

A Bayesian Framework for Sensory Adaptation

Norberto M. Grzywacz

nmg@bmsr.usc.edu

Department of Biomedical Engineering, University of Southern California, Los Angeles, CA, 90089-1451, U.S.A.

Rosario M. Balboa

rbalboa@ski.org

Departamento de Biotecnología, Universidad de Alicante, Apartado de Correos 99, 03080 Alicante, Spain

Adaptation allows biological sensory systems to adjust to variations in the environment and thus to deal better with them. In this article, we propose a general framework of sensory adaptation. The underlying principle of this framework is the setting of internal parameters of the system such that certain prespecified tasks can be performed optimally. Because sensorial inputs vary probabilistically with time and biological mechanisms have noise, the tasks could be performed incorrectly. We postulate that the goal of adaptation is to minimize the number of task errors. This minimization requires prior knowledge of the environment and of the limitations of the mechanisms processing the information. Because these processes are probabilistic, we formulate the minimization with a Bayesian approach. Application of this Bayesian framework to the retina is successful in accounting for a host of experimental findings.

1 Introduction ---

One of the most important properties of biological sensory systems is adaptation. This property allows these systems to adjust to variations in the environment and thus to deal better with it. Adaptation has been studied in many biological sensory systems (Thorson & Biederman-Thorson, 1974; Laughlin, 1989). Are there general principles that guide biological adaptation? A host of theoretical principles has been proposed in the literature, and many of them have been applied to the visual system (Srinivasan, Laughlin, & Dubs, 1982; Atick & Redlich, 1992; Field, 1994). Recently, we showed that many of these principles cannot account for spatial adaptation in the retina (Balboa & Grzywacz, 2000a). We then proposed that the principle underlying this type of adaptation is the setting of internal parameters of the retina such that it can perform optimally certain prespecified visual tasks (Balboa & Grzywacz, 2000b). Key to this proposal was that the visual infor-

mation is probabilistic (Yuille & Bülthoff, 1996; Knill, Kersten, Mamassian, 1996; Knill, Kersten, & Yuille, 1996) and retinal mechanisms noisy (Rush-ton, 1961; Fuortes & Yeadle, 1964; Baylor, Lamb, & Yau, 1979). Hence, any optimization would have to follow Bayesian probability theory (Berger, 1985). For this purpose, the visual system would have to store knowledge about the statistics of the visual environment (Carlson, 1978; van Hateren, 1992; Field, 1994; Zhu & Mumford, 1997; Balboa & Grzywacz, 2000b, 2000c; Balboa, Tyler, & Grzywacz, 2001; Ruderman & Bialek, 1994), and about its own mechanisms and neural limitations.

In this article, we extend these principles to formalize a general framework of sensory adaptation. This framework begins with the Bayesian formulas but encodes tasks and errors through terms akin to the loss function (Berger, 1985), which quantifies the cost of making certain sensory decisions. As an example, we will show how to apply the framework to the spatial adaptation of retinal horizontal cells. The framework and some of the results appeared previously in abstract form (Balboa & Grzywacz, 1999).

2 A Bayesian Framework of Adaptation

2.1 Rationale. The basic architecture of the new framework of adaptation comprises three stages (see Figure 1). The first, which we call the *stage of preprocessing*, transforms the sensory input into the internal language of the system. The second stage, *task coding*, transforms the output of the stage of preprocessing to a code that the system can use directly to perform its tasks.¹ We propose that the tasks are the estimation of desired attributes from the world, which we call the *task-input attributes*. This estimation would be performed by specially implemented recovery functions (task-coding functions). Finally, the third stage estimates the typical error in performing the tasks, that is, the discrepancy between the task-input attributes and their estimates from the task-coding functions. The key phrase here is *typical error*. Because the system does not know what the actual environmental attributes are and can only estimate them, it cannot know what the actual error is at every instant. Hence, the best the system can aim at is a statistically typical estimation of error for the possible ensemble of inputs. Such an estimate is possible only if one assumes that the environment is not random, which means that the statistics of the input to the system at a given time has some consistency with the statistics of the input at subsequent times. For instance, a forest at night produces mostly dark images and thus defines a temporally consistent environment for the visual system. In turn, the same forest during the day produces bright images and represents a different, consistent environment from the one defined at night.

¹ We postulate a stage of preprocessing separate from a stage of task coding, because the output of the stage of preprocessing can be used for multiple tasks. For example, the output of retinal bipolar cells can carry contrast information, which could be useful for motion-, shape-, and color-based tasks.

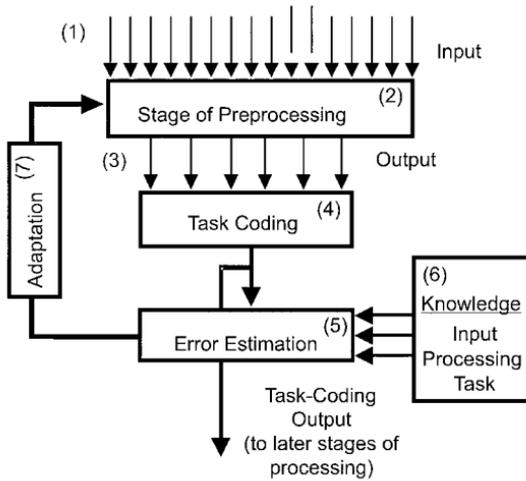


Figure 1: Schematic of the new framework of adaptation. A multivariate input impinges on the system. The input is first processed by the stage of preprocessing, whose output is then fed to a task-coding apparatus. The goal of this apparatus is prepare the code for the system to achieve the desired task. The apparatus has two outputs: one to fulfill the task and another to an error estimation center. This center integrates the incoming information to evaluate the environment and then selects the appropriate prior knowledge from the world to use. In addition, the error estimation center uses knowledge about the stage of preprocessing and the tasks to be performed. With this knowledge, the system estimates the mean task-performance error, not just for the current input but for all possible incoming inputs. The system then sends an adaptation signal to the stage of processing so that this stage modifies its internal parameters to minimize the error.

