

Research report

Multimodality in the superior colliculus: an information theoretic analysis

Paul Patton^a, Kamel Belkacem-Boussaid^a, Thomas J. Anastasio^{a,b,*}

^aBeckman Institute, University of Illinois at Urbana/Champaign, Urbana, IL, USA

^bDepartment of Molecular and Integrative Physiology, University of Illinois at Urbana/Champaign, Urbana, IL, USA

Accepted 2 May 2001

Abstract

The deep superior colliculus (DSC) integrates multisensory input and triggers an orienting movement toward the source of stimulation (target). It would seem reasonable to suppose that input of an additional modality should always increase the amount of information received by a DSC neuron concerning a target. However, of all DSC neurons studied, only about one half in the cat and one-quarter in the monkey were multimodal. The rest received only unimodal input. Multimodal DSC neurons show the properties of multisensory enhancement, in which the neural response to an input of one modality is augmented by input of another modality, and of inverse effectiveness, in which weaker unimodal responses produce a higher percentage enhancement. Previously, we demonstrated that these properties are consistent with the hypothesis that DSC neurons use Bayes' rule to compute the posterior probability that a target is present given their stochastic sensory inputs. Here we use an information theoretic analysis of our Bayesian model to show that input of an additional modality may indeed increase target information, but only if input received from the initial modality does not completely reduce uncertainty concerning the presence of a target. Unimodal DSC neurons may be those whose unimodal input fully reduces target uncertainty and therefore have no need for input of another modality. © 2002 Elsevier Science B.V. All rights reserved.

Theme: Sensory systems

Topic: Multisensory

Keywords: Multisensory integration; Superior colliculus; Bayes' rule; Information theory; Probability theory

1. Introduction

The deep layers of the superior colliculus (DSC) are a major site of multisensory integration in the vertebrate brain [28]. The DSC receives inputs from the visual, auditory, and somatosensory systems, which are derived from cortical and subcortical brain regions [8,11,26]. Its neurons are organized topographically according to the location in sensory space of their receptive fields [13,18,20]. The DSC integrates converging inputs and triggers orienting movements of the eyes, head, body, or pinnae of the ears towards the source of stimulation (target) [10,19,24,26,35]. It would seem reasonable to

suppose that the integration of inputs of multiple sensory modalities should always increase the gain in information concerning a target. Although some individual DSC neurons receive input of two or three sensory modalities, nearly half in the cat and nearly three-quarters in the monkey receive input of only one modality [15,17,19,33]. It is not clear why many DSC neurons are unimodal despite the availability of multimodal input. The purpose of this paper is to describe an information-theoretic analysis that provides a possible explanation for this puzzling experimental finding.

Multisensory DSC neurons show multisensory enhancement (MSE), in which the response to a stimulus of one modality is augmented by a stimulus of another modality [12,14–17,22,31,34]. They also exhibit the property of inverse effectiveness, in which smaller unimodal responses are associated with higher percentages of MSE [17,32]. We have proposed a functional interpretation of MSE based on

*Corresponding author. Beckman Institute, 405 North Mathews Avenue, Urbana, IL 61801, USA. Tel.: +1-217-2442-895; fax: +1-217-2445-180.

E-mail address: tja@uiuc.edu (T.J. Anastasio).

the notion that DSC neurons use Bayes' rule to compute the posterior probability that a target is present given stochastic sensory input [1]. We showed that the posterior probability of a target, given inputs of two modalities, behaved in a manner analogous to that described for MSE, including the phenomenon of inverse effectiveness. Here we will extend our analysis of multisensory integration using information theory [6]. We will quantify the information gain due to various types of random inputs, and show why we might expect to observe multisensory integration for some inputs but not for others.

2. Information, uncertainty, and multisensory integration

We postulate that a DSC neuron uses Bayes' rule to compute the posterior probability that a target is present in its receptive field given sensory inputs from the environment. The target is modeled as random variable T having two states: $T=1$ when a target is present in the receptive field of the DSC neuron, and $T=0$ when a target is absent. The prior probability of a target $P(T=t)$ is the target probability unconditioned by sensory input. Our Bayesian model assumes that the nervous system can represent $P(T=t)$, which it might learn through experience. At any given time, we assume that a target is more probably absent than present in the receptive field of an individual DSC neuron. For our computations, we arbitrarily assign $P(T=0) = 0.9$ and $P(T=1) = 0.1$. These values were chosen to illustrate general principles that apply for a wide range of plausible parameters. Without sensory input, an organism could be completely certain about the target only if it is always absent or always present. There is uncertainty associated with any other prior distribution. We can use information theory to quantify this uncertainty.

Consider any random variable X having possible values x_i ($i = 1, 2, \dots, n$). The information content of the event $X = x_i$ is defined as $I(x_i) = -\log_2 P(X = x_i)$. Information is measured in units of bits when computed using base 2 logarithms. For our target prior distribution, $I(T=1) = 3.32$ bits and $I(T=0) = 0.15$ bits. Note that less probable events contain more information. The mean over n of $I(x_i)$ is known as the entropy of X and is given by:

$$H(X) = -\sum_{i=1}^n P(X = x_i) \log_2 P(X = x_i) \quad (1)$$

where $P(X = x_i) \log_2 P(X = x_i)$ is defined to be 0 when $P(X = x_i) = 0$. $H(X)$ is a measure of our uncertainty about the value of X , and is maximal when $P(X = x_i) = 1/n$ for all i . For our target prior distribution, the unconditional target uncertainty is computed as $H(T) = 0.47$ bits.

Sensory inputs supply information to a DSC neuron, and thereby reduce its uncertainty about the presence of a target in its receptive field. Such inputs can be regarded as

stochastic, both because of the inherent randomness in the firing patterns of neurons, and because of the varying stimulus effectiveness of the target. Therefore, as for the target, the sensory inputs are modeled as random variables. We will first consider a neuron receiving input of a single modality V (visual). V may assume n specific integer values v_i , where each value represents the number of spikes fired per unit time ($v_1 = 0, v_2 = 1, \dots, v_n = r$), and r is the maximum firing rate considered. We will assume a firing range of 0 to $r=30$, with $n=31$. If the time unit chosen is 250 ms, then this range yields plausible values for sensory input firing rates [1].

Likelihood distributions, which specify the conditional probabilities with which each possible firing rate v_i occurs, must also be assigned. The Bayes' rule model is valid for likelihoods having any probability density distribution. The Poisson distribution will be considered first because it requires the fewest assumptions while still reasonably approximating neuronal firing-rate distributions [9,23]. Different Poisson distributions will be assigned for the case where V is firing spontaneously in the absence of a target [$P(V = v_i | T = 0)$] and where it is driven by the presence of a target [$P(V = v_i | T = 1)$]. A Poisson distribution is specified by only one free parameter: the mean λ . Different means are therefore assigned for the spontaneous (λ_0) and stimulus driven (λ_1) cases, and the needed likelihood distributions can then be expressed as:

$$\begin{aligned} P(V = v_i | T = 0) &= \frac{\lambda_0^{v_i} e^{-\lambda_0}}{v_i!}, \quad \text{and} \\ P(V = v_i | T = 1) &= \frac{\lambda_1^{v_i} e^{-\lambda_1}}{v_i!} \end{aligned} \quad (2)$$

Since DSC neurons are excited by sensory input, it is always the case that $\lambda_0 < \lambda_1$. Poisson likelihood distributions for representative spontaneous and stimulus driven inputs are shown in Fig. 1A. Curves are presented for a spontaneous Poisson mean $\lambda_0 = 4$, and for driven means of $\lambda_1 = 5$ and $\lambda_1 = 11$.

