

Adaptive spike coding

Adrienne Fairhall¹ and William Bialek²

Department of Molecular Biology¹ and Department of Physics²,
Princeton University, Princeton NJ 08544

¹ Corresponding author: phone 609.258.7008, fax 609.258.1028

INTRODUCTION

The meaning of any signal that we receive from our environment is modulated by the context within which it appears. Our interpretation of colour, a spoken phoneme, or a patch of luminance depends critically on its context. While ‘context’ may be a rather abstract notion, it is reasonable in a range of situations to take it to mean the statistical ensemble in which the signal is embedded. Interpreting a message requires both registering the signal itself and knowing something about this statistical ensemble. The relevant temporal or spatial ensemble depends on the task. The context may be highly local; we interpret appropriately gradations of light and dark in a scene where local brightness typically varies over orders of magnitude (VISUAL CODING, NATURAL SCENES AND FEATURE DETECTION). For tasks such as decision-making, the relevant statistics may reflect complex descriptions of the world accumulated over long periods.

Neural representations at every level of information processing should be similarly modulated by context. Information theoretically, this has measurable advantages: representations which appropriately take into account the statistical properties of the incoming signal are more efficient. Since the ’50’s it has been suggested that efficiency is a design principle of the nervous system, allowing neurons to transmit more useful information with their limited dynamic range (see OPTIMAL SENSORY ENCODING; INFORMATION THEORY AND VISUAL PLASTICITY). Thus, one expects that learning the context and implementing this knowledge through coding strategy is inherent in the formation of representations.

Such adjustments occur over a wide range of timescales. Through the genetic code, species adapt to environmental changes over many generations. In a single individual, learning, im-

plemented through neural plasticity, continues throughout life in response to experience of the world; perceptual learning is stored even at low levels of neural information processing (see SOMATOTOPY, PLASTICITY OF SENSORY MAPS). Here, we will discuss even more rapid changes: neural adaptation, which we take to mean reversible change in the response properties of neurons on short timescales.

Since Adrian's first observations of adaptation in spiking neurons, it had been suggested that adaptation serves a useful function for information processing, preventing a neuron from continuing to transmit redundant information and increasing its responsiveness to new stimuli. Within the simplified picture of a neuron as a combination of linear filtering followed by a threshold, or a decision rule for spiking, either or both of the two components— the filter and the threshold function— may be adaptive functions of the input, and both may implement the goal of increasing information transmission. We will discuss both of these possibilities.

Neurons in every sensory modality have been shown to have adaptive properties, and the mechanisms governing various types of adaptation have been at least partially explored (Torre et al., 1995). Here we will discuss adaptation as the simplest form of learning and memory. We describe recent experiments which explicitly aim to link the phenomenology of adaptive spike coding to its functional relevance, in particular to improved information transmission. A common feature of adaptation is the existence of multiple timescales. In examining mechanisms, we concentrate on recent work suggesting that the long timescales retaining short-term memory can be generated through single cell properties.

ADAPTIVE CODING

Adaptation of neural firing rate to stationary stimuli has been seen in all modalities of the primary sensory system. In the visual system, photoreceptors adapt to light level, and retinal ganglion cells show rapid contrast gain control. The tradeoffs and information processing gains due to adaptation in insect eyes, relevant also for vertebrate retina, are discussed in (Laughlin, 1989). In the somatosensory system, mechanoreceptors have been classified into four main types of cells, three of which are distinguished by the timescales of their adaptation (rapidly and slowly adapting), and these timescales in part determine the cells' function: slowly adapting cells are implicated in the perception of spatial form and texture, while the experience of flutter and of motion is mediated by rapidly adapting cells (Johnson, 2001). Thus, the dynamics of adaptation can determine a neuron's functional role.

Adaptation is not limited to primary receptors. In visual cortex, V1 neurons show contrast adaptation which is thought to occur entirely at the level of cortex. The motion aftereffect, a familiar phenomenon whereby following exposure to motion in one direction, the visual field appears to move in the opposite direction, is thought to be due to adaptation of direction-sensitive neurons in visual cortex.

