0 = \hat{\alpha}(\alpha_0 - 1)$. We note that the entropy is non-monotonic in α_0; we restrict our analysis to values $\alpha_0 > 2$ where both the mean and the variance of the prior are well-defined, and the entropy is monotonic. We again consider 100 evenly-spaced values of the entropy between 0 and 0.7. For each value of entropy, we compute the value α_0 that elicits a given entropy.

For each combination of prior uncertainty and surprise, we computed the posterior either over the location parameter, or over the scale parameter. We then computed the squared difference between the average value of the prior and the average value of the posterior ($(\hat{\mu}_{t-1} - \hat{\mu}_t(x_t^*))^2$ in the case of location

estimation, and $(\hat{\alpha}_{t-1}^2/2 - \hat{\alpha}_t(x_t^*)^2/2)^2$ in the case of scale estimation), and we used this squared difference as a measure of the impact of a single stimulus observation on the observer's estimate of location or scale.

C.5 Generating spike rasters

Fig. 6A showed simulated spike rasters for an encoding scheme with limited neural response levels. To generate these rasters, a stimulus sample x_t was randomly drawn from the true stimulus distribution $p(x_t|\theta_t)$. This stimulus was then mapped onto one of $N = 4$ neural response levels. Each response level was assigned a binary spike pattern from the set $\{[00], [10], [01], [11]\}$, where 1 or 0 correspond to presence or absence of a spike, respectively. Patterns were assigned to response levels $\{y_t^i\}$ according to the estimated probability, $p(y_t^i|\hat{\theta}_{t-1})$, that a particular level would be used to encode incoming stimuli. In this way, the pattern with the fewest spikes ([00]) was assigned to the response level with the highest estimated probability, and the pattern with the most spikes ([11]) was assigned to the level with the lowest estimated probability. This strategy (called "entropy coding") achieves the shortest average encoding of the input by using the fewest number of spikes [34]. We simulated spike patterns for 800 cycles of the probe environment using the set of response levels optimized for inference or stimulus reconstruction.

C.6 Computing metamer probabilities

We estimated the probability of a metamer as a function of the alignment between the true state of the environment θ and the observer's prediction $\vec{\theta}$. We say that two stimuli x_t^1 and x_t^2 are metamers (i.e., they are indistinguishable to the observer) if in the process of encoding they become mapped on the same neural response level y^M (i.e., $y_t^1 = y_t^2 = y^M$). The probability of a metamer, $p(y_t^1 = y_t^2|\theta_t, \hat{\theta}_{t-1})$, depends on both the true and predicted states of the environment. We numerically estimated this probability for a mean-switching environment in the low state ($\theta = \theta^L$). We generated 100 values of $\hat{\theta}_{t-1}$, evenly spaced between θ^L and θ^H. For each value of $\hat{\theta}_{t-1}$, we drew 100000 pairs of samples from the stimulus distribution $p(x_t|\theta_t = \theta^L)$. We encoded each stimulus by mapping it onto the corresponding response level y_t (using an encoder with 8 response levels, optimized as described in Section C.3.1). If both stimuli in the pair were mapped on the same response level, we counted the trial as a metamer. The total probability of a metamer was computed as the proportion of all trials that resulted in metamers.

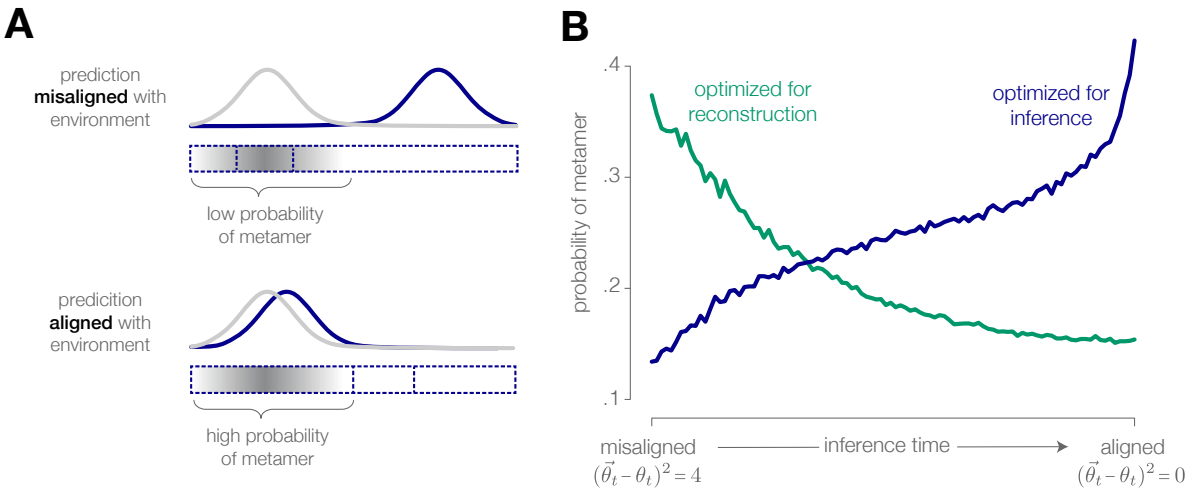

Figure 9: **Physically different stimuli become indistinguishable to an adapted system optimized for inference. (A)** (Schematic) If the observer's prediction is misaligned with the true state of the environment, the stimuli that are most likely to occur in the environment are mapped onto different neural responses. The probability of a metamer—defined by the set of stimuli that give rise to the same neural response—is low. Conversely, when the prediction is closely aligned with the environment, the probability of a metamer is high. **(B)** When an encoding is optimized for inference, the probability of a metamer increases as the observer's prediction aligns with the environment, because an increasing fraction of stimuli are mapped onto the same neural response. The opposite trend is observed when the encoding is optimized for stimulus reconstruction (illustrated for a discretized encoder with eight response levels operating in a mean-switching environment).

