

Gauging sensory representations in the brain

Giedrius T. Buračas and Thomas D. Albright

The stream of information that enters a sensory system is a product of the ecological niche of an organism and the way in which the information is sampled. The most salient characteristic of this sensory stream is the rich temporal structure that is caused by changes in the environment and self motion of sensors (for example, rapid eye or whisker movements). In recent years, substantial progress has been made in understanding how such rapidly varying stimuli are represented in the responses of sensory neurons of a large variety of sensory systems. The crucial observation that has emerged from these studies is that individual action potentials convey substantial amounts of information, which permits the discrimination of rapidly varying stimuli with high temporal precision.

Trends Neurosci. (1999) 22, 303–309

COGNITIVE NEUROSCIENCE hinges on the doctrine that the brain represents the world in patterns of neuronal activity. Recently, this ‘representationalist’ credo was forcefully restated by Blakemore and Movshon¹. ‘The task of sensory systems is to provide a faithful representation of biologically relevant events in the external environment... These representations are [rich] because they contain representations of objects, states, and events that are abstracted from the primitive sensory signals; they are [simple] because they represent the distillation of the vast quantities of raw measurement information offered to the central nervous system by each sensory surface.’ As early as four decades ago Lettvin *et al.*² anticipated this credo in the classical paper ‘What the frog’s eye tells the frog’s brain’, which phrased the neuronal sensitivities of amphibian tectum neurons in terms of biological, or ecological, relevance of the represented stimuli.

The advent of the system-theoretic approach in the 1960s with its stimuli motivated by system identification and signal-filtering considerations, however, pushed the ecological relevance of the stimuli aside as a poorly defined idea. Needless to say, the system-theoretic approach has proven to be of great value in deciphering many important aspects of early sensory processing, albeit at the price of missing the key question: how do sensory systems function when confronted with stimuli that exist in their ecological niches? Probably the most profound difference between conventional impoverished laboratory stimuli and ecologically relevant stimuli is that the parameters of the former stimuli are kept constant for extended periods of time (≥ 1 s), while the latter stimuli are in permanent flux caused by changes either in the environment or the observer. Recent experiments involving the stimulation of neurons in insect, amphibian and mammalian sensory systems on ecologically relevant time scales revealed not only that response properties were predictable from responses to conventional constant stimuli, but also some exciting hitherto overlooked properties, which, thus, highlighted the importance of the temporal factor in sensory signaling.

This article gives a cursory update on the advances in studies of sensory representations that vary on ecologically relevant time scales (reviewed previously by

Bialek and Rieke³). These advances have been stimulated by novel information-theoretic approaches that can be used to gauge the ‘richness’ and ‘simplicity’ of neuronal representations of rapidly varying stimuli^{4–7}, and by accumulation of neurophysiological data from a wide variety of species.

From spikes to representations

Before embarking on a survey of recent developments in mapping neuronal representations of ‘eco-relevant’ stimuli, an explanation of the basic assumptions employed in these studies is necessary. The information-theoretic approaches referred to in the previous paragraph aim to provide a rigorous evaluation of the statistical relationship between a stimulus set and neuronal response. One might argue that the conventional mapping of tuning curves and the measuring of neuronal sensitivities by means of various selectivity indices are directed towards the same goal. Indeed, these measures have proved to be of great value when neurons exhibit smooth tuning curves with a single maximum (preferred stimulus values) and the response-probability distributions can be satisfactorily characterized by the first two moments (see also Ref. 8). No obvious extensions of these measures appear to be available for time-varying stimuli. More importantly, these measures are descriptions of a neuronal signal as a function of stimulus. The idea of neuronal ‘representation’ of a stimulus, however, invokes the inverse relationship: that of stimulus as a function of neuronal response. Indeed, from the organism’s perspective, only the latter relationship is meaningful as, unlike an experimenter, who is pre-occupied with describing neuronal responses to stimuli drawn from a domain of some sensory feature space, an organism is routinely engaged in the inverse task: inference of stimulus value from a neuronal response⁹.

The commonplace measure that comes closest to capturing this type of inference is the discriminability analysis [that is, receiver operating characteristics (ROC)] derived from the signal-detection theory (see Box 1), which evaluates how well an ‘ideal’ observer would tell apart two alternative stimulus values by only looking at a response of a single neuron. While this method provides an explicit measure of discrimination performance, which is very useful in comparisons of

Giedrius T. Buračas and Thomas D. Albright are at the Howard Hughes Medical Institute, Salk Institute for Biological Studies, La Jolla, CA 92037, USA.