While the Poisson distribution suits a simple form of the model, we will also consider the Gaussian distribution as representing its most general form. Input to a DSC neuron from a single sensory modality is likely to arrive over many sensory afferents. According to the central limit theorem, a distribution that is the sum of many individual random variables that are independent and have finite variance will approximate a Gaussian distribution, regardless of the forms of the distributions of the individual variables. A random variable with a Gaussian distribution best models an input that is combined from multiple sources. Two free parameters are needed to specify a Gaussian distribution: the mean μ and the variance σ^2 . The needed Gaussian likelihood distributions can be expressed as:

$$P(V = v_i | T = 0) = \frac{1}{\sigma_0 \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{v_i - \mu_0}{\sigma_0} \right)^2 \right],$$

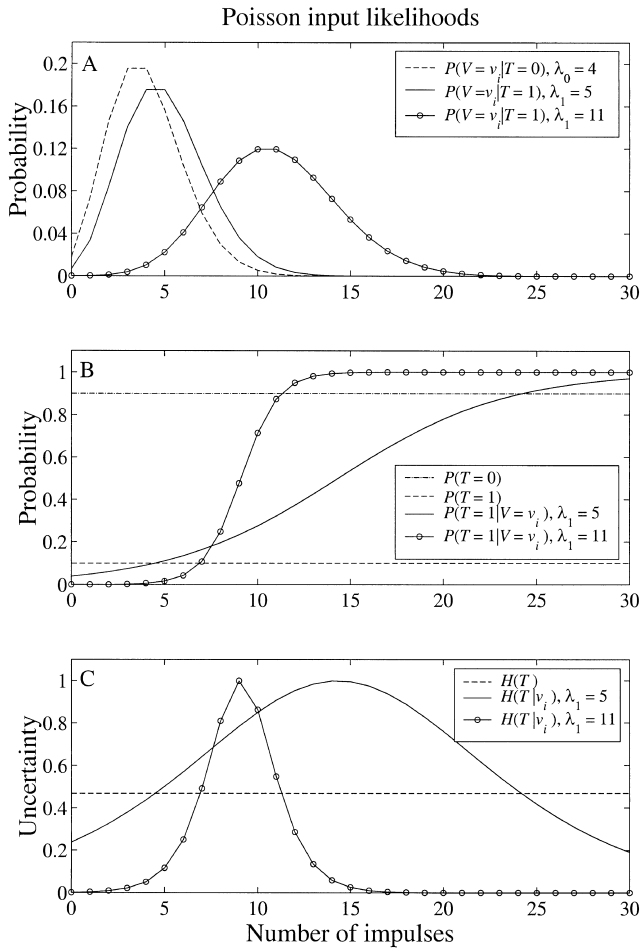


Fig. 1. (A) One spontaneous [$P(V = v_i | T = 0)$] and two stimulus driven [$P(V = v_i | T = 1)$] input likelihoods are plotted as input firing rate varies from 0 to 30. Input likelihoods are modeled as Poisson distributions (Eq. (2)). Spontaneous (λ_0) and driven (λ_1) likelihood means are as indicated. (B) Posterior target probability (Eq. (5)) is plotted for a unimodal sensory input V with firing rate varying from 0 to 30. The spontaneous and driven likelihoods are Poisson distributed as shown in (A). Two cases are considered. The spontaneous mean $\lambda_0 = 4$ in both cases. In the first case [$P(T = 1 | V = v_i)$, $\lambda_1 = 5$] the spontaneous and driven input distributions are strongly overlapping, and in the second [$P(T = 1 | V = v_i)$, $\lambda_1 = 11$] they are well separated. The prior probability of a target [$P(T = 1)$] and of its absence [$P(T = 0)$] are plotted for comparison. (C) Conditional target uncertainty $H(T|v_i)$ (Eq. (4)) is plotted as sensory input V varies from 0 to 30. The spontaneous and driven likelihoods are Poisson distributed as shown in (A). Two cases are considered. The spontaneous mean $\lambda_0 = 4$ in both cases. In the first case [$H(T|v_i)$, $\lambda_1 = 5$] the spontaneous and driven input distributions are strongly overlapping, and in the second [$H(T|v_i)$, $\lambda_1 = 11$] they are well separated. The unconditional uncertainty $H(T)$ is plotted for comparison.

and

$$P(V = v_i | T = 1) = \frac{1}{\sigma_1 \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{v_i - \mu_1}{\sigma_1} \right)^2 \right] \quad (3)$$

As for the Poisson case, means are assigned for the spontaneous (μ_0) and stimulus driven (μ_1) Gaussian cases, with $\mu_0 < \mu_1$. Representative likelihoods are plotted in Fig.

2A, for $\mu_0 = 4$, $\mu_1 = 5$, and $\mu_1 = 11$. For a Poisson distribution the mean and the variance are equal. For consistency, the Gaussian variances are set equal to the means (i.e., $\mu_0 = \sigma_0^2$ and $\mu_1 = \sigma_1^2$). It should be noted that, while the Poisson distribution is discrete, the Gaussian distribution is continuous. Informational measures on the two distribution types are not directly comparable.

Having defined sensory input V as a random variable, we can consider how the information provided by such an input might reduce the uncertainty of a DSC neuron about the presence of a target. To do so we compute $H(T|v_i)$, the conditional target uncertainty (or entropy) given input $V = v_i$. Based on Eq. (1), $H(T|v_i)$ can be stated as:

