

ANNALS OF THE NEW YORK ACADEMY OF SCIENCES

Issue: *The Year in Cognitive Neuroscience***Efficient coding and the neural representation of value**Kenway Louie¹ and Paul W. Glimcher^{1,2,3}¹Center for Neural Science, ²Department of Psychology, ³Department of Economics, New York University, New York, New York

Address for correspondence: Kenway Louie, Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, NY 10003. klouie@cns.nyu.edu

To survive in a dynamic environment, an organism must be able to effectively learn, store, and recall the expected benefits and costs of potential actions. The nature of the valuation and decision processes is thus of fundamental interest to researchers at the intersection of psychology, neuroscience, and economics. Although normative theories of choice have outlined the theoretical structure of these valuations, recent experiments have begun to reveal how value is instantiated in the activity of neurons and neural circuits. Here, we review the various forms of value coding that have been observed in different brain systems and examine the implications of these value representations for both neural circuits and behavior. In particular, we focus on emerging evidence that value coding in a number of brain areas is context dependent, varying as a function of both the current choice set and previously experienced values. Similar contextual modulation occurs widely in the sensory system, and efficient coding principles derived in the sensory domain suggest a new framework for understanding the neural coding of value.

Keywords: decision making; context dependence; reward; efficient coding; neuroeconomics

Introduction

The activity in many brain regions is sensitive to reward, a modulation that can reflect not just value but processes such as sensation, motivation, and attention. Parallel work on humans and animals has begun to elucidate the different functional roles of areas representing value itself: responses linked to the values of specific actions in decision-related areas, activity that represents and perhaps stores action-independent value in the frontal cortices, and value-related teaching signals in subcortical regions that guide learning. These investigations provide details about the mechanism of value representation, and growing evidence is revealing subtle but important consequences resulting from these implementations of value representation in neural activity. For example, although normative models of choice assume that the values of options or goods are evaluated in an absolute manner, independent of other available alternatives, the neural representation of value has been shown to depend significantly on choice context. This finding not only carries implications for behavior but also suggests that the neural encod-

ing of value reflects key features of efficient encoding systems first identified in the visual and auditory systems. In this review, we consider the different forms of context-dependent value representation observed in the brain and compare them to well-known contextual effects in sensory processing.

We begin by reviewing the different forms of value representation observed in the brain, with a focus on results from primate electrophysiology. We then describe the ways that neural value coding depends on both spatial and temporal value context. In the following section, we describe the different ways that spatial and temporal context affect sensory processing. Finally, we briefly review the tenets of the *efficient coding hypothesis* as it applies to sensory systems and address how these context-dependent value representations can be viewed within the efficient coding hypothesis.

Value-related activity in the brain

To efficiently interact with its environment, an organism must be able to predict the consequences of actions and choose the best of possible alternative

options. Value, as a quantification of the expected rewards or costs associated with any choice or action, is thus critical to the decision-making process. This fundamental relationship between value and choice is expressed explicitly in economic theory, which defines the *expected utility* of an object only from an analysis of the choices a decision-maker makes between that object and other options.^{1,2} In this regard, economic theories respect the fact that, for example, a given chooser might view 10 apples as less than 10 times as good as one apple if that is what the subject's choices reveal. In neuroscientific studies of reward and decision making, however, the experimental parameter manipulated is typically an objective quantity such as the number of apples or the amount of liquid reward delivered to the animal, and it is this objective quantity that neurobiologists have typically hypothesized is encoded in the nervous system. Motivated by economic models of choice, a growing number of neuroscientific studies have demonstrated that it is in fact the subjective rather than objective value of rewards that best correlates with reward-related activity in the brain (for more information, see Refs. 3 and 5).

For an organism facing an uncertain and dynamic world, however, optimal behavior requires more than just using value to guide decision making: values must be learned through interaction with the environment, and these values must be stored and updated over time for subsequent use. Consistent with this idea, evidence from neurophysiological studies indicate that value-related neural activity is observed widely throughout the brain and there is growing evidence that activity in these value-related areas can be broadly grouped into three categories according to function: action selection, value storage, and learning. In line with the idea that value is an intrinsic part of the decision process, many of the brain systems involved in action selection and decision making are modulated by value. There is accumulating evidence that frontal regions, notably the ventromedial prefrontal and orbitofrontal cortices, encode the value of options independent of the actions required to achieve them; such a representation is consistent with these areas being involved in the storage and recall of values. Finally, neural circuits that respond near the time of reward receipt are systematically modulated by value in a manner consistent with a role in learning. We briefly review

each of these in turn, with a focus on the primate electrophysiology literature.

Multiple stages of the action-selection process are significantly modulated by reward, with neural activity covarying with the value of actions in these stages. For example, in the primate visuo-saccadic system, value-coding activity is observed at both cortical and subcortical levels of processing. Early in the sensorimotor processing pathway, the activity of neurons in the posterior parietal cortex varies monotonically with the subjective value of the reward associated with a saccade.^{6–10} This influence of value extends through the oculomotor pathway to brain areas more closely tied to saccade execution, with reward expectation modulating activity in the frontal eye fields,¹¹ supplementary eye fields,¹² and the superior colliculus.^{13,14} Reward-related activity is also readily observed in the basal ganglia, where the activity of striatal neurons reflects the expectation of reward.^{15–17} When reinforcement learning models are fit to behavioral data in dynamic tasks, many striatal neurons encode the derived trial-by-trial action values.^{18,19} In these saccade-related brain areas, neurons retain their spatial selectivity, and the influence of value acts primarily like a modulation of gain. Although most of the accumulated evidence pertains to the oculomotor system, similar valuation signals are likely to exist in other effector systems, such as that controlling reach.^{18,20,21}

Although decision making can be viewed as purely a process of action selection, values can also be associated with choice options in a manner independent of motor action. Such a “goods-based” system provides flexibility, allowing an animal to enact a decision independent of simple stimulus–response associations. A recent experiment by Padoa-Schioppa and Assad suggests that neurons in area 13 of the primate orbitofrontal cortex (OFC) can encode exactly this kind of goods-based value.²² Employing a binary choice task that required animals to choose between varying amounts of two different juice types, the authors determined the relative value of the two juice rewards. In contrast to the decision- and motor-related areas mentioned previously, they found that OFC activity was generally insensitive to the spatial configuration of stimuli or the required motor action. However, many of these neurons showed one or more of three particular responses linked to the choice options: *offer type* responses, which varied with the value

(or amount) of one of the offered options; *chosen value* responses, which varied with the value of the selected option; and *taste* responses, which varied in a binary fashion, depending on which juice type was selected. Of these different factors, chosen value most closely resembles an economic, goods-based notion of value; it is a subjective quantity, dependent on an individual animal's valuation of the two juices, and represents value independent of the identity of the particular chosen option.

Unlike the action-value representations appearing in effector-specific circuits, value signals in OFC appear to be present throughout the time course of the choice process, both before and after the receipt of reward, suggesting that they may inform rather than directly participate in action selection. However, such value coding may play a role in economic decision making independent of motor response.²³ This demonstration of a goods-based, subjective value representation aligns with previous neurophysiological reports of OFC responses to reward expectation.^{24–27} In human neuroimaging studies, a growing number of studies have reported activity related to reward expectation and subjective value in the ventromedial prefrontal cortex (vmPFC; see Kable and Glimcher⁴ and Rangel and Hare⁵ for recent reviews).^{4,5} Thus, it appears that frontal circuits in humans convey a unitary representation of action-independent reward value as well. Similar action-independent value representations may occur in related brain regions, as both the OFC and the vmPFC are heavily interconnected with structures implicated in reward processing, such as the cingulate cortex, the amygdala, and the hippocampus;²⁸ many neurons in the amygdala, for example, encode both the positive and negative values of conditioned stimuli or states.^{29,30}

Other brain areas show reward-related activity that is more consistent with a role in learning rather than in the expression or storage of value. Such a system is critical because animals must, in the real world, update the value of choice options through continuing experience. Neurons in the midbrain dopaminergic system are now known to encode a teaching signal well-suited to updating stored value estimates. These neurons influence a large number of brain structures involved in motivation and goal-directed behavior, with projections targeting the nucleus accumbens and frontal cortex via the mesolimbic and mesocortical pathways and the striatum

via the nigrostriatal pathway. Multiple lines of evidence have long suggested that dopamine is involved in the processing of reward: drugs of abuse, such as cocaine, nicotine, and amphetamines, indirectly or directly increase the action of dopamine, and dopaminergic pathways are among the most effective locations for the placement of intracranial electrical self-stimulation electrodes.^{31,32}

Two developments in the 1990s clarified the role of dopamine in reward processing in a fundamental manner. First, in a series of electrophysiological studies of primates, Schultz and colleagues demonstrated that dopamine neurons show a phasic response to appetitive rewards, such as a small piece of apple or quantity of juice.^{33–35} Significantly, when visual stimuli are associated with rewards through a classical conditioning paradigm, the phasic dopaminergic response to reward delivery diminishes; instead, dopamine neurons respond at the presentation of the conditioned, predictive stimulus. This transition of neural response from the reward delivery to the conditioned stimulus mirrors the transfer of the animals' behavioral reactions. Importantly, these responses also carry information about the timing of expected rewards; when an expected reward delivery is omitted, dopaminergic neurons show a marked suppression of activity at the time when reward would have occurred.³⁶

In the second of these critical advances, Montague and colleagues used computational learning theory to provide a theoretical framework for understanding these dopaminergic neuronal responses^{37,38} (for a more complete review of these advances, see Glimcher³⁹). Learning based on prediction errors, in which an organism makes a prediction and learns contingent on errors in that prediction, is a process central to adaptation and learning rules in both computer science and psychology.^{40,41} Building on earlier studies of octopaminergic neurons in the honeybee brain, these authors suggested that mesencephalic dopamine neurons encode an error prediction that provides a dynamic signal of the difference between the expected amount of reward and the actual reward. Montague and colleagues proposed that this dopaminergic reward prediction error (RPE) signal drives learning via a temporal difference (TD) algorithm, a type of method in reinforcement learning first introduced to computer science by Sutton and Barto.^{42,43} In TD learning, the learned *value function* represents, essentially, the

(time-discounted) sum of all expected future rewards associated with any given set of actions. In each iterative learning experience, this value function is updated by a quantity determined by the RPE and the *learning rate* (a model parameter that adjusts the strength with which unexpected rewards update the estimates of future expected rewards). Examination of the primate electrophysiological results in the context of reinforcement learning suggested that dopaminergic responses display the fundamental quantitative and qualitative characteristics of this theoretical RPE signal.^{44–46} Thus, unlike the parietal or orbitofrontal responses described previously, dopaminergic responses are related to value but do not appear to encode value itself, but rather a teaching signal used to update value representations elsewhere in the brain.