C.7 Measuring speed and accuracy of inference

Figure 6D-E compared the accuracy and speed of inference across different encoding schemes and environments. Accuracy was computed separately for the high and low states ($\theta = \theta^H$, and $\theta = \theta^L$, respectively) using the posterior P_t^L. For each time point, we first computed the average value of P_t^L across many cycles of the probe environment (500 cycles for discretization, and 800 cycles for filtering and thresholding, corresponding to the average trajectories of $\hat{\theta}$ shown in Fig. 3-5).

If the observer's estimate is accurate, P_t^L should be close to 1 when the environment is in the low state, and $(1 - P_t^L)$ should be close to 1 when the environment is in the high state. We therefore computed the time-averaged values $\langle P_t^L \rangle_t$ and $\langle (1 - P_t^L) \rangle_t$ to measure the accuracy in the low and high states, respectively. Time-averages were computed over the final 10 timesteps in the high or low state, respectively, corresponding to the adapted portion of the inference process.

Speed was computed separately for downward versus upward switches in the environment by measuring the number of time samples required for the posterior to stabilize after a switch. We used the time-averaged value $\langle P_t^L \rangle_t$ (again averaged over the final 10 timesteps) as a measure of the final value of the posterior in the low state. We then counted the number of timesteps after a switch downward before the posterior came within 0.05 of this value, and we used the inverse of this time as a measure of speed. We computed the speed of response to an upward switch in an analogous manner.

C.8 Natural image simulation

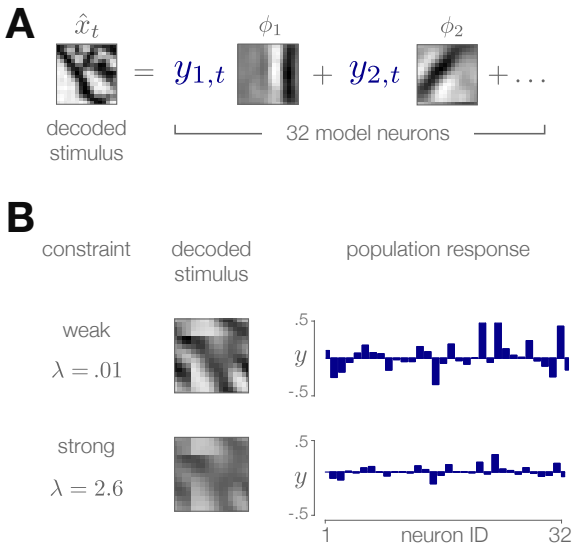

Figure 10: **Sparse coding model of natural image patches (A)** Each reconstructed image patch $\hat{x}_t$ is represented as a linear combination of basis functions ϕ_i multiplied by sparse coefficients $\{y_{i,t}\}$. **(B)** Small values of the sparsity constraint λ (upper row) result in a more accurate reconstruction of the stimulus, but require a stronger population response to encode. Large values of λ (bottom row) result in a weaker population response at the cost of decreased quality of the stimulus reconstruction.

Fig. 7 illustrated a model visual inference task performed on natural images. Within this model, the encoder implemented a sparse encoding of individual image patches [25] using 32 basis functions ϕ_i. The basis functions were chosen to minimize the following cost function:

$$\left(\vec{x}_t - \sum_i y_{i,t}\phi_i\right)^2 + \lambda \sum_i |y_{i,t}|, \tag{47}$$

where $\vec{x}_t$ is an image patch, and y_i are the responses of neurons in a population. The first terms imposes a cost on reconstruction error between the image patch $\vec{x}_t$ and the reconstructed patch $\hat{x}_t = \sum y_{i,t}\phi_i$. The second term imposes a penalty for large population responses $|y_i|$. The parameter λ imposes a constraint on the fidelity of the encoding by controlling the overall sparsity of the population response.

The set of basis functions was trained on 50000 image patches of size 16×16 pixels. Image patches were drawn randomly from the van Hateren database [90]. During training, the value of the sparsity parameter λ was set to 0.1.

A sparse representation $\vec{y}_t$ was inferred for each image patch $\vec{x}_t$ via gradient-descent on the cost function in Eq. 47 [25]. An image reconstruction $\hat{x}_t$ was computed from the sparse representation (Fig. 10A). The reconstructed patch was contrast normalized by dividing each pixel value by the standard deviation across the set of pixel values. The normalized image patch was projected onto four curvature filters C_j, resulting in four curvature coefficients $v_{j,t}$. Curvature filters were hand-designed to bear coarse, qualitative resemblance to curvature-selective receptive fields in V2. The set of four curvature coefficients was used to update the posterior distribution over variance, analogous to the Bayesian estimation of variance described in Section C.4.

Image areas 1 (low curvature) and 2 (high curvature) in Fig. 7 were chosen to be 200×200 pixels in size. For illustrative purposes, they were selected to generate a relatively large difference in the variance of curvature filters, which would require a substantial update of the Bayesian estimate. During all

simulations, the mean of the prior (corresponding to the observer's point estimate $\hat{\theta}_{t-1}$) was fixed to 5.3, equal to the variance of filter outputs in image area 1.

To numerically compute the impact of a stimulus on the estimate as a function of observer's uncertainty (prior variance) and centered surprise (Fig. 7D), a set of 5000 image patches was drawn randomly from image area 2. Image patches were then sorted according to their centered-surprise and divided into 5 groups that uniformly spanned the range of centered surprise in the set. The variance of the prior was chosen to be one of 5 equally-spaced values between 0.018 and 0.18. For each value of prior variance and for each group of stimuli with a given centered-surprise, we computed the change in the observer's estimate before and after incorporating the population response $\vec{y}_t$: $(\hat{\theta}_{t-1} - \hat{\theta}_t(\vec{y}_t))^2$.

We used a similar approach to numerically compute the inference error as a function of the sparsity parameter λ and the centered surprise (Fig. 7F). We chose 5 equally-spaced values of λ between 0.1 and 10. We then randomly drew 5000 image patches from image area 2. Image patches were again sorted according to their centered surprise and were divided into 5 groups that uniformly spanned the range of centered surprise in the set. We then computed the average inference error for each value of λ and each stimulus group. An analogous procedure was used to determine the inference error as a function of the sparsity parameter λ and the observer's uncertainty (Fig. 7 E).

