

An invariance principle for maintaining the operating point of a neuron

TERRY ELLIOTT, XUTAO KUANG, NIGEL R. SHADBOLT, &
KLAUS-PETER ZAUNER

*Department of Electronics and Computer Science, University of Southampton,
Southampton, UK*

(Received 26 February 2008; revised 15 April 2008; accepted 1 June 2008)

Abstract

Sensory neurons adapt to changes in the natural statistics of their environments through processes such as gain control and firing threshold adjustment. It has been argued that neurons early in sensory pathways adapt according to information-theoretic criteria, perhaps maximising their coding efficiency or information rate. Here, we draw a distinction between how a neuron's preferred operating point is determined and how its preferred operating point is maintained through adaptation. We propose that a neuron's preferred operating point can be characterised by the probability density function (PDF) of its output spike rate, and that adaptation maintains an invariant output PDF, regardless of how this output PDF is initially set. Considering a sigmoidal transfer function for simplicity, we derive simple adaptation rules for a neuron with one sensory input that permit adaptation to the lower-order statistics of the input, independent of how the preferred operating point of the neuron is set. Thus, if the preferred operating point is, in fact, set according to information-theoretic criteria, then these rules nonetheless maintain a neuron at that point. Our approach generalises from the unimodal case to the multimodal case, for a neuron with inputs from distinct sensory channels, and we briefly consider this case too.

Keywords: *Neuronal adaptation, operating point, gain control, multisensory integration*

Introduction

Neurons early in sensory pathways are believed to adapt their responses to the statistics of their inputs in order to maximise their coding efficiency, output

Correspondence: T. Elliott, Department of Electronics and Computer Science, University of Southampton, Highfield, Southampton S017 1BJ, UK. Tel: +44 (0)23 8059 6000. Fax: +44 (0)23 8059 9179. E-mail: te@ecs.soton.ac.uk

entropy or information rate (Atteave 1954; Barlow 1961; Laughlin 1981; Atick, 1992; van Hateren 1992; DeWeese 1996; Dan et al. 1996; Baddeley et al. 1997; Smirnakis et al. 1997; Wainwright 1999; Brenner et al. 2000; Fairhall et al. 2001; Maravall et al. 2007). Other adaptive strategies have also been proposed for neurons later in sensory pathways (see, e.g. Carandini and Ferster 2000; Pena and Konishi 2002; Ringach and Malone 2007). Despite the diversity of functional roles of neurons, and however their preferred operating points may be established over evolutionary or developmental timescales, it is possible that there exist adaptive principles, based on input statistics, that enable neurons to maintain their preferred operating points without explicit reference to their functional roles.

We suggest one such principle. We define the operating point of a neuron by the cumulative distribution function (CDF) or probability density function (PDF) of its output spike rate. The operating point therefore embodies information about both the input–output transfer function of the neuron and the statistics of the environment from which the input is drawn. The operating point is not simply the transfer function, nor simply the environment, but both, united in the PDF of the neuron’s output spike rate. We propose, in particular, the adaptive principle that a neuron adapts its transfer function in order to keep its output PDF invariant, or as invariant as possible, under changes in its input statistics. Such invariance would ensure that a neuron remains at its preferred operating point, regardless of how that point is set. One consequence of this view is that when a neuron’s operating point is set by a principle such as maximum information rate, if a neuron can maintain an invariant output PDF, then it automatically remains at the point determined by that principle.

In this article we develop a model of neuronal adaptation in which adaptation ensures that a neuron’s output PDF remains invariant, or approximately invariant, under changes in its input statistics. For simplicity, we consider only sigmoidal transfer functions. We initially consider a unimodal neuron with input from a single sensory channel, and derive rules for threshold adaptation and gain control that seek to maintain an invariant output PDF. These rules are independent of how the preferred operating point of the neuron is set, but we also consider how the operating point may be set by a maximum entropy principle. The structure of our model permits generalisation to multimodal neurons receiving input from multiple, distinct sensory channels, and we extend our rules to this multimodal case. This extension necessitates the introduction of separate gains for each modality. Having developed the underlying approach to adaptation, we present several examples of a model neuron functioning at different preferred operating points and adapting to changing input statistics, both unimodal and bimodal. Finally, we discuss our approach, its merits and limitations and possible future work.

An invariance principle for adaptation

We first consider a neuron with an input from only a single sensory channel. Then we extend our approach to a neuron with inputs from multiple, distinct sensory channels.

Unimodal inputs

Suppose that a neuron's unimodal input is determined by the (univariate) random variable X with PDF $f_X(x)$ and CDF $F_X(x)$, where X , without loss of generality, also may absorb any input noise processes. Under the assumption of certain regularity conditions, X is uniquely determined by its moments, so that its moment generating function, $M_X(t)$, exists and its characteristic function, $\phi_X(t)$, is analytic on the real line (Feller 1967).¹ We denote the moments of X by m_i , with mean $\mu = m_1$, and the central moments by μ_i , with variance $\sigma^2 = \mu_2$. We denote the transfer or output function of the neuron by $r = r(x)$ for input x drawn from the distribution X . The output is then a random variable, denoted by R , and its CDF and PDF are given by

$$F_R(r) = F_X(x(r)), \quad (1)$$

$$f_R(r) = \frac{dx(r)}{dr} f_X(x(r)), \quad (2)$$

respectively, where $x(r)$ is the inverse of $r(x)$, i.e. $x(r(x')) = x'$. If the response is bounded, so that $r \in [0, s]$, and if $r(x) = sF_X(x)$, then $f_R(r) = 1/s$, so that R is uniform on $[0, s]$, and hence has maximum entropy, corresponding to Laughlin's result (Laughlin 1981).

Popular choices of response function in the visual processing literature are the hyperbolic ratio function,

$$r(x) = s \frac{x^{4\gamma\theta/s}}{\theta^{4\gamma\theta/s} + x^{4\gamma\theta/s}}, \quad (3)$$

where θ determines the point of semi-saturation, $r(\theta) = s/2$, and γ the gain at semi-saturation; and the linear model with rectification,

$$r(x) = \gamma[x - \theta]_+, \quad (4)$$

where $[\]_+$ denotes the positive part, θ is here the threshold for response onset, and γ the gain. For simplicity, however, we will instead use the sigmoidal response function,

$$r(x) = \frac{s}{2} \left[1 + \tanh \frac{2\gamma}{s} (x - \theta) \right], \quad (5)$$

where θ and γ are, as for the hyperbolic ratio function, the semi-saturation constant and the gain, respectively. For γ large, θ determines the transition from no response to saturated response, so we refer to it as the threshold. We employ the sigmoidal response function because it is defined for $x < 0$ (unlike the hyperbolic ratio function for most parameter choices) and everywhere differentiable (unlike the linear rectified response).

We shall regard the saturation value s as fixed, since it is determined by the maximum firing rate of a neuron, which is limited by a neuron's refractory period. The sigmoidal response function thus endows a neuron with two parameters, the threshold and the gain, which it can change in an attempt to keep a neuron's output PDF invariant, or as invariant as possible. Clearly, then, for an arbitrary input distribution X , perfect adaptation is in general impossible, since a sigmoidal

response $r(x)$ lacks the potentially infinite number of degrees of freedom required to adapt to the potentially infinite number of independent degrees of freedom in all the moments, m_i , of X . Furthermore, R cannot adopt the uniform, maximum entropy distribution, given the upper bound s on the response, except in the particular case that X obeys the logistic distribution, with CDF $F_X(x) = (1/2) \times [1 + \tanh(2\gamma/s)(x - \theta)]$, satisfying Laughlin's condition, $r(x) = sF_X(x)$. The choice of a two-parameter response function thus imposes limitations on the extent to which a model neuron can adapt to its input statistics. Although this may appear unsatisfactory, of course any finite-parameter response function will be so limited. Unless we hold to the view that a neuron's response function may be set arbitrarily, which seems unlikely, then such limitations are unavoidable. We could improve the adaptability of a model neuron by considering more complicated, multi-parameter response functions, but we have selected the two-parameter, sigmoidal function precisely because of its simplicity, so that the underlying approach, of adaptation to input statistics, is not obscured by unnecessary complexity.