To review, we have outlined a broad framework loosely categorizing value-related neural activity into action selection, value representation, and value updating. This conceptual organization provides a useful heuristic with which to systematically consider the various functional roles of value-related activity and contextual modulation that we discuss later. However, it is important to note that this heuristic can serve only as a very general guide, and some experimental findings do not fit neatly within this simple framework. One important issue that we have oversimplified concerns the relative roles of action- and goods-based selection and their integration in decision making. One possibility that has been raised by a number of authors is that all decisions are implemented at the stage of action selection, when information about option value (also called goods or stimulus value) is combined with information about motor costs into action values to guide choice.⁵ Alternatively, it has been suggested that many or all decisions may be made in an abstract goods space independent of motor implementation and occur solely in brain regions such as OFC.²³ Finally, goods-based and action-based representations may coexist and decisions might be accomplished in either framework depending on the nature of any given task.⁴⁷ A better understanding of the precise mechanism of valuation and decision making in these kinds of choices will be critical for predicting the effects of the different kinds of contextual modulation observed in these areas and discussed later.

Another important area of simplification concerns the fact that many of these brain regions may

play a role in other functional processes, and these alternative signals may confound our interpretation of the role of these areas in valuation and decision making. For example, value itself often co-varies with other important behavioral quantities, such as motivation and attention.⁴⁸ This is a particularly relevant issue in the parietal cortex, where lateral intraparietal area (LIP) neurons are known to be modulated by visual salience and the allocation of spatial attention.^{49,50} Such findings have led some to suggest that activity in the parietal cortex reflects the allocation of spatial attention and plays little or no direct role in decision making.⁵¹ Given the tight behavioral correlation between attention and decisions about eye and arm movements, we take a middle ground in this review, suggesting that both quantities are represented. Indeed, the observation that attention and decision are for the most part behaviorally inseparable has led us to suggest elsewhere that it is unlikely that these two quantities are fully orthogonalized in the nervous system.

Finally, we note that some brain areas may be involved in more than one of these stages in the processing of value information. For example, in an oculomotor foraging task, some neurons in the striatum are correlated with the value of a specific directional saccade (action value), but others are instead correlated with the value of the saccade that is chosen, regardless of direction.^{18,19} These chosen value neurons cannot guide selection: their activity is contingent on the selected action, and they are most active immediately before or persistently after reward delivery. However, because such activity represents the expected outcome of any action, it may play a role in updating stored representations of value (for example, the difference between chosen value and obtained reward can be used as a teaching signal).

Contextual modulation in value representation

In normative models of decision making, such as those in neoclassical economics, a fundamental assumption is that options are evaluated in an absolute manner, and that the values assigned to goods or actions are stable, stationary quantities.^{52,53} Under this kind of value representation, a decision-maker will have a complete preference order over all possible choice options, and will always choose the highest valued option from a set of possible alternatives. However, growing behavioral evidence

indicates that choice behavior in both animals and humans is often context dependent, for example, varying depending on the composition of the choice set or exposure to previous priming situations. Choice behavior is also to a certain degree stochastic, varying somewhat randomly from moment to moment and trial to trial. These behavioral data suggest that choice processes rely on a noisy and comparative form of evaluation, driven by a relative representation of value dependent on both spatial and temporal context. Below, we review the emerging neurophysiological evidence for context-dependent effects on the activity of value coding areas.

Spatial context and value coding

Many computational theories of action selection and decision making require a representation of the value of individual actions, from which a single action is then selected. However, the precise relationship between value and neural firing rates is not known. Although the concept of utility in economic models of choice is unattached to a particular unit of measure (that is, utility is *ordinal* not *cardinal and unique*), the neural representation of value is instantiated via actual spiking rates, which are necessarily fully cardinal unique values. As a result, many different possible neural representations of value will be consistent with a given set of choice data (and a given ordinal ranking); for example, systems whose value representations are linear transforms of one another (e.g., $V_1 = 10$, $V_2 = 20$, and $V_3 = 30$ spikes/sec vs. $V_1 = 50$, $V_2 = 100$, and $V_3 = 150$ spikes/sec) would produce identical behavioral choice preferences. Thus, behaviorally generated models of value only provide limited constraints on how neural systems represent values.

Many economic models thus do not distinguish between representations coding absolute value, where action values are modulated strictly by the value of the target option, and those coding relative value, in which action values are normalized to the value of all available options. In the animal decision-making literature, however, there is a strong historical precedent for the idea that value is represented in a fractional manner:

$$FV_1 = \frac{V_1}{V_1 + V_2},$$

where FV_1 is the fractional value of option 1, and V_1 and V_2 are the values of the two options in the choice set. In studies of foraging behavior in

pigeons, Herrnstein established that the relative response probability for a given option was intricately tied to the relative rate of reinforcement.⁵⁴ This became the basis of the influential *matching law*, which proposes that the fraction of choices an animal allocates to a given option will match the fraction of rewards earned from that option. Under this proposal, the primary determinant of choice behavior is the relative (fractional) value of rewards.

Does the brain represent action values in absolute terms, independent of the other available options, or in relative terms? In visuomotor areas like the parietal cortex, neurons display visual and motor selectivity that coincides in space, suggesting that they link sensory and motor information during decision making. Thus, a given neuron in LIP represents a specific action—a saccadic eye movement to its response field (RF). In different paradigms, decision-related LIP activity is modulated by a number of behaviorally relevant factors that affect the choice of saccade, including accumulated motion evidence, target color, temporal information, and probabilistic cues.^{55–58} One unifying hypothesis for these various correlates is that they control LIP activity by influencing the subjective value of the associated saccades. The representation of the value of a saccade (action value) is seen as a modulation of the action-selective activity, a finding now demonstrated in a number of studies.^{6–8,10,59} Growing experimental evidence suggests that these value signals reflect the relative rather than absolute value of a given saccade. LIP activity correlates with the value of the response field target when that quantity is varied alone. However, if the value of both targets in a two-target task are varied, LIP responses depend on both the response field target value (V_{in}) and the extra-response field target value (V_{out}), consistent with a fractional, relative reward representation.^{7,59} For example, Rorie *et al.* recorded LIP neurons while monkeys performed a classic perceptual discrimination task, which requires animals to judge the motion direction of a noisy motion-dot stimulus.⁵⁹ LIP neurons display characteristic decision-related activity in this paradigm,^{55,60} with firing rate increases paralleling the accumulating sensory evidence for a given saccade. When Rorie *et al.* manipulated the values associated with the two choices, they found that LIP activity reflected the absolute value of the RF target, consistent with previous studies. More significantly, they found that LIP neurons also

reflected the *relative* value of the RF target, with higher firing rates for a given RF target reward if the other, extra-RF target was associated with a low-value reward.

How is a relative value representation constructed in decision circuits? If relative value representation is considered in terms of spatial context, activity driven by the response field target value is suppressed by the value of other targets situated outside the RF. In sensory cortices, stimuli outside the classical receptive field can nonetheless significantly modulate neuronal activity driven by receptive field stimulation. Many of these extra-classical effects are characterized by models in which response is specified by the sensory properties of the stimulus inside the receptive field, divided by the weighted sum of the sensory properties of stimuli both outside and inside the receptive field.⁶¹ For example, the output of a cell in the visual cortex can be described as

$$R_i \propto \frac{A_i}{\sigma^2 + \sum_j A_j},$$

where A_i is the driving input of the cell in question, the summation in the denominator is taken over inputs A_j to a large population of similar neurons, and σ^2 is an empirical semisaturation constant. This *divisive normalization* mechanism is widely found in the visual cortex and explains nonlinear phenomena in the striate and extra-striate cortex^{62–65} as well as object-driven normalization in the ventral visual stream.⁶⁶ Furthermore, divisive normalization produces an efficient coding of natural signals^{67,68} and may underlie the attentional modulation of neural responses,⁶⁹ suggesting that it may be a canonical computational algorithm in the cortex.⁷⁰ If the parietal cortex employs an analogous functional architecture, then a similar form of divisive normalization may underlie the representation of saccade value in the LIP.