The principle underlying our framework for adaptation is the setting of internal parameters to adjust the system to the environment such that the system can perform optimally certain prespecified tasks (Balboa & Grzywacz, 2000a, 2000b). This optimization means that the goal of the system is to minimize the expected number of errors in the task-coding process, adjusting the free parameters of adaptation at the stage of preprocessing such that the estimated error is as low as possible. Therefore, optimality implies that the system has knowledge about the environment and about the mechanisms and limitations of the system itself (for instance, noisiness, sluggishness, and small number of information channels) (see Figure 1). These kinds of knowledge could come from evolution, development, or learning.²

² Although people often refer to these three processes as forms of adaptation, we are not addressing them in this article. We mention them here only as a means through which biological systems can attain knowledge useful for sensory adaptation.

2.2 Framework. This section expresses mathematically each element of the framework illustrated in Figure 1 and emphasized in section 2.1. Let \bar{I} be the input in Figure 1 (process labeled 1 in the figure) and $H_1(\bar{I}), H_2(\bar{I}), \dots, H_N(\bar{I})$ the relevant task-input attributes to extract from \bar{I} . For instance, in section 3, these attributes will be things like contrasts and positions of occluding borders in a visual image. The output (the process labeled 3) of the stage of processing (process labeled 2), \bar{O} , is transformed by the task-coding stage (the process labeled 4) to $R_{H_1}(\bar{O}), R_{H_2}(\bar{O}), \dots, R_{H_N}(\bar{O})$. These functions are estimates of the values of the task-input attributes. To evaluate the error (the process labeled 5), E_i , in each of these estimates, one must measure the discrepancy between $H_i(\bar{I})$ and $R_{H_i}(\bar{O})$. The system cannot know exactly what this discrepancy is, since it does not have access to the input, only to its estimates. However, the system can estimate the expected amount of error. To do so, the system must have previous knowledge (the process labeled 6) about the probable \bar{I} , about the task, and about the stage of processing that produces \bar{O} . Adaptation (the process labeled 7) would be the setting of the stage-of-preprocessing parameters (A) such that the error is minimal over the ensemble of inputs.³ The error to be minimized is

$$E(A) = \sum_{i=1}^N E_i(A). \quad (2.1)$$

The error, $E_i(A)$, can be defined using statistical decision theory (Berger, 1985). We begin by defining a loss function $\mathcal{L}(\bar{I}, \bar{R}_{H_i})$ (where $\bar{R}_{H_i} = R_{H_i}(\bar{O})$), which measures the cost of deciding that the i th attribute is \bar{R}_{H_i} given that the input is \bar{I} . Using the loss function, statistical decision theory defines the Bayesian expected loss as

$$l_A(\bar{R}_{H_i}) = \int_{\bar{I}} P_A(\bar{I}|\bar{R}_{H_i}) \mathcal{L}(\bar{I}, \bar{R}_{H_i}), \quad (2.2)$$

where $P_A(\bar{I}|\bar{R}_{H_i})$ is the conditional probability, with the parameters of adaptation set to A , that the correct interpretation of the input is \bar{I} given that the response of the system is \bar{R}_{H_i} . The Bayesian expected loss is the mean loss

³ The parameters of adaptation can be thought of as a representation of the environment. Here, we choose to define the framework through the parameters of adaptation instead of those of the environment (for instance, time of the day) for two reasons. First, our framework is a mathematical formulation for the role of adaptation. Second, there may not be a one-to-one correspondence between the parameters of the environment and those of adaptation. If we had chosen to parameterize our framework through the environment, then the framework would be similar to the Bayesian process of prior selection (Berger, 1985). In prior selection, one chooses the correct prior distribution to use before performing a task. One defines the various possible priors by special hyperparameters.

of the system given that response. We define the error as the mean Bayesian expected loss over all possible responses, that is,

$$E_i(A) = \int_{\vec{R}_{H_i}} P_A(\vec{R}_{H_i}) l_A(\vec{R}_{H_i}). \quad (2.3)$$

That we minimize the sum of these errors (see equation 2.1) is the same as using conditional Bayes' principle of statistical decision theory (Berger, 1985). In simple words, we choose the "action" that minimizes the loss.