Suppose that a neuron's initial response is defined by the parameter set $\{\theta, \gamma\}$, giving rise to the random output variable R with PDF $f_R(r)$ for the initial random input variable X . Suppose that the input statistics then change, giving a new input random variable X' , inducing a new response parameter set $\{\theta', \gamma'\}$ and output random variable R' with PDF $f_{R'}(r)$. Since we wish to make the output PDF as invariant as possible under the change $X \rightarrow X'$, an obvious strategy would be to define a functional metric or distance \mathcal{D} and explicitly find θ' and γ' such that $\mathcal{D}(f_R, f_{R'})$ is minimised. A popular candidate for \mathcal{D} would be the relative entropy or Kullback–Leibler divergence between f_R and $f_{R'}$. Although explicit minimisation of $\mathcal{D}(f_R, f_{R'})$ is, mathematically speaking, a natural approach, it is potentially very expensive, computationally speaking. Therefore, we prefer to seek an alternative approach for determining θ' and γ' , one that may be computationally simpler for a neuron to implement, but accept that it may not be statistically optimal under all circumstances.

Since our model neuron possesses only two degrees of freedom in θ and γ , it is reasonable to assume that these two parameters are modified in order to accommodate the most significant variations in X . Defining the Z -score as usual by $Z = (X - \mu)/\sigma$, the Chebyshev inequality,

$$P[|Z| \geq k] \leq \frac{1}{k^2}, \quad (6)$$

where k is any positive number, suggests that the lowest two moments represent good targets for adaptive processes. We therefore propose that the parameters θ and γ should thus be adapted to accommodate changes in the mean, μ (Barlow and Mollon 1982) and the standard deviation, σ (Meister and Berry 1999; Smirnakis et al. 1997; Brenner et al. 2000; Fairhall et al. 2001), with the higher-order moments m_3, \dots playing second fiddle to m_1 and m_2 (Bonin et al. 2006; but see Kvale and Schreiner 2004). More degrees of freedom in the response function would allow adaptation to moments higher than the second.

Consider an input PDF of the particular form

$$f_X(x) = \frac{1}{\sigma} g\left(\frac{x - \mu}{\sigma}\right), \quad (7)$$

where, by this, we mean that the dependence of f_X on μ and σ appears *only* through the combination $(x - \mu)/\sigma$ in f_X , together with the overall scale factor of $1/\sigma$, and that g does not depend on any other parameters. The normal distribution is an example of such a PDF. Since

$$x(r) = \theta - \frac{s}{2\gamma} \log\left(\frac{s-r}{r}\right), \quad (8)$$

and

$$\frac{dx(r)}{dr} = \frac{1}{4\gamma} \frac{s^2}{r(s-r)}, \quad (9)$$

we then have

$$f_R(r) = \frac{1}{4\gamma\sigma} \frac{s^2}{r(s-r)} g\left(\frac{\theta - \mu}{\sigma} - \frac{s}{4\gamma\sigma} \log\left(\frac{s-r}{r}\right)\right). \quad (10)$$

We see immediately that $f_R(r) = f_{R'}(r) \forall r \in [0, s]$, and hence $\mathcal{D}(f_R, f_{R'}) \equiv 0$ for any metric \mathcal{D} , provided that

$$\frac{\theta - \mu}{\sigma} = \frac{\theta' - \mu'}{\sigma'}, \quad (11)$$

$$\gamma\sigma = \gamma'\sigma'. \quad (12)$$

Equation 11 implies that a neuron adjusts the threshold θ to keep the Z -score of θ invariant with respect to the input statistics. This ensures that a neuron's mean output is independent of its mean input, where this mean output is determined by a neuron's preferred operating point. Equation 12 then sets the output gain in inverse proportion to the input standard deviation, so that larger (smaller) standard deviations correspond to smaller (larger) gains. This ensures that the dynamic range of a neuron's input is mapped onto an invariant dynamic range of its output, where this output dynamic range is again set by a neuron's preferred operating point, keeping the output dynamic range constant. This relationship between the gain of a neuron and the standard deviation of the input is a well-established experimental observation (see, e.g. Kvale and Schreiner 2004; Bonin et al. 2006; Maravall et al. 2007; and references therein).

Many standard distributions may be written in the form defined by Equation 7, including the normal distribution, the exponential, Laplace, doubly exponential distributions and the logistic distribution, and in general an infinity of forms for the function g is available. For this class of input distributions, if a neuron sets θ and γ according to

$$\theta = \mu + \Theta\sigma, \quad (13)$$

$$\gamma = \Gamma \frac{s}{\sigma}, \quad (14)$$

where the constants Θ and Γ determine the preferred operating point of a neuron, then a neuron remains at this operating point, with $f_R(r)$ exactly invariant in response to changes in the mean and variance of the input distribution X . We have made explicit a factor of s in Equation 14 since scaling the response range s will scale

the response gain γ by the same factor. With s made explicit, Γ will then be independent of s and scale-free.

For a completely general form of input distribution $f_X(x)$, Equations 13 and 14 will not suffice to maintain response PDF invariance. However, we propose that a neuron nonetheless sets θ and γ according to Equations 13 and 14. First, these rules are simple to implement. Second, they result in exact invariance for a large class of input distributions including the normal distribution. Third, they ensure adaptation to the lowest-order moments, μ and σ , of a distribution, albeit at the price of possible non-invariance due to the higher-order moments of completely general distributions. We might expect, in any event, that a neuron should not adapt perfectly to higher-order moments (but see Kvale and Schreiner 2004), since such moments contain information of importance to learning algorithms such as independent component analysis (Hyvarinen et al. 2001).

The confounding influence of the higher-order moments m_3, \dots on the invariance of $f_R(r)$ in the presence of a general input distribution $f_X(x)$ can be confirmed by transforming to the Z variable,

$$\begin{aligned} f_R(r) &= \frac{dz(r)}{dr} f_Z(z(r)) \\ &= \frac{1}{4\gamma\sigma} \frac{s^2}{r(s-r)} f_Z\left(\frac{\theta - \mu}{\sigma} - \frac{s}{4\gamma\sigma} \log\left(\frac{s-r}{r}\right)\right). \end{aligned} \quad (15)$$

Of course, Equation 15 is identical to Equation 10 with the PDF of the Z -transformed input f_Z replacing g . However, for a general form of f_Z , Equation 15 is not invariant under Equations 13 and 14 because the higher-order moments of Z are not invariant under changes in μ and σ . To see this, we write $f_Z(z)$ to exhibit explicitly all its moments,

$$f_Z(z) = f_Z(z; 0, 1, \tilde{m}_3, \tilde{m}_4, \dots), \quad (16)$$

where 0 and 1 are the first and second moments of Z , by definition, and the i th moment of Z , \tilde{m}_i , is just the i th central moment of X divided by σ^i , $\tilde{m}_i = \mu_i/\sigma^i$. Making explicit the dependence of f_R on θ and γ , we may then write

$$f_R(r; \theta, \gamma) = \frac{1}{4\gamma\sigma} \frac{s^2}{r(s-r)} f_Z\left(\frac{\theta - \mu}{\sigma} - \frac{s}{4\gamma\sigma} \log\left(\frac{s-r}{r}\right); 0, 1, \tilde{m}_3, \dots\right), \quad (17)$$

and under a change of statistics, $\mu \rightarrow \mu'$, $\sigma \rightarrow \sigma'$, $\tilde{m}_i \rightarrow \tilde{m}'_i$ and adaptive changes $\theta \rightarrow \theta'$, $\gamma \rightarrow \gamma'$, we have

$$f_R(r; \theta', \gamma') = \frac{1}{4\gamma'\sigma'} \frac{s^2}{r(s-r)} f_Z\left(\frac{\theta' - \mu'}{\sigma'} - \frac{s}{4\gamma'\sigma'} \log\left(\frac{s-r}{r}\right); 0, 1, \tilde{m}'_3, \dots\right). \quad (18)$$

Now, if a neuron implements Equations 13 and 14 in an attempt to maintain an invariant output PDF, we have

$$f_R(r; \theta, \gamma) = \frac{1}{4\Gamma} \frac{s}{r(s-r)} f_Z\left(\Theta - \frac{1}{4\Gamma} \log\left(\frac{s-r}{r}\right); 0, 1, \tilde{m}_3, \dots\right), \quad (19)$$

$$f_R(r; \theta', \gamma') = \frac{1}{4\Gamma} \frac{s}{r(s-r)} f_Z\left(\Theta - \frac{1}{4\Gamma} \log\left(\frac{s-r}{r}\right); 0, 1, \tilde{m}'_3, \dots\right), \quad (20)$$

where Θ and Γ are constants as shown above, so that

$$\begin{aligned}
 & f_R(r; \theta, \gamma) - f_R(r; \theta', \gamma') \\
 &= \frac{1}{4\Gamma} \frac{s}{r(s-r)} \sum_{i=3}^{\infty} (\tilde{m}_i - \tilde{m}'_i) \frac{\partial f_Z}{\partial \tilde{m}_i} \left(\Theta - \frac{1}{4\Gamma} \log\left(\frac{s-r}{r}\right); 0, 1, \xi_3, \xi_4, \dots \right), \quad (21)
 \end{aligned}$$

for some values of $\xi_i \in (\tilde{m}_i, \tilde{m}'_i)$. Hence, the attempt to maintain the invariance of f_R is contaminated by the possible changes in the higher-order moments of the Z -transformed input distribution.