Simulated population responses (Fig. 7E-F) were generated by selecting a random subset of 45 image patches with a given centered surprise, or specified values of uncertainty. Image patches were then encoded with a sparsity value of either $\lambda = 0.1$ or $\lambda = 10$ (corresponding to the inference errors marked with red and white circles). 40 images patches were encoded with the higher value of λ, and 5 image patches were encoded with the lower value of λ. For illustrative purposes, the image patches were arranged such that the first and last 20 patches corresponded to high values of λ values (white), while the middle 5 patches correspond to low values of λ (red). High and low lambda values were matched to generate similar average inference error.

Centered-surprise was computed for each image patch $\vec{x}_t$ as follows:

$$\sum_j \left| H[p(v_{j,t}|\hat{\theta}_{t-1})] + \log p(v_{j,t}|\hat{\theta}_{t-1}) \right| \tag{48}$$

where $H(v_{j,t}|\hat{\theta}_{t-1}) = \frac{1}{2}\log(2\pi e\hat{\theta}_{t-1}^2)$ is the entropy of the Gaussian distribution of curvature coefficients given the prior estimate $\hat{\theta}_{t-1}$.

References

[1] D. Kersten and P. Schrater, *Pattern inference theory: A probabilistic approach to vision.* na, 2002.

[2] D. H. Brainard, P. Longere, P. B. Delahunt, W. T. Freeman, J. M. Kraft, and B. Xiao, "Bayesian model of human color constancy," *Journal of Vision*, vol. 6, no. 11, pp. 10–10, 2006.

[3] Y. Weiss, E. P. Simoncelli, and E. H. Adelson, "Motion illusions as optimal percepts," *Nature Neuroscience*, vol. 5, no. 6, pp. 598–604, 2002.

[4] B. J. Fischer and J. L. Peña, "Owl's behavior and neural representation predicted by bayesian inference," *Nature Neuroscience*, vol. 14, no. 8, pp. 1061–1066, 2011.

[5] W. Młynarski, "The opponent channel population code of sound location is an efficient representation of natural binaural sounds," *PLOS Computational Biology*, vol. 11, no. 5, p. e1004294, 2015.

[6] D. Kersten, P. Mamassian, and A. Yuille, "Object perception as bayesian inference," *Annual Review of Psychology*, vol. 55, pp. 271–304, 2004.

[7] J. Fiser, P. Berkes, G. Orbán, and M. Lengyel, "Statistically optimal perception and learning: from behavior to neural representations," *Trends in Cognitive Sciences*, vol. 14, no. 3, pp. 119–130, 2010.

[8] R. P. Rao, B. A. Olshausen, and M. S. Lewicki, *Probabilistic models of the brain: Perception and neural function.* MIT press, 2002.

[9] R. Coen-Cagli, A. Kohn, and O. Schwartz, "Flexible gating of contextual influences in natural vision," *Nature neuroscience*, vol. 18, no. 11, p. 1648, 2015.

[10] B. Wark, A. Fairhall, and F. Rieke, "Timescales of inference in visual adaptation," *Neuron*, vol. 61, no. 5, pp. 750–761, 2009.

[11] M. DeWeese and A. Zador, "Asymmetric dynamics in optimal variance adaptation," *Neural Computation*, vol. 10, no. 5, pp. 1179–1202, 1998.

[12] M. R. Nassar, R. C. Wilson, B. Heasly, and J. I. Gold, "An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment," *Journal of Neuroscience*, vol. 30, no. 37, pp. 12366–12378, 2010.

[13] T. Lochmann, U. A. Ernst, and S. Deneve, "Perceptual inference predicts contextual modulations of sensory responses," *Journal of Neuroscience*, vol. 32, no. 12, pp. 4179–4195, 2012.

[14] S. Deneve, "Bayesian spiking neurons i: inference," *Neural Computation*, vol. 20, no. 1, pp. 91–117, 2008.

[15] S. B. Laughlin, R. R. d. R. van Steveninck, and J. C. Anderson, "The metabolic cost of neural information," *Nature Neuroscience*, vol. 1, no. 1, pp. 36–41, 1998.

[16] P. Mehta and D. J. Schwab, "Energetic costs of cellular computation," *Proceedings of the National Academy of Sciences*, vol. 109, no. 44, pp. 17978–17982, 2012.

[17] V. Balasubramanian, D. Kimber, and M. J. Berry II, "Metabolically efficient information processing," *Neural Computation*, vol. 13, no. 4, pp. 799–815, 2001.

[18] J. J. Harris, R. Jolivet, and D. Attwell, "Synaptic energy use and supply," *Neuron*, vol. 75, no. 5, pp. 762–777, 2012.

[19] D. Attwell and S. B. Laughlin, "An energy budget for signaling in the grey matter of the brain," *Journal of Cerebral Blood Flow & Metabolism*, vol. 21, no. 10, pp. 1133–1145, 2001.

[20] W. B. Levy and R. A. Baxter, "Energy efficient neural codes," *Neural Computation*, vol. 8, no. 3, pp. 531–543, 1996.

[21] S. B. Laughlin and T. J. Sejnowski, "Communication in neuronal networks," *Science*, vol. 301, no. 5641, pp. 1870–1874, 2003.

[22] A. M. Hermundstad, J. J. Briguglio, M. M. Conte, J. D. Victor, V. Balasubramanian, and G. Tkačik, "Variance predicts salience in central sensory processing," *eLife*, vol. 3, p. e03722, 2014.

[23] H. B. Barlow, "Possible principles underlying the transformations of sensory messages," 1961.

[24] J. Van Hateren, "A theory of maximizing sensory information," *Biological Cybernetics*, vol. 68, no. 1, pp. 23–29, 1992.