From standard Bayesian analysis, one can obtain a particularly useful form of the error in equation 2.3 by using Bayes' theorem,⁴ that is,

$$P_A(\vec{I}|\vec{R}_{H_i}) = \frac{P_A(\vec{R}_{H_i}|\vec{I})P(\vec{I})}{P_A(\vec{R}_{H_i})}. \quad (2.4)$$

By substituting equation 2.4 for $P_A(\vec{I}|\vec{R}_{H_i})$ in equation 2.2 and then plugging the result in equation 2.3, one gets

$$E_i(A) = \int_{\vec{I}} \int_{\vec{R}_{H_i}} P_A(\vec{R}_{H_i}|\vec{I})P(\vec{I})\mathcal{L}(\vec{I}, \vec{R}_{H_i}). \quad (2.5)$$

One can give intuitive and practical interpretations of the probability terms in this equation. The first term, the likelihood function $P_A(\vec{R}_{H_i}|\vec{I})$, embodies the knowledge about the sensory mechanisms; the second, $P(\vec{I})$, is the prior knowledge about the input. In other words, $P_A(\vec{R}_{H_i}|\vec{I})$ tells how the system responds when the stimulus is \vec{I} . Because the response is composed of the stage of preprocessing and the task-coding functions, it is sometimes useful to unfold equation 2.5 as

$$E_i(A) = \int_{\vec{I}} \int_{\vec{O}} \int_{\vec{R}_{H_i}} P_A(\vec{R}_{H_i}|\vec{O})P_A(\vec{O}|\vec{I})P(\vec{I})\mathcal{L}(\vec{I}, \vec{R}_{H_i}).$$

In this equation, we divided the knowledge of the sensory process into $P_A(\vec{R}_{H_i}|\vec{O})$ and $P_A(\vec{O}|\vec{I})$. The former reflects the task, while the latter reflects the computations performed in the stage of preprocessing. In this article, we assume for simplicity that the task-coding functions have no noise and thus use the simpler form in equation 2.5.

⁴ We write $P(\vec{I})$ in this equation as if this probability function does not depend on the parameters of adaptation. However, as discussed in note 3, these parameters can be thought of as a representation of the environment. Hence, it would have been more correct to insert the subindex A in $P(\vec{I})$ in equation 2.4. The reason that we avoid doing that is that for some readers, it might be confusing having a notation suggesting that the distribution of images depends on adaptation. We preferred not to use the subindex A and to think of $P(\vec{I})$ as a prior distribution of the current environment (see section 2.1).

3 An Application of the Framework: Spatial Retinal Adaptation

3.1 Some Task-Input Attributes Proposed for the Retina. We previously proposed that retinal horizontal cells are part of a system to extract from images the position of occluding borders, contrast at occluding borders, and intensity away from them (Balboa & Grzywacz, 2000b). To quantify these image attributes (the functions H_i of equation 2.5) we begin by defining contrast as $|\nabla I(\vec{r})/I(\vec{r})|$ (see Balboa & Grzywacz, 2000b, 2000c, for a justification).⁵ From this definition, the functions H_i may be

$$\begin{aligned}
 H_1(I(\vec{r})) &= P_O \left(\left| \frac{\nabla I(\vec{r})}{I(\vec{r})} \right| \right) \\
 H_2(I(\vec{r})) &= \left| \frac{\nabla I(\vec{r})}{I(\vec{r})} \right| \\
 H_3^{n+1}(I(\vec{r})) + H_3(I(\vec{r})) &= I(\vec{r}) G(\bar{I}), \tag{3.1}
 \end{aligned}$$

where $0 \leq P_O(|\nabla I(\vec{r})/I(\vec{r})|) \leq 1$ is the probability that a point with the given contrast has an occluding border and $0 < G(\bar{I}) \leq 1$ is a prespecified, positive, decreasing function of the mean intensity, \bar{I} . Consequently, H_1 , H_2 , and H_3 quantify the likelihood of a border at position \vec{r} , contrast, and intensity, respectively. In this application, H_1 is made to depend only on $\frac{|\nabla I(\vec{r})}{I(\vec{r})}$ for simplicity; one may introduce other variables in more complex models. The quantification of intensity by H_3 is indirect, but we can justify this quantification from independent results. The function G transforms the right-hand side of the definition of H_3 (see equation 3.1) to a compressed function of I . The polynomial form of H_3 causes it to be a fractional power law of the right-hand side. Hence, H_3 is a compressed function of intensity, allowing the retina to deal with the wide range of natural intensities (Shapley, 1997; Smirnakis, Berry, Warland, Bialek, & Meister, 1997; Adorjan, Piepenbrock, & Obermayer, 1999).

3.2 Prior Knowledge of Images. Many literature studies address the prior knowledge on the distribution of natural images (Carlson, 1978; van Hateren, 1992; Field, 1994; Zhu & Mumford, 1997; Ruderman & Bialek, 1994; Balboa & Grzywacz, 2000b, 2000c). These studies focus on certain statistical moments of the images and thus cannot tell what $P(\bar{I})$ is. What these

⁵ For notational simplicity, we assume that the input is mapped to the output of the outer plexiform layer (OPL) in a one-to-one manner. (In this layer, the photoreceptor synapses serve as input, the horizontal cells serve as interneurons, and the bipolar cells serve as output.) Thus, we will use \vec{r} to indicate position in both the input and the output. Alternatively, one could use subscripts to indicate the discrete positions of receptive-field centers (Srinivasan et al., 1982; Balboa & Grzywacz, 2000b).

studies reveal is that the moments obey certain regularities, and thus not all possible images occur naturally. The moments of interest here are those specified by equation 3.1, that is, the task-input attributes. For simplicity, the approach of this article is to assume first that $P(\vec{I})$ is constant over the range of naturally occurring images and zero outside it. Then we use a sample of natural images to estimate the occurrence of the attributes in images. (By the last assumption, the images in the sample have equal probability among themselves and the same probability as any image outside the sample.)