In a recent study, we explicitly examined the influence of alternative option values on LIP activity.⁷¹ To confirm a relative rather than absolute coding of value, we first quantified LIP responses in a two-target task, in which the RF target value was held constant and the extra-RF target value was explicitly varied. Consistent with other reports,^{7,59} we found that the activity of both single neurons and the population is inversely related to the value of the alternative target, suggesting that action values

are coded relative to the choice context. To more precisely quantify the nature of this modulation, we recorded additional LIP neurons in a three-target task in which we systematically varied the number of saccade targets and their values, enabling us to test the divisive normalization model. Monkeys fixated a central cue and were presented with either one, two, or three targets, each of which was associated with a different magnitude of water reward. After target presentation, monkeys were subsequently instructed to select one of the presented targets: a medium reward target situated in the RF or either the small or large reward targets placed outside the RF, typically in the opposite hemifield. Each trial consisted of one of seven possible target arrays, presented randomly and with equal probability (three single target, three dual target, and one triple target trial). Each randomized target array provided a unique combination of value associated with the target in the RF and values available outside the RF (Fig. 1A), allowing us to quantify the relationship between target value (V_{in}) and value context (V_{out}).

We found two primary effects of the value context, defined by the instantaneous choice set, on LIP responses (Fig. 1B). First, consistent with qualitative reports in two target tasks, activity elicited by target onset in the RF is modulated by the value of the alternatives, with larger V_{out} magnitudes leading to greater suppression. Second, activity when no RF target is present is suppressed in a context dependent manner, with larger V_{out} values driving activity further below baseline activity levels. Analogous to extra-classical modulation in the early visual cortex, both of these effects are driven by the value of targets that themselves do not drive the recorded neuron. Significantly, when we performed a quantitative model comparison with other possible relative reward representations including the fractional value of Herrnstein, $V_{in}/(V_{in}+V_{out})$, and a differential value model, $V_{in}-V_{out}$, contextual value modulation was clearly best explained by a divisive normalization-based model:

$$R_i \propto \frac{V_i + \beta}{\sigma^2 + \sum_j V_j},$$

where the activity of a neuron R_i is dependent on both the value of the target in its RF V_i and the sum over all available target values V_j (the empirical parameter β models suppression below

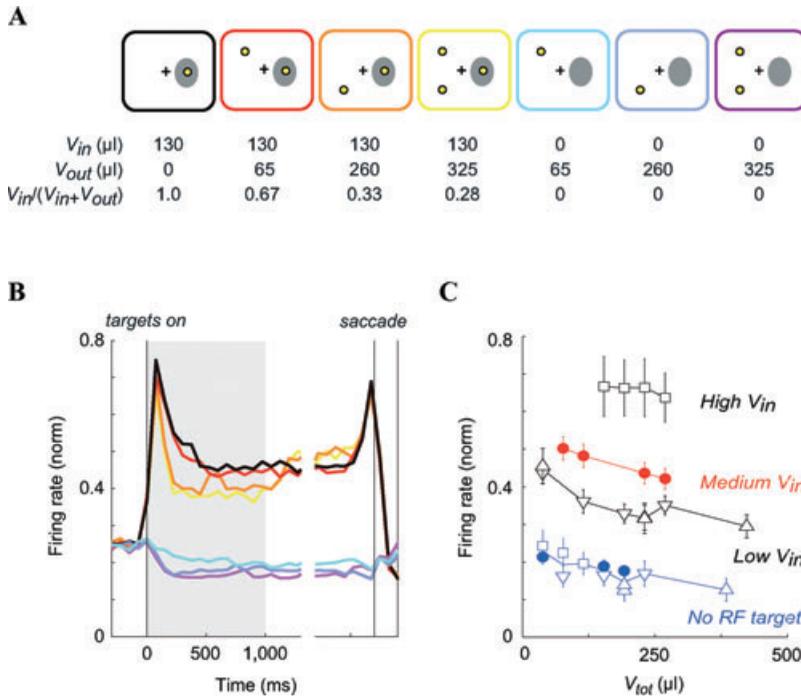


Figure 1. Spatial context dependence in LIP value coding. (A) Different value conditions in an oculomotor saccade task. Monkeys were presented with a target array of one, two, or three peripheral targets associated with different reward magnitudes. The value of the RF target was constant, whereas the value context varied with the number and reward magnitude of extra-RF targets. (B) Population parietal neuron activity. The value context (extra-RF target value) modulates LIP activity both in the presence and absence of the target in the RF. (C) Neural coding of value and value context. Increasing the value of the RF target increases LIP activity (different lines). Increasing the value of extra-RF targets suppresses LIP activity (connected lines). Together, these data suggest that LIP activity encodes a relative reward representation incorporating both target value and value context. (Adapted from Louie *et al.*⁷¹)

a baseline rate). Significantly, the implementation of relative value through divisive normalization suggests a functional linkage to contextual modulation in sensory systems, which use an analogous normalization algorithm over sensory inputs.

Temporal context and value coding

A fundamental question about the neural representation of value is how such coding changes with behavioral context; in other words, is value coding relative or absolute? The results described previously suggest that at least some parietal circuits involved in decision making reflect a normalization process across the available choice options, but it is not yet entirely clear how this parietal circuit fits into the larger network of areas involved in decision making. A number of brain regions in addition to the parietal cortex display decision-related activity during the choice process, such as the basal ganglia, frontal eye fields,

supplementary eye fields, and superior colliculus. A remaining open question is whether normalization is a general aspect of value coding at all stages of action selection. We do know, however, that neural activity in the monkey dorsal premotor cortex during reach decisions is modulated by the relative value of reach targets,⁷² indicating that value normalization may be a general phenomenon. Thus, one important issue is how contextual value coding varies in different brain areas performing different value-related processing. It is reasonable to expect that action value-based gain control, presumably occurring in an instantaneous manner when the options are presented, may be a feature specific to areas involved in decision making, and that value-sensitive brain regions subserving other functions—such as value storage—employ different value coding strategies.

Given the demonstration that orbitofrontal neurons encode a goods-based representation of value, Padoa-Schioppa and Assad examined whether those

value representations are dependent on the other available rewards in a choice situation.⁷³ As in their original demonstration of value coding by OFC neurons,²² monkeys chose between varying amounts pairs of juices (A:B, B:C, C:A) that could be ranked by relative preference order (when offered in equal amounts, $A > B > C$). Trials with different pairwise reward offers were randomly interleaved. Monkeys in this task displayed transitivity, for example, choosing 1A over 1C if they chose 1A over 1B and 1B over 1C, indicating that the different rewards could be compared on a common value scale, enabling the examination and comparison of the different neural value representations. As before, the authors found three general types of response, which they termed *offer value* (the presented value of a specific reward type), *chosen value* (the value of the selected option in a given trial, regardless of type), and *taste* (received reward type). When the value-specific responses (offer and choice value) were examined, they did not depend on the specific pair of rewards offered: for example, the activity of a neuron encoding *offer value A* had the same linear relationship between firing rate and offered amount of A whether the other available reward was B or C. The authors concluded that OFC responses are invariant to the menu of choice options, and do not reflect the relative preference ranking of the possible rewards.

This menu invariance appears at first glance to contradict earlier results by Tremblay and Schultz, describing relative reward preference in OFC neurons.²⁵ In that study, monkeys performed a spatial delayed-response task in which they were presented with a stimulus that predicted which of two possible rewards would be delivered at the end of the trial; a single trial consisted of one stimulus that was associated with a specific liquid or food reward. The task was conducted in blocks of trials, with two different stimuli and their respective associated rewards employed in a given block. Of the OFC neurons active in this task, many showed reward-related responses, responding in a phasic manner to the instruction stimulus or reward delivery and in a sustained activation preceding reward. These responses often showed greater activation for one kind of reward over others but did not differentiate between left and right instructions or different cues indicating the same reward, suggesting that they reflected information about the predicted reward.

Notably, in contrast to the Padoa-Schioppa and Assad findings, OFC neurons in the Tremblay and Schultz task encoded a value representation that was relative rather than absolute. In separate choice trials, the authors established the relative behavioral preferences between each of three different reward pairs (A:B, B:C, and C:A). The majority (40 of 65) of reward sensitive neurons showed reward responses that were dependent on the block context, with different activation for a given predicted reward contingent on which other reward was available in the block. For example, in a monkey that preferred reward A to reward B and reward B to reward C, this kind of response would show low activation in a B trial when both A and B were offered in a block but high activation in a B trial when B and C were paired. Note that these differential responses occurred in trials with identical visual stimuli and rewards; only the larger context of which reward was available in other trials varied. These results suggest that OFC neurons encode a subjective, context-dependent value driven by relative preference rather than absolute, unchanging properties of the reward itself.

How can menu invariance and relative reward preferences both occur in orbitofrontal neurons? A distinct difference between the two experiments is that the different pairwise combinations of rewards were presented randomly interleaved in the study reporting menu invariance, whereas the study showing relative reward preference presented reward pairs in blocks. If the orbitofrontal value representation adapts to the recent history of received rewards, relative reward coding over time would appear as differential adaptation occurs to the rewards in different blocks of contiguous trials. When the choice context changes rapidly with randomly interleaved pairwise rewards, the local distribution will appear almost identical to the global distribution if the integration time for adaptation is sufficiently larger, and value representations will appear absolute, or invariant.

Additional work by Padoa-Schioppa indicates that the results described in the two preceding studies can indeed be reconciled by accounting for the temporal dynamics of how different reward possibilities were presented.⁷⁴ This study consisted primarily of a reanalysis of two large datasets presented previously, including data from the study demonstrating menu invariance. In these experiments,

animals chose between different amounts of two (or in some cases, three) types of juice rewards. The distribution of possible reward sizes for a given juice type were fixed for each neuron, but varied across neurons. For example, one neuron may have been recorded with B rewards ranging from 0 to 2 (in equivalent units of juice A, determined by behavior), whereas a separate neuron was recorded with B rewards ranging from 0 to 10. To examine value-based adaptation, the authors examined whether, across the population of OFC neurons, firing rates depended on the value range (an example of local value distribution).

