Why aren't all deep superior colliculus neurons multisensory? A Bayes' ratio analysis

HANS COLONIUS
Oldenburg University, Oldenburg, Germany

and

ADELE DIEDERICH
International University Bremen, Bremen, Germany

Multisensory neurons in the deep superior colliculus (SC) show response enhancement to cross-modal stimuli that coincide in time and space. However, multisensory SC neurons respond to unimodal input as well. It is thus legitimate to ask why not all deep SC neurons are multisensory or, at least, develop multi-sensory behavior during an organism's maturation. The novel answer given here derives from a signal detection theory perspective. A Bayes' ratio model of multisensory enhancement is suggested. It holds that deep SC neurons operate under the Bayes' ratio rule, which guarantees optimal performance—that is, it maximizes the probability of target detection while minimizing the false alarm rate. It is shown that optimal performance of multisensory neurons vis-à-vis cross-modal stimuli implies, at the same time, that modality-specific neurons will outperform multisensory neurons in processing unimodal targets. Thus, only the existence of both multisensory and modality-specific neurons allows optimal performance when targets of one or several modalities may occur.

In order to survive in a possibly hostile environment with stimuli occurring at various positions in space and time, an individual must constantly discriminate between signals relevant for action planning (*targets*) and signals that demand no immediate response (*nontargets*). Separate sensory channels process stimuli from different modalities, but in order to succeed in the discrimination task, the individual must integrate the information coming from different modalities in an optimal way. In an orienting task, saccadic reaction time (the time to initiate an eye movement) to a target stimulus is facilitated when the stimulus is composed of several modalities, rather than a single one. This multisensory interaction effect is larger than predicted by a simple probability summation mechanism, and it is most conspicuous if the stimulus elements coincide in time and space (e.g., Colonius & Diederich, 2004; Frens, Van Opstal, & Van der Willigen, 1995; for a review, see Diederich & Colonius, 2004).

Evidence for multisensory interaction at the neural level has been documented in a series of studies on the superior colliculus (SC), a midbrain structure involved in the control of eye movements (Stein & Meredith, 1993). Multisensory neurons in the deep SC of anesthetized cats (Meredith & Stein, 1986a, 1986b; Stein, Magalhães-Castro, & Kruger, 1976) and monkeys (Wallace, Wilkin-son, & Stein, 1996) show an enhanced response to partic-ular combinations of visual, auditory, and tactile stimuli paralleling the spatial–temporal rules observed in behav-ioral studies. Similar results have recently been obtained for recordings from unanesthetized cats by Wallace, Meredith, and Stein (1998) and from the awake behaving monkey by Bell, Corneil, Meredith, and Muñoz (2001) and by Frens and Van Opstal (1998).

Although the multisensory integration properties of deep SC neurons facilitate the detection of cross-modal targets, not all deep SC neurons are multisensory. In the cat, about half of deep SC neurons are multisensory; in the monkey, only about one quarter. Moreover, the per-centage of neurons responsive to visual–auditory stimuli is about twice that of neurons responsive to visual–somatosensory stimulation, whereas about 9% have been found to respond to all three modalities in the cat and 6% in the monkey (Wallace & Stein, 1996).[1] Given that multi-sensory neurons do respond to unimodal input, this raises the issue of why not all SC neurons are multisensory or, at least, develop multisensory response behavior during an organism's maturation? There have been several at-tempts to answer this question, but none of them appears entirely satisfactory as of yet. This article suggests that one possible explanation for the existence of modality-specific SC neurons can be given from a signal detection theory perspective.

As was mentioned above, an organism confronted with the task of discriminating targets from nontargets not only should keep up a high rate of detecting targets but, at the same time, must strive to minimize *false alarms* to

irrelevant nontarget stimuli. This can be achieved following a simple decision rule, the so-called *Bayes' ratio* (BR) rule: Given the evidence from the afferent sensory input, decide that a target is present (rather than a non-target) if the probability for a target is greater than the probability for a nontarget; otherwise, vote for a non-target to be present. Most important, it can be shown that any optimally adapted system maximizing the number of correct decisions (i.e., hits and correct rejections) has to follow this decision rule (cf. Egan, 1975). We postulate that the neural response is proportional to the probability to respond "Yes" computed under the BR decision rule. Although both bimodal (visual–auditory) and unimodal targets exist, the organism is not required to distinguish between different types of targets, but only between targets and nontargets. From these simple assumptions, it is demonstrated that optimal performance with both cross-modal and unimodal targets cannot be accomplished with multisensory SC neurons alone. In fact, it turns out that optimal performance of multisensory neurons vis-à-vis cross-modal stimuli implies, at the same time, that modality-specific neurons will outperform multisensory neurons in processing unimodal targets.

For simplicity, we will consider only the bimodal (visual–auditory) case here, although an extension to three or more modalities can be developed. The present model extends the approach of Colonius and Diederich (2002), in that cross-modal targets are distinguished from unimodal ones, which makes possible an answer to the title question.