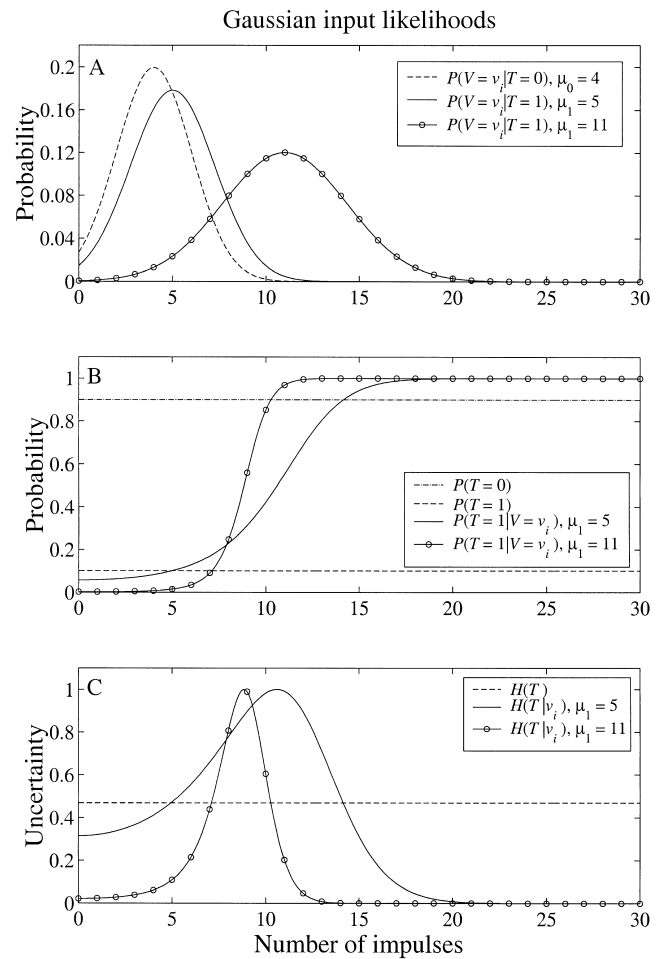


Fig. 2. (A) One spontaneous [$P(V = v_i | T = 0)$] and two stimulus driven [$P(V = v_i | T = 1)$] input likelihoods are plotted as in (Fig. 1A). Input likelihoods are modeled as Gaussian distributions (Eq. (3)). Spontaneous (μ_0) and driven (μ_1) means are as indicated. (B) Unimodal posterior target probability (Eq. (5)) is plotted for two representative cases as in (Fig. 1B). Input likelihoods are Gaussian. The spontaneous mean is $\mu_0 = 4$ in both cases. Driven means are as indicated. The prior probability of a target [$P(T = 1)$] and of its absence [$P(T = 0)$] are plotted for comparison. (C) Conditional target uncertainty or entropy $H(T|v_i)$ (Eq. (4)) is plotted for two representative cases as in (Fig. 1C). Input likelihoods are Gaussian. The spontaneous mean is $\mu_0 = 4$ in both cases. Driven means are as indicated. The unconditional uncertainty $H(T)$ is plotted for comparison.

$$H(T|v_i) = - [P(T=0|V=v_i) \log_2 P(T=0|V=v_i) + P(T=1|V=v_i) \log_2 P(T=1|V=v_i)] \quad (4)$$

The target can assume one of only two states: present ($T=1$) or absent ($T=0$). We earlier computed the target uncertainty $H(T)$ in the absence of sensory input, using Eq. (1) and the unconditional (prior) probabilities of these two states. In order to compute the conditional uncertainty of the target using Eq. (4), we must compute the conditional probabilities of the two target states given sensory input. The conditional (posterior) probability that a target is present [$P(T=1|V=v_i)$] or that it is absent [$P(T=0|V=v_i)$] can be expressed using Bayes' rule [2,7]:

$$P(T=0|V=v_i) = \frac{P(V=v_i|T=0)P(T=0)}{P(V=v_i)},$$

$$P(T=1|V=v_i) = \frac{P(V=v_i|T=1)P(T=1)}{P(V=v_i)} \quad (5)$$

Since $T=0$ and $T=1$ are mutually exclusive events, we can use the principle of total probability [21] to express $P(V=v_i)$ as:

$$P(V=v_i) = P(V=v_i|T=1)P(T=1) + P(V=v_i|T=0)P(T=0) \quad (6)$$

The posterior probability of a target, $P(T=1|V=v_i)$, can now be computed using Eq. (5) and the input likelihood distributions defined above.

The ability of a sensory input to reduce the uncertainty of a DSC neuron concerning a target may be related to the difference between the spontaneous and driven likelihood distributions of the input. To illustrate this, we consider two cases, one in which these distributions are strongly overlapping ($\lambda_0=4$, $\lambda_1=5$, and $\mu_0=4$, $\mu_1=5$), and another in which they are well separated ($\lambda_0=4$, $\lambda_1=11$, and $\mu_0=4$, $\mu_1=11$). The effects of overlapping and well-separated likelihoods are illustrated in Fig. 1 for Poisson and in Fig. 2 for Gaussian distributed inputs. Posterior target probabilities for the overlapping and well-separated cases are plotted in Figs. 1B and 2B for Poisson and Gaussian distributed inputs, respectively. In each plot, the prior probabilities of the target [$P(T=0)$ and $P(T=1)$] are also indicated. As input V increases from 0 to 30, it changes from being most likely spontaneous to most likely driven by a target, and the posterior probability that a target is present rises sigmoidally from nearly 0 to nearly 1. Note that the sigmoid is shallow for the case where the distributions are strongly overlapping, and is steeper for the case where they are more separated. At the point where the spontaneous and driven distributions cross, the input is just as likely to be spontaneous as to be driven, and the posterior target probability is equal to the prior target probability. We postulate that the firing rates of DSC neurons encode such posterior probabilities [1]. An ideal

Bayesian decision-maker should always decide that $T=1$ whenever $P(T=1|V=v_i) > P(T=0|V=v_i)$, and decide $T=0$ otherwise [7]. In the flat portions of the sigmoid, where the posterior probability of a target is nearly 0 or nearly 1 and target uncertainty is low, such a procedure has a very high probability of producing correct decisions regarding T .

Target uncertainty (or entropy) is plotted for the Poisson and Gaussian cases in Figs. 1C and 2C, respectively. The dashed line indicates the unconditional uncertainty $H(T)$ computed using Eq. (1). Sensory input V can have the effect of reducing uncertainty about the presence of a target. The conditional target uncertainty $H(T|v_i)$ is nearly 0 for those portions of the plot corresponding to the flat portions of the probability sigmoid. At low input values where the sigmoid is flat, sensory input provides a near certainty that a target is absent. At high input values where again the sigmoid is flat, there is a near certainty that a target is present. The sloping portion of the sigmoid, however, corresponds to a region of uncertainty concerning the target. For firing rates v_i such that $P(T=1|V=v_i) > P(T=1)$ and $P(T=1|V=v_i) < P(T=0)$, the conditional uncertainty $H(T|v_i)$ exceeds the unconditional uncertainty $H(T)$. Within this region, an ideal Bayesian decision-maker would have a probability of between 0.1 and 0.5 of deciding erroneously. When $P(T=1|V=v_i) = P(T=0|V=v_i) = 0.5$, any judgement about the presence of a target is just as likely to be wrong as to be right. At this point uncertainty is maximal, and $H(T|v_i) = 1$ bit, which represents complete uncertainty concerning binary random variable T . Note that the zone of uncertainty is larger for the case of strongly overlapping spontaneous and driven input distributions, and is smaller for the more separated case. These effects are generally similar for Poisson and Gaussian input distributions. However, direct comparisons on informational measures between the two distribution types are not valid (see above).