Box 1. Prolegomena of information theory

Shannon's mutual information^{a,b} can be considered as a generalized measure of stimulus discriminability. The following thought experiment underscores the analogy between this measure and the conventional measure of signal discriminability – the receiver operating characteristics (ROC) of the signal detection theory, which is applicable in the context of two-alternative-forced-choice tasks. Imagine an observer who relies on responses of two perfect neurons, the visual neuron is capable of perfect discrimination between red and yellow, and the tactile neuron is capable of perfect discrimination between apple and cherry shapes. Our stimulus set consists of four possible combinations of fruit and color (Fig. 1A). If all stimuli have equal probability of appearance ($P = 1/4$), then this stimulus set offers $H(\text{stimulus}) = H(\text{color, shape})$, that is, $\log_2 4 = 2$ bits of potential information, or in Shannon's terms – uncertainty (also called 'entropy'). If the observer obtains only a response r_{shape} of the 'shape neuron' to a randomly drawn stimulus, all uncertainty associated with shape is resolved, but uncertainty about color remains: $H(\text{color} | r_{\text{shape}}) = \log_2 2$, that is, 1 bit. The gained total mutual information is $I(\text{stimulus}; r_{\text{shape}}) = H(\text{stimulus}) - H(\text{color} | r_{\text{shape}})$, which equals 1 bit (Fig. 1A, bottom). Hence, the mutual information is equal to the amount by which uncertainty about the stimulus is reduced after observing a response^{a,b}. Note that one bit is equivalent to the perfect discrimination between two possibilities, which corresponds to a ROC area of 1 (100% correct discrimination). Likewise, $I(\text{stimulus}; r_{\text{shape}}) = 0$ corresponds to chance performance (ROC = 0.5, that is, 50% of correct discriminations). If the shape neuron is noisy and the distributions of responses to the two shapes are unimodal, mutual information, like ROC, increases as the distance between the medians of the two distributions grows (Fig. 1B, left). If these distributions are bimodal, as in Fig. 1B (right), the ROC measure breaks down, while the mutual information reflects the best possible discriminability of the two stimuli correctly, which decreases as the overlap of the two distributions increases (Fig. 1B, right). Mutual information becomes especially valuable when judging discriminability of large stimulus sets.

In general, the mutual information can be written as a function of both the conditional and marginal stimulus and response probability distributions:

$$I(r;x) = H(x) - H(x|r) \\ = - \sum_x p(x) \log_2 p(x) - \sum_r p(r) \sum_x p(x|r) \log_2 p(x|r)$$

Here $p(x)$ is the *a priori* distribution of the stimulus (s) parameter x , and $p(x|r)$ is the *a posteriori* conditional distribution of the parameter x for a given response r . This equation implies that in order to evaluate the conditional entropy $H(x|r)$, one first has to convert back from the neuronal

response distribution, to stimulus probability distribution by means of Bayes rule (see Fig. 1C):

$$p(x|r) = c \cdot p(r|x)p(x)$$

where $c = 1/p(r)$. Estimation of the above equation is a simple task for constant-stimulus sets as their size only rarely exceeds 2^7 stimuli (for an equiprobable stimulus set of this size the uncertainty is 7 bits). By contrast, time-varying stimuli might offer uncertainty rates of 100 bits/s, which implies that the stimulus ensemble could be composed of 2^{100} ($\sim 10^{30}$) different stimuli of 1 s duration! These rich stimulus ensembles are much closer to the 'real-world' stimuli, but they require a special approach, such as optimal linear reconstruction (reconstruction method), which was developed by Bialek and colleagues^{c,d}. This method can be used to estimate a lower bound on information about the stimulus parameter x .

Recently, a number of researchers have produced an alternative method for the estimation of information rates, which is based on the fact that mutual information is symmetrical with respect to stimulus and response uncertainties (Fig. 1D):

$$I(r;s) = H(s) - H(s|r) = H(r) - H(r|s) \\ = H(r) - \sum_s p(s) H(r|s = s)$$

The symbol s represents the stimulus as a function, probably stochastic, of parameters x and y : $s = f(x,y)$. The rightmost expression suggests that one can evaluate information rates by using two types of stimuli: (1) for the evaluation of the first term $H(r)$ – (total spike-train uncertainty) one needs a large range of stimuli drawn randomly from a given stimulus ensemble, because these stimuli have to test the limits of response variability; and (2) the second term $H(r|s)$ reflects the variability of the response when the stimulus is fixed, and can be evaluated from responses to a typical stimulus instance repeated many times (for details, see Ref. e). As this method (the 'direct method') is not limited by the quality of stimulus reconstruction, it can yield, given a sufficient amount of data, an accurate estimate of information that is associated with the stimulus ensemble. In addition, because this method does not use information about stimulus parameters, it evaluates the total information that a neuron carries about all aspects of the stimulus.

References

- a Shannon, C.E. (1948) *Bell Sys. Tech. J.* 27, 379–423
- b Cover, T.M. and Thomas, J.A. (1991) *Elements of Information Theory*, Wiley
- c Bialek, W. and Rieke, F. (1992) *Trends Neurosci.* 15, 428–433
- d Rieke, F. et al. (1997) *Spikes*, MIT Press
- e Strong, S.P. et al. (1998) *Phys. Rev. Lett.* 80, 197

neuronal responses with responses of the whole animal (see, for example, Ref. 10), the main limitation of this method is that it does not generalize readily to cases when stimulus sets are large or infinite, as in the case of continuous stimuli, or when discriminant functions are highly nonlinear or not specified at all (or both). The natural generalization of this discriminability measure appears to be Shannon's mutual information, which is a well-understood measure that captures stimulus–response correlations of all orders and has a number of very appealing properties (see Box 1 for a direct comparison of ROC and mutual information). Most importantly, this measure addresses the question that is of key interest to most neurophysiologists: how much information about my stimulus set does the neur-

onal response carry? Furthermore, Shannon information offers tools for tracing information about what is being computed and what information is being discarded as the raw sensory information proceeds to downstream cortical areas.