[25] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, p. 607, 1996.

[26] S. Laughlin, "A simple coding procedure enhances a neuron's information capacity," *Zeitschrift für Naturforschung c*, vol. 36, no. 9-10, pp. 910–912, 1981.

[27] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," *arXiv preprint physics/0004057*, 2000.

[28] S. E. Palmer, O. Marre, M. J. Berry, and W. Bialek, "Predictive information in a sensory population," *Proceedings of the National Academy of Sciences*, vol. 112, no. 22, pp. 6908–6913, 2015.

[29] W. S. Geisler, J. Najemnik, and A. D. Ing, "Optimal stimulus encoders for natural tasks," *Journal of Vision*, vol. 9, no. 13, pp. 17–17, 2009.

[30] J. Burge and W. S. Geisler, "Optimal speed estimation in natural image movies predicts human performance," *Nature Communications*, vol. 6, 2015.

[31] A. L. Fairhall, G. D. Lewen, W. Bialek, and R. R. d. R. van Steveninck, "Efficiency and ambiguity in an adaptive neural code," *Nature*, vol. 412, no. 6849, pp. 787–792, 2001.

[32] C. Robert, *The Bayesian choice: from decision-theoretic foundations to computational implementation.* Springer Science & Business Media, 2007.

[33] K. P. Murphy, "Conjugate bayesian analysis of the gaussian distribution," *def*, vol. 1, no. $2\sigma 2$, p. 16, 2007.

[34] T. M. Cover and J. A. Thomas, *Elements of information theory.* John Wiley & Sons, 2012.

[35] R. C. Wilson, M. R. Nassar, and J. I. Gold, "A mixture of delta-rules approximation to bayesian inference in change-point problems," *PLOS Computational Biology*, vol. 9, no. 7, p. e1003150, 2013.

[36] A. E. Radillo, A. Veliz-Cuba, K. Josić, and Z. P. Kilpatrick, "Evidence accumulation and change rate inference in dynamic environments," *Neural Computation*, 2017.

[37] A. Veliz-Cuba, Z. P. Kilpatrick, and K. Josic, "Stochastic models of evidence accumulation in changing environments," *SIAM Review*, vol. 58, no. 2, pp. 264–289, 2016.

[38] P. Sterling and S. Laughlin, *Principles of neural design.* MIT Press, 2015.

[39] W. Bialek, F. Rieke, R. de Ruyter van Steveninck, and D. Warland, "Spikes: exploring the neural code," *MIT. Roddey, JC, Girish, B., & Miller, JP (2000). Assessing the performance of neural encoding models in the presence of noise. Journal of Computational Neuroscience*, vol. 8, no. 95, p. 112, 1997.

[40] J. Gjorgjieva, M. Meister, and H. Sompolinsky, "Optimal sensory coding by populations of on and off neurons," *bioRxiv*, p. 131946, 2017.

[41] V. Balasubramanian and M. J. Berry, "A test of metabolically efficient coding in the retina," *Network: Computation in Neural Systems*, vol. 13, no. 4, pp. 531–552, 2002.

[42] W. Bialek, F. Rieke, R. R. d. R. van Steveninck, and D. Warland, "Reading a neural code," in *Advances in Neural Information Processing Systems*, pp. 36–43, 1990.

[43] T. O. Sharpee, H. Sugihara, A. V. Kurgansky, S. P. Rebrik, M. P. Stryker, and K. D. Miller, "Adaptive filtering enhances information transmission in visual cortex," *Nature*, vol. 439, no. 7079, pp. 936–942, 2006.

[44] T. O. Sharpee, K. I. Nagel, and A. J. Doupe, "Two-dimensional adaptation in the auditory forebrain," *Journal of Neurophysiology*, vol. 106, no. 4, pp. 1841–1861, 2011.

[45] M. V. Srinivasan, S. B. Laughlin, and A. Dubs, "Predictive coding: a fresh view of inhibition in the retina," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 216, no. 1205, pp. 427–459, 1982.

[46] W. E. Vinje and J. L. Gallant, "Sparse coding and decorrelation in primary visual cortex during natural vision," *Science*, vol. 287, no. 5456, pp. 1273–1276, 2000.

[47] M. Weliky, J. Fiser, R. H. Hunt, and D. N. Wagner, "Coding of natural scenes in primary visual cortex," *Neuron*, vol. 37, no. 4, pp. 703–718, 2003.

[48] M. R. DeWeese and A. M. Zador, "Binary coding in auditory cortex," in *Advances in Neural Information Processing Systems*, pp. 117–124, 2003.

[49] B. A. Olshausen and D. J. Field, "Sparse coding of sensory inputs," *Current Opinion in Neurobiology*, vol. 14, no. 4, pp. 481–487, 2004.

[50] D. L. Rathbun, D. K. Warland, and W. M. Usrey, "Spike timing and information transmission at retinogeniculate synapses," *Journal of Neuroscience*, vol. 30, no. 41, pp. 13558–13566, 2010.

[51] L. M. Miller, M. A. Escabı, and C. E. Schreiner, "Feature selectivity and interneuronal cooperation in the thalamocortical system," *Journal of Neuroscience*, vol. 21, no. 20, pp. 8136–8144, 2001.

[52] J. Freeman, C. M. Ziemba, D. J. Heeger, E. P. Simoncelli, and J. A. Movshon, "A functional and perceptual signature of the second visual area in primates," *Nature neuroscience*, vol. 16, no. 7, pp. 974–981, 2013.

[53] M. Ito and H. Komatsu, "Representation of angles embedded within contour stimuli in area v2 of macaque monkeys," *Journal of Neuroscience*, vol. 24, no. 13, pp. 3313–3324, 2004.

[54] "Visual processing of informative multipoint correlations arises primarily in v2," *eLife*, vol. 4, p. e06604, 2015.

[55] R. P. Rao and D. H. Ballard, "Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects," *Nature Neuroscience*, vol. 2, no. 1, pp. 79–87, 1999.