We can now examine the effects of input from a second sensory modality on target uncertainty. We consider two modalities V (visual) and A (auditory). This particular combination of modalities was chosen arbitrarily, and our analysis applies with equal validity for any other combination. For the case of a neuron receiving input from these two modalities, the joint conditional target uncertainty $H(T|v_i, a_j)$ can be expressed based on Eq. (1) as:

$$H(T|v_i, a_j) = -P(T=0|V=v_i, A=a_j) \log_2 P(T=0|V=v_i, A=a_j) - P(T=1|V=v_i, A=a_j) \log_2 P(T=1|V=v_i, A=a_j) \quad (7)$$

To perform this computation, we must first compute the joint conditional probability that a target is present [$P(T=1|V=v_i, A=a_j)$] or that it is absent [$P(T=0|V=v_i, A=$

a_j)] in terms of the likelihood and prior distributions. These are again given by Bayes' rule as:

$$\begin{aligned}
 P(T = 1|V = v_i, A = a_j) &= \frac{P(V = v_i, A = a_j|T = 1)P(T = 1)}{P(V = v_i, A = a_j)}, \quad \text{and} \\
 P(T = 0|V = v_i, A = a_j) &= \frac{P(V = v_i, A = a_j|T = 0)P(T = 0)}{P(V = v_i, A = a_j)} \quad (8)
 \end{aligned}$$

To simplify the analysis, we will assume that V and A are conditionally independent given T . The joint likelihoods of $V = v_i$ and $A = a_j$ can then be computed as:

$$\begin{aligned}
 P(V = v_i, A = a_j|T = 1) &= P(V = v_i|T = 1)P(A = a_j|T = 1), \quad \text{and} \\
 P(V = v_i, A = a_j|T = 0) &= P(V = v_i|T = 0)P(A = a_j|T = 0) \quad (9)
 \end{aligned}$$

Similarly, the unconditional joint probability that $V = v_i$ and $A = a_j$ can be computed using the principle of total probability as:

$$\begin{aligned}
 P(V = v_i, A = a_j) &= \\
 &P(V = v_i|T = 1)P(A = a_j|T = 1)P(T = 1) \\
 &+ P(V = v_i|T = 0)P(A = a_j|T = 0)P(T = 0) \quad (10)
 \end{aligned}$$

To compute the posterior probability of a target given bimodal input, we will assume that A , like V , assumes integer values over the range from $a_1 = 0$ to $a_{31} = r = 30$. For simplicity we consider the special case where $i = j$, with $V = v_i$ and $A = a_j$ varying together over this range. All other combinations are possible, and the results are qualitatively similar so long as V and A increase together. We will further assume, for simplicity, that the two modalities have likelihood distributions possessing the same spontaneous and driven means. For the case where the spontaneous and driven input likelihood distributions strongly overlap ($\lambda_0 = 4$, $\lambda_1 = 5$ and $\mu_0 = 4$, $\mu_1 = 5$) we have seen that the posterior probability sigmoid is shallow and the zone of uncertainty is wide for the unimodal case. Figs. 3A and 4A show the effect on posterior target probability of input of a second modality for overlapping likelihoods that are Poisson and Gaussian distributed, respectively. The bimodal (joint) posterior target probabilities [$P(T = 1|V = v_i, A = a_i)$] are co-plotted with the unimodal posterior probabilities presented earlier [$P(T = 1|V = v_i)$]. Note that for both the Poisson and Gaussian likelihoods the slope of the probability sigmoid is substantially higher in the bimodal than in the unimodal case.

The joint conditional uncertainty for the bimodal case [$H(T|v_i, a_j)$] and the conditional uncertainty for the unimodal case [$H(T|v_i)$] are plotted in Figs. 3C and 4C.

Note that the zone of uncertainty is considerably narrower for the bimodal than for the unimodal case. With Poisson input likelihoods, the bimodal zone of uncertainty is about half as large as it is for the unimodal case. With Gaussian input likelihoods it is about 60% as large. Input from a second modality can greatly decrease uncertainty about the target.

Fig. 3B and D provide an example in which spontaneous and driven Poisson likelihoods are better separated ($\lambda_0 = 4$ and $\lambda_1 = 11$). A similar example using Gaussian likelihoods ($\mu_0 = 4$ and $\mu_1 = 11$) is illustrated in Fig. 4B and D. Posterior probabilities are plotted in panel B of each figure, and conditional uncertainties in panel D. As already noted for the single modality case, the probability sigmoid $P(T = 1|V = v_i)$ associated with the well-separated input distributions is already quite steep, and the zone of uncertainty associated with $H(T|v_i)$ is already quite narrow. Addition of a second modality results in a steeper sigmoid for $P(T = 1|V = v_i, A = a_i)$, and a narrower zone of uncertainty for $H(T|v_i, a_j)$. As for the previous case, addition of a second, well separated input reduces the zone of uncertainty to roughly half of its unimodal width for the Poisson case, and roughly 60% of its unimodal width for the Gaussian case. However, in absolute terms, the benefit of adding a second modality is not as large for the well-separated input likelihoods (Figs. 3D and 4D) as it was for strongly overlapping input likelihoods considered previously (Figs. 3C and 4C).

3. Likelihood distributions and the benefits of multimodality

Sensory input to a DSC neuron can supply information concerning a target and thereby reduce target uncertainty. We have seen that when the spontaneous and driven input likelihoods are similar to one another, there is a wide range of input firing rates for which target uncertainty is large. Input from a second modality can reduce this range by as much as half. When the spontaneous and driven input likelihoods are more different from one another, the size of the range of uncertainty is smaller. The effect of adding input from a second modality under such circumstances is of less consequence. We conclude from this analysis that inputs with similar spontaneous and driven likelihoods reduce uncertainty less than inputs with different spontaneous and driven likelihoods. It appears that inputs with similar or different spontaneous and driven likelihood distributions are relatively ambiguous or unambiguous, respectively. We now wish to generalize these conclusions for a range of plausible Poisson and Gaussian input likelihoods.

The amount by which spontaneous and driven likelihood distributions differ from one another can be quantified using the relative entropy, or Kullback–Leibler difference [6]. The Kullback–Leibler difference (D) between the