The information rates that are cited in subsequent sections for time-varying stimuli are calculated using methods ('direct' and 'reconstruction' methods) that are merely special cases of Shannon's formula (Box 1) for mutual information applied to time-varying signals. The key modification in these special cases is that one has to deal with the probability distribution of spike-train vectors rather than simply a probability distribution of the scalar spike count over a particular fixed time window. This modification could be costly with

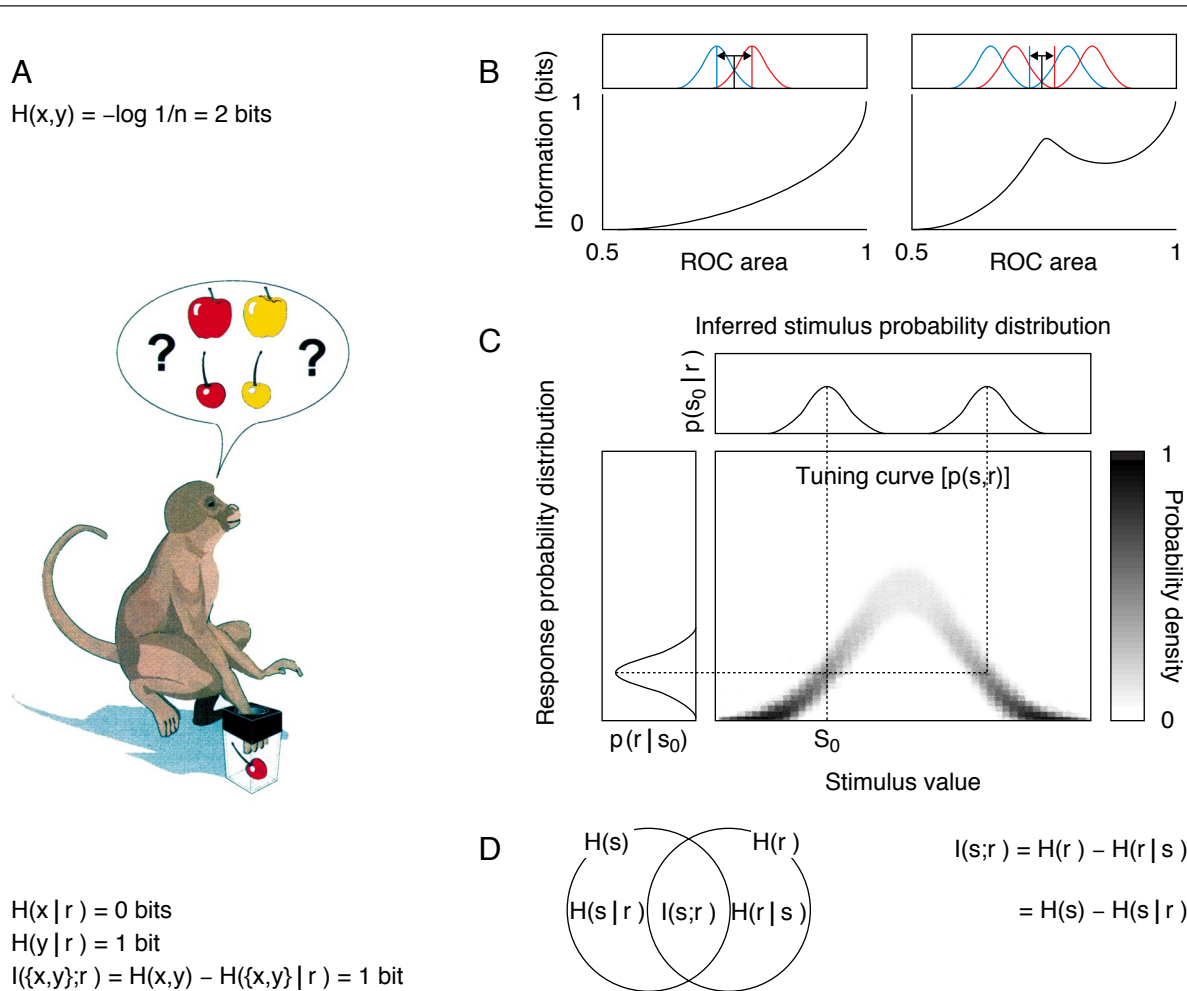


Fig. 1. Mutual information is the reduction of uncertainty after observing a signal. (A) The stimulus set consists of four different stimuli. Each stimulus appears with equal probability of 0.25. Hence, the initial stimulus uncertainty $[H(x,y)]$ is 2 bits. The subject possesses neurons that convey the signal, r , about the color (x) and about the shape (y), which permits a perfect discrimination of the feature. After observing (feeling) the stimulus shape the observer is left with uncertainty about the unobserved feature (cartoon). The gained information $[I(s;r)]$ is the difference between the initial uncertainty and uncertainty after the signal was observed. Observing the responses of both neurons conveys 2 bits about the stimulus s . This is equivalent to a perfect discrimination among all $2^{L(s;r)} = 4$ stimuli. (B) The receiver operating characteristics (ROC) and mutual information are monotonically related when probability distributions are unimodal (left), but not when they are bi- or multi-modal (right). The top panels plot the distributions of responses to two different stimuli. Vertical lines represent the criterion value for optimal discrimination by a linear ideal observer (ROC). The bottom curves show the relationship between ROC and mutual information as the distance between the modes of the two distributions vary from zero to infinity. (C) Calculation of mutual information from stimulus uncertainties involves reconstruction of the stimulus probability distribution. s_0 is the parameter of a presented stimulus. The probability distribution of neuronal responses associated with this stimulus $[p(r|s_0)]$ is plotted on the left. It can be obtained from the joint stimulus and response probability $[p(s,r)]$ depicted in the center. The top panel shows the stimulus probability distribution $[p(s_0|r)]$ inferred from the response to stimulus, s_0 . (D) Mutual information is symmetrical with respect to stimulus and response uncertainties. The circles are diagrams that depict the uncertainty of the stimulus and response sets $[H(s)$ and $H(r)$, respectively]. The area of their intersection corresponds to the mutual information. Hence, the mutual information $[I(s;r)]$ can be estimated from the total response uncertainty $[H(r)]$ (that is, the range of possible neuronal responses) and the uncertainty $[H(r|s)]$ that remains after presenting a specific stimulus (noise).

respect to data needed for obtaining reliable estimates of mutual information, but it permits the rigorous treatment of the temporal aspect of stimulus-response correlations. Understanding details of these methods is not necessary for subsequent sections in this article. More-detailed descriptions of the temporal reconstruction method and the general stimulus-reconstruction approach are covered in Refs 3 and 11, respectively; and the direct method is discussed in detail in Refs 12 and 13.