3.3 Knowledge of Neural Processes. To specify the knowledge of the neural processing, one has to describe what is known about the OPL (see note 5), define the task-coding functions (see Figure 1), and then calculate $P_A(\vec{R}_{H_i}|\vec{I})$. Unfortunately, the outputs of the task-coding functions have complex, nonlinear dependencies on \vec{I} . Therefore, it is hard to provide complete analytical formulas for $P_A(\vec{R}_{H_i}|\vec{I})$. In Balboa and Grzywacz (2000b), we diminish the need for such formulas through some reasonable approximations.

A model for the OPL is what we use for the stage-of-preprocessing box of Figure 1. This model is based on a wealth of literature (for a review, see Dowling, 1987) and was presented elsewhere (Balboa & Grzywacz, 2000b). The basic equation of the model is

$$B(\vec{r}) = \frac{T(\vec{r})}{1 + (h_A(\vec{r}) * B(\vec{r}))^n}, \quad (3.2)$$

where $T(\vec{r})$ and $B(\vec{r})$ are the intensity-dependent input and output of the synapse of the photoreceptor, $*$ stands for convolution, $h_A(\vec{r})$ is the lateral-inhibition filter (which can change depending on the state A of adaptation), and $n \geq 1$ is a constant. The variable $B(\vec{r})$ can be pre- or postsynaptic, since we assume here for simplicity that the photoreceptor-bipolar synapse is linear. It is also postulated that $h_A(\vec{r})$ is positive ($h_A(\vec{r}) > 0$), isotropic (if $|\vec{r}| = |\vec{r}'|$, then $h_A(\vec{r}) = h_A(\vec{r}')$), and normalized ($\int h_A(\vec{r}) = 1$). The normalization assumption is made for simplicity, because it was shown that the results are invariant with the integral of h_A (Balboa & Grzywacz, 2000b).

One can think of $T(\vec{r})$ as the current generated by phototransduction. Hence, a simple way to introduce the phototransduction adaptation is to define a gain function $G(\vec{I})$ such that $T(\vec{r}) = I(\vec{r})G(\vec{I})$ (see equation 3.1). In this case, equation 3.2 becomes

$$B(\vec{r}) = \frac{I(\vec{r}) G(\vec{I})}{1 + (h_A(\vec{r}) * B(\vec{r}))^n}. \quad (3.3)$$

Defined as in this equation, $B(\vec{r})$ provides a straightforward signal for the desired attributes specified in equation 3.1. If one stimulates the retina

with a full-field illumination and disregards noise, then in steady state, equation 3.3 becomes

$$\bar{B} = \frac{\bar{I}G(\bar{I})}{1 + \bar{B}^n}, \quad (3.4)$$

which resembles the H_3 term in equation 3.1. Moreover, mathematical analysis (derivations not shown) shows that if one stimulates the retina with an edge and disregards noise, then in steady state, one gets to a good approximation at the edge:

$$\left| \frac{\nabla B(\vec{r})}{B(\vec{r})} \right| = \left| \frac{\nabla I(\vec{r})}{I(\vec{r})} \right|. \quad (3.5)$$

In other words, luminance contrast is proportional to the contrast in the bipolar cell responses.

3.4 Task-Coding Functions. The role of the task-coding functions is to estimate the task attributes (H_i) from the output of the OPL. By comparing equation 3.4 to the H_3 term of equations 3.1, one sees that B and H_3 have the same dependence on light intensity. Therefore, to recover a compressed version of light intensity from the bipolar signals, all that the task-coding stage has to do is to read them directly. Comparison of equation 3.5 to the H_2 term of equation 3.1 yields a similar conclusion. The task-coding stage can compute the illumination contrast directly from the bipolar signals. To do so, this mechanism only has to compute the gradient of the bipolar cell responses and divide it by the responses themselves. Furthermore, assume that the function P_O from the first term of equation 3.1 is known (Balboa & Grzywacz, 2000c). In this case, the task-coding stage can extract the border-position attribute from the contrast signal in the bipolar cells. For given adaptation settings (A), we can express these conclusions in the following equation:

$$\begin{aligned} R_{H_1}(\vec{r}) &= P_O \left(\left| \frac{\nabla B(\vec{r})}{B(\vec{r})} \right| \right) \\ R_{H_2}(\vec{r}) &= \left| \frac{\nabla B(\vec{r})}{B(\vec{r})} \right| \\ R_{H_3}(\vec{r}) &= B(\vec{r}). \end{aligned} \quad (3.6)$$

3.5 Retinal Loss Functions. We now must specify the cost of the system extracting these image attributes incorrectly. The first attribute (H_1) quantifies the likelihood that there is a border in a given position. A good loss

function for this attribute should penalize missing true occluding borders. (We worry less about false borders being detected, since the central visual system has heuristics to eliminate false borders, such as false borders not giving rise to long, continuous contours; Field, Hayes, & Hess, 1993; Kovacs & Julesz, 1993; Pettet, McKee, & Grzywacz, 1997.) For the contrast attribute (H_2), the loss function should increase as the contrast diverges from truth, especially at points likely to contain a border. In contrast, a good loss function for the third attribute (H_3) should penalize discrepancies of intensity estimation at points unlikely to be occluding borders.