This contamination may arise from a number of sources. For a general distribution f_Z , the higher-order moments \tilde{m}_3, \dots will in general change when μ and σ are changed. Although Equations 13 and 14 attempt to accommodate changes in μ and σ , these changes will leak into the higher-order moments. However, for the specific form of distribution $f_Z(z) = g(z)$ discussed above, the higher-order moments \tilde{m}_3, \dots are independent of μ and σ and hence are constants. Thus, for the form $f_Z = g$, the invariance of f_R is exact precisely because the higher-order moments are, by construction, also invariant. Another source of ‘‘contamination’’ could arise from distributions in which some of the higher-order moments are independent degrees of freedom that may be freely changed in an attempt to probe a neuron’s ability to adapt to higher-order moments (Kvale and Schreiner 2004; Bonin et al. 2006). In this case, of course, the non-invariance of f_R is explicit and direct, rather than implicit and indirect.

Despite the breakdown in exact invariance of $f_R(r)$ in the case of a completely general input distribution $f_X(x)$, we see from Equation 21 that the magnitude of the change in the output PDF is controlled by the magnitude of the change in the higher-order moments. Unless an input distribution exhibits a high degree of sensitivity to its higher-order moments, we would thus not expect the non-invariance of f_R to be too severe. Of course, it is always possible to construct counter-examples to such arguments, but we might expect naturally-occurring input statistics to be reasonably well-behaved (Simoncelli and Olshausen 2001).

In the foregoing, we have allowed the possibility that X may contain noise from the environment, from signal transduction processes, and from neuronal transmission, but we have not considered the possibility that the output PDF could be contaminated by noise. In fact, under a simple additive model of output noise, in which the actual output distribution R is the sum of the noiseless output R_0 and a noise source N , $R = R_0 + N$, so that

$$\begin{aligned}
 f_R(r) &= \int dr_0 f_{R_0}(r_0) f_N(r - r_0) \\
 &= \int dx f_X(x) f_N(r - r(x)), \quad (22)
 \end{aligned}$$

where f_N is the PDF of the noise process, it is a simple matter to show that the rules in Equations 13 and 14 follow directly from a Z -transformation under the integral sign. The adaptation rules in Equations 13 and 14 are therefore also valid under the assumption of additive output noise.

So far, we have been concerned with how a neuron maintains its preferred operating point in the face of changing input statistics by adapting its threshold θ and gain γ so as to keep its output PDF f_R invariant, or approximately invariant.

In our approach, the operating point of a neuron is set by the two constants Θ and Γ in Equations 13 and 14. Thus, we have not been concerned with how Θ and Γ are set by a neuron. We now consider the possibility that Θ and Γ are set so as to maximise the entropy of the output distribution R , given the sigmoidal transfer function in Equation 5.

The output entropy $S[R]$ is defined by

$$S[R] = - \int dr f_R(r) \log f_R(r), \quad (23)$$

from which we have

$$S[R] = - \int dx f_X(x) \log \left[\frac{dx}{dr(x)} f_X(x) \right]. \quad (24)$$

Writing

$$\gamma(x) = \frac{dr(x)}{dx}, \quad (25)$$

the instantaneous gain of the neuronal response $r(x)$ for input x , with $\gamma(\theta) \equiv \gamma$ at semi-saturation, we then have

$$S[R] = S[X] + \langle \log \gamma(x) \rangle_X, \quad (26)$$

where $S[X]$ is the entropy of the input distribution and $\langle \rangle_X$ means an average over the distribution X . If the response function is characterised by a set of adjustable parameters p_i , so that $r(x) = r(x; p_i)$, then we maximise the entropy $S[R]$ with respect to these parameters by evaluating the derivatives

$$\frac{\partial S[R]}{\partial p_i} = \left\langle \frac{1}{\gamma(x)} \frac{\partial \gamma}{\partial p_i}(x; p_i) \right\rangle_X \quad (27)$$

and setting them to zero. The derivative $\partial \gamma(x; p_i) / \partial p_i$ is the sensitivity of the instantaneous gain of the response function to the parameter p_i . For the sigmoidal function in Equation 5, we have

$$\gamma(x) = \gamma \operatorname{sech}^2 \frac{2\gamma}{s}(x - \theta), \quad (28)$$

depending on the two parameters θ and γ , from which we obtain the two conditions

$$\frac{\partial S[R]}{\partial \theta} = \frac{4\gamma}{s} \int dx f_X(x) \tanh \frac{2\gamma}{s}(x - \theta) = 0, \quad (29)$$

$$\frac{\partial S[R]}{\partial \gamma} = \frac{1}{\gamma} - \frac{4}{s} \int dx x f_X(x) \tanh \frac{2\gamma}{s}(x - \theta) = 0, \quad (30)$$

and hence values of θ and γ that maximise $S[R]$. For the given values of μ and σ associated with input X , we can then fix Θ and Γ from Equations 13 and 14. Of course, in principle adaptation could take the form of determining θ and γ directly from Equations 29 and 30, and thus of always being optimal, in the sense of achieving maximum output entropy. However, if Θ and Γ are fixed once, perhaps over evolutionary or developmental timescales, Equations 13 and 14 then afford an arguably computationally easier method of remaining at, or near, maximum output entropy.

Multimodal inputs

Many neurons, both subcortical and cortical, receive input not from just one sensory modality, but often from two or three distinct sensory modalities. For example, neurons in the deep layers of the superior colliculus can receive and integrate visual, auditory and somatosensory input. Each modality will be associated with its own intrinsic statistical parameters, and these may vary independently of those in other modalities. Moreover, different input sources may exhibit time-dependent correlations. A multimodal neuron may therefore be expected to exhibit adaptation to all its various input sources. Then, it is natural to consider extending the principles of adaptation in unimodal neurons to multimodal neurons.

Therefore, we now turn to the case in which a neuron receives input from at least two different sensory channels. For simplicity, we restrict to bimodal neurons, but our results generalise to the full, multimodal case. We therefore consider two input channel distributions, X_1 and X_2 with PDFs $f_{X_1}(x_1)$ and $f_{X_2}(x_2)$ and means μ_1 and μ_2 and variances σ_1^2 and σ_2^2 , respectively. The correlation coefficient between the two channels X_1 and X_2 is defined to be ρ . In general, of course, a joint PDF $f_{X_1, X_2}(x_1, x_2)$ defines the joint input distribution, with $f_{X_1}(x_1)$ and $f_{X_2}(x_2)$ being the marginal distributions.

For a unimodal input, we defined the response function in Equation 5 so that the argument of the tanh function is, up to factors, just $\gamma(x - \theta)$. We then derived rules for adapting θ and γ to the input statistics based on the invariance of the output PDF $f_R(r)$. We have not discussed how a neuron instantiates its threshold θ and gain γ , and we have also not discussed the implementation mechanism leading to adaptive changes in θ and γ . Many mechanisms are implicated in adaptation to input statistics, both at the single neuron level and the circuit level (Sanchez-Vives et al. 2000; Rieke 2001; Baccus and Meister 2002; Chance et al. 2002; Kim and Rieke 2003; Shu et al. 2003; Dean et al. 2005; Ingham and McAlpine 2005; Arganda et al. 2007). If adaptation occurs at the single neuron level, then information about the stimulus mean and standard deviation must be available locally. If adaptation to stimulus statistics occurs for a neuron with multimodal inputs, then locality demands that the statistics pertaining to a single input channel are available only at that channel's synapses onto the neuron, or at least at the local dendritic level, rather than the whole neuron level, at which presumably all the separate statistics for the individual input channels become merged and therefore lost.