Neuronal representations of time-varying stimuli

Until recently, beliefs about how information in a neuronal response is encoded were influenced by an early observation that multiple presentations of identical stimuli yielded highly variable numbers of action

potentials¹⁴. In order to obtain a good estimate of average firing rates in response to a particular stimulus, the duration of stimulus presentation had to be extended to lengths in the order of 1 second. Implanted in this manipulation is the assumption that information about exact spike arrival times is largely irrelevant.

A decade ago this assumption was challenged by a series of experiments that exploited a newly developed method to estimate how well sensory systems could discriminate between stimuli that varied at rates commonly found under natural conditions (see, for example, Ref. 3). This method, based on the organism's perspective, reconstructed the stimulus solely from the observed spike train. The results of experiments, which employed this stimulus-reconstruction method, demonstrated

TABLE 1. Information rates and coding efficiency in a number of species

Neural system and species	Information rate (bits/spike)	Information rate (bits/s)	Coding efficiency ^a	Best timing precision and variance:mean	Refs
Constant stimulus^b					
Cat retinal ganglion cells	~0.04–0.10	0.4–0.8 ^c	–	–	15
Primary visual cortex (V1) of rhesus monkeys	–	0.62 ^d	–	–	16
Middle temporal area (MT) of rhesus monkeys	0.025	0.89 ± 0.29	–	~1.3	12
Inferior temporal area (IT) of rhesus monkeys	~0.13	0.9 ^d	–	–	17–19
Hippocampus of rhesus monkeys	~0.18 ^e	0.32 (maximum = 1.2) ^f	–	–	20
Variable stimulus: reconstruction method					
H1 motion-sensitive neuron of a fly	0.75	64	30%	–	21
Frog auditory afferents	0.66	23	11%	~2 ms	22
Vibratory receptors of the bullfrog sacculus	2.6	155	50–60%	–	22
Crickets mechanoreceptors	0.6–3.2	75–294	50–60%	~0.4 ms	22,23
Salamander retinal ganglion cells	1.9	3.7 (up to 10 for a population of >10 cells)	26% (>79% for >10 cells)	–	24
The MT of anesthetized rhesus monkeys	~0.65	6.7 (maximum = 12.3)	–	2–4 ms	25
The MT of alert rhesus monkeys	0.6	5.5	<30%	–	12
Variable stimulus: direct method					
H1 motion-sensitive neuron of a fly	2.43	80	50%	1.5–3 ms; <0.1	13
Salamander and rabbit retinal ganglion cells	3.7	16.3	59%	≥0.70 ms; ≥0.05	26
The MT of alert rhesus monkeys	1.5	12 (maximum = 29)	Up to 45%	<2 ms; ~1.4	12

^aEvaluated for the bin size used in estimating the information rate.

^bInformation rates are based on spike counts in time windows of duration ≥400 ms.

^cThe information rate is given in bits/stimulus. The stimulus was a 5 ms flash but the integration time was a few hundred milliseconds.

^dEstimated for the stimulus presentation period of 400 ms and the first three principal components of the average firing rate.

^eObtained by dividing 1.2 by the reported firing rate of 6.5 spikes in 0.5 s (see Ref. 19).

^fEstimated in a period of 500 ms.

Abbreviation: ‘–’, not known.

that sensory neurons of invertebrates and amphibians could faithfully track stimuli that varied on an ecologically relevant time scale [~30 ms (Ref. 9)] with high efficiency. Surprisingly, the stimulus could be estimated with a reliability that approached physical limits, despite the fact that its rate of change exceeded that of spike generation. As these neurons did not produce enough spikes in sufficiently short intervals for firing rate to be judged, a code based on mean firing rate could be ruled out for these particular nervous systems (see Table 1). More recently, experiments with rapidly varying stimuli were performed in tiger salamander and rabbit retinas (Fig. 1A). The results were comparable: in all of these systems individual spikes were strikingly informative (conveying in the range of 1–3 bits of information about the stimulus), suggesting that this might be a property shared by all sensory systems (see Table 1; see also Ref. 15). It remained to be seen, however, whether spikes recorded from the mammalian neocortex would be capable of matching the discrimination power of retinal spikes upstream.