A loss function for the attribute H_1 that meets these conditions is

$$\mathcal{L}(\bar{I}, \bar{R}_{H_1}) = \left(\int_{\bar{r}} H_1(I(\bar{r})) (1 - R_{H_1}(\bar{r})) \times (R_{H_1}(\bar{r}) - H_1(I(\bar{r})))^{2k} \right)^{1/2}, \quad (3.7)$$

where k is an integer. This equation penalizes missing borders. The term $H_1(I(\bar{r}))$ enforces computations on input borders, while the term $(1 - R_{H_1}(\bar{r}))$ matters only if the system missed a border at \bar{r} . Because $0 \leq H_1(I(\bar{r}))$, $R_{H_1}(\bar{r}) \leq 1$ (see equations 3.1 and 3.6), large values of k force the loss function to penalize clear-cut errors. Such errors are $R_{H_1}(\bar{r}) \approx 0$ and $H_1(I(\bar{r})) \approx 1$. In contrast, results like $R_{H_1}(\bar{r}) = H_1(I(\bar{r})) = 0.5$, which are not in error but would contribute to equation 3.7 if $k = 0$ would not do so if $k \gg 1$. In general, this equation works as a counter of the number of errors instead of as a measure of their magnitudes. This equation is similar to the standard 0 – 1 loss function (Berger, 1985).

For the attribute H_2 , a loss function meeting the conditions above is

$$\mathcal{L}(\bar{I}, \bar{R}_{H_2}) = \left(\int_{\bar{r}} H_1(I(\bar{r})) (R_{H_2}(\bar{r}) - H_2(I(\bar{r})))^2 \right)^{1/2}. \quad (3.8)$$

This equation also enforces computations on borders as equation 3.7. The H_2 term in equation 3.8 expresses contrast discrepancies in absolute terms, since the visual system is sensitive to contrast. This equation is similar to a standard type of loss function called the squared-error loss (Berger, 1985) and to the L_2 metric of functional analysis (Riesz & Nagy, 1990).

Finally, a good loss function for the attribute H_3 is

$$\mathcal{L}(\bar{I}, \bar{R}_{H_3}) = \left(\int_{\bar{r}} (1 - H_1(I(\bar{r}))) \left(\frac{R_{H_3}(\bar{r}) - H_3(I(\bar{r}))}{H_3(I(\bar{r}))} \right)^2 \right)^{1/2}, \quad (3.9)$$

which computes errors away from borders (the $1 - H_1(I(\bar{r}))$ term). The H_3 ratio of this equation shows that intensity discrepancies are measured as

a noise-to-signal ratio. Such a ratio allows the expression of the noise in a dimensionless form and is the same type of error measurement that successfully explained spatial adaptation in Balboa and Grzywacz (2000b).

4 Results

Elsewhere (Balboa & Grzywacz, 2000b), we showed that our retinal application of the general framework of adaptation can account correctly for the spatial adaptation in horizontal cells. This application can also explain the division-like mechanism of inhibitory action (see equation 3.3 and Balboa & Grzywacz, 2000b) and linearity at low contrasts (Tranchina, Gordon, Shapley, & Toyoda, 1981). Moreover, the application is consistent with the psychophysical Stevens' power law of intensity perception (Stevens, 1970). This section extends the validity of this application by showing that it can account for three new retinal effects.

The first new retinal effect that this application of the framework can account for is the shrinkage of the extent of lateral inhibition for stimuli that are not spatially homogeneous. Experiments by Reifsnider and Tranchina (1995) show that background contrast reduces the OPL lateral spread of responses to superimposed stimuli (see also Smirnakis et al., 1997). Karamans, Haak, Habraken, and Spekreijse (1996) have neurally realistic computer simulations of the OPL, demonstrating a similar reduction. Figure 2 shows that our model could account for such reductions. Furthermore, this figure illustrates a qualitative prediction of our model: it can account for this reduction when the retina is stimulated by checkerboards. The smaller the squares in the checkerboard are, the less extent the predicted lateral inhibition has. The reason is that with smaller squares, this extent must decrease to prevent as much as possible the straddling of multiple borders by lateral inhibition. If this straddling occurs, the model makes errors in border localization.

Because an important task postulated in the retinal application of the general framework is the measurement of contrast, we argued for a division-like model of lateral inhibition (see section 3.3). Figure 3 shows that a consequence of this model is the disappearance of lateral inhibition at low background intensities. The inhibition emerges only at high intensities, with its threshold decreasing as n (the cooperativity of inhibition) increases. The disappearance of lateral inhibition at low background intensities is well known, having been demonstrated both physiologically (Barlow, Fitzhugh, & Kuffler, 1957; Bowling, 1980) and psychophysically (Van Ness & Bouman, 1967). Functionally, it is not too hard to understand why horizontal cell inhibition disappears in the model. As intensity falls, horizontal cell responses also do so. This causes the conductance change in the photoreceptor synapse (the B^n term in equation 3.4) to be smaller than the resting conductance (the 1 in that equation). The conductance changes specified by the model also lead to a quantitative prediction about the modulation of the inhibitory strength

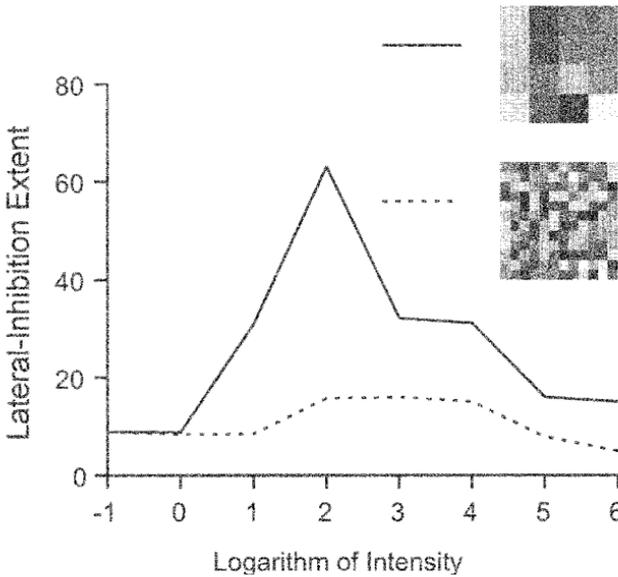


Figure 2: Dependence of lateral inhibition extent on the spatial structure of the image. The stimuli (checkerboard images) are those that appear in the insets. The curves display a bell-shape behavior, with the lateral inhibition extent increasing as the size of the squares in the checkerboard increases.