[56] J. Homann, S. A. Koay, A. M. Glidden, D. W. Tank, and M. J. Berry, "Predictive coding of novel versus familiar stimuli in the primary visual cortex," *bioRxiv*, p. 197608, 2017.

[57] T. S. Lee and D. Mumford, "Hierarchical bayesian inference in the visual cortex," *JOSA A*, vol. 20, no. 7, pp. 1434–1448, 2003.

[58] G. Marsat, A. Longtin, and L. Maler, "Cellular and circuit properties supporting different sensory coding strategies in electric fish and other systems," *Current opinion in neurobiology*, vol. 22, no. 4, pp. 686–692, 2012.

[59] A.-M. M. Oswald, M. J. Chacron, B. Doiron, J. Bastian, and L. Maler, "Parallel processing of sensory input by bursts and isolated spikes," *Journal of Neuroscience*, vol. 24, no. 18, pp. 4351–4362, 2004.

[60] P. Sabourin and G. S. Pollack, "Behaviorally relevant burst coding in primary sensory neurons," *Journal of Neurophysiology*, vol. 102, no. 2, pp. 1086–1091, 2009.

[61] G. Marsat and G. S. Pollack, "A behavioral role for feature detection by sensory bursts," *Journal of Neuroscience*, vol. 26, no. 41, pp. 10542–10547, 2006.

[62] G. Schwartz, S. Taylor, C. Fisher, R. Harris, and M. J. Berry II, "Synchronized firing among retinal ganglion cells signals motion reversal," *Neuron*, vol. 55, no. 6, pp. 958–969, 2007.

[63] T. Meyer and C. R. Olson, "Statistical learning of visual transitions in monkey inferotemporal cortex," *Proc. Natl. Acad. Sci. USA*, vol. 108, no. 48, pp. 19401–19406, 2011.

[64] A. H. Bell, C. Summerfield, E. L. Morin, N. J. Malecek, and L. G. Ungerleider, "Encoding of stimulus probability in macaque inferior temporal cortex," *Current Biology*, vol. 26, no. 17, pp. 2280–2290, 2016.

[65] M. Halassa, J. Siegle, J. Ritt, J. Ting, G. Feng, and C. Moore, "Selective optical drive of thalamic reticular nucleus generates thalamic bursts and cortical spindles," *Nature Neuroscience*, vol. 10, pp. 1118–1120, 2011.

[66] N. A. Lesica and G. B. Stanley, "Encoding of natural scene movies by tonic and burst spikes in the lateral geniculate nucleus," *Journal of Neuroscience*, vol. 24, no. 47, pp. 10731–10740, 2004.

[67] P. Gill, S. M. Woolley, T. Fremouw, and F. E. Theunissen, "What's that sound? auditory area clm encodes stimulus surprise, not intensity or intensity changes," *Journal of neurophysiology*, vol. 99, no. 6, pp. 2809–2820, 2008.

[68] G. Beckers and M. Gahr, "Large-scale synchronized activity during vocal deviance detection in the zebra finch auditory forebrain," *Journal of Neuroscience*, vol. 32, no. 31, pp. 10594–10608, 2012.

[69] P. Sabourin and G. Pollack, "Temporal coding by populations of auditory receptor neurons," *Journal of Neurophysiology*, vol. 103, pp. 1614–1621, 2010.

[70] Z. D. Zhang and M. J. Chacron, "Adaptation to second order stimulus features by electrosensory neurons causes ambiguity," *Scientific Reports*, vol. 6, p. 28716, 2016.

[71] O. Mazor and G. Laurent, "Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons," *Neuron*, vol. 48, no. 4, pp. 661–673, 2005.

[72] J. H. McDermott, M. Schemitsch, and E. P. Simoncelli, "Summary statistics in auditory perception," *Nature Neuroscience*, vol. 16, no. 4, pp. 493–498, 2013.

[73] W. S. Geisler, "Contributions of ideal observer theory to vision research," *Vision Research*, vol. 51, no. 7, pp. 771–781, 2011.

[74] E. P. Simoncelli, "Optimal estimation in sensory systems," *The Cognitive Neurosciences, IV*, pp. 525–535, 2009.

[75] W. J. Ma, J. M. Beck, P. E. Latham, and A. Pouget, "Bayesian inference with probabilistic population codes," *Nature Neuroscience*, vol. 9, no. 11, pp. 1432–1438, 2006.

[76] X.-X. Wei and A. A. Stocker, "A bayesian observer model constrained by efficient coding can explain'anti-bayesian'percepts," *Nature Neuroscience*, vol. 18, no. 10, pp. 1509–1517, 2015.

[77] D. Ganguli and E. P. Simoncelli, "Efficient sensory encoding and bayesian inference with heterogeneous neural populations," *Neural Computation*, 2014.

[78] T. Putzeys, M. Bethge, F. Wichmann, J. Wagemans, and R. Goris, "A new perceptual bias reveals suboptimal population decoding of sensory responses," *PLOS Computational Biology*, vol. 8, no. 4, p. e1002453, 2012.

[79] B. N. Lundstrom, M. H. Higgs, W. J. Spain, and A. L. Fairhall, "Fractional differentiation by neocortical pyramidal neurons," *Nature Neuroscience*, vol. 11, no. 11, pp. 1335–1342, 2008.

[80] K. Koch, J. McLean, M. Berry, P. Sterling, V. Balasubramanian, and M. A. Freed, "Efficiency of information transmission by retinal ganglion cells," *Current Biology*, vol. 14, no. 17, pp. 1523–1530, 2004.

[81] T. O. Sharpee, "Computational identification of receptive fields," *Annu. Rev. Neurosci*, vol. 36, pp. 103–20, 2013.

[82] B. A. Brinkman, A. I. Weber, F. Rieke, and E. Shea-Brown, "How do efficient coding strategies depend on origins of noise in neural circuits?," *PLoS computational biology*, vol. 12, no. 10, p. e1005150, 2016.