Studies of the informativeness of neocortical neurons that were conducted using constant stimuli yielded about 0.5–2 bits/s routinely, suggesting that these neurons could discriminate between only 2–4 stimulus values (Table 1). In effect, an average neocortical spike, in contrast to its non-cortical relatives, was deemed to be worth very little – only about 0.01–0.2 bits/spike. It has long been known, however, that neurons in the visual cortex, especially those in the middle temporal

area (MT), respond promptly to rapid variations in stimulus motion (see, for example, Refs 28,29). Encouraged by this evidence and inspired by the results from non-mammalian neural systems, two laboratories^{12,25,30} recently embarked upon experiments designed to test the limits of the informativeness of motion-sensitive neocortical neurons by applying stimuli that varied on an ecologically relevant 30–300 ms time scale. Indeed, while saccadic eye movements occur at the rate of 2–5 per second (and, hence, have a duration of ~300 ms), hyperacuity experiments suggest that temporal differences that are less than 30 ms are perceived as simultaneous³¹. In addition, there are neurons in the MT that are tuned to such speeds (for example, ~250°/s) that optimal stimuli pass their receptive fields (which, for example, have a diameter of 8°) in about 30 ms (Ref. 32). When the MT neurons of primate (rhesus monkey) visual cortex were presented with these stimuli and the stimulus-reconstruction approach was applied, the net worth of a neocortical spike increased by nearly two orders of magnitude to ~0.6 bits/spike. This was true, despite the fact that one laboratory used anesthetized monkeys and the linear stimulus-reconstruction method²⁵, and another used alert monkeys and a nonlinear stimulus-reconstruction method¹². A casual inspection of the response histogram in the bottom panel of Fig. 1B reveals that spikes recorded from the MT are very well timed when the neuron is challenged by a rapidly varying stimulus. It is this faithful tracking of the stimulus characteristics in time

that endowed the MT neurons, much like retinal ganglion neurons of Fig. 1A, with high information content (nearly one bit per spike), despite the fact that the spike-count variability elicited by time-varying stimuli was as high (variance:mean = 1.4; 10 ms and 100 ms time windows used) as that elicited by a stimulus moving at a constant rate (variance:mean = 1.3).

Neurosemantics: content in spikes

What is the content of the information conveyed by these informative spikes? When asked to make fine discrimination about the exact stimulus value (the ‘what’ of the stimulus), the spikes of the neocortical neurons failed miserably¹² but became very informative when challenged with time-varying stimuli. Indeed, the spikes produced in response to varying stimuli were generated with high temporal precision (0.9–5.0 ms; cf. timing precisions of other neural systems in Table 1), which supports the idea that most of the information carried by spikes is about stimulus timing. The same preference of ‘when’ over ‘what’ was reported for salamander retinal ganglion cells²⁷ and motion-sensitive neurons of the fly¹³, suggesting that this preference might reflect a fundamental principle that underlies the design of sensory systems. Indeed, this preference for ‘when’ over ‘what’ is a strategy that maximizes transmitted information³³. For example, if a neuron doubles the precision with which it represents stimulus orientation (stimulus ‘what’), information increases by only one bit (that is, information grows logarithmically with message precision). By contrast, if the same neuron doubles its timing precision, which effectively halves its transmission time, the information rate doubles (information grows linearly with transmission speed). This biased allocation of channel bandwidth in favor of timing is especially useful for sensory systems that rely on large populations of neurons to convey the signal: one can readily increase the precision of the stimulus estimate by simply pooling more neurons (for a review of population coding see Ref. 11). By contrast, the temporal precision of responses cannot be increased easily in this way.

In conclusion, when sensory systems confront stimuli that vary on ecologically relevant time scales, individual neuronal spikes convey in the range of 1–3 bits of information. It appears that the information conveyed is primarily about the timing of abrupt variations in the stimulus.

Gauging efficiency of neuronal representations

The high information rates in most of the tested neural systems could have resulted solely from the fact that transmission speed was being maximized. The other factor that could contribute to high information rates is the use of an efficient code to represent information (see Box 2; Refs 9,22). Indeed, the coding efficiencies as estimated from information rates obtained using the reconstruction method, were rather high for all tested systems (11–60%), meaning that a large percent of variability in the spike trains was caused by stimulus rather than noise. This was also true for neocortical neurons (up to 45%), despite their comparatively high trial-to-trial variability.

The reconstruction method discussed above addresses information about a specific stimulus parameter. Recently, a number of researchers have developed algorithms for calculating the total information available