Since we now wish to discuss the possibility of adaptation of a multimodal neuron to the separate statistics of its different input channels, we will therefore write the response function $r(x_1, x_2)$, for a bimodal neuron, in the form

$$r(x_1, x_2) = \frac{s}{2} \left\{ 1 + \tanh \frac{2}{s} [\gamma_1(x_1 - \theta_1) + \gamma_2(x_2 - \theta_2)] \right\}, \quad (31)$$

where the ‘‘thresholds’’ θ_1 and θ_2 and the ‘‘gains’’ γ_1 and γ_2 are now specific to each input distribution X_1 and X_2 and permit adaptation according to only locally-available information about each input. Of course, θ_i and γ_i are no longer neuronal thresholds and gains in the conventional, unimodal sense, but we retain this nomenclature because of the clear analogy with the unimodal case.

Writing $U = \gamma_1 X_1 + \gamma_2 X_2$, $\gamma_U \equiv 1$ and $\theta_U = \gamma_1 \theta_1 + \gamma_2 \theta_2$, the bimodal response function in Equation 31 can be reduced to an effective unimodal response function,

$$r(u) = \frac{s}{2} \left[1 + \tanh \frac{2\gamma_U}{s} (u - \theta_U) \right], \quad (32)$$

where the effective neuronal input u takes values from the effective input distribution U . The mean and variance of U are given by

$$\mu_U = \gamma_1 \mu_1 + \gamma_2 \mu_2, \quad (33)$$

$$\sigma_U^2 = \gamma_1^2 \sigma_1^2 + \gamma_2^2 \sigma_2^2 + 2\gamma_1 \sigma_1 \gamma_2 \sigma_2 \rho, \quad (34)$$

respectively, where the correlation coefficient ρ appears in the expression for σ_U^2 . The output PDF $f_R(r)$ is therefore given by

$$f_R(r) = \frac{du(r)}{dr} f_U(u(r)), \quad (35)$$

where $u(r)$ is the inverse function of $r(u)$, given by an equation analogous to Equation 8 under the replacements $\theta \rightarrow \theta_U$ and $\gamma \rightarrow \gamma_U$, and $f_U(u)$ is the PDF of the effective unimodal input U , determined from the joint PDF $f_{X_1 X_2}(x_1, x_2)$. Since $\gamma_U \equiv 1$, $f_R(r)$ can be written simply as

$$f_R(r) = \frac{s^2}{4r(s-r)} f_U \left(\theta_U - \frac{s}{4} \log \frac{s-r}{r} \right). \quad (36)$$

This reduction allows us to write down, by analogy, the rules according to which the separate channel thresholds and gains should be set to ensure the (perhaps approximate) invariance of $f_R(r)$ according to our previous results.

Adapting to the effective, unimodal input distribution U , and hence to the actual, bimodal input distributions X_1 and X_2 is then achieved, according to our earlier rules in Equations 13 and 14, by setting

$$\theta_U = \mu_U + \Theta \sigma_U, \quad (37)$$

$$\gamma_U = \Gamma \frac{s}{\sigma_U}, \quad (38)$$

where the constants Θ and Γ determine, as usual, the preferred operating point of a neuron. We therefore have

$$\gamma_1(\theta_1 - \mu_1) + \gamma_2(\theta_2 - \mu_2) = \Theta \Gamma s, \quad (39)$$

$$\gamma_1^2 \sigma_1^2 + \gamma_2^2 \sigma_2^2 + 2\gamma_1 \sigma_1 \gamma_2 \sigma_2 \rho = \Gamma^2 s^2. \quad (40)$$

Equation 40 defines, in general, an ellipse (or ellipsoid for the multimodal case) on which solutions γ_1 and γ_2 of this equation exist. The semi-major and semi-minor axes of this ellipse are set by σ_1 and σ_2 , as well as by ρ . Thus, a general point on the ellipse will be influenced by both σ_1 and σ_2 , and hence the solution for γ_1 , say, will be influenced by σ_2 . However, we have argued that information about the σ_i should only be available locally, at or near the site of the synapses associated with the inputs X_i . Hence, the gain associated with, say,

input X_1 , γ_1 , should not be influenced, non-locally, by the statistics of X_2 . We can achieve this by insisting that

$$\gamma_i \sigma_i = \Gamma_i s, \quad (41)$$

where the Γ_i are constants and, as usual, an overall scale is made explicit. In order to ensure that each modality is mapped onto the same output dynamic range, we set all these constants equal, so that $\Gamma_i = \Gamma$, $\forall i$. Then, from Equation 40, we have $\Gamma' = \Gamma/\sqrt{2(1+\rho)}$. Thus, the separate channel gains should be set according to

$$\gamma_i = \frac{\Gamma}{\sqrt{2(1+\rho)}} \frac{s}{\sigma_i}, \quad (42)$$

providing solutions of Equation 40 satisfying the principle of locality. The channel-specific gains γ_i are therefore fixed locally, by the factor of $1/\sigma_i$, but the overall scaling also depends on the correlation coefficient ρ . For the possible influence of correlations on gain control, although not in a multimodal context, see Sharpee et al. (2006) and Lesica et al. (2007). For the general, multimodal case, with n distinct channels, a sum of all $(1/2)n(n-1)$ separate correlation coefficients between all distinct pairs of channels appear in the denominator in Equation 42. The overall scaling is a global factor affecting all gains equally, and this global modification of the local gains could be achieved at the whole neuron level, rather than at the local input level. Nonetheless, Equation 42 does require knowledge of the correlation coefficients. We shall discuss this later.

Turning to Equation 39, since $\theta_U \equiv \gamma_1 \theta_1 + \gamma_2 \theta_2$, we have

$$\theta_U = \gamma_1 \mu_1 + \gamma_2 \mu_2 + \Theta \Gamma s. \quad (43)$$

This equation defines the combined threshold θ_U and not θ_1 and θ_2 separately. However, we defined the θ_i separately only for convenience, for the analogy to the unimodal case. We see from Equation 31 that only the combination θ_U is real, with θ_1 and θ_2 having no independent meaning. θ_U may be regarded as the actual threshold of the neuron, although in the multimodal case it does not specify the semi-saturation point, and Equation 43 sets this threshold uniquely.

In summary, in the bimodal case, the response function of the neuron is given by

$$r(x_1, x_2) = \frac{s}{2} \left[1 + \tanh \frac{2}{s} (\gamma_1 x_1 + \gamma_2 x_2 - \theta_U) \right], \quad (44)$$

and if the separate channel gains γ_i are adapted locally according to Equation 42 and the pseudo-threshold θ_U is adapted according to Equation 43, then the output PDF $f_R(r)$ remains invariant, or approximately so. Thus, Equations 42 and 43 provide a means for a bimodal neuron to adapt to the changing input statistics of two distinct input channels, when these channels' statistics vary either separately or simultaneously. These results generalise directly to a multimodal neuron with more than two separate input channels.

Examples of adaptation in model neurons

First, we consider the application of our adaptation rules to the case of a purely unimodal neuron and show examples of adaptation, both perfect and imperfect,

to changes in the input mean and standard deviation. We then consider a bimodal neuron and discuss, in particular, adaptation to changes in the correlation coefficient, which is not an available form of adaptation in the unimodal case.

Unimodal inputs

We demonstrate the ability of the adaptation rules in Equations 13 and 14 to accommodate changes in the input statistics. We consider one example of an input distribution that allows perfect adaptation, and one example demonstrating only approximate invariance of $f_R(r)$.

Logistic distribution. We first consider the logistic or sech-squared distribution, defined by the PDF

$$f_X(x) = \frac{\pi}{4\sigma\sqrt{3}} \operatorname{sech}^2 \frac{\pi}{2\sqrt{3}} \left(\frac{x - \mu}{\sigma} \right), \quad (45)$$

where μ and σ^2 are its mean and variance, respectively. Clearly $f_X(x)$ has a form for which $f_Z(z) = g(z)$, with $g(z) = (\pi/4\sqrt{3}) \operatorname{sech}^2(\pi/2\sqrt{3})z$, so perfect adaptation of θ and γ to changing μ and σ is possible. We consider the logistic distribution for two reasons. First, it is frequently employed as an alternative to the normal distribution, in order to simplify analysis (Johnson et al. 1995). Second, as mentioned above, it is the only input distribution for which it is possible, in the presence of the sigmoidal response function in Equation 5, to generate an output distribution R that is uniform on $[0, s]$, and hence has the maximum entropy distribution, given the saturation constraint.