with intensity. Because the phototransduction adaptation gain ($G(\bar{I})$) can be estimated from the photoreceptor literature (Hamer, 2000; Fain, Matthews, Cornwall, & Koutalos, 2001), equation 3.4 can be solved. With \bar{B} in hand, one can then predict how the inhibitory strength increases with intensity and test dependence experimentally.

The division-like mechanism of lateral inhibition also has implications for the model Mach bands. A property of Mach bands is that there is an asymmetry of positive and negative Mach bands at an intensity border (Fiorentini & Radici, 1957; Ratliff, 1965); the Mach band is larger on the positive side than it is on the negative side. This is a well-known Mach band effect without a good explanation until now. Figure 4 shows that this Mach band asymmetry occurs in our model. This figure also shows that the asymmetry is more prominent as the parameter n increases. From our model, it is not difficult to understand this edge asymmetry. Although the inhibition coming from the positive side of the edge is strong, if the intensity at the negative side is near zero, then the response there will be near zero (see the $I(\vec{r})$ term in equation 3.3). In contrast, the intensity on the positive side of the edge is high, and thus the effect of lateral inhibition can work there, producing a large Mach band.

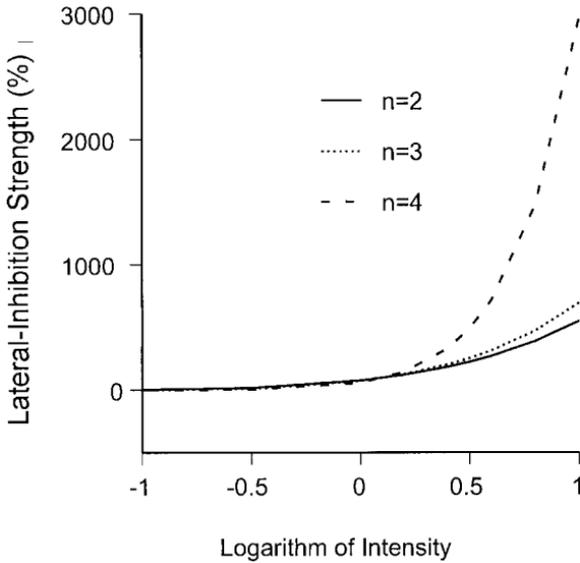


Figure 3: Lateral inhibition strength as a function of intensity. These curves are parametric on n (see equation 3.2) and assume $G(\bar{I}) = 1$, and their inhibitory strength is defined as $B^n / (1 + B^n)$. As the intensity falls, the inhibition disappears, as indicated by its strength going to zero. The threshold for the emergence of inhibition as a function of intensity falls as n increases.

5 Discussion

5.1 Adaptation as a Bayesian Process. Ours is not the first Bayesian framework of sensory processing. Other theories have proposed Bayesian mechanisms for perception (for instance, Yuille & Bülthoff, 1996; Knill, Kersten, & Mamassian, 1996; Knill, Kersten, & Yuille, 1996). Those theories begin with the output of neurons or filters, and then ask how to interpret the environment most correctly. What is new in our framework is not the use of the Bayesian framework but the computation of the best internal state for the system. In other words, the new framework is concerned not with interpreting the current input but rather with setting the parameters of the system such that future interpretations are as correct as possible. Another novel aspect of the new framework is the emphasis on the loss function. Past theories tended to make their perceptual decision based on the maximal probability of interpreting the input correctly given some internal data (Yuille & Bülthoff, 1996; Knill, Kersten, & Mamassian, 1996; Knill, Kersten, & Yuille, 1996). The use of loss functions means that the most probable interpretation of the input may not be sufficient. Sometimes errors made by not picking less probable interpretations are more costly.

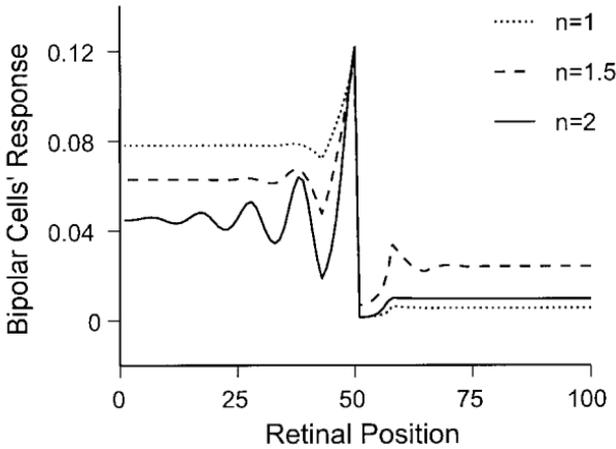


Figure 4: Mach band asymmetries in the retinal application of the framework. The stimulus was an edge at position 50 (arbitrary units) and of intensities 1.8 and 0.1 at positions lower and higher than 50, respectively. The responses of the bipolar cells were simulated parametric on n (see equation 2.7), with $G(\bar{I}) = 1$, and the filter h being flat and with a radius of 15. The Mach band was larger at the high-intensity side of the edge than at the low-intensity side. This Mach band asymmetry became more prominent as n increased.