Adaptation to a distribution

Understanding the significance of adaptation for information processing requires going beyond fixed stimuli. Recently, studies have focused on adaptation to the stimulus *distribution*. This approach is necessary to characterise coding information theoretically: the evaluation of coding strategy requires considering the entire ensemble of inputs and outputs. In Smirnakis et al. (1997), retinal ganglion cells were stimulated with dynamic movies of flickering light intensity, where the mean light level was fixed, but the variance switched periodically from one value to

another. The spike rate of the neurons showed typical adaptive behaviour, Fig. 1: following an increase in variance, the firing rate initially increased, but gradually returned to a considerably lower level; a decrease in variance led to a sudden dip in firing rate, with eventual recovery.

The experiments of Smirnakis et al. (1997) consider only firing rate. However, the timing of single spikes can convey a great deal of information about the stimulus. In the visual system of the fly, in particular the motion-sensitive identified neuron H1 in the fly's lobula plate, much is understood about single spike coding, providing an excellent opportunity to study the effects of adaptation in detail.

H1 responds to a simple stimulus, wide-field horizontal motion. The neuron is characterised by its input/output relation $P(\text{spike}|s)$, the probability of a spike given the projection s of the dynamic stimulus onto a relevant feature, determined by reverse correlation.

When the system has reached steady state through exposure to a zero mean, white noise velocity stimulus with a given variance σ^2 , its input/output relation is measured. The resulting curves, measured for a range of values of the variance, are shown in Fig. 2. Clearly the input/output relation is not a fixed property of the system, but adapts to the distribution of inputs. Indeed, it does so in such a way that the stimulus appears to be measured in units of its standard deviation; when the curves are replotted with the stimulus normalised by its RMS value, they superimpose. Thus a scale factor λ multiplying the stimulus, and thus matching the dynamic range of the response to the distribution of the inputs, is a degree of freedom for the system. The value of λ chosen by the system achieves a maximum of information transmission (Brenner et al., 2000).

This is a simple form of learning: the system gauges the standard deviation of the signal, and modifies its response properties to adjust its dynamic range to the range of inputs. The adjustment must take some time, as the new distribution must be sampled from examples. This sets fundamental physical and statistical limits for the system's estimate of the current variance. We can examine the timescale for learning (Fairhall et al, 2001) by, as in the retina experiments described above, switching periodically between two distributions. The firing rate shows the same pattern of adaptation seen in the experiments of Smirnakis et al. (1997), but this need not correspond to the timescale for adjustment of the input/output relation. Indeed, it was found that the scale factor adjusts much more rapidly than the relaxation time of the rate—on the order of 100 ms, compared with several seconds. This short timescale is consistent with the limits imposed by estimates of noise from the photoreceptors. One can verify that the dynamic adaptation of the input/output relation maintains information transmission through the system by computing how much information one can extract from the spikes about the stimulus (see SENSORY CODING AND INFORMATION TRANSMISSION and Fairhall et al. (2001)). The information rate recovers on comparably short timescales.

For the decoder, a potential drawback of adaptive coding is *ambiguity*: it is necessary to know the context in order to interpret the signal correctly. Thus, information about the context must be conveyed independently. While this information might be carried by other neurons in the network, here the information about the ensemble is carried simultaneously by the same spike train: it can be read off, either through the rate (taking into account the delays due to the slow relaxation) or, more accurately, through the statistics of *spike time differences* (Fairhall et al, 2001). Thus for the code of H1, spikes carry multiple meanings: in absolute timing, as precise markers of single stimulus events, and in relative timing, as indicators of the stimulus ensemble.

Multiple timescales

The slow relaxation of the rate appears to be related to a commonly observed property of adapting primary sensory neurons: a power law decay of the firing rate r , $r \sim t^{-\alpha}$. More generally, in the case above, the rate is close to the *fractional derivative* of the logarithm of the stimulus variance. For each frequency ω , fractional differentiation shifts the frequency component by a constant phase, and scales each component by ω^α , where α is a power less than 1. Some of the properties of a fractional differentiator are illustrated in Fig. 3. Several examples of a power law decay of the rate following a step change in stimulus amplitude were collected by Thorson and Biederman-Thorson (1974) (see Fig. 4) and more have since been observed; examples include various invertebrate mechanoreceptors and photoreceptors, mammalian carotid-sinus baroreceptors and cat retinal ganglion cells.

While we have already noted a separation of timescales in the adaptation of the input/output relation compared with the rate, this type of adaptation on its own signals the existence of many timescales. Power-law scaling implies the lack of a typical timescale, or the presence of multiple timescales. Fractional differentiation is non-local; the response at time t_0 is increasingly affected by times $t \ll t_0$. This is a linear “memory” mechanism.