Since $f_R(r) = 1/s$ has maximum entropy on the bounded interval $[0, s]$, we deduce that Equations 29 and 30 are satisfied when

$$\theta = \mu, \quad (46)$$

$$\gamma = \frac{\pi}{4\sqrt{3}} \frac{s}{\sigma}, \quad (47)$$

from which we see that the constants Θ and Γ satisfy $\Theta = 0$ and $\Gamma = (\pi/4\sqrt{3}) \approx 0.4534$. The scaling of the gain γ with the range of the response s was anticipated earlier, giving rise to the scale-free definition of the constant Γ . This can be seen directly from Equations 29 and 30.

In Figure 1, we consider two different operating points. The first corresponds to the maximum entropy operating point, defined by $\Theta = 0$ and $\Gamma = \pi/4\sqrt{3}$ (Figure 1A), and the second corresponds to a representative example of an operating point away from maximum entropy, which we take to be defined by $\Theta = -0.5$ and $\Gamma = 0.4$ (Figure 1B). In both cases, we assume that the neuron has adapted to an input mean $\mu = 10$ and standard deviation $\sigma = 2$, and use Equations 13 and 14 to determine the induced thresholds and gains. In Figure 1(A), we have a threshold $\theta = 10$ and gain $\gamma = s\Gamma/2$. We set $s = 10$ without loss of generality throughout. In Figure 1(B), we have threshold $\theta = 9$ and gain $\gamma = 2$. The PDFs of the corresponding output distributions are shown, defining the preferred operating points of the neuron in both cases. When the input standard deviation is changed from $\sigma = 2$ to $\sigma = 1$ without any concomitant neuronal adaption, the output PDFs




Figure 1. Adaptation of a unimodal neuron with input defined by a logistic distribution. Two different operating points are considered, defined by (A) $\Theta = 0$ and $\Gamma = \pi/4\sqrt{3}$, corresponding to a maximum entropy output PDF, and (B) $\Theta = -0.5$ and $\Gamma = 0.4$, corresponding to an operating point away from maximum entropy. In each case, the neuron is initially adapted to an input mean $\mu = 10$ and standard deviation $\sigma = 2$ (PDF a). The standard deviation is then changed to $\sigma = 1$ (PDFs b–e). Output PDFs are shown without any concomitant changes in θ and γ , showing how the output PDFs move away from their preferred operating points (PDF b). Output PDFs are also shown as θ and γ are moved to their target values for $\sigma = 1$ (PDFs c–e). (A) The values of γ are: $\gamma = 2.27$ (PDF a); $\gamma = 2.27$ (PDF b); $\gamma = 2.83$ (PDF c); $\gamma = 3.40$ (PDF d); $\gamma = 3.97$ (PDF e). In all cases, $\theta = 10$, a constant independent of σ because $\Theta = 0$. The PDF corresponding to the target value of $\gamma = 4.53$ for $\sigma = 1$ is not shown, because it is identical to PDF a. (B) The values of θ and γ are: $\theta = 9$, $\gamma = 2$ (PDF a); $\theta = 9$, $\gamma = 2$ (PDF b); $\theta = 9.125$, $\gamma = 2.5$ (PDF c); $\theta = 9.25$, $\gamma = 3$ (PDF d); $\theta = 9.375$, $\gamma = 3.5$ (PDF e). The PDF for $\theta = 9.5$, $\gamma = 4$, the target values for $\sigma = 1$, is again not shown.

move away from the preferred operating point. To return the PDFs back to the preferred operating points, we must set θ and γ according to Equations 13 and 14. Incrementing θ and γ from their initial values at $\sigma = 2$ to their target values at $\sigma = 1$ shows how the output PDFs return to the preferred operating points. We see that adaptation is in this case perfect, and, crucially, does not depend on the details of how the preferred operating point is set. In particular, when the preferred operating point is initially set according to a maximum entropy principle, as shown in Figure 1(A), adaptation restores the output PDF to the maximum entropy

distribution in precisely the same manner that it returns the output PDF to a non-optimal preferred operating point, as shown in Figure 1(B). The use of Equations 13 and 14 thus avoids an explicit recomputation of the maximum entropy point following a change in the input statistics.

Convolved exponential distribution. We now consider an input distribution for which perfect adaptation is not possible, defined by the PDF

$$f_X(x) = \frac{\exp(-x/\lambda_+) - \exp(-x/\lambda_-)}{\lambda_+ - \lambda_-}, \quad x \geq 0. \tag{48}$$

This is the distribution of a variable X defined as the sum of two independent, exponentially-distributed variables X_{\pm} with different means λ_{\pm} , respectively, so that $X = X_+ + X_-$. The mean and variance of X are given by $\mu = \lambda_+ + \lambda_-$ and $\sigma^2 = \lambda_+^2 + \lambda_-^2$, respectively. Regarding μ and σ as the fundamental parameters, we invert these relations to obtain

$$\lambda_{\pm} = \frac{1}{2} \left(\mu \pm \sqrt{2\sigma^2 - \mu^2} \right). \tag{49}$$

For this distribution to exist, we must have $\sigma < \mu < \sqrt{2}\sigma$, so that the mean can only take values defined in a range set by the standard deviation, and *vice versa*. Writing $q = \mu / \sigma$, so that $1 < q < \sqrt{2}$, and transforming to the Z variable, we have

$$f_Z(z) = \frac{\exp\left[\frac{-2(z+q)}{q + \sqrt{2-q^2}}\right] - \exp\left[\frac{-2(z+q)}{q - \sqrt{2-q^2}}\right]}{\sqrt{2-q^2}}. \tag{50}$$

We see that $f_Z(z)$ does not take the form $g(z)$ considered earlier, because the ratio q , depending on μ and σ , appears in $f_Z(z)$. Therefore, adapting θ and γ according to Equations 13 and 14 will not, in general, result in an exactly invariant output PDF $f_R(r)$.

We set $\mu = 10$ and take a value $\sigma = 8.5$ [approximately midway in the allowed range of σ of $(5\sqrt{2}, 10)$] and use Equations 29 and 30 to determine the values of θ and γ corresponding to the maximum entropy output PDF for the convolved exponential input PDF. For this input distribution, Equations 29 and 30 cannot be solved analytically, so we perform direct numerical searches for the values of θ and γ that induce the maximum entropy output distribution for given input mean and standard deviation. For $\mu = 10$ and $\sigma = 8.5$, we find $\theta \approx 8.68$ and $\gamma \approx 0.60$, corresponding to values $\Theta \approx -0.16$ and $\Gamma \approx 0.51$. In Figure 2, we again consider two operating points, one defined by this maximum entropy distribution, and the second, as for the logistic distribution, defined by $\Theta = -0.5$ and $\Gamma = 0.4$. We show the induced output PDFs when the thresholds and gains are set according to the input statistics, defining the preferred operating point. Also shown are how these PDFs move away from the preferred operating points when σ changes to $\sigma = 7.5$ without compensating neuronal adaptation. However, we see that the PDFs do not return to the preferred operating points when the thresholds and gains are restored to the target values induced by Equations 13 and 14 when $\sigma = 7.5$.

Despite this expected, general non-invariance of the output PDF $f_R(r)$ for convolved exponential input, it is natural to wonder whether adaptation




Figure 2. Attempted adaptation of a unimodal neuron with input defined by a convolved exponential distribution. (A) Fixing $\mu = 10$ and $\sigma = 8.5$ results in $\theta = 8.68$ and $\gamma = 0.60$ for a maximum entropy output distribution, inducing the values $\Theta = -0.16$ and $\Gamma = 0.51$, which we take as defining the preferred operating point. The second operating point is defined by: (B) $\Theta = -0.5$ and $\Gamma = 0.4$, corresponding to a preferred operating point away from maximum entropy when $\mu = 10$ and $\sigma = 8.5$. The format of this figure is otherwise identical to Figure 1, except that we change σ from $\sigma = 8.5$ (PDF a) to $\sigma = 7.5$ (PDFs b–e). (A) The values of θ and γ are: $\theta = 8.68$ and $\gamma = 0.60$ (PDF a, the preferred operating point); $\theta = 8.68$, $\gamma = 0.60$ (PDF b); $\theta = 8.73$, $\gamma = 0.63$ (PDF c); $\theta = 8.79$, $\gamma = 0.65$ (PDF d); $\theta = 8.84$, $\gamma = 0.68$, the target values for $\sigma = 7.5$ (PDF e). (B) The values of θ and γ are: $\theta = 5.75$ and $\gamma = 0.47$ (PDF a, the preferred operating point); $\theta = 5.75$, $\gamma = 0.47$ (PDF b); $\theta = 5.92$, $\gamma = 0.49$ (PDF c); $\theta = 6.08$, $\gamma = 0.51$ (PDF d); $\theta = 6.25$, $\gamma = 0.53$, the target values for $\sigma = 7.5$ (PDF e).