One straightforward model of range adaptation is that the firing rate range is adjusted to match the range of possible values. Under this model, the slope of the relationship between firing rate and value would decrease as the possible value range increases, and the high end of the value range should be represented by the same firing rate in different value-range conditions (Fig. 2, top). When the mean population firing rates were examined in this manner, segregated by value range, OFC activity showed a clear adaptation to the locally experienced range of values, for both offer value and chosen value responses (Fig. 2, bottom). Consistent results were observed when individual neurons were recorded under both low-range and high-range conditions, indicating that range adaptation is not an artifact of averaging across the population. Thus, value representation in OFC is independent of the immediate context (menu invariant) but dependent on the local temporal context (range adapting), in contrast to the trial-by-trial normalization observed in areas representing action value.

Although this temporal adaptation can be framed in terms of the range of available values, there are many characteristics that describe the local temporal distribution; for example, in the study described previously, the maximum, mean, and standard deviation of the value distribution varied along with the range. In a recent experiment, Kobayashi *et al.* examined how orbitofrontal neurons adapt their firing rates to reward distributions with different standard deviations but identical means.⁷⁵ When individual neurons were exposed to three possible liquid rewards with either a narrow distribution (low standard deviation) or a wide distribution (high standard deviation) of volumes, approximately a quarter of the neurons displayed adaptive coding, with

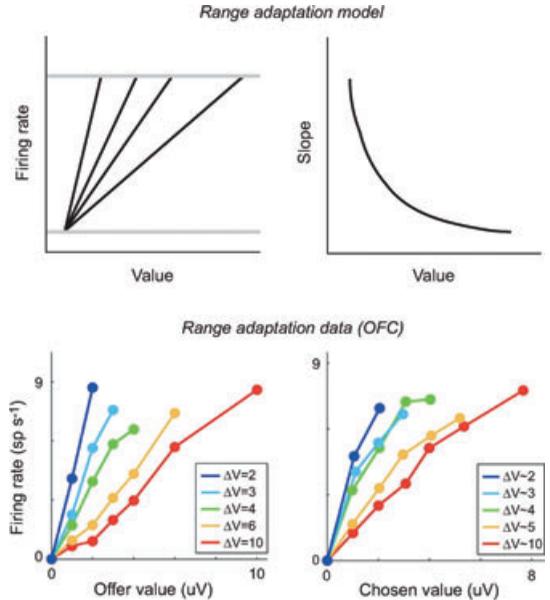


Figure 2. Temporal context dependence in OFC value coding. Top: Simple model of range adaptation in value coding neurons. The key assumption is that the range of neural activity is constant across different behavioral value conditions. Bottom: Range adaptation in orbitofrontal neuron activity. The two panels show average OFC activity in two different types of value-coding neurons, color-coded by the range of experienced values (plotted as normalized unit value). Population OFC activity adapts to the range of possible values, indicating that such activity is sensitive to the temporal value context. (Adapted from Padoa-Schioppa.⁷⁴)

steeper response slopes to the narrow range of rewards, an effect that was significant in the population response. This adaptive coding is of particular interest because it allows the full dynamic range of neural responses to be employed in representing both the narrow and wide distributions. In the light of these effects of value range and standard deviation, the results of Tremblay and Schultz can probably best be interpreted as an adaptation of orbitofrontal firing rates to the mean reward values available in different trial blocks. Importantly, this finding is compatible with the divisive normalization algorithm described previously and observed throughout the cortex, with normalization occurring across time rather than space. These results reinforce the idea that value coding can be adapted to multiple aspects of the (temporally) local probability distribution of values, though the exact parameters of the value distribution that influence

adaptation requires further study (e.g., by exploring the effects of higher-order moments like skew).

This kind of adaptive response to the local value statistics may generalize to regions other than the cortex. In a recent study of the dopamine system, widely believed to carry an error prediction signal for updating value representations, Tobler *et al.* examined how recent rewards affect the activity of midbrain dopamine neurons.⁷⁶ Monkeys were presented with three different conditioned stimuli, each of which predicted one of two possible rewards that occurred with equal probability. Consistent with coding the difference between outcome and the reward predicted by the cue, dopaminergic activity at the time of reward receipt always increased with the larger reward and decreased with the smaller reward. However, rather than a strict prediction error representing the absolute difference between predicted and received reward, dopamine responses adapted such that better rewards always elicited the same increase in activity regardless of the absolute prediction error. This again represents a form of gain control where the sensitivity of dopaminergic responses is adjusted to the range of possible prediction errors.

Functional implications of contextual coding

The results reviewed previously suggest that contextual modulation plays an important role in determining the neural coding of value in multiple brain circuits. Although research is just beginning to document such effects, it appears that the nature and extent of contextual influence in a given brain area may be closely tied to its functional role in valuation and decision making. Consistent with a role in storing value information independent of action selection, OFC value coding adapts to the temporal context but appears independent of spatial context. In contrast, consistent with a role in action selection, LIP neurons are strongly influenced by the spatial context of the available choice alternatives. One important open question is how these different forms of contextual modulation are combined in the decision process. For example, the role of the temporal context in decision circuits like the parietal cortex remains unknown; action selection areas may inherit temporally adapted value signals from the frontal cortex, receive value information from nonadapting brain areas and encode value independent of the temporal context, or apply a different

temporal weighting function to such nonadapted value signals. Future experimental work will be necessary to fully explore the nature of these different spatial and temporal context effects.

One important and largely unexplored question is the functional consequence of contextual value coding. According to normative, rational theories of choice—such as those in economics—decisions between any pair of options should be independent of the context in which the choice is made.^{53,77} For example, the relative preference between any two options should be independent of the presence or value of other alternatives, a property known as *independence from irrelevant alternatives*.⁵² However, a large body of behavioral evidence indicates that both animal and human choosers are sensitive to both spatial and temporal forms of context. In trinary choice studies, adding a third low-valued option changes the relative preference between two high-valued options in species ranging from insects to birds to humans.^{78–81} Experiments have documented a number of such phenomena in humans, which rely on the alternatives differing in two attribute dimensions, with the effect dependent on the relationship between alternatives in two-dimensional attribute space. A related effect is evident in the so-called *paradox of choice*: despite the rational prediction that more options increase welfare, choosers facing larger choice sets are more likely to select the default option, defer choosing, experience regret, and exhibit inconsistent choice behavior.^{82–84}

These behavioral context effects suggest that biological decision making employs some form of comparative valuation, but the mechanism underlying such phenomena remains unknown. One might well hypothesize that relative value coding provides a possible link between decision-making circuits and context-dependent valuation. Contextual modulation, for example normalization in the spatial domain, can significantly alter the relative distance between the mean firing rates that represent different actions. Consider a chooser selecting from three options, two high-value target items and a low-valued distracter item. Under a relative value coding system, like that described previously in the parietal cortex, the mean firing rates representing the values of each option will be divisively scaled by the total value of all alternatives. Accordingly, higher-valued distracters will decrease the distance between the neural representations coding for the values of the

target items. Noise, or variability, will critically influence these representations and the choices they produce in such systems, driving increasingly stochastic choice behavior as cardinal representations (neural firing rates) are affected by context. Importantly, variability is an inherent feature of spiking neuron activity: at fixed levels of input, neurons generate action potentials in a stochastic manner, which leads to well-characterized variability in spike counts in repeated measurements.^{85–87} Given the increasing evidence for both spatial and temporal forms of contextual modulation in value coding, understanding the interaction between noisy neural systems, context dependence, and stochastic choice behavior remains a key area of future research.

Contextual modulation in sensory coding

As described previously, there is increasing evidence that the nature of value representation in the brain is dependent on both spatial and temporal context. These responses suggest that the neural mechanisms of value coding are more complex than the classical, behavioral concept of value derived from theories of choice like economics and decision theory. To gain a better understanding of how and why such contextual effects may arise, we turn next to the electrophysiological study of sensory systems, focusing on the visual system where there is well-established literature on the effects of context on both perception and electrophysiological responses.⁸⁸ These spatial and temporal context effects are closely linked to the statistics of the natural environment, indicating that sensory circuits are tailored to the signals they are likely to encounter and that contextual modulation may play an important role in information processing.

Spatial context in visual coding

At the psychological level, spatial context produces robust and well-documented effects on perception. Figure 3A illustrates simultaneous contrast, one of the simplest and most powerful examples of spatial context. The two smaller squares are colored the same shade of gray and have the same *luminance* (actual amount of light traveling from the object to the eye), but because of the effect of the surrounding areas the square with the dark surround imparts a higher level of *brightness* (perceived luminance). Spatial context can strongly affect the perception of a wide number of visual features, in-

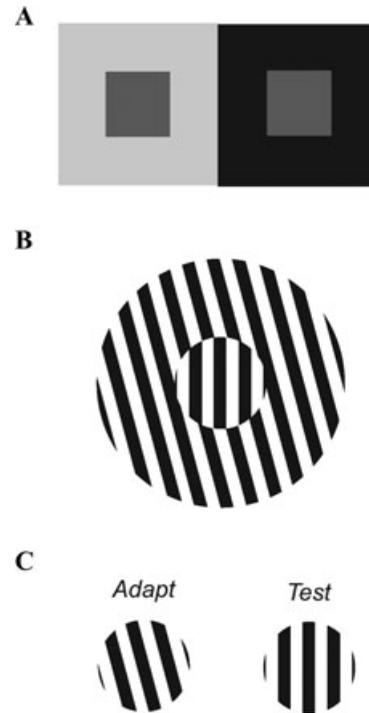


Figure 3. Spatial and temporal context effects in visual perception. (A) Simultaneous contrast. The two small squares are an identical shade of gray, but the surrounding spatial contexts drive a differential perception of brightness. (B) Tilt illusion. The presence of a spatial context tilted 15 degrees counterclockwise drives a perceived orientation of the center region that is tilted clockwise from the true vertical orientation. (C) Tilt after-effect. To induce this effect, fixate the adaptation stimulus for 30 sec and then shift fixation to the test stimulus. As in the tilt illusion, the presence of the temporal context (counterclockwise adaptation stimulus) induces a perceived orientation that is repulsed from the context (shifted clockwise).

cluding brightness, motion, orientation, and object recognition, and may underlie high-level processes like perceptual filling in, figure-ground segregation, and contour detection.⁸⁹ Figure 3B shows the tilt illusion, which represents a spatial contextual effect on orientation perception. Here, the surrounding bars oriented 15° counterclockwise away from vertical drives the perceived orientation of the center patch of vertical bars clockwise relative to the true vertical orientation.