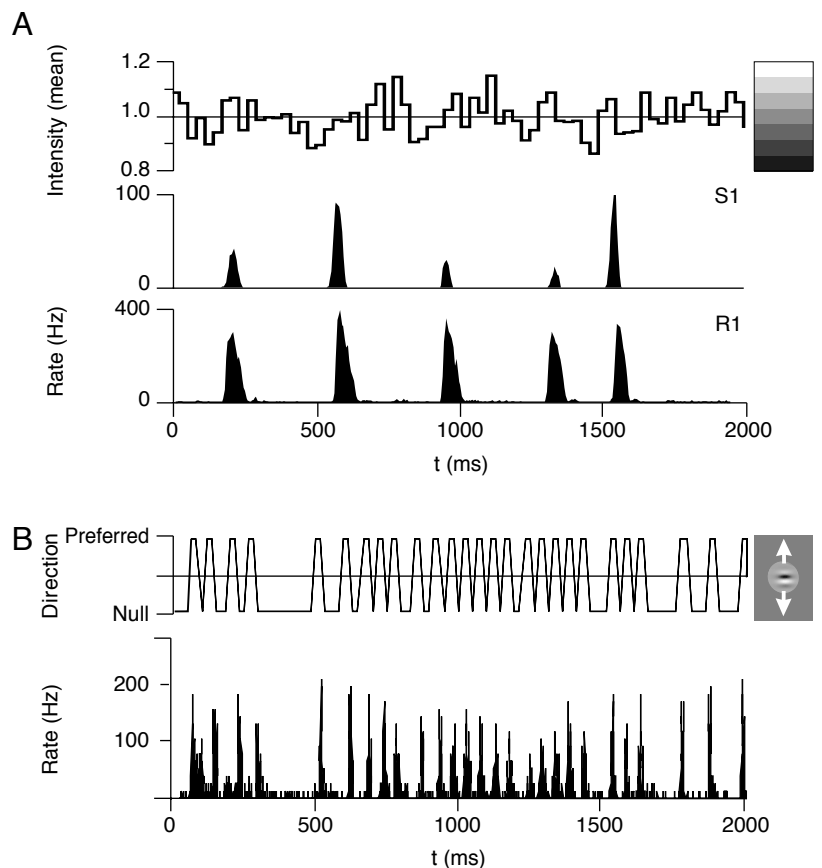


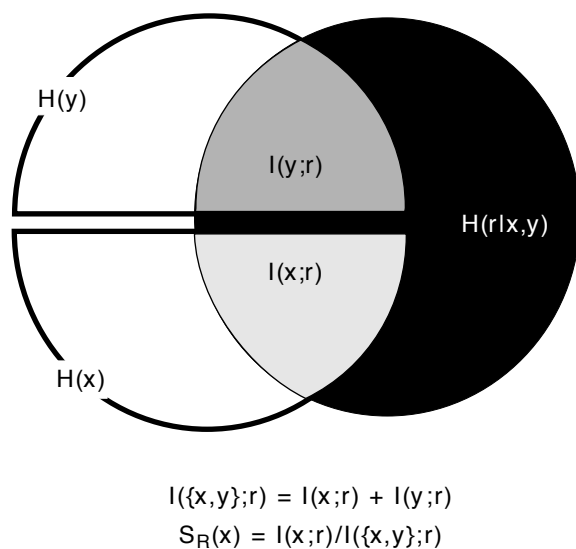
Fig. 1. The high temporal precision response of sensory neurons to stimuli varying on ecologically relevant time scales. (A) Electrical responses of salamander (S1) and rabbit (R1) retinal ganglion cells (lower panels) to rapid full-field variations of luminance (top panel). These neurons responded discretely to specific stimulus events. **(B)** Responses of middle temporal area neurons of alert monkeys (bottom panel) to rapid direction reversals of a moving Gabor patch (top panel). The direction of motion reversed repeatedly after short unpredictable periods of time. Neuronal responses to direction reversals were timed with high precision. (A) reproduced, with permission, from Ref. 27 and (B) reproduced, with permission, from Ref. 12.

in a spike train^{4–6}. This ‘direct’ method has two advantages over the reconstruction method: (1) it does not suffer from information loss that might be caused by an imperfect reconstruction algorithm and it is capable of providing a rather accurate estimate of information rate (provided that there are enough data); and (2) it makes no assumptions about what stimulus parameter is encoded in the spike train. When applied to the fly¹³ and monkey data¹² these two properties conspire to yield information rates that exceed the rates obtained using the reconstruction method substantially (see Table 1). It seems unlikely that the twofold increase in information rate reflects only the recovery of information lost by imperfect stimulus reconstruction, as the stimulus employed is simple (could take on only one of two motion values) and readily reconstructed. A more likely possibility is that the twofold improvement in the information rate was at least partially caused by the fact that the activity of these neurons conveys information not only about direction of motion but also about other stimulus properties, such as spatial phase.

In order to quantify the degree to which a neuronal response captures a particular stimulus parameter, this article proposes a metric for representational specificity S_R , which is a ratio of the information about the stimulus parameter in question (obtained by reconstruction method) relative to total information (obtained by the direct method; see Box 2). According to this metric, at

Box 2. Information transfer versus computation

At neuroscientific gatherings it is common to hear that information theory, which was designed for solving communication problems, is irrelevant for neuroscience, because the brain is not designed simply to transfer information, but rather to perform some sort of computations that discard irrelevant information. Every computation, however, including neural computations, uses some information that is transferred to the site of computation^a. In addition, information theory can be used to evaluate the efficiency of neuronal computations. For example, a computational system (that is, a brain area) that extracts a parameter x (for example, direction of motion) from the raw data $s(x)$, can be evaluated in terms of its coding efficiency E_c , which is a ratio of information about the parameter, x [$I(x;r)$], and the total uncertainty in the neuronal response [$H(r)$]: $E_c = I(x;r)/H(r)$. Here r is the signal produced by the computational system. In addition, it is proposed



that the efficiency with which a certain parameter x is computed from the stimulus $s(x)$ can be measured by representational specificity S_R in the following way. If the information about one parameter x [$I(x;r)$], is calculated (for example, by means of the reconstruction method), and then the total information (using the direct method) about the stimulus $I(s;r) = I(\{x,y\};r)$ is calculated, then the representational specificity, $S_R(x)$, for the parameter x equals $I(x;r)/I(\{x,y\};r)$. $I(s;r) = I(\{x,y\};r)$ and is written merely to reflect the fact that information about a stimulus can be decomposed into information about the parameter x and all other sources of information y (Fig. 1). Note that, because the reconstruction method estimates the lower bound for information, when using this method only a lower bound can be obtained for $S_R(x)$.

Representational specificity $S_R(x)$ might also be used in evaluating invariance of neuronal representations: when the stimulus is allowed to vary along an additional dimension y , then $S_R(x) = I(x;r)/I(s(x,y);r)$ can be used to quantify the degree to which the representation of the parameter x is invariant with respect to parameter y .