Such adaptation is particularly interesting both because it is so prevalent, and because it may have an important role in optimising information transmission. Fractional-differentiationlike behaviour is observed in fly photoreceptors, and there, the exponent of the fractional differentiator appears to be matched to the spectrum of natural stimuli (van Hateren and Snippe, 2001). Thus the effect of the transformation is to whiten the spectrum of natural signals. As many natural stimuli have power-law characteristics, it is intriguing to speculate that fractional differentiation

at the sensory periphery may be a general neural mechanism for whitening input statistics.

MECHANISMS

Adaptation requires retaining memory of activity over extended timescales. These long timescales can arise from a number of sources. Intracellular calcium concentration has been identified as playing an important role in information processing, acting as a slowly-changing “integrator” of activity. Other forms of adaptation— particularly the power-law like behaviour discussed above— are also likely to be a property of single cells rather than of the network. Recent biophysical studies show that membrane dynamics can have long timescales that retain memory of the history of stimulation/activity over hundreds of seconds (Marom, 1998). This could be brought about either by the modification of intrinsic properties, or by intrinsic properties that have built-in long timescales through *state-dependent inactivation* (Marom, 1998; Turrigiano et al., 1996).

Calcium as an integrator of activity

Each spike introduces a roughly constant amount of calcium into the cell through voltage-dependent Ca^{2+} channels. The Ca^{2+} concentration then decays slowly. Thus $[\text{Ca}^{2+}]$ can be modelled as a leaky integrator of activity, with a decay timescale of ~ 100 ms. This calcium signal can allow activity-dependent regulation of subsequent neural activity through the modification of conductances (ACTIVITY-DEPENDENT REGULATION OF NEURONAL CONDUCTANCES).

Recent evidence indicates that single-cell properties may contribute to contrast adaptation in cortex (Sanchez-Vives et al., 2000). Previous work has shown that contrast adaptation is asso-

ciated with hyperpolarization of the membrane potential in cat area 17 neurons. By stimulating the neurons directly with injected current, effects similar to contrast adaptation are seen (though less dramatically than to real visual input). This suggests that these effects can be induced through the modulation of intrinsic cell properties; the activation of Ca^{2+} and Na^+ -dependent potassium conductances is indicated.

State-dependent channel dynamics

In some cases, the relevant dynamics may be due to the complex behaviour of the channels themselves. Recently it has become clear that the dynamics of inactivation provide the membrane with the possibility for extended history dependence (Marom, 1998).

A simplified picture of the gating of voltage-gated ion channels is a three-state scheme:



where channels can be either closed (C), open (O) or inactivated (I). Generally, the transition between closed and open is voltage-dependent and rapid, on the order of the duration of an action potential. The transition between open and inactivated, on the other hand, is voltage-independent and can have very long timescale dynamics. Intriguingly, studies *in vitro* show that some sodium channel types have inactivation rates that scale with the duration of the input (Marom, 1998), providing timescales up to several minutes. The precise mechanism underlying this large variety of timescales is not yet well-understood; it is hypothesised that the system cascades through a multiplicity of inactivation states. Earlier theoretical work has shown that the coupling of many states leads to a scaling relation between the duration of activity and the rate of recovery from inactivation.

In a step closer to a realistic preparation, the dynamic clamp was applied to cultured stomatogastric ganglion neurons to add an effective slowly inactivating potassium current (Turrigiano et al., 1996). As had been observed previously, this produced long delays to firing during depolarisation, and an increase in excitability with a timescale much longer than the duration of the input. Further, the slow channel dynamics produced a long-lasting effect on the firing properties of the neuron.

In vivo, the contribution of slowly inactivating sodium channels to power-lawlike adaptation has been suggested. Mechanosensory neurons in the cockroach femoral tactile spine have been shown to display power-law adaptation. From intracellular measurements, Basarsky and French (1991) found that the spike rate adaptation is due to successive slowing of the recovery of the membrane potential between spikes. Previous work had demonstrated that calcium channel blockers or blockers of Ca^{2+} -activated K^+ channels did not reduce adaptation, while modifying sodium channel inactivation did.

These mechanisms might be seen as primitives for short-term “learning and memory”.