One positive aspect of using a Bayesian framework is that it forces us to state explicitly the assumptions of our models. In particular, one must state the knowledge that the system has about the environment and about its mechanisms, and state the tasks that the system must perform. How does one go about specifying tasks? We believe that tasks are chosen by both what the systems need to do and can do. For instance, although it would be lovely for a barnacle to perform recognition tasks, this animal has only ten photoreceptors. (It has four photoreceptors in its only median eye—Hayashi, Moore, & Stuart, 1985; Oland & Stuart, 1986—and three photoreceptors in each of its two lateral eyes—Krebs & Schaten, 1976; Oland & Stuart, 1986.) Nevertheless, its visual system is sufficiently good to allow the animal to hide in its shell when a predator approaches. Therefore, one should not assume automatically that the goal of the early sensory system is to maximize transmission of information (Haft & van Hemmen, 1998; Wainwright, 1999). Furthermore, the tasks that a system performs might be intimately coupled to its hardware, against what was argued by Marr (1982). The system must perform computation, and its hardware makes certain types of computations easier than others.

5.2 Limitations of the Framework. Elsewhere, we discuss the limitations of our retinal application of the framework (Balboa & Grzywacz, 2000b); here, we focus on problems related to the framework itself.

Perhaps the most serious problem with the framework is that to apply equation 2.5, one must have $P_A(\bar{R}_{H_i}|\bar{I})$, $P(\bar{I})$, and $\mathcal{L}(\bar{I}, \bar{R}_{H_i})$, that is, the prior knowledge and the loss function. We assume that (but do not state how) the system would attain these things through evolution, development, and learning. A complete framework would have to specify the algorithms and mechanisms to attain the prior knowledge and the loss function. An even more serious problem is how to apply the framework when the environment changes (see note 4).

5.3 The Importance of Understanding Tasks in Neurobiology. We have provided one example of how to apply the new framework for sensory adaptation—namely, to spatial adaptation in horizontal cells. Besides specifying the neural processes, the example had to formalize the retinal tasks. As we showed here and elsewhere (Balboa & Grzywacz, 2000a, 2000b), the choice of the tasks can have a large impact on the behavior of the system. For instance, the assumed horizontal cell functions were border localization, contrast estimation, and intensity estimation. One could have argued that all of this may just be a consequence of contrast sensitivity (Shapley, 1997). However, requiring maximization of contrast sensitivity has different consequences from requiring optimization of border location. Responding to any edge is not the same as encoding its position with precision. The requirement of encoding position correctly is essential to obtain the correct behavior for the spatial adaptation of horizontal cells (Balboa & Grzywacz, 2000b). Hence, we argue that to understand the behavior of ensembles of nerve cells, it is not sufficient to figure out their neurobiological processes. One must also carefully study their information processing tasks.

Acknowledgments

We thank Alan Yuille and Christopher Tyler for critical comments on the manuscript and Joaquín de Juan for support and many discussions in the early phases of this project. We also thank the Smith-Kettlewell Institute, where we performed a large portion of the work. This work was supported by National Eye Institute Grants EY08921 and EY11170 to N.M.G.

References

- Adorjan, P., Piepenbrock, C., & Obermayer, K. (1999). Contrast adaptation and infomax in visual cortical neurons. *Rev. Neurosci.*, *10*(3–4), 181–200.
- Atick, J. J., & Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Comp.*, *4*, 196–210.
- Balboa, R. M., & Grzywacz, N. M. (1999). Biological evidence for an ecological-based theory of early retinal lateral inhibition. *Invest. Ophthalmol. Vis. Sci.*, *40*, S386.
- Balboa, R. M., & Grzywacz, N. M. (2000a). The role of early retinal lateral inhibition: More than maximizing luminance information. *Visual Neurosci.*, *17*, 77–89.