Reference

^a Cover, T.M. and Thomas, J.A. (1991) *Elements of Information Theory*, Wiley

Fig. 1. The relationship between uncertainties [$H(x)$ and $H(y)$] of independent stimulus parameters (x and y) and specific information a neural response (r) conveys about those parameters. This relationship suggests an index for the degree to which a neuron represents a particular parameter (for example, x) [known as the representational specificity $S_R(x)$], which is a fraction of total information dedicated to a specific parameter x . The split circle on the left depicts stimulus entropies associated with two independently varying parameters x and y . The uncertainty $H(r|x,y)$ (black), measures variability in response r , which is not correlated with stimulus parameters. Information about parameters x and y is indicated by the two gray regions [$I(x;r)$ and $I(y;r)$, respectively], which in this case are independent and, hence, their entropies are additive.

least 45% of the total information conveyed by spike trains recorded from individual neurons in the MT reflects direction of motion. Hence, in the MT, the representation of direction, a 'simple' representation of a single stimulus parameter, is 'rich,' because at least a half of the bandwidth of the neuronal signal in this area is dedicated to this parameter. For retinal ganglion cells the S_R of luminance variations is very close to that of the MT (51%; see Refs 24,26). The lower bound of the representational specificity is even higher in the direction-selective fly H1 neuron ($\geq 80\%$), which is understandable in view of the fact that, in contrast to the MT, the fly has very few neurons available to estimate direction of motion.

Progression of information abstraction in sensory systems

These newly developed information-theoretic methods allow one to quantify precisely the degree to which neuronal representations of sensory information become more abstract, by selective loss of irrelevant information, as one proceeds from the periphery to the CNS. This approach has been taken in recent studies of the electric fish *Eigenmannia*³⁴. Peripheral electrosensors, which sense the structure of the environment reflected in the electric field, were found to be able to convey

large quantities of raw information (in terms of signal-to-noise ratio) about the rapid variation of the electric field. By contrast, pyramidal neurons of the first electro-sensory central nucleus (lateral-line lobe) were able to discriminate high-level features with high efficiency. According to the thesis presented, if explicitly computed, S_R should grow for abstract features and not for low-level features, as more and more low-level information is discarded and abstract features are computed (see Box 2).

Mechanisms that support high information rates

The high information rates observed in a wide spectrum of species testify that encoding of sensory information could be very rapid and is characterized by a fast modulation of firing rate. König *et al.*³⁵ argued that, on the basis of computer simulations, this fast modulation is possible if stimuli cause synchronous synaptic inputs and the neurons involved operate using coincidence detection. This hypothesis is supported by recent experiments that used synthetic synaptic-input currents injected into cortical sensory neurons in brain slices³⁶ (cf. Ref. 37). This synchronous mode of signal encoding might explain high timing precision observed in the primate MT (Refs 29,30), which is as far as five synapses away from photoreceptors.

Indeed, local field potentials (LFPs), which index local-population activity, appear to covary with single-unit responses in the MT, in a manner that is consistent with the synchronous stimulus-locked activity needed to propagate precise timing information to downstream neurons¹². Interestingly, increasing the contrast of a flickering stimulus increased the timing precision²⁷, which suggests that stronger stimuli might synchronize synaptic inputs more efficiently.

The other key factor, which determines the information rate, is trial-to-trial response variability. While for primate MT neurons this variability was found to be consistent with that predicted by the nonhomogenous Poisson process model¹² (that is, a model of maximum randomness), both the fly H1 neurons¹³ and retinal ganglion cells²⁷ exhibited variability, which was below that predicted by this Poisson model. Berry and Meister have found that the trial-to-trial variability is reduced for stimuli that drive cells more robustly, probably because for stronger synaptic inputs the frequency of spike generation approaches limits imposed by the refractory period, which, thus, causes regular firing²⁶. In conclusion, current experimental evidence suggests that, when sensory neurons are stimulated by rapidly varying stimuli, the observed high information rates result from synergy between the synchronous synaptic input and finite refractory period.

Unanswered questions

A parallel line of research on temporal aspects of neuronal coding has explored the hypothesis that static stimuli could be encoded in the temporal distribution of neuronal responses^{16,17}. Recently McClurkin and colleagues³⁸ have reported that different temporal neuronal-response waveforms are used to encode color and shape in visual-cortex areas V1, V2 and V4. Indeed, afferents to these areas possess a wide range of response latencies, which could, in principle, support the encoding of static features in time³⁹. In view of the data indicating that temporal precision and response variability varies with stimulus intensity^{26,40}, however, it will be important to determine whether this hypothetical temporal encoding remains invariant for time-varying stimuli.

Most of the work reviewed in this article has addressed the information conveyed by single neurons. How will information rates be affected by pooling responses from many neurons? The results of studies that addressed this question for retinal ganglion cells^{24,41} suggest that response pooling could have non-trivial effects: as the pool of neurons increases, the contribution of each neuron to the stimulus reconstruction changes in a complicated way. Furthermore, noise correlations between neurons^{11,15}, in contrast to conventional wisdom, might increase signal-to-noise ratio and information rates.

Most importantly, the measurements of information rates detailed in this article only address how much information is available at a given processing level. It remains to be seen how this information about rapidly varying stimuli metamorphoses as one progresses to higher tiers of visual-processing hierarchy and how it is used in guiding behavior. Thus, while many important questions still confront us, the approach reviewed here frames the discourse and provides a novel means for investigating sensory processing under conditions that approximate the richness of normal sensory experience.