MODELLING

Historically, attempts to model adaptation have considered the process to involve a dynamic threshold. More recently, modelling approaches take a functional perspective on the outcome of adaptation, and propose algorithms whereby the conductances may adjust to provide the cell with desirable properties such as approximately constant activity (see **ACTIVITY-DEPENDENT REGULATION OF NEURONAL CONDUCTANCES**). Closer to our earlier discussion, Stemmler and Koch (1999) derive a learning rule for conductances which maximises

the mutual information between input and output, where the output is taken to be the neuron's firing rate. The learning rule adjusts conductances at every new presentation of the stimulus, subject to biologically plausible constraints. Under this learning rule, a realistic conductance-based model neuron was indeed able to learn a changing distribution and adjust its firing statistics accordingly. The timescales treated were order of magnitudes longer than those observed experimentally in (Fairhall et al., 2001) and predicted theoretically from statistical considerations. Experimental evidence is still required to determine whether such a model is realistic.

As noted, many adaptation processes in sensory receptors follow a power-law relaxation. Assuming that most elementary processes involve a single time scale, with exponential dynamics, Thorson and Biederman-Thorson (1974) proposed that power laws may arise from a superposition of many elementary processes with a wide range of timescales. From the definition of the gamma function,

$$t^{-\alpha} = \frac{1}{\Gamma(\alpha)} \int_0^{\infty} dr r^{\alpha-1} e^{-rt}, \quad (2)$$

a power law may be generated by a weighted sum of exponentials with a range of timescales. This distribution was considered to be generated through geometric factors, such as the inhomogeneous distribution of elements within the receptor.

This model has merited some scepticism due to the requirements both for a continuous distribution of timescales and for these to be present in the appropriate proportions. It has been noted that power-law-like behaviour results from much less stringent conditions: the superposition of only a few exponentials can produce a power law over the decade or two normally available to experiment. However, recent experimental advances outlined in the previous section may provide a better underpinning for the derivation of power-law adaptation from membrane mechanisms.

ADAPTATION OF RECEPTIVE FIELDS

As noted in the Introduction, a neuron can be modelled as a combination of feature extraction (linear filtering) and a nonlinear decision function (or threshold). While we have discussed the effects of adaptation on the nonlinear decision function, adaptation can also affect the feature that causes the neuron to spike: the receptive field can depend upon the ensemble of inputs. While this had been frequently observed in work on invertebrate vision, recent experiments demonstrate analogous results for cortical receptive fields. Sceniak et al. (1999) show that the extent of spatial summation implemented by neurons in V1 depends adaptively on contrast; this has parallels in the adaptation of filters in retina (Laughlin, 1989). Theunissen et al. (2000) found that the spatiotemporal receptive fields of neurons in auditory cortex showed a strong dependence on the stimulus ensemble. This is a natural consequence of neural nonlinearity, but such a dependence is also necessary for optimal information processing.

DISCUSSION

The ubiquity of adaptation throughout the nervous system should be proof of its fundamental importance. While the phenomenology of adaptation, particularly to constant stimuli, has been extensively explored, recent experimental and theoretical approaches have made contact with the principles of information theory in order to evaluate adaptive coding. For fly motion-sensitive neurons, it was found that the coding strategy of the system adapts rapidly and continuously to track dynamic changes in the statistics of the stimulus.

We have discussed here a variety of mechanisms that may implement adaptive coding at the level of single cells. While it is likely that systems will implement such important behaviour at

many levels, it is appealing that the simplest elements of neural computation have the power to carry out dynamic aspects of information processing.

REFERENCES

Basarsky, T. and French, A., 1991, Intracellular measurements from a rapidly adapting sensory neuron. J. Neurophysiol. 65 49-56.

Brenner, N., Bialek, W. and de Ruyter van Steveninck, R., 2000, Adaptive rescaling maximizes information transmission. Neuron 26, 695-702.

Fairhall, A.L., Lewen, G., Bialek, W. and de Ruyter van Steveninck, R. R., 2001, Efficiency and ambiguity in an adaptive neural code. Nature 412, 787-792.

Johnson, K.O., 2001, The roles and functions of cutaneous mechanoreceptors. Curr. Opin. Neurobiol. 11(4):455-61.**

Laughlin, S.B., 1989, The role of sensory adaptation in the retina. J. Exp. Biol. 146:39-62.**

Marom, S., 1998, Slow changes in the availability of voltage-gated ion channels: effects on the dynamics of excitable membranes. J. Membrane Biol. 161, 105-113.**

Sanchez-Vives, M., Nowak, L. and McCormick, D., 2000, Membrane mechanisms underlying contrast adaptation in cat area 17 *in vivo*. J. Neurosci. 20(11):4267-4285.