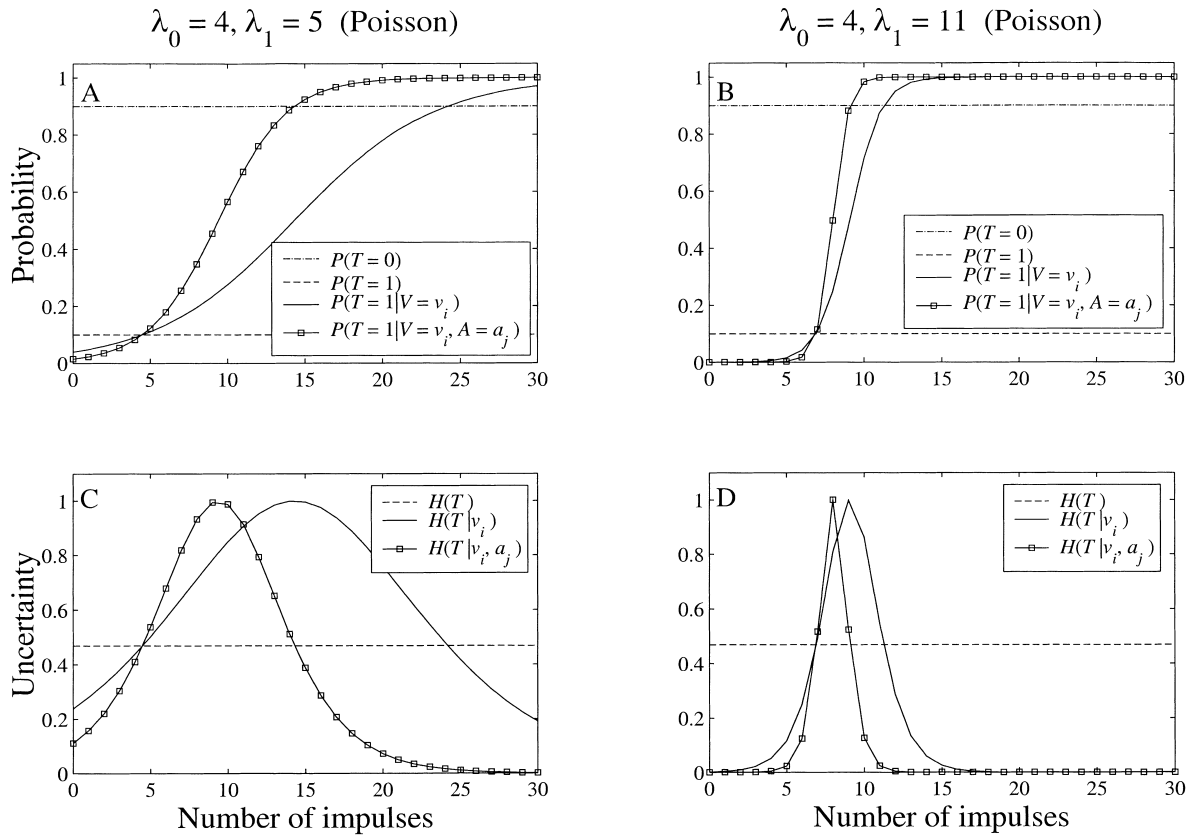


Fig. 3. (A) Posterior target probability is plotted as sensory input firing rate varies from 0 to 30 in the unimodal case for V only [$P(T = 1|V = v_i)$] (Eq. (5)), and in the bimodal case [$P(T = 1|V = v_i, A = a_j)$] (Eq. (8)) with inputs V and A varying together. Input likelihoods are closely spaced Poisson distributions with the indicated means. The prior target probabilities $P(T = 1)$ and $P(T = 0)$ are shown for comparison. (B) Posterior target probability is plotted in the unimodal and bimodal cases as in (A). Input likelihoods are widely separated Poisson distributions with the indicated means. (C) Conditional target uncertainty is plotted as a function of input firing rate in the case of unimodal sensory input V [$H(T|v_i)$] (Eq. (4)), and in the case of bimodal sensory inputs V and A varying together [$H(T|v_i, a_j)$] (Eq. (7)). Input likelihoods are closely spaced Poisson distributions with the indicated means. The unconditional uncertainty $H(T)$ is plotted for comparison. (D) Conditional target uncertainty is plotted in the unimodal and bimodal cases as in (C). Input likelihoods are widely spaced Poisson distributions with the indicated means.

spontaneous and driven likelihoods in the unimodal case is:

$$D(P(V = v_i|T = 0)||P(V = v_i|T = 1)) = \sum_{i=1}^n P(V = v_i|T = 0) \log_2 \frac{P(V = v_i|T = 0)}{P(V = v_i|T = 1)} \quad (11)$$

For the bimodal case, the Kullback–Leibler difference between the spontaneous and driven likelihoods is:

$$D(P(V = v_i, A = a_j|T = 0)||P(V = v_i, A = a_j|T = 1)) = \sum_{i=1}^n \sum_{j=1}^n P(V = v_i|T = 0)P(A = a_j|T = 0) \times \log_2 \frac{P(V = v_i|T = 0)P(A = a_j|T = 0)}{P(V = v_i|T = 1)P(A = a_j|T = 1)} \quad (12)$$

We use D to quantify the difference between the spontaneous and driven likelihood distributions as the

difference between their means is increased. Eqs. (11) and (12) were evaluated while the spontaneous likelihood mean was fixed at 4, and the driven likelihood mean was varied over the range from 4 to 30. To allow for the larger means, firing rates from 0 to $r = 50$ were assumed for both modalities. Computations were made for both Poisson and Gaussian likelihoods (Fig. 5A and B). This and subsequent computations involving the Gaussian distribution utilized a discrete approximation, with a bin size of one spike. The results displayed in all subsequent panels (Fig. 5C–F) use the same range of distribution means (4–30) and firing rates (0–50). The two distributions overlap extensively when the spontaneous and driven means are close in value, and D is small. D increases monotonically as the means of the two distributions grow more widely separated. Note that D is twice as large for the bimodal as for the unimodal case.

The Kullback–Leibler difference D provides a means of quantifying the difference between the spontaneous and