[83] E. Doi and M. S. Lewicki, "A simple model of optimal population coding for sensory systems," *PLoS computational biology*, vol. 10, no. 8, p. e1003761, 2014.

[84] E. T. Jaynes, *Probability theory: The logic of science*. Cambridge university press, 2003.

[85] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, "Information bottleneck for gaussian variables," *Journal of Machine Learning Research*, vol. 6, no. Jan, pp. 165–188, 2005.

[86] D. Strouse and D. J. Schwab, "The deterministic information bottleneck," *arXiv preprint arXiv:1604.00268*, 2016.

[87] J. Rubin, N. Ulanovsky, I. Nelken, and N. Tishby, "The representation of prediction error in auditory cortex," *PLoS Computational Biology*, vol. 12, no. 8, p. e1005058, 2016.

[88] I. M. Park and J. W. Pillow, "Bayesian efficient coding," *bioRxiv*, p. 178418, 2017.

[89] M. Chalk, O. Marre, and G. Tkacik, "Towards a unified theory of efficient, predictive and sparse coding," *bioRxiv*, p. 152660, 2017.

[90] J. H. van Hateren and D. L. Ruderman, "Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 265, no. 1412, pp. 2315–2320, 1998.

APPENDIX I

Here, we provide a detailed derivation of the decoding distribution for the filtering encoder (described in Section B.5.2).

To simplify Eq. 31, we rewrite the first Gaussian as a function of $\alpha_t x_t$ (for notational simplicity, we will write $\sigma'^2 = (1-\alpha_t)^2 \hat{\sigma}_{t-1}^2$:

$$\begin{aligned}
\mathcal{N}(*) &= \mathcal{N}\left(y_t; \alpha_t x_t + (1-\alpha_t)\hat{\mu}_{t-1}, \sigma'^2\right) \\
&= \frac{1}{\sqrt{2\pi\sigma'^2}} \exp\left(-\frac{(y_t - (\alpha_t x_t + (1-\alpha_t)\hat{\mu}_{t-1}))^2}{2\sigma'^2}\right) \\
&= \frac{1}{\sqrt{2\pi\sigma'^2}} \exp\left(-\frac{(\alpha_t x_t - (y_t - (1-\alpha_t)\hat{\mu}_{t-1}))^2}{2\sigma'^2}\right) \\
&= \mathcal{N}\left(\alpha_t x_t; y_t - (1-\alpha_t)\hat{\mu}_{t-1}, \sigma'^2\right)
\end{aligned} \tag{49}$$

We can now pull out the factor of α_t (again, for notational simplicity, we will write $\mu' = y_t - (1-\alpha_t)\hat{\mu}_{t-1}$):

$$\begin{aligned}
\mathcal{N}(*) &= \mathcal{N}\left(\alpha_t x_t; \mu', \sigma'^2\right) \\
&= \frac{1}{\sqrt{2\pi\sigma'^2}} \exp\left(-\frac{(\alpha_t x_t - \mu')^2}{2\sigma'^2}\right) \\
&= \frac{1}{\alpha_t\sqrt{2\pi\sigma'^2/\alpha^2}} \exp\left(-\frac{\alpha_t^2(x_t - \mu'/\alpha_t)^2}{2\sigma'^2}\right) \\
&= \frac{1}{\alpha_t\sqrt{2\pi\sigma''^2}} \exp\left(-\frac{(x_t - \mu'')^2}{2\sigma''^2}\right) \\
&= \frac{1}{\alpha_t}\mathcal{N}\left(x_t; \mu'', \sigma''^2\right)
\end{aligned} \tag{50}$$

where $\mu'' = \mu'/\alpha_t = (y_t - (1-\alpha_t)\hat{\mu}_{t-1}))/\alpha_t$ and $\sigma''^2 = \sigma'^2/\alpha_t^2 = (1-\alpha_t)^2\hat{\sigma}_{t-1}^2/\alpha_t^2$. Eq. 49 can now be written as a Gaussian over x_t:

$$\mathcal{N}(*) = \frac{1}{\alpha_t}\mathcal{N}\left(x_t; (y_t - (1-\alpha_t)\hat{\mu}_{t-1})/\alpha_t, (1-\alpha_t)^2\hat{\sigma}_{t-1}^2/\alpha_t^2\right) \tag{51}$$

This allows us to combine the two distributions in Eq. 31:

$$\begin{aligned}
p(x_t|y_t, \vec{\theta}_t) &= \frac{1}{Z(y_t, \hat{\theta}_{\tau<t})}\frac{1}{\alpha_t}\mathcal{N}\left(x_t; \mu_A, \sigma_A^2\right)\mathcal{N}(x_t; \mu_B, \sigma_B^2) \\
&= \frac{1}{Z(y_t, \hat{\theta}_{\tau<t})}\frac{1}{\alpha_t}\frac{1}{\sqrt{2\pi(\sigma_A^2+\sigma_B^2)}}\exp\left(-\frac{(\mu_A-\mu_B)^2}{2(\sigma_A^2+\sigma_B^2)}\right)\mathcal{N}\left(x_t; \frac{\sigma_B^2\mu_A + \sigma_A^2\mu_B}{\sigma_A^2+\sigma_B^2}, \frac{\sigma_A^2\sigma_B^2}{\sigma_A^2+\sigma_B^2}\right) \\
&= \frac{f(y_t, \hat{\theta}_{\tau<t})}{Z(y_t, \hat{\theta}_{\tau<t})}\mathcal{N}\left(x_t; \frac{\sigma_B^2\mu_A + \sigma_A^2\mu_B}{\sigma_A^2+\sigma_B^2}, \frac{\sigma_A^2\sigma_B^2}{\sigma_A^2+\sigma_B^2}\right)
\end{aligned} \tag{52}$$

where:

$$\begin{aligned}
\mu_A &= (y_t - (1-\alpha_t)\hat{\mu}_{t-1})/\alpha_t \\
\mu_B &= \hat{\mu}_{t-1} \\
\sigma_A^2 &= (1-\alpha_t)^2\hat{\sigma}_{t-1}^2/\alpha_t^2 \\
\sigma_B^2 &= \hat{\sigma}_{t-1}^2
\end{aligned} \tag{53}$$