- Balboa, R. M., & Grzywacz, N. M. (2000b). The minimal-local asperity hypothesis of early retinal lateral inhibition. *Neural Comp.*, *12*, 1485–1517.
- Balboa, R. M., & Grzywacz, N. M. (2000c). Occlusions and their relationship with the distribution of contrasts in natural images. *Vision Res.*, *40*, 2661–2669.
- Balboa, R. M., Tyler, C. W., & Grzywacz, N. M. (2001). Occlusions contribute to scaling in natural images. *Vision Res.*, *41*, 955–964.
- Barlow, H. B., Fitzhugh, R., & Kuffler, S. W. (1957). Change of organisation in the receptive fields of the cat's retina during dark adaptation. *J. Physiol.*, *137*, 338–354.
- Baylor, D. A., Lamb, T. D., & Yau, K.-W. (1979). Responses of retinal rods to single photons. *J. Physiol.*, *288*, 613–634.
- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. New York: Springer-Verlag.
- Bowling, D. B. (1980). Light responses of ganglion cells in the retina of the turtle. *J. Physiol.*, *209*, 173–196.
- Carlson, C. R. (1978). Thresholds for perceived image sharpness. *Phot. Sci. and Eng.*, *22*, 69–71.
- Dowling, J. E. (1987). *The retina: An approachable part of the brain*. Cambridge, MA: Belknap Press, Harvard University Press.
- Fain, G. L., Matthews, H. R., Cornwall, M. C., & Koutalos, Y. (2001). Adaptation in vertebrate photoreceptors. *Physiol. Rev.*, *81*, 117–151.
- Field, D. J. (1994). What is the goal of sensory coding? *Neural Comp.*, *6*, 559–601.
- Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local “association field.” *Vision Res.*, *33*(2), 173–193.
- Fiorentini, A., & Radici, T. (1957). Binocular measurements of brightness on a field presenting a luminance gradient. *Atti. Fond. Giorgio Ronchi*, *12*, 453–461.
- Fuortes, M. G. F., & Yeadle, S. (1964). Probability of occurrence of discrete potential waves in the eye of the *Limulus*. *J. Physiol.*, *47*, 443–463.
- Haft, M., & van Hemmen, J. L. (1998). Theory and implementation of infomax filters for the retina. *Network*, *9*(1), 39–71.
- Hamer, R. D. (2000). Analysis of CA^{++} -dependent gain changes in PDE activation in vertebrate rod phototransduction. *Mol. Vis.*, *6*, 265–286.
- Hateren, J. H. van (1992). Theoretical predictions of spatiotemporal receptive fields of fly LMC's, and experimental validation. *J. Comp. Physiol. A*, *171*, 157–170.
- Hayashi, J. H., Moore, J. W., & Stuart, A. E. (1985). Adaptation in the input-output relation of the synapse made by the barnacle's photoreceptor. *J. Physiol.*, *368*, 179–195.
- Kamermans, M., Haak, J., Habraken, J. B., & Spekreijse, H. (1996). The size of the horizontal cell receptive fields adapts to the stimulus in the light adapted goldfish retina. *Vision Res.*, *36*, 4105–4119.
- Knill, D. C., Kersten, D., & Mamassian, P. (1996). Implications of a Bayesian formulation of visual information for processing for psychophysics. In D. C. Knill & R. Whitman (Eds.), *Perception as Bayesian inference* (pp. 239–286). Cambridge: Cambridge University Press.

- Knill, D. C., Kersten, D., & Yuille, A. (1996). Introduction: A Bayesian formulation of visual perception. In D. C. Knill & R. Whitman (Eds.), *Perception as Bayesian inference* (pp. 1–21). Cambridge: Cambridge University Press.
- Kovacs, I., & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in figure-ground segmentation. *Proc. Natl. Acad. Sci. USA*, *90*, 7495–7497.
- Krebs, W., & Schaten, B. (1976). The lateral photoreceptor of the barnacle, *Balanus eburneus*: quantitative morphology and fine structure. *Cell Tissue Res.*, *168*, 193–207.
- Laughlin, S. B. (1989). The role of sensory adaptation in the retina. *J. Exp. Biol.*, *146*, 39–62.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Oland, L. A., & Stuart, A. E. (1986). Pattern of convergence of the photoreceptors of the barnacle's three ocelli onto second-order cells. *J. Neurophysiol.*, *55*, 882–895.
- Pettet, M. W., McKee, S. P., & Grzywacz, N. M. (1997). Constraints of long range interactions mediating contour detection. *Vision Res.*, *38*, 865–879.
- Ratliff, F. (1965). *Mach bands: Quantitative studies on neural networks in the retina*. San Francisco: Holden-Day.
- Reifsnider, E. S., & Tranchina, D. (1995). Background contrast modulates kinetics and lateral spread of responses to superimposed stimuli in outer retina. *Vis. Neurosci.*, *12*, 1105–1126.
- Riesz, F., & Nagy, B. S. Z. (1990). *Functional analysis*. New York: Dover.
- Ruderman, D. L., & Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Phys. Rev. Letter*, *73*, 814–817.
- Rushton, W. A. H. (1961). *The intensity factor in vision*. In W. D. McElroy & H. B. Glass (Eds.), *Light and life* (pp. 706–722). Baltimore, MD: Johns Hopkins University Press.
- Shapley, R. (1997). Retinal physiology: Adapting to the changing scene. *Curr. Biol.*, *7*(7), R421–R423.
- Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W., & Meister, M. (1997). Adaptation of retinal processing to image contrast and spatial scale. *Nature*, *386*(6620), 69–73.
- Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: A fresh view of inhibition in the retina. *Proc. R. Soc. Lond. B*, *216*, 427–459.
- Stevens, S. S. (1970). Neural events and the psychophysical law. *Science*, *170*, 1043–1050.
- Thorson, J., & Biederman-Thorson, M. (1974). Distributed relaxation processes in sensory adaptation. *Science*, *183*, 161–172.
- Tranchina, D., Gordon, J., Shapley, R., & Toyoda, J. (1981). Linear information processing in the retina: A study of horizontal cell responses. *Proc. Natl. Acad. Sci. USA*, *78*, 6540–6542.
- Van Ness, F. L., & Bouman M. A. (1967). Spatial modulation transfer in the human eye. *J. Opt. Soc. Am.*, *57*, 401–406.
- Wainwright, M. J. (1999). Visual adaptation as optimal information transmission. *Vision Res.*, *39*, 3960–3974.

- Yuille, A., & Bülthoff, H. H. (1996). Bayesian decision theory and psychophysics. In D. C. Knill & R. Whitman (Eds.), *Perception as Bayesian inference* (pp. 123–161). Cambridge: Cambridge University Press.
- Zhu, S. C., & Mumford, D. (1997). Prior learning and Gibbs reaction-diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, *19*, 1236–1250.

Received November 13, 2000; accepted May 22, 2001.