Do the perceptual effects of context reflect a contextual modulation of neural responses? Neurons in most parts of the visual system respond to visual stimuli in a restricted portion of visual space termed the *classical receptive field* (CRF), a feature that applies to neurons from the earliest stage

of visual processing in the retina up to high-level visual cortical areas. This area is typically defined as the portion of the visual field in which a stimulus can elicit spiking activity; by definition, stimuli outside this area do not elicit a spiking response. However, there is a large body of evidence that neurons are modulated by stimuli falling outside the boundaries of the classical receptive field, in an area called the *extra-classical receptive field* (eRF) or *surround*. Many neurons show a differential response to the combination of a stimulus in the cRF and a stimulus in the surround compared to the cRF stimulus alone, an effect that is typically suppressive but can include facilitation as well.

Such contextual interactions are widespread in the visual pathway, and extra-classical effects on neural responses are observed at multiple levels of processing, from retina to lateral geniculate nucleus (LGN) to visual cortical areas, suggesting that contextual modulation may be a fundamental feature of sensory processing. Extra-classical modulation was reported by Hubel and Wiesel in their initial pioneering description of the primary visual cortex (V1), in which they found certain cells (which they called hypercomplex) that were tuned to the length of a bar stimulus, with firing rate increasing with length up to a certain magnitude but attenuating to longer bars. This *end-inhibition* is now known to be one of multiple examples of suppressive modulation driven by stimuli outside the cRF.⁹⁰ For example, many V1 neurons show selectivity for the orientation of stimuli within their receptive fields, with a unimodal tuning curve peaked at the optimal stimulus orientation. However, these neurons exhibit a strong contextual modulation termed *iso-orientation suppression*, in which cRF activity is most strongly suppressed by surround stimuli of the same orientation that optimally drives the cRF. Spatial context effects also affect higher levels of visual processing, such as the motion-sensitive neurons of the macaque middle temporal area (area MT). These extrastriate neurons are selective for the direction and speed of motion stimuli in their cRF, with unimodal direction tuning curves similar to V1 selectivity for orientation. Analogous to iso-orientation suppression in V1, stimuli in the surround drive a direction-selective modulation, producing a marked suppression when motion in the surround is the same as motion in the center.

Although spatial contextual modulation appears to be a fundamental feature of visual processing, the anatomic basis and influence on perception of such processes is still an area of active investigation. The mechanisms underlying these spatial modulatory effects may be diverse: although lateral inhibition mediated by intra-area horizontal connections is the standard explanation for surround suppression in V1,⁹¹ feedback projections from higher cortical areas⁹² or feedforward inheritance of surround suppression from the LGN⁹³ may also play a role. There is also a diversity of surround interactions, which can occur with either fast or slower dynamics and drive selectivity changes such as changes in tuning-curve width or shifts in the preferred direction. However, regardless of the underlying mechanism, sensory processing appears to be organized in a manner in which the neural representation of a given feature is a function of the spatial context in which it appears. The ubiquity of center-surround organization and contextual interactions may be related to the inherent structure present in the environment, as this form of processing is well suited to efficiently code natural stimuli⁹⁴ (see section “Efficient coding in sensory systems”). The influence of space represents an interaction between the neural representations of different inputs that appear simultaneously, and provides a framework to understand how decision areas encode the value-related activity of multiple choice options, which we explore later.

Temporal context in visual coding

The environment is dynamic, and one of the critical problems faced by the sensory system is processing a constant stream of changing stimuli. The effects of spatial context described previously reflect the interaction of different stimulus features at an instant in time, but stimuli also have a temporal context—the input stimuli in the recent past. The effect of temporal context is referred to by the general term *adaptation*, which describes the response to a sustained presentation of a stimulus (or stimulus distribution).

Like spatial contextual modulation, adaptation is a well-described phenomenon in both the perception and neurophysiology literature. An everyday example of adaptation is the ability to see at different levels of illumination, which is driven by adaptation to ambient illumination in the retina.⁹⁵ As an

observer moves from a high to a low illumination environment, for example carrying a newspaper from the sunlit outdoors into a darkened room, the perceived brightness of both the dark letters and the gray background remains stable. This adaptation to the local luminance allows the visual system to function over the vast range of possible light levels in the world despite the limited dynamic range of neural firing rates.

The visual system also adapts to a number of higher order visual features beyond simple luminance, though the link between such effects and functional benefits are not as obvious as that for luminance adaptation. There are clear perceptual adaptation effects to features including contrast (relative illumination), orientation, motion, spatial frequency, and even complex objects like faces (for a recent review, see Kohn⁹⁶). Figure 3C illustrates the tilt after-effect, a prominent example of orientation adaptation. Fixating the counterclockwise adaptation stimulus on the left for 30 sec and then shifting fixation to the target stimulus on the right induces a perception that the vertically oriented bars occur tilted clockwise, rotated away from the adapting orientation (a repulsive shift). Interestingly, the tilt after-effect provides a temporal counterpart to the spatial tilt illusion discussed previously; in the tilt illusion, the target and context occur at the same time but separated in space, whereas here the stimuli occur colocalized in space but separated in time.

At the mechanistic level, visual adaptation produces a diverse array of changes in the response of neurons in the visual system. Many of the earliest studies on adaptation studied the effects of presenting two different levels of a stimulus feature, a paradigm that investigates adaptation to the mean of the stimulus distribution. When V1 neurons are exposed to different levels of ambient contrast, the responses shift to encode higher levels of contrast, indicating a decrease in sensitivity.⁹⁷ Adaptation to higher-contrast stimuli induces larger reductions in sensitivity, maintaining the neuronal dynamic range close to the average of the recently experienced average contrast. Similar to spatial iso-orientation suppression, suppression of V1 responses is stronger if adapted to stimuli in the preferred versus the opposite or orthogonal orientation.⁹⁶ Adaptation to mean responses suppresses activity in a number of visual areas, including MT, V4, and the inferotemporal cortex.

However, when considering the temporal context, the average value of a stimulus is only one way of characterizing the distribution of recent stimuli. A system that responds to the local context of a dynamic environment will be influenced by the shape of the stimulus distribution if the timescale of integration is short relative to the timescale of changes. A number of studies have shown that visual areas can adapt to higher order statistics of local stimulus distributions, such as the variation in stimulus feature. In a study of vertebrate retinal ganglion cells, Smirnakis *et al.* presented random stimuli drawn from distributions with the same mean intensity but differing variances.⁹⁸ They found that retinal neurons adapted their responses to the width of the intensity distributions, an effect driven by recent sensory experience. When this temporal modulation is examined in the form of a linear kernel, which shows the average effect of a stimulus as a function of time from presentation, there is a clear temporally weighted dependency on recent intensity. Similar adaptations to stimulus variance have been demonstrated in other visual areas, such as LGN⁹⁹ and V1,¹⁰⁰ as well as in other modalities including somatosensation in rats¹⁰¹ and audition in songbirds.¹⁰² As in spatial contextual modulation, the effects of adaptation are likely mediated by heterogeneous mechanisms that differ by both locale of neuronal adaptation and timescale. Nevertheless, adaptation is a widespread feature of sensory processing, suggesting that neural circuits have evolved to respond to both the spatial and temporal statistics of the environment.

Efficient coding in sensory systems

Given the vast number of possible environmental inputs and the finite amount of neural hardware and metabolic energy possibly devoted to sensory processing, it is natural to assume that sensory systems evolved toward functioning as efficiently as possible. Although there are many possible definitions of efficiency, one enduring and influential proposal for a general principle of sensory system function is the *efficient coding hypothesis*. The fundamental idea is that sensory systems adapt their responses to the regularities of their input, and employ knowledge about these regularities to increase the amount of transmitted information at any given time. This approach relies heavily on the work of Shannon, who developed a quantitative theory of

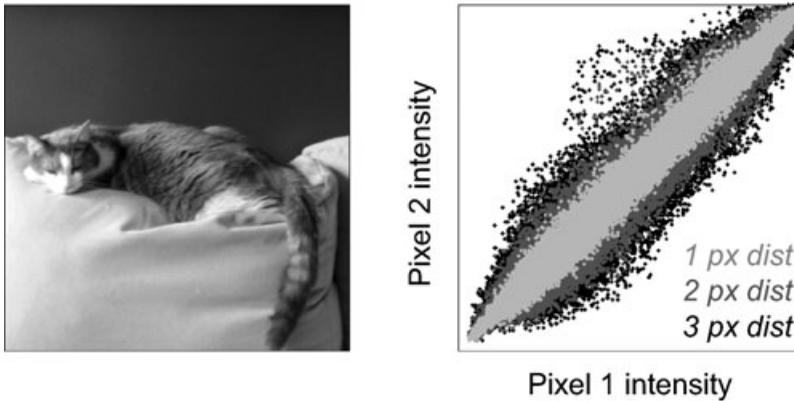


Figure 4. Structure in the sensory world. Sensory stimuli in the environment, such as the image of the cat, display significant statistical structure. For example, the luminance value of nearby pixels in the image are significantly correlated, an effect that exists for even nonadjacent pixels. Statistical structure in the sensory environment extends beyond simple two-point correlations, for example, to stereotyped spatial frequency characteristics. Neural systems can improve their coding efficiency by accounting for and reducing such information redundancy.