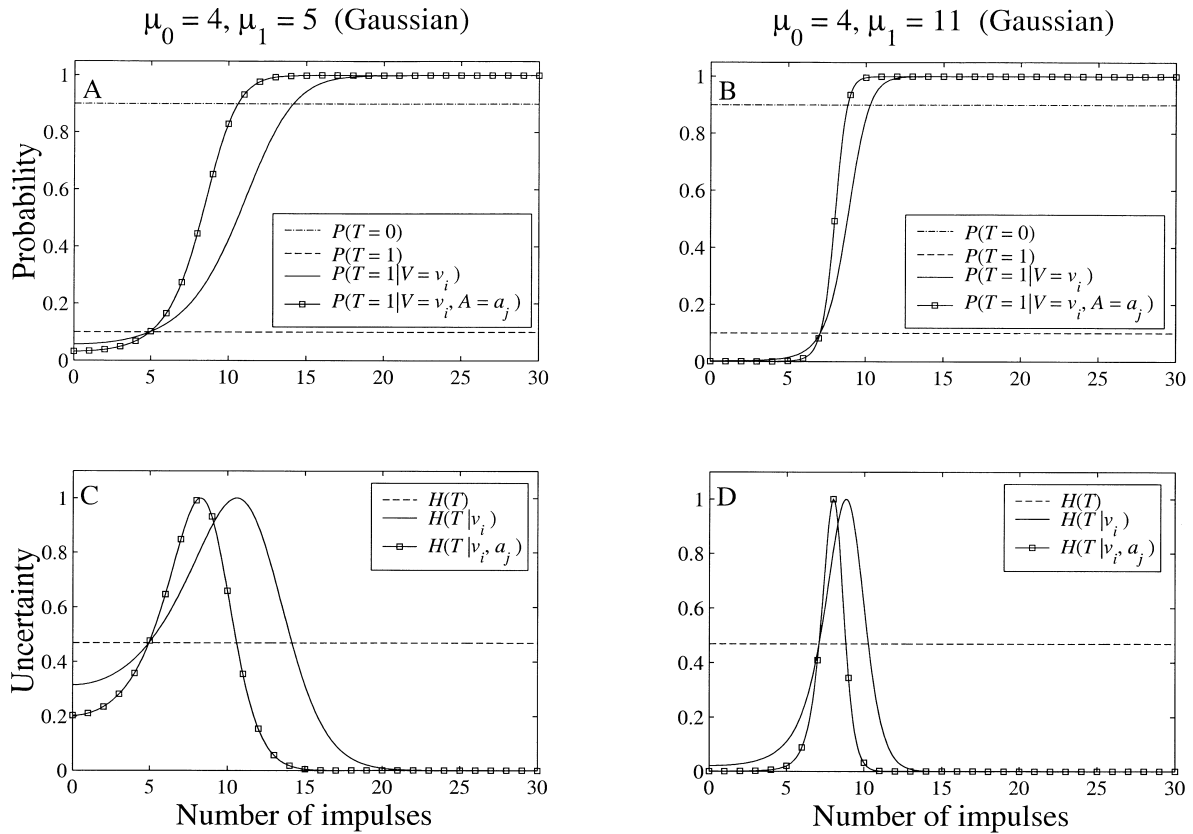


Fig. 4. (A) Posterior target probability is plotted as sensory input firing rate varies from 0 to 30 in the unimodal case for V only [$P(T=1|V=v_i)$] (Eq. (5)), and in the bimodal case [$P(T=1|V=v_i, A=a_j)$] (Eq. (8)) with V and A varying together. Input likelihoods are closely spaced Gaussian distributions with the indicated means. The prior target probabilities $P(T=1)$ and $P(T=0)$ are shown for comparison. (B) Posterior target probability is plotted in the unimodal and bimodal cases as in (A). Input likelihoods are widely separated Gaussian distributions with the indicated means. (C) Conditional target uncertainty is plotted as a function of input firing rate in the case of unimodal sensory input V [$H(T|v_i)$] (Eq. (4)), and in the case of bimodal sensory inputs V and A varying together [$H(T|v_i, a_j)$] (Eq. (7)). Input likelihoods are closely spaced Gaussian distributions with the indicated means. The unconditional uncertainty $H(T)$ is plotted for comparison. (D) Conditional target uncertainty is plotted in the unimodal and bimodal cases as in (C). Input likelihoods are widely spaced Gaussian distributions with the indicated means.

driven likelihood distributions of an input. As such, it should be related to the ability of an input to reduce uncertainty and thereby provide information concerning a target. For comparison with D we need a measure of the uncertainty of the target conditioned not at an individual value of an input, but over the entire range of possible values of an input. For the case of unimodal input V , this measure is provided by the average conditional uncertainty $H(T|V)$. This measure is best understood by comparison with the uncertainty measures already introduced.

$H(T)$ is a measure of uncertainty about the presence of a target, given only the target prior probability distribution. $H(T|v_i)$ is a measure of uncertainty about the presence of a target, given sensory input of some particular value $V=v_i$. For a given spontaneous and driven likelihood distribution, $H(T|v_i)$ varies systematically with variation in firing rate v_i . $H(T|V)$ is a measure of the average uncertainty (or entropy) across the entire range of values v_i that V may assume. $H(T|V)$ is the average of $H(T|v_i)$ over a range of values of $V=v_i$, and can be computed as:

$$\begin{aligned}
 H(T|V) &= -\sum_{i=1}^n P(v_i)H(T|v_i) \\
 &= -\sum_{i=1}^n P(v_i) \left[\sum_{T=0}^1 P(T|v_i) \log_2 P(T|v_i) \right] \\
 &= -\sum_{i=1}^n P(v_i) [P(T=0|v_i) \log_2 P(T=0|v_i) \\
 &\quad + P(T=1|v_i) \log_2 P(T=1|v_i)] \quad (13)
 \end{aligned}$$

For bimodal input V and A , the joint average conditional uncertainty can be computed as:

$$\begin{aligned}
 H(T|V, A) &= -\sum_{i=1}^n \sum_{j=1}^n P(v_i, a_j)H(T|v_i, a_j) \\
 &= -\sum_{i=1}^n \sum_{j=1}^n P(v_i, a_j) \left[\sum_{T=0}^1 P(T|v_i, a_j) \log_2 P(T|v_i, a_j) \right]
 \end{aligned}$$