Because the function $f(y_t, \hat{\theta}_{\tau<t})$ does not depend on x_t, we can trivially obtain $Z(y_t, \hat{\theta}_{\tau<t})$ by integrating over x_t (as given by Eq. 7):

$$\begin{aligned} Z(y_t, \hat{\theta}_{\tau<t}) &= \int dx_t p(y_t|x_t, \hat{\theta}_{\tau<t}) p(x_t|\hat{\theta}_{\tau<t}) \\ &= f(y_t, \hat{\theta}_{\tau<t}) \int dx_t \, \mathcal{N}\left(x_t; \frac{\sigma_B^2 \mu_A + \sigma_A^2 \mu_B}{\sigma_A^2 + \sigma_B^2}, \frac{\sigma_A^2 \sigma_B^2}{\sigma_A^2 + \sigma_B^2}\right) \\ &= f(y_t, \hat{\theta}_{\tau<t}) \end{aligned} \tag{54}$$

The remaining terms in Eq. 52 are given by:

$$\begin{aligned} \sigma_A^2 + \sigma_B^2 &= \left(\frac{1 - 2\alpha_t + 2\alpha_t^2}{\alpha_t^2}\right) \hat{\sigma}_{t-1}^2 \\ \sigma_A^2 \sigma_B^2 &= \left(\frac{(1-\alpha_t)^2}{\alpha_t^2}\right) \hat{\sigma}_{t-1}^4 \\ \frac{\sigma_A^2 \sigma_B^2}{\sigma_A^2 + \sigma_B^2} &= \left(\frac{(1-\alpha_t)^2}{1 - 2\alpha_t + 2\alpha_t^2}\right) \hat{\sigma}_{t-1}^2 \\ \sigma_B^2 \mu_A + \sigma_A^2 \mu_B &= \left(\frac{\alpha_t y_t - (1-\alpha_t)(2\alpha_t - 1)\hat{\mu}_{t-1}}{\alpha_t^2}\right) \hat{\sigma}_{t-1}^2 \\ \frac{\sigma_B^2 \mu_A + \sigma_A^2 \mu_B}{\sigma_A^2 + \sigma_B^2} &= \frac{\alpha_t y_t - (1-\alpha_t)(2\alpha_t - 1)\hat{\mu}_{t-1}}{1 - 2\alpha_t + 2\alpha_t^2} \end{aligned} \tag{55}$$

Putting everything together, the final form of Eq. 31 becomes:

$$p(x_t|y_t, \hat{\theta}_{\tau<t}) = \mathcal{N}\left(x_t; \frac{\alpha_t y_t - (1-\alpha_t)(2\alpha_t - 1)\hat{\mu}_{t-1}}{1 - 2\alpha_t + 2\alpha_t^2}, \left(\frac{(1-\alpha_t)^2}{1 - 2\alpha_t + 2\alpha_t^2}\right) \hat{\sigma}_{t-1}^2\right) \tag{56}$$

For $\frac{1}{2} \leq \alpha_t \leq 1$, we can see that: $0 \leq (1-\alpha_t)(2\alpha_t - 1) \leq \frac{1}{8}$ and $\frac{1}{2} \leq (1 - 2\alpha_t + 2\alpha_t^2) \leq 1$.

APPENDIX II

Here we provide a detailed derivation of the entropy of the output of filtering encoder (described in Section B.5.3).

To compute $H(y_t, y_{t+1}|\hat{\theta}_{\tau<t})$, we assume that the encoder has access to the history of estimates $\hat{\theta}_{\tau<t}$, and that it uses the most recent estimate $\hat{\theta}_{t-1}$ as an approximate prediction of future states (i.e., $\hat{\theta}_{t-1} \approx \vec{\theta}_t \approx \vec{\theta}_{t+1}$).

For reference, the entropy of a normal distribution is:

$$H[\mathcal{N}(x; \mu, \sigma^2)] = \frac{1}{2} \log_e(2\pi e \sigma^2) \tag{57}$$

We want to compute $H(y_t, y_{t+1}|\hat{\theta}_{\tau<t})$:

$$H(y_t, y_{t+1}|\hat{\theta}_{\tau<t}) = H(y_{t+1}|y_t, \hat{\theta}_{\tau<t}) + H(y_t|\hat{\theta}_{\tau<t}) \tag{58}$$

where $y_t = \alpha_t x_t + (1-\alpha_t)x_{t-1}$ is the output of the encoder, and $\alpha_t \in [0.5, 1]$ is the filtering coefficient.

To compute each of the terms in Eq. 58, we need to compute $p(y_t|\hat{\theta}_{\tau<t})$ and $p(y_{t+1}|y_t, \hat{\theta}_{\tau<t})$. The first of these distributions is given by:

$$p(y_t|\hat{\theta}_{\tau<t}) = \mathcal{N}\left(y_t; \alpha_t\hat{\mu}_{t-1} + (1-\alpha_t)\hat{\mu}_{t-1}, \alpha_t^2\hat{\sigma}_{t-1}^2 + (1-\alpha_t)^2\hat{\sigma}_{t-1}^2\right), \tag{59}$$

whose entropy is given by:

$$H(y_t|\hat{\theta}_{\tau<t}) = \frac{1}{2}\log_e\left(2\pi e(\alpha_t^2\hat{\sigma}_{t-1}^2 + (1-\alpha_t)^2\hat{\sigma}_{t-1}^2)\right). \tag{60}$$

The second of these distributions can be written as:

$$p(y_{t+1}|y_t, \hat{\theta}_{\tau<t}) = \int dx_t p(y_{t+1}|x_t, \hat{\theta}_{\tau<t})p(x_t|y_t, \hat{\theta}_{\tau<t}) \tag{61}$$