information fundamental to the field of communication.¹⁰³ Attneave applied these ideas to perception, suggesting that a guiding principle for sensory systems is the statistically efficient representation of available information.¹⁰⁴ Extending this idea to the neural level, Barlow proposed that the goal of early neurons in sensory processing is to remove the redundancy in the input stimuli.¹⁰⁵

These approaches are motivated by the fact that signals arising from the natural environment are highly structured.⁹⁴ Such structural regularities imply informational redundancy because an observer with knowledge about part of a signal can predict other parts of the signal with greater than chance probability. Consider the processing of visual information, from which most studies of natural statistics and empirical evidence for neural efficient coding have arisen. As initially pointed out by Attneave,¹⁰⁴ there is a significant degree of redundancy in natural visual images because of correlation in both the spatial and temporal domains. For example, as shown in Figure 4, if the responses of a pair of pixels separated by a fixed distance are examined across all such pairs in a natural image, this activity will be highly correlated. This kind of spatial correlation is clearly evident upon even a casual examination of a natural image, and such structural redundancies underlie current image compression and transmission technologies. These correlations reflect the underlying smoothness of natural images in both space and time: luminance primarily changes gradually

(with the exception of sharp transitions at edges), line segments vary as contours, and visual inputs change smoothly with time. Additional structure is evident when one considers the spatial and temporal dimensions together.¹⁰⁶ These statistical regularities constrain the images a visual system is likely to encounter to a tiny fraction of the set of all possible images, and visual circuits must be tuned to this probable subset in order to represent this information efficiently.

How is efficient coding evident in neural responses? At the level of single neurons, efficient coding requires that the input–output function be adjusted so that the entire response range is employed to represent the stimulus distribution.¹⁰⁷ For example, under the constraints of a maximum firing rate and finite precision, efficient neurons should employ all activity levels equally in response to the distribution (Fig. 5). If the input–output function sensitivity is set too low, high levels of the stimulus feature will be indistinguishable as the response function saturates; if the sensitivity is set too high, low levels of the stimulus feature cannot drive responses. One early demonstration of the precise correspondence between activity and natural stimulus statistics is Laughlin’s work on contrast-sensitive neurons in the fly compound eye.¹⁰⁸ Much like the model depicted in Figure 5, the contrast-response function of these neurons replicates the curve that transforms the probability distribution of natural contrasts into a flat response distribution,

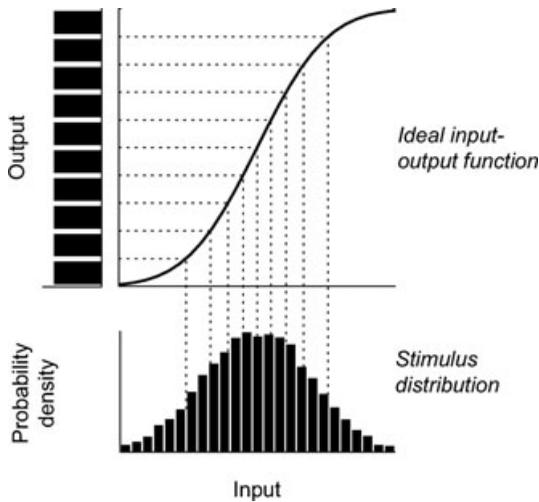


Figure 5. Efficient coding in sensory systems. For a given distribution of sensory characteristics in the world (bottom), an efficient neural input–output function produces an output (top) that equally uses all possible levels of neural activity. Such a function matches the greatest sensitivity of neural responses to the most probable stimuli in the environment, and has been demonstrated in neurophysiological data from sensory systems. (Adapted from Laughlin, 1981.¹⁰⁸)

thus matching neural activity to the environmental statistics. When the activity of multiple neurons are considered together, the efficient encoding hypothesis requires that the joint encoding of a stimulus should reflect both optimal responses in individual neurons and efficiency across the set of neurons. For example, to maximize efficiency and reduce redundancy, neural responses should be independent of one another (decorrelated), and a given stimulus should involve only a small fraction of the available neurons (sparse).

Recent work suggests that contextual effects in visual processing like surround modulation and adaptation serve to implement efficient coding representations. Theoretically, retinal and LGN center-surround structures implement a spatial decorrelation of outputs¹⁰⁹ (termed a *whitening* of the spectrum, in the frequency domain) and the nonlinear interactions mediated by cortical eRF interactions increase the sparseness of the output representations.¹¹⁰ Efficient coding has been proposed to explain the effect of context on neural responses, like the orientation tuning curve changes driven by surround stimuli.^{88,111,112} Empirically, when natural stimuli drive the surround of V1 neurons, responses are decorrelated and show a more efficient

sparse representation compared to cRF stimulation alone.^{113,114} To explore how contextual modulation increases the efficiency of sensory processing, Schwartz and Simoncelli explored the possible computational mechanisms linking the two. Specifically, they examined *divisive normalization*, a gain-control mechanism that characterizes contextual effects like surround suppression and contrast gain control. They found that models incorporating divisive normalization increase the independence of neural responses and allow for efficient encoding of natural visual and auditory signals.⁶⁷

Modulation by temporal context can also serve to improve the efficiency of sensory processing. The fundamental principle of efficient coding is that a sensory system is adjusted to the specific statistics of the natural environment from which it encodes and transmits information. However, if a sensory system is hard-wired to only the global, long-term average statistics of the world, it cannot efficiently transmit information if the short-term, local statistics vary. Experimental studies such as those reviewed previously suggest that sensory systems adapt to not only the mean of stimulus distributions, but to high-order statistics such as the variance. When coding efficiency is quantified, adaptation rescales neuronal input–output functions in a manner that maximizes the transmission of information.^{115,116}

An efficient coding framework for context-dependent value encoding

How can sensory context effects and an efficient coding framework illuminate our understanding of the representation of value in neural circuits? The efficient coding hypothesis proposes that sensory systems reduce the redundancy in natural signals (by increasing the efficiency of their output responses) in order to maximize the information that can be transmitted through limited capacity channels. Unlike sensory information, value is not a product of the environment alone, but a subjective construct determined by both external information and the state of the animal. Thus, value coding systems face both external constraints given by the statistics of reward distributions in the natural world and internal constraints governed by factors such as physiological status and metabolic needs. However, research in foraging theory shows that animals in the real world behave in manners distinct and highly suited to the reward structure in their environmental niches.⁵³

This suggests that the number of possible value distributions in the environment is constrained, and that the neural circuits driving behavior are tuned to this environmental structure.

We hypothesize here that value systems in the brain adopt coding strategies specific to their functional requirements, a feature that specifies the different forms of context-dependent value representation. In particular, we hypothesize that it is the statistics of the value distributions that each circuit encounters that will structure the form of the value representations. For action-value coding in decision-related systems, value representations arise simultaneously during action selection when a choice must be made between multiple options. For value systems involved in value and representation and storage, different value representations may be activated at different times depending on the environmental requirements. These response patterns are analogous to the manner in which the spatial and temporal aspects of incoming sensory signals are represented, and contextual effects in sensory processing offer a framework in which to explore the analogous context-dependent value representations.

As reviewed previously, the neural representation of action value appears to be instantiated in a relative manner, dependent on the other action values available at the time. This value representation is a spatial form of contextual modulation: analogous to extra-classical RF effects in the visual system, modulation of a given neuron is driven by stimuli that themselves do not drive the spiking activity. Although the function of the early sensory system is to efficiently transmit information about stimulus features, a decision circuit must select the best option and discriminate between the values of the possible choices. In terms of efficient coding, the simultaneous options should be represented optimally in neurons coding single actions and across the set of active neurons coding the choice set.

For a single neuron, divisive normalization acts as a means of compressive gain control even when a single option is presented alone. Because reward amounts are potentially limitless in the real world, a gain control mechanism operates to transform this wide range into the limited dynamic range of neural firing rates. Analogous to mechanisms like the one Laughlin described in the fly eye, divisive normalization may adjust the value input–output function

to efficiently encode the distribution of possible values. Seen in this light, the shape of the gain control observed in area LIP may reflect an underlying value distribution that is significantly skewed, with most possible options occurring at low values and a long tail of rare, higher values. The degree to which such compression increases the efficiency of the representation is difficult to quantify without knowledge about the natural statistics of rewards. Although ecological reward distributions have been studied for some animals and niches, the distribution of natural rewards in higher order primates, particularly over evolutionary timescales, remains unknown.