according to Equations 13 and 14 is better than not adapting θ and γ at all. To this end, we determine the extent to which adaptation according to Equations 13 and 14 is able to track the maximum entropy distribution governed by the solutions of Equations 29 and 30, and compare this to the deviation induced in the absence of adaptation of θ and γ to changes in μ and σ . Thus, here, we regard the maximum entropy distribution as the preferred operating point of a neuron, and determine the deviations from this preferred operating point when either θ and γ are imperfectly adapted according to Equations 13 and 14 or θ and γ are not adapted at all. For $\mu = 10$ and $\sigma = 7.2$, the maximum entropy operating point, from the solution of Equations 29 and 30 is set by $\theta \approx 9.10$ and

$\gamma \approx 0.66$, inducing values $\Theta \approx -0.13$ and $\Gamma \approx 0.48$. We increase σ up to its upper limit of $\mu = 10$ and for each value of σ , we compute three different output entropies. First, we determine the new, maximum output entropy values of θ and γ from Equations 29 and 30, and the associated value of $S_{\max}[R]$. In this case, Θ and Γ are not fixed, but determined functionally. Second, we determine the entropy, $S_{\text{adapt}}[R]$, of the output PDF when θ and γ are instead adapted according to Equations 13 and 14, so that Θ and Γ are held constant. Finally, we calculate the output PDF entropy, $S_{\text{no}}[R]$ when θ and γ are held constant, so that there is no adaptation to the changing standard deviation. The results are shown in Figure 3. We see that adapting θ and γ according to Equations 13 and 14, although not perfect, tracks the maximum entropy distribution very closely, while the non-adapted distribution deviates from the target distribution to a greater extent. It is worthwhile commenting that the maximum output entropy here is not constant precisely because the output distribution cannot be made exactly invariant. If we were to display a similar graph to this for the logistic distribution, however, $S_{\max}[R]$ would remain constant as σ is varied, and the perfect adaptation present in that case would ensure that $S_{\text{adapt}}[R] = S_{\max}[R]$ for all values of σ provided that equality is established for any one value of σ .

Multimodal inputs

We have seen that the multimodal input case can be reduced, mathematically-speaking, to the unimodal input case by considering the effective input $U = \gamma_1 X_1 + \gamma_2 X_2$ and adapting the channel-specific gains γ_i and the pseudo-threshold θ_U according to Equations 42 and 43, respectively. Our presentation of the above results for the unimodal case therefore, in general, completely characterises the multimodal results too. We thus discuss only one example of the

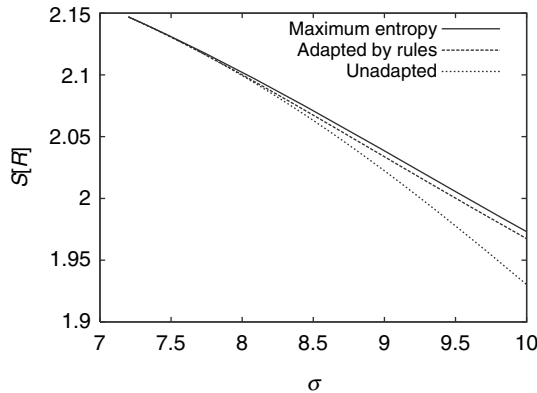


Figure 3. Tracking the maximum entropy output PDF by adaptation according to Equations 13 and 14, as σ varies, for a unimodal neuron with input drawn from a convolved exponential distribution. Here the preferred operating point of the neuron is defined functionally, as the maximum entropy output PDF. Adaptation according to Equations 13 and 14 in this case is imperfect, but tracks the maximum entropy distribution closely. In contrast, not adapting the output PDF quickly shifts the output PDF away from the preferred, maximum entropy operating point.

invariance of $f_R(r)$ in the presence of multimodal inputs for illustrative purposes. Novel to the multimodal case is the possibility of adaptation to the correlation coefficients between different modalities, so we focus on this case specifically.

As perhaps the simplest, non-trivial example of a bivariate distribution, we suppose that the two channel inputs are drawn from a bivariate normal distribution with joint PDF $f_{X_1, X_2}(x_1, x_2)$ given by

$$f_{X_1, X_2}(x_1, x_2) = \frac{1}{2\pi} \frac{1}{\sqrt{\det \Sigma}} \exp\left[-\frac{1}{2}(\underline{x} - \underline{\mu})^T \Sigma^{-1}(\underline{x} - \underline{\mu})\right], \quad (51)$$

where $\underline{x}^T = (x_1, x_2)$, the superscript T denoting the transpose, $\underline{\mu}^T = (\mu_1, \mu_2)$ and Σ is the covariance matrix,

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_1\sigma_2\rho \\ \sigma_1\sigma_2\rho & \sigma_2^2 \end{pmatrix}, \quad (52)$$

where $\det \Sigma$ is its determinant. It is easy to see that the distribution of the effective, unimodal input U is normal, with mean and variance given by Equations 33 and 34, respectively. In order to determine the operating point corresponding to the maximum entropy output distribution in the presence of a bivariate normal input distribution, we must calculate the values of Θ and Γ determining this point for a univariate normal input distribution.

A univariate normal input is defined by the standard PDF

$$f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{x - \mu}{\sigma}\right)^2\right], \quad (53)$$

with mean μ and variance σ^2 , respectively. As with the logistic distribution, the normal distribution can clearly be written in the form $f_Z(z) = g(z)$, permitting perfect adaptation. Unlike the logistic distribution, the maximum entropy output distribution R , given the sigmoidal response function in Equation 5, must be determined explicitly from the evaluation of the integrals in Equations 29 and 30. For the first integral, we have that

$$\frac{\partial S[R]}{\partial \theta} = \frac{1}{\sqrt{2\pi}} \frac{4\gamma}{s} \int_{-\infty}^{\infty} dz e^{-z^2/2} \tanh \frac{2\gamma\sigma}{s} \left(z - \frac{\theta - \mu}{\sigma}\right). \quad (54)$$

Since $\exp(-z^2/2)$ is even around $z=0$, the integral can be made to vanish when we choose the argument of the tanh function so that the tanh function is odd around $z=0$. This is possible only if $\theta = \mu$. Indeed, it is clear that for any input distribution that is symmetric about its mean, we must set $\theta = \mu$ in order to satisfy Equation 29, which implies that $\Theta \equiv 0$. With Equation 29 satisfied, Equation 30 reduces to

$$\Gamma \int_{-\infty}^{\infty} dz z e^{-z^2/2} \tanh 2\Gamma z = \sqrt{\frac{\pi}{8}}, \quad (55)$$

from which we must determine the solution for Γ numerically, resulting in $\Gamma \approx 0.4372$ for the maximum entropy output PDF. Notice the similarity between this value of Γ for the normal distribution and the value $\Gamma = (\pi/4\sqrt{3}) \approx 0.4534$ for the logistic distribution, confirming the utility of

replacing a normal distribution by a logistic distribution for the purposes of analytical tractability (Johnson et al. 1995).

This examination of the unimodal, univariate normal input case suffices to understand the bimodal, bivariate normal input case. In the bimodal case, adaptation of γ_1 and γ_2 according to Equation 42 and θ_U according to Equation 43 leads to the exact invariance of $f_R(r)$ under changes in all the bivariate input statistics, including the correlation coefficient ρ . Since the correlation coefficient does not appear in the unimodal case studied earlier, we examine, in particular, adaptation to changes in ρ only, with the means μ_1 and μ_2 and the variances σ_1^2 and σ_2^2 held constant.

In Figure 4, we as usual consider two different operating points, the first (Figure 4A) corresponding to the maximum entropy distribution, with $\Theta \equiv 0$ and $\Gamma \approx 0.4372$ for a univariate normal distribution, and the second (Figure 4B) defined as usual by $\Theta = -0.5$ and $\Gamma = 0.4$. We fix $\mu_1 = 10$, $\sigma_1 = 2$ as for the logistic distribution above, and then fix $\mu_2 = 8$ and $\sigma_2 = 4$ as representative values. We select an initial correlation coefficient of $\rho = -0.5$. The preferred operating points then determine γ_1 , γ_2 and θ_U via Equations 42 and 43. We then decrease ρ to $\rho = -0.9$. Without adaptive changes in the γ_i and θ_U , the output PDFs moves away from the preferred operating points. As the γ_i and θ_U are restored to their induced values according to Equations 42 and 43, the output PDFs $f_R(r)$ return to the preferred operating points. Adaptation to changes in the correlation coefficient in this bivariate normal input case is perfect.