Noting that $p(y_{t+1}|x_{t+1}, x_t, \hat{\theta}_{\tau<t}) = \delta(y_{t+1} - (\alpha_t x_{t+1} + (1-\alpha_t)x_t))$, the first term in the integral in Eq. 61 is given by:

$$\begin{aligned} p(y_{t+1}|x_t, \hat{\theta}_{\tau<t}) &= \int dx_{t+1}\, \delta(y_{t+1} - (\alpha_t x_{t+1} + (1-\alpha_t)x_t))\mathcal{N}\left(x_{t+1}; \hat{\mu}_{t-1}, \hat{\sigma}_{t-1}^2\right) \\ &= \frac{1}{\alpha_t}\mathcal{N}\left(\frac{(y_{t+1} - (1-\alpha_t)x_t)}{\alpha_t}; \hat{\mu}_{t-1}, \hat{\sigma}_{t-1}^2\right) \\ &= \frac{1}{(1-\alpha_t)}\mathcal{N}\left(x_t; \frac{(y_{t+1} - \alpha_t\hat{\mu}_{t-1})}{(1-\alpha_t)}, \frac{\alpha_t^2}{(1-\alpha_t)^2}\hat{\sigma}_{t-1}^2\right) \end{aligned} \tag{62}$$

The second term in the integral in Eq. 61 is given by:

$$p(x_t|y_t, \hat{\theta}_{\tau<t}) = \mathcal{N}\left(x_t; \frac{\alpha_t y_t - (1-\alpha_t)(2\alpha_t - 1)\hat{\mu}_{t-1}}{\alpha_t^2 + (1-\alpha_t)^2}, \left(\frac{(1-\alpha_t)^2}{\alpha_t^2 + (1-\alpha_t)^2}\right)\hat{\sigma}_{t-1}^2\right) \tag{63}$$

Combining the two terms, we have:

$$\begin{aligned} p(y_{t+1}|x_t, \hat{\theta}_{\tau<t})p(x_t|y_t, \hat{\theta}_{\tau<t}) &= \frac{1}{(1-\alpha_t)}\mathcal{N}\left(x_t; \mu_A, \sigma_A^2\right)\mathcal{N}\left(x_t; \mu_B, \sigma_B^2\right) \\ &= \frac{1}{(1-\alpha_t)}\frac{1}{\sqrt{2\pi(\sigma_A^2 + \sigma_B^2)}}\exp\left(-\frac{(\mu_A - \mu_B)^2}{2(\sigma_A^2 + \sigma_B^2)}\right) \\ &\mathcal{N}\left(x_t; \frac{\sigma_B^2\mu_A + \sigma_A^2\mu_B}{\sigma_A^2 + \sigma_B^2}, \frac{\sigma_A^2\sigma_B^2}{\sigma_A^2 + \sigma_B^2}\right) \end{aligned} \tag{64}$$

where

$$
\begin{aligned}
\mu_A &= \frac{(y_{t+1} - \alpha_t \hat{\mu}_{t-1})}{(1-\alpha_t)} \\
\mu_B &= \frac{\alpha_t y_t - (1-\alpha_t)(2\alpha_t - 1)\hat{\mu}_{t-1}}{\alpha_t^2 + (1-\alpha_t)^2} \\
\sigma_A^2 &= \frac{\alpha_t^2}{(1-\alpha_t)^2}\hat{\sigma}_{t-1}^2 \\
\sigma_B^2 &= \frac{(1-\alpha_t)^2}{\alpha_t^2 + (1-\alpha_t)^2}\hat{\sigma}_{t-1}^2
\end{aligned}
\tag{65}
$$

Putting these terms back into the integral in Eq. 61 gives:

$$
\begin{aligned}
p(y_{t+1}|y_t, \hat{\theta}_{\tau<t}) &= \frac{1}{(1-\alpha_t)} \frac{1}{\sqrt{2\pi(\sigma_A^2 + \sigma_B^2)}} \exp\left(-\frac{(\mu_A - \mu_B)^2}{2(\sigma_A^2 + \sigma_B^2)}\right) \\
&= \mathcal{N}\left(y_{t+1}; \alpha_t \hat{\mu}_{t-1} + \frac{(1-\alpha_t)(\alpha_t y_t - (1-\alpha_t)(2\alpha - 1)\hat{\mu}_{t-1})}{1 - 2\alpha + 2\alpha_t^2}, \alpha_t^2 \hat{\sigma}_{t-1}^2 + \frac{(1-\alpha_t)^4}{1 - 2\alpha_t + 2\alpha_t^2}\hat{\sigma}_{t-1}^2\right)
\end{aligned}
\tag{66}
$$

The conditional entropy $H(y_{t+1}|y_t, \hat{\theta}_{\tau<t})$ is determined by the variance in this distribution:

$$
H(y_{t+1}|y_t, \hat{\theta}_{\tau<t}) = \frac{1}{2}\log_e\left(2\pi e\left(\alpha^2 \hat{\sigma}_{t-1}^2 + \frac{(1-\alpha)^4}{1 - 2\alpha + 2\alpha^2}\hat{\sigma}_{t-1}^2\right)\right) \tag{67}
$$

Combining the two entropy terms in Eqs. 60 and 67, we get:

$$
\begin{aligned}
H(y_t, y_{t+1}|\hat{\theta}_{\tau<t}) &= H(y_{t+1}|y_t, \hat{\theta}_{\tau<t}) + H(y_t|\hat{\theta}_{\tau<t}) \\
&= \frac{1}{2}\log_e\left(4\pi^2 e^2\left(\alpha_t^2 \hat{\sigma}_{t-1}^2 + \frac{(1-\alpha_t)^4}{1 - 2\alpha_t + 2\alpha_t^2}\hat{\sigma}_{t-1}^2\right)\left(\alpha_t^2 \hat{\sigma}_{t-1}^2 + (1-\alpha_t)^2 \hat{\sigma}_{t-1}^2\right)\right)
\end{aligned}
\tag{68}
$$