However, like sensory systems, decision systems must process multiple representations at the same instant in time. For neurons encoding the value of multiple actions, contextual modulation may serve to adjust the range of neuronal firing rates across the set of neurons to the value distribution of the choice set. Consider a decision system whose goal is to distinguish the higher valued of two rewards separated by a small amount when the values are low versus when they are high. In addition to providing a compression that keeps outputs within the dynamic range of single neurons, a relative value representation dependent on the total reward available adjusts the gain across the entire set of active neurons. Significantly, the divisive normalization computation that precisely characterizes relative value representation in the LIP also underlies spatial context effects in visual processing, raising the possibility that divisive normalization represents a canonical cortical computation that drives efficient gain control. This form of instantaneous normalization across active neurons may also apply to other brain areas that simultaneously represent the value of stimuli or actions. For example, areas that serve as saliency or attentional maps, representing simultaneous information from locations spanning the environment, may also implement a divisive normalization.⁶⁹

In the temporal domain, contextual effects in sensory processing have an analogue in the adaptation of neurons coding economic value in the OFC that reflects the proposed role of orbitofrontal neurons in the storage of value information. Unlike action-value coding neurons in the LIP or premotor cortex, OFC neurons are menu invariant and are unmodulated by the presence of other choice options. However, like neurons in the visual system that adapt to statistics of the stimulus feature

distribution, OFC responses show modulation by values encountered over a longer timescale. This adaptation is sensitive to multiple distribution statistics, including the mean, range, and variance of the recent value signals. Intuitively, adaptation effects improve the efficiency of coding by adjusting the input–output function of value neurons to the appropriate local statistics of input values. For example, when presented with reward distributions with identical means but different variances, OFC neurons adapt their responses so that their reward sensitivity slopes align with the probable environmental rewards. Using mutual information theory to quantify the ability of neurons to discriminate rewards, Kobayashi *et al.* showed that adaptive neurons in OFC preserve the amount of encoded information regardless of the input statistics.⁷⁵ Thus, adaptive processes in value storage areas may produce a more efficient neural representation of value for use in downstream decision processes.

Caveats

There remains much that is unknown about valuation systems in the brain and their relationship to the environment. Unlike natural sensory signals, the distribution of values in the natural environment is not nearly as easily defined or measured, particularly if one wishes to examine the statistics of values over an evolutionary timescale. It is comparatively easy to extract the statistics of the sensory environment, but the statistics of value will always be subject to the interaction between an organism and its environment. Experimenters are beginning to address this issue in the laboratory by constraining the statistics of local values, an approach particularly suited to the study of adaptive processes.

The parallels between contextual effects in value coding and sensory processing reviewed previously may reflect a shared functional architecture (like divisive normalization), a shared design principle (dependence on natural statistics), or both. Although the principle of efficient coding in sensory processing provides an attractive framework to examine contextual effects in value representation, it is important to note that the sensory system and valuation networks have different functional goals. Information theory approaches representation as purely a problem of transmission, concerned with maximizing the amount of information in the signal while reducing redundancy. This approach has reliably

characterized many of the early stages of sensory processing, in which the primary goal is to transmit information about the environment to the rest of the brain, but higher order brain areas are likely to have functional goals beyond strict transmission. Ultimately, the selection pressure for all neural systems, including sensory ones, is not maximizing the efficient representation of information *per se* but maximizing the survival of the animal. Given the high metabolic cost of neural systems, however, evolution should favor an optimal use of resources, regardless of whether this is implemented as efficiency in representation, as proposed by the efficient-coding hypothesis, or as another constraint. It is likely that efficient coding will not be the only principle through which to understand the design of value systems, but it provides an attractive starting point to examine the nature of value processing.

Conflicts of interest

The authors declare no conflicts of interest.

References

1. Von Neumann, J. & O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press. Princeton, NJ.
2. Samuelson, P.A. 1947. *Foundations of Economic Analysis*. Harvard University Press. Cambridge.
3. Kable, J.W. & P.W. Glimcher. 2007. The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**: 1625–1633.
4. Kable, J.W. & P.W. Glimcher. 2009. The neurobiology of decision: consensus and controversy. *Neuron* **63**: 733–745.
5. Rangel, A. & T. Hare. 2010. Neural computations associated with goal-directed choice. *Curr. Opin. Neurobiol.* **20**: 262–270.
6. Platt, M.L. & P.W. Glimcher. 1999. Neural correlates of decision variables in parietal cortex. *Nature* **400**: 233–238.
7. Dorris, M.C. & P.W. Glimcher. 2004. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* **44**: 365–378.
8. Sugrue, L.P., G.S. Corrado & W.T. Newsome. 2004. Matching behavior and the representation of value in the parietal cortex. *Science* **304**: 1782–1787.
9. Seo, H., D.J. Barraclough & D. Lee. 2009. Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. *J. Neurosci.* **29**: 7278–7289.
10. Louie, K. & P.W. Glimcher. 2010. Separating value from choice: delay discounting activity in the lateral intraparietal area. *J. Neurosci.* **30**: 5498–5507.
11. Ding, L. & O. Hikosaka. 2006. Comparison of reward modulation in the frontal eye field and caudate of the macaque. *J. Neurosci.* **26**: 6695–6703.

12. So, N.Y. & V. Stuphorn. 2010. Supplementary eye field encodes option and action value for saccades with variable reward. *J. Neurophysiol.* **104**: 2634–2653.
13. Dorris, M.C. & D.P. Munoz. 1998. Saccadic probability influences motor preparation signals and time to saccadic initiation. *J. Neurosci.* **18**: 7015–7026.
14. Ikeda, T. & O. Hikosaka. 2003. Reward-dependent gain and bias of visual responses in primate superior colliculus. *Neuron* **39**: 693–700.
15. Hollerman, J.R., L. Tremblay & W. Schultz. 1998. Influence of reward expectation on behavior-related neuronal activity in primate striatum. *J. Neurophysiol.* **80**: 947–963.
16. Kawagoe, R., Y. Takikawa & O. Hikosaka. 1998. Expectation of reward modulates cognitive signals in the basal ganglia. *Nat. Neurosci.* **1**: 411–416.
17. Cromwell, H.C. & W. Schultz. 2003. Effects of expectations for different reward magnitudes on neuronal activity in primate striatum. *J. Neurophysiol.* **89**: 2823–2838.
18. Samejima, K. *et al.* 2005. Representation of action-specific reward values in the striatum. *Science* **310**: 1337–1340.
19. Lau, B. & P.W. Glimcher. 2008. Value representations in the primate striatum during matching behavior. *Neuron* **58**: 451–463.
20. Musallam, S. *et al.* 2004. Cognitive control signals for neural prosthetics. *Science* **305**: 258–262.
21. Hori, Y., T. Minamimoto & M. Kimura. 2009. Neuronal encoding of reward value and direction of actions in the primate putamen. *J. Neurophysiol.* **102**: 3530–3543.
22. Padoa-Schioppa, C. & J.A. Assad. 2006. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**: 223–226.
23. Padoa-Schioppa, C. 2011. Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* **34**: 333–359.
24. Thorpe, S.J., E.T. Rolls & S. Maddison. 1983. The orbitofrontal cortex: neuronal activity in the behaving monkey. *Exp. Brain Res.* **49**: 93–115.
25. Tremblay, L. & W. Schultz. 1999. Relative reward preference in primate orbitofrontal cortex. *Nature* **398**: 704–708.
26. Wallis, J.D. & E.K. Miller. 2003. Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur. J. Neurosci.* **18**: 2069–2081.
27. Roesch, M.R. & C.R. Olson. 2004. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* **304**: 307–310.
28. Ongur, D. & J.L. Price. 2000. The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cereb. Cortex* **10**: 206–219.
29. Paton, J.J. *et al.* 2006. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* **439**: 865–870.
30. Belova, M.A., J.J. Paton & C.D. Salzman. 2008. Moment-to-moment tracking of state value in the amygdala. *J. Neurosci.* **28**: 10023–10030.
31. Wise, R.A. 1996. Addictive drugs and brain stimulation reward. *Annu. Rev. Neurosci.* **19**: 319–340.
32. Wise, R.A. 2004. Dopamine, learning and motivation. *Nat. Rev. Neurosci.* **5**: 483–494.
33. Romo, R. & W. Schultz. 1990. Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements. *J. Neurophysiol.* **63**: 592–606.
34. Ljungberg, T., P. Apicella & W. Schultz. 1992. Responses of monkey dopamine neurons during learning of behavioral reactions. *J. Neurophysiol.* **67**: 145–163.
35. Mirenowicz, J. & W. Schultz. 1994. Importance of unpredictability for reward responses in primate dopamine neurons. *J. Neurophysiol.* **72**: 1024–1027.
36. Hollerman, J.R. & W. Schultz. 1998. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat. Neurosci.* **1**: 304–309.
37. Montague, P.R. *et al.* 1995. Bee foraging in uncertain environments using predictive hebbian learning. *Nature* **377**: 725–728.
38. Montague, P.R., P. Dayan & T.J. Sejnowski. 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* **16**: 1936–1947.
39. Glimcher, P.W. 2011. Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proc. Natl. Acad. Sci. USA* **108**(Suppl 3): 15647–15654.
40. Kalman, R.E. 1960. A new approach to linear filtering and prediction problems. *J. Basic Eng. Trans. ASME* **82**: 35–45.
41. Dickinson, A. 1980. *Contemporary Animal Learning Theory*. Cambridge University Press. Cambridge [Eng.]; New York.
42. Sutton, R.S. & A.G. Barto. 1981. Toward a modern theory of adaptive networks: expectation and prediction. *Psychol. Rev.* **88**: 135–170.
43. Sutton, R.S. & A.G. Barto. 1998. *Reinforcement Learning: An Introduction*. MIT Press. Cambridge, MA.
44. Schultz, W., P. Dayan & P.R. Montague. 1997. A neural substrate of prediction and reward. *Science* **275**: 1593–1599.
45. Waelti, P., A. Dickinson & W. Schultz. 2001. Dopamine responses comply with basic assumptions of formal learning theory. *Nature* **412**: 43–48.
46. Bayer, H.M. & P.W. Glimcher. 2005. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* **47**: 129–141.
47. Glimcher, P.W. 2011. *Foundations of Neuroeconomic Analysis*. Oxford University Press. New York.
48. Maunsell, J.H. 2004. Neuronal representations of cognitive state: reward or attention? *Trends Cogn. Sci.* **8**: 261–265.
49. Gottlieb, J.P., M. Kusunoki & M.E. Goldberg. 1998. The representation of visual salience in monkey parietal cortex. *Nature* **391**: 481–484.
50. Bisley, J.W. & M.E. Goldberg. 2003. Neuronal activity in the lateral intraparietal area and spatial attention. *Science* **299**: 81–86.
51. Bisley, J.W. & M.E. Goldberg. 2010. Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* **33**: 1–21.
52. Luce, R.D. 1959. *Individual Choice Behavior: A Theoretical Analysis*. Wiley. New York.