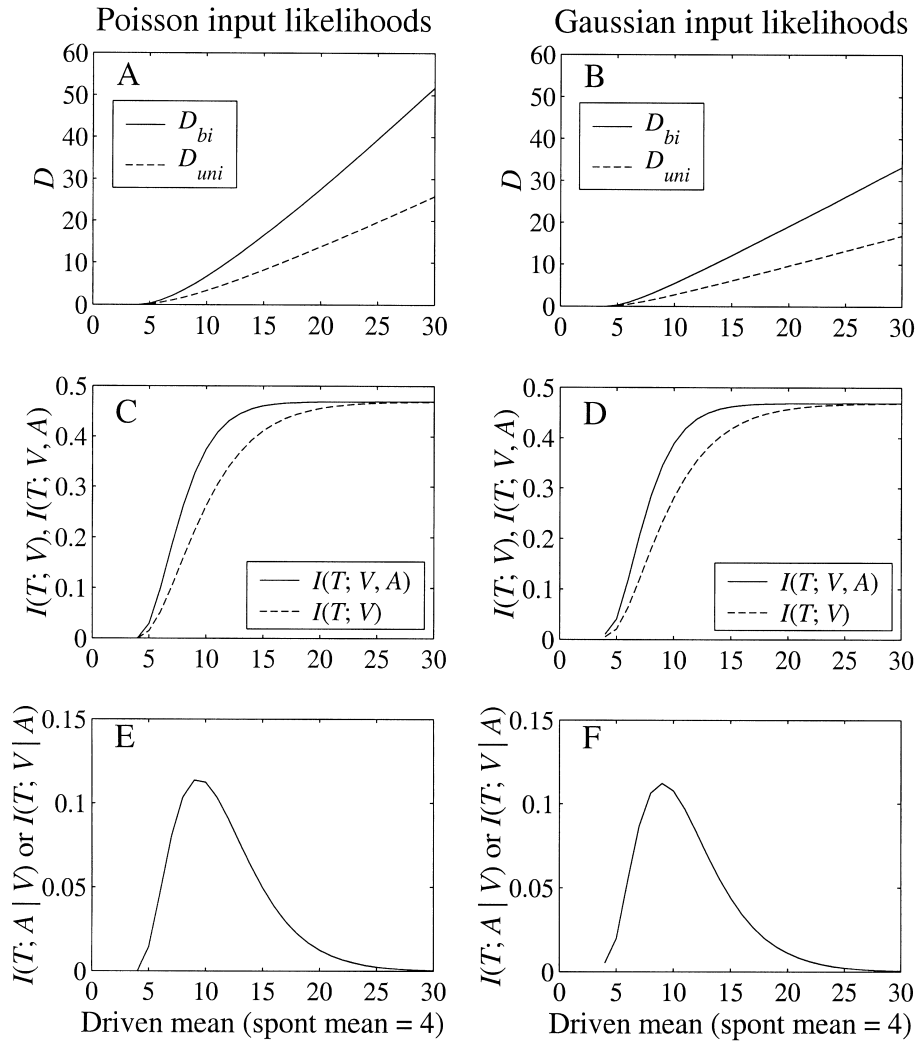


Fig. 5. (A) Kullback–Leibler distance (D) between the spontaneous and driven likelihoods is plotted versus the mean of the driven likelihoods, which varies from 4 to 30. The spontaneous mean is fixed at 4. Likelihoods are modeled as Poisson distributions. D for a unimodal input V was computed using Eq. (11) [$D_{uni} = D(P(V = v_i|T = 0)||P(V = v_i|T = 1))$]. D for bimodal inputs V and A was computed using Eq. (12) [$D_{bi} = D(P(V = v_i, A = a_j|T = 0)||P(V = v_i, A = a_j|T = 1))$]. For simplicity, means for both modalities (V and A) are equal. (B) Kullback–Leibler distance plotted using Gaussian likelihoods. Means and plotting conventions are as described for (A). Likelihood variances were set equal to the means. (C) The mutual information between the target and sensory input is plotted versus the driven mean. For both modalities, the spontaneous mean is fixed at 4, and the driven mean varies from 4 to 30. Likelihoods are Poisson distributed. Mutual information for the unimodal case is computed using Eq. (15) [$I(T; V) = H(T) - H(T|V)$], and for the bimodal case using Eq. (16) [$I(T; V, A) = H(T) - H(T|V, A)$]. (D) Mutual information plotted using Gaussian likelihoods. Means and plotting conventions are as described for (C). Likelihood variances were set equal to the means. (E) The amount of information gained (conditional mutual information) by the addition of a second input modality [$I(T; A|V) = H(T|V) - H(T|V, A)$] is plotted versus the driven mean for Poisson distributed likelihoods. Because the likelihoods for V and A are set equal for simplicity, $I(T; A|V) = I(T; V|A)$ where $I(T; V|A) = H(T|A) - H(T|V, A)$. The spontaneous mean is held fixed at 4, and the driven mean varies from 4 to 30. (F) Conditional mutual information is plotted for Gaussian likelihoods. Means and plotting conventions are as described for (E). Variances were set equal to the means.

$$\begin{aligned}
 &= - \sum_{i=1}^n \sum_{j=1}^n P(v_i, a_j) [P(T = 0|v_i, a_j) \log_2 P(T = 0|v_i, a_j) \\
 &\quad + P(T = 1|v_i, a_j) \log_2 P(T = 1|v_i, a_j)] \quad (14)
 \end{aligned}$$

The needed unimodal and bimodal conditional probabilities can be computed using Eqs. (5) and (8).

If a sensory input can indeed provide information concerning a target, then it should reduce the uncertainty concerning the target. In other words, if an input does

provide target information, then the average target uncertainty conditioned on input should be less than the unconditional target uncertainty. The gain in information due to sensory input can then be quantified as the difference between the unconditional and conditional uncertainties. For unimodal input V , the average information gain, also known as mutual information, is computed as:

$$I(T; V) = H(T) - H(T|V) \quad (15)$$

Likewise, for bimodal input V and A , the average information gain is given by:

$$I(T; V, A) = H(T) - H(T|V, A) \quad (16)$$

We can now quantify the average information gained in the bimodal and unimodal cases for inputs of varying D . To do so we evaluate Eqs. (15) and (16) over the same range of spontaneous and driven likelihood parameters used for computing D above. The information gain (Fig. 5C and D) is zero when the spontaneous and driven means are equal. The average reduction in uncertainty becomes greater as the separation between the spontaneous and driven means increases (i.e., as D increases). Average information gain increases faster in the bimodal $[I(T; V, A)]$ than in the unimodal $[I(T; V)]$ case. In both cases, however, average information gain reaches a plateau equal to $H(T)$. At that point, the input has completely reduced the uncertainty associated with the target. Further increasing the separation between the spontaneous and driven likelihood means produces further increases in D , but cannot further increase the amount of information that the inputs provide concerning the target.

The average reduction in uncertainty produced by two inputs can exceed that produced by one alone, but only over a limited range of D . To determine how much information is, on average, gained through integration of a second modality, the difference between the average conditional and joint conditional uncertainty, or conditional mutual information $I(T; A|V) = H(T|V) - H(T|V, A)$, was computed. The conditional mutual information $I(T; A|V)$ is the average amount of information gained about T due to A when V has already been taken into consideration. Because we assume, for simplicity, that V and A have the same input likelihoods, $I(T; A|V) = I(T; V|A)$. The conditional mutual information is shown in Fig. 5E and F. The greatest average gains are realized for an intermediate range of driven means. For the most ambiguous likelihoods, further information could be gained by integrating input of a third modality (not shown).

4. Discussion

Multisensory integration in the DSC has been the subject of extensive neurophysiological investigation [15,17,19,27–31,34]. We have earlier hypothesized that a DSC neuron uses Bayes' rule to compute the posterior probability that a target is present in its receptive field, given stochastic sensory inputs [1]. Bayes' rule is a fundamental concept in probability theory. It has found wide application and underlies all modern systems for probabilistic inference in artificial intelligence [5,7,25]. The idea that neural responses encode stimulus probabilities has been proposed previously by Barlow [3,4], however Barlow's proposal involved classical statistical

inference and not Bayesian decision theory. The significance of our model is that it provides, for the first time, a rigorous method of exploring the possible functional significance of multisensory integration in the DSC. In past work, we have shown that both multisensory enhancement and inverse effectiveness arise as natural consequences of the Bayesian model [1]. We have here extended this earlier work by performing an information theoretic analysis of the model.

It might seem obvious that multisensory integration should increase the information received by a DSC neuron concerning a target in its receptive field. However, many DSC neurons are not multisensory. Respectively, in cat and monkey, 46% and 73% of DSC neurons are unimodal [33]. Fewer than 50% are bimodal and fewer than 10% are trimodal in either species [33]. Information theoretic analysis of our model provides an explanation for the otherwise puzzling fact that some DSC neurons are unimodal despite the availability of multisensory input.

A sensory neuron indicates the presence of a target by the difference between its spontaneous and stimulus driven discharge distributions (input likelihoods). If the spontaneous and driven likelihoods are the same, (i.e., $D=0$) then the input is completely ambiguous. Such an input would provide no information to a DSC neuron concerning the presence of a target. The target uncertainty would remain equal to the unconditional uncertainty $H(T)$. Inputs for which the spontaneous and driven distributions are different are not completely ambiguous. When $D>0$, the posterior probability of a target forms a sigmoid when plotted against sensory input firing rate. The sloping portion of the sigmoid is a region of high conditional uncertainty. When the input is relatively ambiguous and D is small, the sigmoid is shallow and the zone of high uncertainty is large. The average information gain provided by an ambiguous input is small. As the driven mean increases, the input becomes less ambiguous. The sigmoid becomes steeper, the zone of uncertainty becomes narrower, and the average information gain provided by the input increases. Average uncertainty regarding the presence of a target decreases as the input likelihoods become more well-separated and so more unambiguous. Eventually, the input becomes sufficiently unambiguous that the average information it provides is maximal, with average uncertainty regarding the target reduced nearly to zero. Further increase in the driven mean results in further increase in D , but little or no further increase in average information gain.