Sceniak, M. P., Ringach, D. L., Hawken, M. J. and Shapley, R., 1999, Contrast's effect on spatial summation by macaque V1 neurons. Nat. Neurosci. 2:733-739.

Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W. and Meister, M., 1997, Adaptation of retinal processing to image contrast and spatial scale. Nature 386, 69-73.

Stemmler, M. and Koch, C., 1999, How voltage-dependent conductances can adapt to maximize

the information encoded by neuronal firing rate. Nat. Neurosci. 2(6):521-527.

Thorson, J. and Biederman-Thorson, M., 1974, Distributed relaxation processes in a sensory adaptation. Science 183, 161-172.**

Theunissen, F., Sen, K., and Doupe, A., 2000, Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. J. Neurosci. 20(6):2315-2331.

Torre, V., Ashmore, J.F., Lamb, T.D. and Menini, A., 1995, Transduction and adaptation in sensory receptor cells. J. Neurosci. 15 12:7757-7763.**

Turrigiano, G., Marder, E., and Abbott, L., 1996, Cellular short-term memory from a slow potassium conductance. J. Neurophysiol. 75(2): 963-968.

van Hateren, J. H. and Snippe, H. P., 2001, Information theoretical evaluation of parametric models of gain control in blowfly photoreceptor cells. Vision Res. 41, 1851-1865.

Figure captions

Fig. 1. a. Firing rate of rabbit retinal ganglion cells in response to a flicker stimulus where the variance of the light intensity I switches periodically in time, as illustrated. Reproduced from *Nature* 386, Smirnakis et al., 69-73. Copyright (1997) Macmillan Publishers Ltd.

Fig. 2. A set of input/output relations relating the probability of spiking to the velocity stimulus, measured for stationary white noise stimuli with different variances. The curves differ only by a scale factor λ multiplying the stimulus. Reprinted from *Neuron* 26, Brenner et al., pp 695-702, copyright (2000), with permission from Elsevier Science.

Fig. 3. Illustration of some properties of a fractional differentiator with exponent $\alpha = 0.3$. a. A step function stimulus leads to a power law decaying rate. In a log-log plot the curve would appear as a straight line with slope $-\alpha$. b. A square wave leads to a similar adaptation curve as shown in Fig. 1.

Fig. 4. Examples of data taken from *Science*, 183, Thorson and Biederman-Thorson, pp 161-172. Copyright (1974), American Association for the Advancement of Science. The four curves show power law adaptation in response to a step increase in stimulus in four different receptors. A. Cockroach leg mechanoreceptor, in response to distortion of the tactile spine on the femur; B. Slit sensillum on the leg of the hunting spider, in response to 1200 Hz sound; C. Slowly adapting stretch receptor of the crayfish; D. Increase of response over light-adapted level of Limulus lateral-eye eccentric cell to an increase in light intensity.




Figure 1: Adaptive Spike Coding




Figure 2: Adaptive Spike Coding

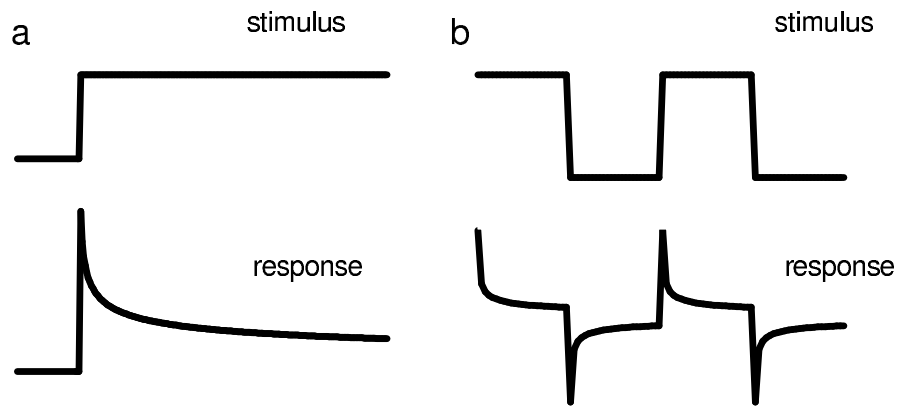


Figure 3: Adaptive Spike Coding




Figure 4: Adaptive Spike Coding