Selected references

- 1 Blakemore, C. and Movshon, J.A. (1995) in *The Cognitive Neurosciences* (Gazzaniga, M., ed.), pp. 235–236, MIT Press
- 2 Lettvin, J.Y. et al. (1959) *Proc. I. R. E.* 47, 1940–1951
- 3 Bialek, W. and Rieke, F. (1992) *Trends Neurosci.* 15, 428–433
- 4 DeWeese, M. (1996) in *Advances in Neural Information Processing* (Vol. 8) (Touretzky, D.S., Mozer, M.C. and Hasselmo, M.E., eds), pp. 281–287, MIT Press
- 5 Stevens, C.F. and Zador, A.M. (1996) in *Advances in Neural Information Processing* (Vol. 8) (Touretzky, D.S., Mozer, M.C. and Hasselmo, M.E., eds), pp. 75–81, MIT Press
- 6 Strong, S.P. et al. (1998) *Phys. Rev. Lett.* 80, 197
- 7 Zador, A. (1998) *J. Neurophysiol.* 79, 1219–1229
- 8 Van Essen, D.C. and DeYoe, E.A. (1995) in *The Cognitive Neurosciences* (Gazzaniga, M., ed.), pp. 393–394, MIT Press
- 9 Rieke, F. et al. (1997) *Spikes*, MIT Press
- 10 Newsome, W.T., Britten, K.H. and Movshon, J.A. (1989) *Nature* 341, 52–54
- 11 Oram, M.W. et al. (1998) *Trends Neurosci.* 21, 259–265
- 12 Buračas, G.T. et al. (1998) *Neuron* 20, 959–969
- 13 de Ruyter van Steveninck R.R. et al. (1997) *Science* 21, 1805–1808
- 14 Adrian, E.D. (1926) *The Basis of Sensation: The Action of the Sense Organs*, W.W. Norton
- 15 Dan, Y. et al. (1998) *Nat. Neurosci.* 1, 501–507
- 16 Richmond, B.J. and Optican, L.M. (1990) *J. Neurophysiol.* 64, 370–380
- 17 Optican, L.M. and Richmond, B.J. (1987) *J. Neurophysiol.* 57, 162–178
- 18 Rolls, E.T. et al. (1997) *J. Comp. Neurosci.* 4, 309–333
- 19 Toveé, M.J. et al. (1993) *J. Neurophysiol.* 70, 640–654
- 20 Rolls, E.T. et al. (1998) *J. Neurophysiol.* 79, 1797–1813
- 21 Bialek, W. et al. (1991) *Science* 252, 1854–1857
- 22 Rieke, F., Warland, D. and Bialek, W. (1993) *Europhys. Lett.* 22, 151–156
- 23 Roddey, J.C. and Jacobs, G.A. (1996) *J. Neurophysiol.* 75, 1365–1376
- 24 Warland, D.K., Reinagel, P. and Meister, M. (1997) *J. Neurophysiol.* 78, 2336–2350
- 25 Bair, W., Cavanaugh, J.R. and Movshon, J.A. (1997) in *Advances in Information Processing Systems* (Vol. 9) (Mozer, M.C., Jordan, M.I. and Petsche, T., eds), pp. 34–40, MIT Press
- 26 Berry, M.J. and Meister, M. (1998) *J. Neurosci.* 18, 2200–2211
- 27 Berry, M.J., Warland, D.K. and Meister, M. (1997) *Proc. Natl. Acad. Sci. U. S. A.* 94, 5411–5416
- 28 Shadlen, M.N. and Newsome, W.T. (1994) *Curr. Opin. Neurobiol.* 4, 569–579
- 29 Bair, W. and Koch, C. (1996) *Neural Comp.* 8, 44–66
- 30 Buračas, G.T. et al. (1996) *Soc. Neurosci. Abstr.* 22, 717
- 31 Leibovic, K.N. (1990) in *Science of Vision* (Leibovic, K.N., ed.), pp. 153–171, Springer-Verlag
- 32 Maunsell, J.H.R. and Van Essen, D.C. (1983) *J. Neurophysiol.* 49, 1127–1167
- 33 Softky, W. (1996) *Neural Networks* 9, 15–24
- 34 Gabbiani, F. et al. (1996) *Nature* 384, 564–567
- 35 König, P., Engel, A.K. and Singer, W. (1996) *Trends Neurosci.* 19, 130–137
- 36 Stevens, C.F. and Zador, A.M. (1998) *Nat. Neurosci.* 1, 210–217
- 37 Shadlen, M.N. and Newsome, W.T. (1998) *J. Neurosci.* 18, 3870–3896
- 38 McClurkin, J.W., Zarbock, J.W. and Optican, L.W. (1996) *J. Neurophysiol.* 75, 496–507
- 39 Golomb, D. et al. (1994) *J. Neurophysiol.* 72, 2990–3003
- 40 Reich, D.S. et al. (1997) *J. Neurophysiol.* 77, 2836–2841
- 41 Meister, M., Lagnado, L. and Baylor, D.A. (1995) *Science* 270, 1207–1210

Acknowledgements

The authors are grateful to M. Berry, D. Chklovskii, M. DeWeese, A. Messinger, O. Rukšenas, G. Stoner, A. Thiele and A. Zador for useful comments on the manuscript, and W. Bair, G. Baranauskas, F. Crick, S. Janušonis, D. Kirvelis, M. Meister, W. Newsome and M. Shadlen for stimulating discussions.

Letters to the Editor

Letters to the Editor concerning articles published recently in *TINS* are welcome. Please mark clearly whether they are intended for publication; the authors of the article referred to are given an opportunity to respond to any of the points made in the Letter. Maximum length: 500 words. The Editor reserves the right to edit Letters for publication.

Please address Letters to:
Dr Gavin Swanson, Editor
Trends in Neurosciences, 68 Hills Road
Cambridge, UK CB2 1LA.