For the general multivariate normal input case with n inputs, there are n means, n variances, and $(1/2)n(n-1)$ correlation coefficients. Adaptation to all n first-order moments is accomplished through just one parameter, the pseudo-threshold θ_U . The n gains γ_i permit adaptation to the variances of the n inputs. Although it is an immediate consequence of the form of Equation 42 in the general case, in which a sum of all the correlation coefficients appears in the denominator, it is nonetheless remarkable that adaptation to all $(1/2)n(n+1)$ second-order moments, consisting of n variances and $(1/2)n(n-1)$ correlation coefficients, is possible with just n gain parameters.

Discussion

We first discuss the general issues surrounding our approach to adaptation, then briefly examine the specific issues regarding adaptation in neurons receiving input from multiple sources.

General considerations

In this article we have proposed an invariance principle for neuronal adaptation to changing input statistics. We have suggested that a neuron seeks to maintain an invariant output spike rate PDF by adjusting its threshold and gain to accommodate changes in the lowest-order moments of an input distribution. For simplicity, we have considered only a sigmoidal transfer function, although the linear region of such a function also provides a fair approximation to the linear, rectified transfer




Figure 4. Adaptation of a bimodal neuron to changes in the correlation coefficient for inputs drawn from a bivariate normal distribution. The input means and standard deviations are set as $\mu_1 = 10$, $\mu_2 = 8$, $\sigma_1 = 2$ and $\sigma_2 = 4$. (A) The preferred operating point is set by $\Theta \equiv 0$ and $\Gamma = 0.4372$, corresponding to a maximum entropy distribution. For an initial correlation coefficient $\rho = -0.5$, the induced pseudo-threshold and gains are set as $\theta_U = 30.60$, $\gamma_1 = 2.19$, $\gamma_2 = 1.09$ (PDF a), defining the preferred PDF. Changing the correlation coefficient to $\rho = -0.9$ while keeping the pseudo-threshold and gains constant shifts the neuron from its preferred PDF (PDF b). Moving the pseudo-threshold and gains to their target values for $\rho = -0.9$ restores the neuron's output PDF to its preferred point: $\theta_U = 40.07$, $\gamma_1 = 2.86$, $\gamma_2 = 1.43$ (PDF c); $\theta_U = 49.52$, $\gamma_1 = 3.54$, $\gamma_2 = 1.77$ (PDF d); $\theta_U = 58.98$, $\gamma_1 = 4.21$, $\gamma_2 = 2.11$ (PDF e). (B) The preferred operating point of the neuron is set by $\Theta = -0.5$ and $\Gamma = 0.4$, moving the neuron away from a maximum entropy distribution. For an initial correlation coefficient $\rho = -0.5$, the induced pseudo-threshold and gains are set as $\theta_U = 26$, $\gamma_1 = 2$, $\gamma_2 = 1$ (PDF a), defining the preferred PDF. Changing the correlation coefficient to $\rho = -0.9$ while keeping the pseudo-threshold and gains constant shifts the neuron from its preferred PDF (PDF b). Moving the pseudo-threshold and gains to their target values for $\rho = -0.9$ restores the neuron's output PDF to its preferred point: $\theta_U = 34.65$, $\gamma_1 = 2.62$, $\gamma_2 = 1.31$ (PDF c); $\theta_U = 43.31$, $\gamma_1 = 3.24$, $\gamma_2 = 1.62$ (PDF d); $\theta_U = 51.96$, $\gamma_1 = 3.85$, $\gamma_2 = 1.93$ (PDF e).

function popular in the visual processing literature. For a class of input distributions, including the normal distribution, we have shown that exact invariance is achieved. For completely general input distributions, however, changes in higher-order moments can contaminate the invariance of the output PDF, resulting in only approximate invariance. We have examined an example of

non-invariance, and shown that adaption according to the lowest-order moments nonetheless tracks the maximum output entropy distribution (assuming response saturation) closely. More generally, given that the output spiking process and any downstream processing will be noisy, maintaining precise invariance is perhaps neither required nor realistic. It likely suffices, for noisy neuronal systems, to maintain the preferred operating point of a neuron only within acceptable bounds, rather than precisely.

We have considered adaption only to the input mean and variance because the sigmoidal response function employed here provides us with two degrees of response freedom, making changes in the input mean and variance suitable targets for adaptive changes in those two parameters. Evidence from neurons in the inferior colliculus suggests that neurons may be able to adapt to moments higher than the second (Kvale and Schreiner 2004), although evidence from neurons in the lateral geniculate nucleus suggests that contrast gain control is sensitive only to the input mean and variance (Bonin et al. 2006). In order to admit adaptation to higher-order moments in our approach, it would be necessary to consider a more general response function characterised by more parameters. Of course, to allow perfect adaptation to the potentially infinite number of independent moments of a completely general input distribution, it would be necessary to consider an essentially arbitrary response function.

The extent to which a neuron can modify the functional relationship between its input spike rate and its output spike rate is, however, unclear. A neuron can modify its threshold and gain, but the input–output mapping is presumably not arbitrarily modifiable. It is therefore unrealistic to assume, in a moment-orientated approach to adaptation, that adaptation to more than a few moments is possible. The same issues, however, apply to information-theoretic approaches. Assuming, for example, that the output response is bounded, so that maximum output entropy is achieved with a uniform output probability distribution, optimal adaptation is achieved by setting a neuron's response function proportional to the CDF of the input distribution (Laughlin 1981). Over evolutionary time, it is possible that such a mapping has been acquired by neurons in the presence of natural stimuli. But can such neurons adapt to a rapid, dramatic change in their input distributions? An affirmative answer requires identical assumptions, in terms of the freedoms assumed to be available in the response function characterisation, as a moment-based view.

A view of adaptation based on moments of course requires that a neuron, or a circuit, can estimate the moments of its input distribution. We have not discussed here how that could be achieved, nor the timescales on which changes in moments would affect a neuron's, or a circuit's, estimates of them. It is a simple matter to build a model based on running estimates of moments requiring access only to the instantaneous input rate, and we will pursue this elsewhere. The higher the moment, however, the more complicated the form for the running estimate, so it is probably unrealistic to assume that a neuron can estimate more than a handful of the lowest-order moments. Equally, however, in the above information-theoretic considerations, estimation of the input CDF will be tightly constrained by the computational resources available to a neuron, or a circuit.

In the above, we have therefore restricted to a consideration of adaptation to the lowest-order moments, and in particular to the mean and variance, of an input distribution, in order to maintain an invariant output distribution. Compared to the

complexity of the full problem, requiring all the moments and an essentially arbitrary response function, restricting to the two lowest-order moments produces results, in terms of adapting the neuronal threshold and gain via Equations 13 and 14, that are easy to derive and simple to implement. Despite the fact that these rules ensure only approximate output PDF invariance for general input distributions, and although some functional properties of neurons may then escape our analysis (e.g. large deviation detection, for which the higher-order moments are clearly critical), we regard this simplicity as a virtue. Under the assumption that a neuron can indeed estimate the mean and variance of its input, Equations 13 and 14 provide an adaptive strategy that does not require a neuron to perform elaborate computations in order to adapt to changes in its input statistics. If we regard neurons not as perfect optimisers but rather as devices of extremely limited computational ability and resources, then adaptation according to Equations 13 and 14 may very well represent a solution for suboptimal, resource-constrained computation.