Our hypothesis is that the function of DSC neurons is to detect the presence of targets in their receptive fields. The results of our simple example illustrate that it is possible for a DSC neuron to receive such informative input from a single modality that it would gain little or no further information about the presence of a target by integrating input of another modality. Unimodal DSC neurons may be those that receive unambiguous unimodal input and re-

liably detect targets while maintaining their modality specificity. On the other hand, a DSC neuron receiving ambiguous input from a single modality could gain appreciably more information about the presence of a target by integrating inputs of one or two additional modalities. Multisensory DSC neurons may be those that combine ambiguous inputs of multiple modalities, and therefore make use of them to reliably detect targets, but they do so at the price of losing their ability to convey modality specific information.

Acknowledgements

We thank Prakash Ishwar, Pierre Moulin, Jesse Reichler, and Liudmila Yafremava for helpful discussions and for comments on the manuscript prior to submission. This work was supported by NSF grants IBN 92-21823 and IBN 00-80789, and a grant from the Critical Research Initiatives of the State of Illinois, all to T.J.A.

References

- [1] T.J. Anastasio, P.E. Patton, K. Belkacem-Boussaid, Using Bayes' rule to model multisensory enhancement in the superior colliculus, *Neural Comp.* 12 (2000) 997–1019.
- [2] D. Appelbaum, *Probability and Information: An Integrated Approach*, Cambridge University Press, Cambridge, 1996.
- [3] H.B. Barlow, Pattern recognition and the responses of sensory neurons, *Ann. N.Y. Acad. Sci.* 156 (1969) 872–881.
- [4] H.B. Barlow, Single units and sensation: a neuron doctrine for perceptual psychology?, *Perception* 1 (1972) 371–394.
- [5] C.M. Bishop, *Neural Networks For Pattern Recognition*, Clarendon Press, Oxford, 1995.
- [6] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, Wiley, New York, 1991.
- [7] R.O. Duda, P.E. Hart, D.G. Stork, *Pattern Classification*, 2nd Edition, Wiley, New York, 2001.
- [8] S.B. Edwards, C.L. Ginsburgh, C.K. Henkel, B.E. Stein, Sources of subcortical projections to the superior colliculus in the cat, *J. Comp. Neurol.* 184 (1979) 309–329.
- [9] F. Gabbiani, C. Koch, Principles of spike train analysis, in: C. Koch, I. Segev (Eds.), *Methods in Neuronal Modeling*, 2nd Edition, MIT Press, Cambridge, MA, 1998, pp. 312–360.
- [10] L.R. Harris, The superior colliculus and movements of the head and eyes in cat, *J. Physiol. (Lond.)* 300 (1980) 367–391.
- [11] J.K. Harting, B.V. Updyke, D.P. Van Lieshout, Corticotectal projections in the cat: anterograde transport studies of twenty-five cortical areas, *J. Comp. Neurol.* 324 (1992) 379–414.
- [12] A.J. King, A.R. Palmer, Integration of visual and auditory information in bimodal neurones in guinea-pig superior colliculus, *Exp. Brain Res.* 60 (1985) 492–500.
- [13] M.A. Meredith, H.R. Clemo, B.E. Stein, Somatotopic component of the multisensory map in the deep laminae of the cat superior colliculus, *J. Comp. Neurol.* 312 (1991) 353–370.
- [14] M.A. Meredith, J.W. Nemitz, B.E. Stein, Determinants of multisensory integration in superior colliculus neurons. I. Temporal factors, *J. Neurosci.* 7 (1987) 3215–3229.
- [15] M.A. Meredith, B.E. Stein, Interactions among converging sensory inputs in the superior colliculus, *Science* 221 (1983) 389–391.
- [16] M.A. Meredith, B.E. Stein, Spatial factors determine the activity of multisensory neurons in cat superior colliculus, *Brain Res.* 365 (1986) 350–354.
- [17] M.A. Meredith, B.E. Stein, Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration, *J. Neurophysiol.* 56 (1986) 640–662.
- [18] M.A. Meredith, B.E. Stein, The visuotopic component of the multisensory map in the deep laminae of the cat superior colliculus, *J. Neurosci.* 10 (1990) 3727–3742.
- [19] M.A. Meredith, M.T. Wallace, B.E. Stein, Visual, auditory and somatosensory convergence in output neurons of the cat superior colliculus: multisensory properties of the tecto-reticulo-spinal projection, *Exp. Brain Res.* 88 (1992) 181–186.
- [20] J.C. Middlebrooks, E.I. Knudsen, A neural code for auditory space in the cat's superior colliculus, *J. Neurosci.* 4 (1984) 2621–2634.
- [21] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd Edition, McGraw-Hill, Boston, MA, 1991.
- [22] C.K. Peck, J.A. Baro, S.M. Warder, Effects of eye position on saccadic eye movements and on the neuronal responses to auditory and visual stimuli in cat superior colliculus, *Exp. Brain Res.* 103 (1995) 227–242.
- [23] F. Rieke, D. Warland, R. de Ruyter van Stevenick, *Spikes: Exploring the Neural Code*, MIT Press, Cambridge, MA, 1997.
- [24] D.A. Robinson, Eye movements evoked by collicular stimulation in the alert monkey, *Vision Res.* 12 (1972) 1795–1808.
- [25] S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, Upper Saddle River, NJ, 1995.
- [26] D.L. Sparks, R. Hartwich-Young, The deep layers of the superior colliculus, in: R.H. Wurtz, M. Goldberg (Eds.), *The Neurobiology of Saccadic Eye Movements*, Vol. 3, Elsevier, Amsterdam, 1989, pp. 213–255.
- [27] B.E. Stein, Neural mechanisms for synthesizing sensory information and producing adaptive behaviors, *Exp. Brain Res.* 123 (1998) 124–135.
- [28] B.E. Stein, M.A. Meredith, *The Merging of the Senses*, MIT Press, Cambridge, MA, 1993.
- [29] B.E. Stein, M.A. Meredith, M.T. Wallace, The visually responsive neuron and beyond: multisensory integration in cat and monkey, *Prog. Brain Res.* 95 (1993) 79–90.
- [30] B.E. Stein, M.T. Wallace, T.R. Stanford, Development of multisensory integration: transforming sensory input into motor output, *Ment. Retard. Dev. Disabil. Res. Rev.* 5 (1999) 72–85.
- [31] M.T. Wallace, M.A. Meredith, B.E. Stein, Multisensory integration in the superior colliculus of the alert cat, *J. Neurophysiol.* 20 (1998) 1006–1010.
- [32] M.T. Wallace, B.E. Stein, Cross-modal synthesis in the midbrain depends on input from cortex, *J. Neurophysiol.* 71 (1994) 429–432.
- [33] M.T. Wallace, B.E. Stein, Sensory organization of the superior colliculus in cat and monkey, *Prog. Brain Res.* 112 (1996) 301–311.
- [34] M.T. Wallace, L.K. Wilkinson, B.E. Stein, Representation and integration of multiple sensory inputs in primate superior colliculus, *J. Neurophysiol.* 76 (1996) 1246–1266.
- [35] R.H. Wurtz, M.E. Goldberg, Activity of superior colliculus in behaving monkey. III. Cells discharging before eye movements, *J. Neurophysiol.* 35 (1972) 575–586.