53. Stephens, D.W. & J.R. Krebs. 1986. *Foraging Theory*. Princeton University Press. Princeton, NJ.
54. Herrnstein, R.J. 1961. Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav.* **4**: 267–272.
55. Shadlen, M.N. & W.T. Newsome. 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**: 1916–1936.
56. Toth, L.J. & J.A. Assad. 2002. Dynamic coding of behaviourally relevant stimuli in parietal cortex. *Nature* **415**: 165–168.
57. Leon, M.I. & M.N. Shadlen. 2003. Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* **38**: 317–327.
58. Yang, T. & M.N. Shadlen. 2007. Probabilistic reasoning by neurons. *Nature* **447**: 1075–1080.
59. Rorie, A.E. *et al.* 2010. Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS One* **5**: e9308.
60. Roitman, J.D. & M.N. Shadlen. 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**: 9475–9489.
61. Heeger, D.J. 1992. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**: 181–197.
62. Carandini, M., D.J. Heeger & J.A. Movshon. 1997. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* **17**: 8621–8644.
63. Britten, K.H. & H.W. Heuer. 1999. Spatial summation in the receptive fields of MT neurons. *J. Neurosci.* **19**: 5074–5084.
64. Cavanaugh, J.R., W. Bair & J.A. Movshon. 2002. Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *J. Neurophysiol.* **88**: 2530–2546.
65. Heuer, H.W. & K.H. Britten. 2002. Contrast dependence of response normalization in area MT of the rhesus macaque. *J. Neurophysiol.* **88**: 3398–3408.
66. Zoccolan, D., D.D. Cox & J.J. DiCarlo. 2005. Multiple object response normalization in monkey inferotemporal cortex. *J. Neurosci.* **25**: 8150–8164.
67. Schwartz, O. & E.P. Simoncelli. 2001. Natural signal statistics and sensory gain control. *Nat. Neurosci.* **4**: 819–825.
68. Valerio, R. & R. Navarro. 2003. Optimal coding through divisive normalization models of V1 neurons. *Network* **14**: 579–593.
69. Reynolds, J.H. & D.J. Heeger. 2009. The normalization model of attention. *Neuron* **61**: 168–185.
70. Carandini, M. & D.J. Heeger. 2012. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**: 51–62.
71. Louie, K., L.E. Grattan & P.W. Glimcher. 2011. Reward value-based gain control: divisive normalization in parietal cortex. *J. Neurosci.* **31**: 10627–10639.
72. Pastor-Bernier, A. & P. Cisek. 2011. Neural correlates of biased competition in premotor cortex. *J. Neurosci.* **31**: 7083–7088.
73. Padoa-Schioppa, C. & J.A. Assad. 2008. The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat. Neurosci.* **11**: 95–102.
74. Padoa-Schioppa, C. 2009. Range-adapting representation of economic value in the orbitofrontal cortex. *J. Neurosci.* **29**: 14004–14014.
75. Kobayashi, S., O. Pinto de Carvalho & W. Schultz. 2010. Adaptation of reward sensitivity in orbitofrontal neurons. *J. Neurosci.* **30**: 534–544.
76. Tobler, P.N., C.D. Fiorillo & W. Schultz. 2005. Adaptive coding of reward value by dopamine neurons. *Science* **307**: 1642–1645.
77. Samuelson, P.A. 1947. *Foundations of Economic Analysis*. Harvard University Press. Cambridge, MA.
78. Tversky, A. 1972. Elimination by aspects: a theory of choice. *Psychol. Rev.* **79**: 281–8.
79. Huber, J., J.W. Payne & C. Puto. 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *J. Consum. Res.* **9**: 90–98.
80. Shafir, S., T.A. Waite & B.H. Smith. 2002. Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behav. Ecol. Sociobiol.* **51**: 180–187.
81. Bateson, M., S.D. Healy & T.A. Hurly. 2003. Context-dependent foraging decisions in rufous hummingbirds. *P. Roy. Soc. Lond. B Bio.* **270**: 1271–1276.
82. Tversky, A. & E. Shafir. 1992. Choice under conflict: the dynamics of deferred decision. *Psychol. Sci.* **3**: 358–361.
83. Iyengar, S.S. & M.R. Lepper. 2000. When choice is demotivating: can one desire too much of a good thing? *J. Pers. Soc. Psychol.* **79**: 995–1006.
84. DeShazo, J.R. & G. Fermo. 2002. Designing choice sets for stated preference methods: the effects of complexity on choice consistency. *J. Environ. Econ. Manag.* **44**: 123–143.
85. Tolhurst, D.J., J.A. Movshon & A.F. Dean. 1983. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* **23**: 775–785.
86. Softky, W.R. & C. Koch. 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* **13**: 334–350.
87. Shadlen, M.N. & W.T. Newsome. 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**: 3870–3896.
88. Schwartz, O., A. Hsu & P. Dayan. 2007. Space and time in visual context. *Nat. Rev. Neurosci.* **8**: 522–535.
89. Albright, T.D. & G.R. Stoner. 2002. Contextual influences on visual processing. *Annu. Rev. Neurosci.* **25**: 339–379.
90. DeAngelis, G.C., R.D. Freeman & I. Ohzawa. 1994. Length and width tuning of neurons in the cat's primary visual cortex. *J. Neurophysiol.* **71**: 347–374.
91. Gilbert, C.D. & T.N. Wiesel. 1985. Intrinsic connectivity and receptive field properties in visual cortex. *Vision Res.* **25**: 365–374.
92. Bair, W., J.R. Cavanaugh & J.A. Movshon. 2003. Time course and time-distance relationships for surround suppression in macaque V1 neurons. *J. Neurosci.* **23**: 7690–7701.
93. Ozeki, H. *et al.* 2004. Relationship between excitation and inhibition underlying size tuning and contextual response

- modulation in the cat primary visual cortex. *J. Neurosci.* **24**: 1428–1438.
94. Simoncelli, E.P. & B.A. Olshausen. 2001. Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**: 1193–1216.
 95. Shapley, R.M. & C. Enroth-Cugell. 1984. Visual adaptation and retinal gain control. *Prog. Ret. Res.* **3**: 263–346.
 96. Kohn, A. 2007. Visual adaptation: physiology, mechanisms, and functional benefits. *J. Neurophysiol.* **97**: 3155–3164.
 97. Ohzawa, I., G. Sclar & R.D. Freeman. 1982. Contrast gain control in the cat visual cortex. *Nature* **298**: 266–268.
 98. Smirnakis, S.M. *et al.* 1997. Adaptation of retinal processing to image contrast and spatial scale. *Nature* **386**: 69–73.
 99. Bonin, V., V. Mante & M. Carandini. 2006. The statistical computation underlying contrast gain control. *J. Neurosci.* **26**: 6346–6353.
 100. Sharpee, T.O. *et al.* 2006. Adaptive filtering enhances information transmission in visual cortex. *Nature* **439**: 936–942.
 101. Maravall, M. *et al.* 2007. Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biol.* **5**: e19.
 102. Nagel, K.I. & A.J. Doupe. 2006. Temporal processing and adaptation in the songbird auditory forebrain. *Neuron* **51**: 845–859.
 103. Shannon, C.E. & W. Weaver. 1949. *The Mathematical Theory of Communication*. University of Illinois Press. Urbana, IL.
 104. Attneave, F. 1954. Some informational aspects of visual perception. *Psychol. Rev.* **61**: 183–193.
 105. Barlow, H.B. 1961. Possible principles underlying the trans-formation of sensory messages. In *Sensory Communication*. W.A. Rosenblith, Ed. MIT Press. Cambridge, MA.
 106. van Hateren, J.H. & D.L. Ruderman. 1998. Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc. Biol. Sci.* **265**: 2315–2320.
 107. Simoncelli, E.P. 2003. Vision and the statistics of the visual environment. *Curr. Opin. Neurobiol.* **13**: 144–149.
 108. Laughlin, S. 1981. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch. C.* **36**: 910–912.
 109. Atick, J.J. 1992. Could information theory provide an ecological theory of sensory processing? *Network* **3**: 213–251.
 110. Olshausen, B.A. & D.J. Field. 1997. Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res.* **37**: 3311–3325.
 111. Muller, J.R. *et al.* 2003. Local signals from beyond the receptive fields of striate cortical neurons. *J. Neurophysiol.* **90**: 822–831.
 112. Felsen, G., J. Touryan & Y. Dan. 2005. Contextual modulation of orientation tuning contributes to efficient processing of natural stimuli. *Network* **16**: 139–149.
 113. Vinje, W.E. & J.L. Gallant. 2000. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**: 1273–1276.
 114. Vinje, W.E. & J.L. Gallant. 2002. Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. *J. Neurosci.* **22**: 2904–2915.
 115. Brenner, N., W. Bialek & R. de Ruyter van Steveninck. 2000. Adaptive rescaling maximizes information transmission. *Neuron* **26**: 695–702.
 116. Fairhall, A.L. *et al.* 2001. Efficiency and ambiguity in an adaptive neural code. *Nature* **412**: 787–792.