Multimodal adaptation

In a purely unimodal context, firing threshold adaptation and gain control are usually considered to be whole-neuron-level processes or circuit-level processes affecting the whole-neuron level (Sanchez-Vives et al. 2000; Rieke 2001; Baccus and Meister 2002; Chance et al. 2002; Kim and Rieke 2003; Shu et al. 2003; Dean et al. 2005; Ingham and McAlpine 2005; Arganda et al. 2007). Considering adaptation in multimodal neurons, however, we have allowed the possibility that a neuron admits separate gains for its distinct sensory inputs. We have argued that there is a local contribution to these gains, set by the standard deviations of the local, channel-specific inputs, and a global, common contribution to the gains, set by the correlation coefficients. Although adaptation processes at the whole-neuron level (for example, by changing the parameters of the spike generation mechanism) could accommodate changes in all the input means and the global change in the gains due to correlations, such processes are not suitable candidates for the local gain control for adaptation to local input standard deviations. It is possible that such local gain control could only be achieved at the circuit-level, with the circuit controlling or modulating either specific afferent inputs or the dendrites on which these afferents synapse. If afferents from different modalities synapse on the same dendrite, however, then it is difficult to conceive how the circuit could control or modulate dendritic properties in order to induce different gains associated with different modalities. Furthermore, the estimation of correlation coefficients is, by definition, a non-local process, since knowledge of the activities in pairs of distinct sensory inputs is required. One way to achieve this would be through the close juxtaposition of different modalities' synapses on a dendrite, thus ruling out local dendritic gain control. On balance, then, we consider that local dendritic gain control is unlikely in a multimodal setting, and that circuit-level processes control both the overall gain of a neuron, in order to allow adaptation to correlations (and means), and the local gains, in order to allow adaptation to variances.

The potential for adaptation in multimodal neurons raises the possibility that adaptation may play an important role in multisensory integration. Perhaps the best-studied form of multisensory integration occurs in the deep layers of the superior colliculus (DSC). Multisensory DSC neurons exhibit the properties of cross-modal

enhancement (CME), associated with inverse effectiveness, and modality-specific suppression (MSS). In CME, the response of a DSC neuron to simultaneous, multimodal stimulation is greater than its response to distinct, unimodal stimulation in only one channel (Meredith and Stein 1986). In MSS, simultaneous stimulation of separate parts of the receptive field in a single sensory channel of a DSC neuron can suppress the response of the neuron (Kadunce et al. 1997). If DSC neurons undergo adaptation to the multivariate statistics of their multisensory inputs, then these neurons could exhibit separate gains for their different sensory inputs. Moreover, these gains should also be sensitive to the correlation coefficients between different sensory modalities. The superior colliculus is the first sensory area in which multiple modalities converge, and thus offers scope for adaptation to correlation coefficients. The experimental opportunities for testing adaptation in DSC neurons are therefore rich, and in future work we shall examine in detail the consequences of adaptation in multisensory neurons for the dynamics of multisensory integration.

Note

- [1] Carleman's condition on the moments is sufficient to ensure uniqueness (Feller 1967). The classic counter-examples to this condition are the Cauchy distribution, which has no moments, and the log-normal distribution, for which $M_X(t)$ does not exist despite all the moments existing.

Acknowledgement

X. Kuang thanks the University of Southampton for the support of a studentship.

Declaration of interest: The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

References

- Arganda S, Guantes R, de Polavieja GG. 2007. Sodium pumps adapt spike bursting to stimulus statistics. *Nature Neuroscience* 10:1467–1473.
- Atick JJ. 1992. Could information theory provide an ecological theory of sensory processing?. *Network: Computation in Neural Systems* 3:213–251.
- Attave F. 1954. Some informational aspects of visual perception. *Psychology Review* 61:183–193.
- Baccus SA, Meister M. 2002. Fast and slow contrast adaptation in retinal circuitry. *Neuron* 36:909–919.
- Baddeley R, Abbott LF, Booth MC, Sengpiel F, Freeman T, Wakeman EA, Rolls ET. 1997. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society of London, Series B* 264:1775–1783.
- Barlow HB. 1961. Possible principles underlying the transformation of sensory messages. In: Rosenblith WA, editor. *Sensory communication*. Cambridge, MA: MIT Press. pp 217–234.
- Barlow HB, Mollon JD. 1982. *The senses*. Cambridge: Cambridge University Press.
- Bonin V, Mante V, Carandini M. 2006. The statistical computation underlying contrast gain control. *Journal of Neuroscience* 26:6346–6353.

- Brenner N, Bialek W, de Ruyter van Steveninck RR. 2000. Adaptive rescaling maximizes information transmission. *Neuron* 26:695–702.
- Carandini M, Ferster D. 2000. Membrane potential and firing rate in cat primary visual cortex. *Journal of Neuroscience* 20:470–484.
- Chance FS, Abbott LF, Reyes AD. 2002. Gain modulation from background synaptic input. *Neuron* 35:773–782.
- Dan Y, Atick JJ, Read RC. 1996. Efficient coding of natural scenes in the lateral geniculate nucleus: Experimental test of a computational model. *Journal of Neuroscience* 16:3351–3362.
- Dean I, Harper NS, McAlpine D. 2005. Neural population coding of sound level adapts to stimulus statistics. *Nature Neuroscience* 8:1684–1689.
- DeWeese M. 1996. Optimization principles for the neural code. *Network: Computation in Neural Systems* 7:325–331.
- Fairhall AL, Lewen GD, Bialek W, de Ruyter van Steveninck RR. 2001. Efficiency and ambiguity in an adaptive neural code. *Nature* 412:787–792.
- Feller W. 1967. An introduction to probability theory and its applications. New York: Wiley.
- Hyvarinen A, Karhunen J, Oja E. 2001. Independent component analysis. New York: Wiley.
- Ingham NJ, McAlpine D. 2005. GABAergic inhibition controls neural gain in inferior colliculus neurons sensitive to interaural time differences. *Journal of Neuroscience* 25:6187–6198.
- Johnson NL, Kotz S, Balakrishnan N. 1995. Continuous univariate distributions. New York: Wiley.
- Kadunc DC, Vaughan JW, Wallace MT, Benedek G, Stein BE. 1997. Mechanisms of within- and cross-modality suppression in the superior colliculus. *Journal Neurophysiology* 78:2834–2847.
- Kim KJ, Rieke F. 2003. Slow Na⁺ inactivation and variance adaptation in salamander retinal ganglion cells. *Journal of Neuroscience* 23:1506–1516.
- Kvale MN, Schreiner CE. 2004. Short-term adaptation of auditory receptive fields to dynamic stimuli. *Journal of Neurophysiology* 91:604–612.
- Laughlin SB. 1981. A simple coding procedure enhances a neuron's information capacity. *Zeitschrift fur Naturforschung. C* 36:910–912.
- Lesica NA, Jin J, Weng C, Yeh C-I, Butts DA, Stanley GB, Alonso J-M. 2007. Adaptation to stimulus contrast and correlations during natural visual stimulation. *Neuron* 55:479–491.
- Maravall M, Petersen RS, Fairhall AL, Arabzadeh E, Diamond ME. 2007. Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biology* 5:0323–0334.
- Meister M, Berry MJ. 1999. The neural code of the retina. *Neuron* 22:435–450.
- Meredith MA, Stein BE. 1986. Visual, auditory and somatosensory convergence on cells in superior colliculus results in multisensory integration. *Journal of Neurophysiology* 56:640–662.
- Pena JL, Konishi M. 2002. From postsynaptic potentials to spikes in the genesis of auditory spatial receptive fields. *Journal of Neuroscience* 22:5652–5658.
- Rieke F. 2001. Temporal contrast adaptation in salamander bipolar cells. *Journal of Neuroscience* 21:9445–9454.
- Ringach DL, Malone BJ. 2007. The operating point of the cortex: Neurons as large deviation detectors. *Journal of Neuroscience* 27:7673–7683.
- Sanchez-Vives MV, Nowak LG, McCormick DA. 2000. Membrane mechanisms underlying contrast adaptation in cat area 17 *in vivo*. *Journal of Neuroscience* 20:4267–4285.
- Sharpee TO, Sugihara H, Kurgansky AV, Rebrik SP, Stryker MP, Miller KD. 2006. Adaptive filtering enhances information transmission in visual cortex. *Nature* 439:936–942.
- Shu Y, Hasenstaub A, Badoual M, Bal T, McCormick DA. 2003. Barrages of synaptic activity control the gain and sensitivity of cortical neurons. *Journal of Neuroscience* 23:10388–10401.
- Simoncelli EP, Olshausen BA. 2001. Natural image statistics and neural representation. *Annual Review of Neuroscience* 24:1193–1216.
- Smirnakis SM, Berry MJ, Warland DK, Bialek W, Meister M. 1997. Adaptation of retinal processing to image contrast and spatial scale. *Nature* 386:69–73.
- van Haterren JH. 1992. A theory maximizing sensory information. *Biological Cybernetics* 68:23–29.
- Wainwright MJ. 1999. Visual adaptation as optimal information transmission. *Vision Research* 39:3960–3974.