The Bayes' Ratio Model of Multisensory Enhancement

The notion that part of the saccadic response duration reflects an on-going decision process in the SC (and other areas, such as the frontal eye fields [FEFs]) has been discussed in many recent contributions (e.g., Krauzlis, Liston, & Carello, 2004). Saccadic reaction times typically lie in the range between 100 and 500 msec (e.g., Muñoz & Schall, 2004). From single-unit studies, it is known that afferent delays to the SC take between 40 and 100 msec, whereas the SC motor command (efferent processing) is carried out within 30–50 msec. "The fact that saccades are not initiated as soon as theoretically possible based on known afferent and efferent delays is presumed to reflect the time necessary to make decisions about where to move the eyes in the event of competing alternatives" (Stanford, 2004, p. 49).

It should be noted that the model proposed here is computational—that is, it purports to describe the decision processes that are being made in generating a saccade—but that its physiological implementation is not at the level of detail found in models that show, for example, how Bayesian probabilities can be computed in a network architecture (Rao, 2004). The BR model is formulated as describing the function of a single deep SC neuron, but an alternative implementation at the level of an entire cell assembly will be considered in the Discussion section.

After introducing some necessary notation, we will present the important concept of the BR and the fundamental decision rule underlying our model.

The Bayes' Ratio Decision Rule

Visual and auditory afferent input to a deep SC neuron are represented by random variables V and A, respectively, taking on integer values 0, 1, 2, ... These values are to be interpreted as the number of neural impulses arriving at the deep SC in a (small) unit interval of time. We first will restrict the development to the unimodal (visual) case.

Let T_{V+} denote the event that a visual target is present, while T_{V-} denotes the occurrence of a (visual) nontarget. Then $P(\text{Yes} \mid T_{V+})$ stands for the conditional probability to decide that a target is present if it is present (*hit*), whereas $P(\text{Yes} \mid T_{V-})$ is the probability to decide that a target is present when in fact only a nontarget is present (*false alarm*). In order to maximize the probability of a correct response,

$$P(\text{C}) = P(\text{Yes} \mid T_{V+}) \, P(T_{V+}) \\ + [1 - P(\text{Yes} \mid T_{V-})] \, P(T_{V-}), \qquad (1)$$

the following decision rule must be adopted in the unimodal visual case (the unimodal decision rule):

If $P(T_{V+} \mid V = v) > P(T_{V-} \mid V = v)$,

then decide "Yes"; otherwise, decide "No."

In the signal detection theory literature, this is known as the max-P(C) rule[2] (cf. Egan, 1975, pp. 20–23). $P(T_{V+} \mid V = v)$ is the (Bayesian) posterior probability that a target is present given afferent input $V = v$. It can be computed via the classic Bayes rule,

$$P(T_{V+} \mid V = v) = \frac{P(V = v \mid T_{V+}) \, P(T_{V+})}{P(V = v)}, \qquad (2)$$

whenever the likelihood $P(V = v \mid T_{V+})$, the a priori target probability $P(T_{V+})$, and the input probability distribution $P(V = v)$ are available.

Interestingly, however, decisions following the above max-P(C) rule can be made without using the Bayes formula (Equation 2) explicitly. In fact, note that the inequality in the max-P(C) rule is equivalent to

$$\frac{P(T_{V+} \mid V = v)}{P(T_{V-} \mid V = v)} = \frac{P(T_{V+})}{P(T_{V-})} \frac{P(V = v \mid T_{V+})}{P(V = v \mid T_{V-})} > 1. \qquad (3)$$

The left-hand side is called *Bayes' ratio* and is equal to the ratio of the a priori probabilities times the likelihood ratio $L(v)$,

$$L(v) = \frac{P(V = v \mid T_{V+})}{P(V = v \mid T_{V-})}.$$

Thus, a Bayesian decision maker can be interpreted as using a form of likelihood ratio testing in which the likelihood ratio is adjusted for the prior probabilities (cf. Wickens, 2002). Under the max-P(C) rule, a "Yes" deci-

sion is made whenever the BR is larger than one. Transforming Equation 3, this rule is equivalent to the following unimodal decision rule:

If $L(v) > \dfrac{P(T_{V-})}{P(T_{V+})}$,

then decide "Yes"; otherwise, decide "No."

Considered as a function of the random variable V, the likelihood ratio $L(V)$ is itself a random variable, and this leads to the conditional probability to decide "Yes," given that a target is present (hit rate):

$$P(\text{Yes} \mid T_{V+}) = P\left(L(V) > \dfrac{P(T_{V-})}{P(T_{V+})} \middle| T_{V+} \right). \qquad (4)$$

Computation of the hit rate requires specification of the probability distributions

$$P(V = v \mid T_{V+})$$

and

$$P(V = v \mid T_{V-}),$$

which we will assume to be Poisson (see below).

The Bayes' Ratio Model

The BR model casts an individual deep SC neuron as a computational unit that processes afferent input and prior target probabilities and that calculates a response in the form of the mean spike rate. Specifically, we postulate that the neural response is proportional to the probability to respond "Yes" computed under the max-P(C) rule. This guarantees a maximum number of correct decisions (i.e., hits and correct rejections) when targets and nontargets are to be discriminated (see Egan, 1975).

Let T_{V+A+}, T_{V+A-}, and T_{V-A+} denote the event that a bimodal target, a unimodal visual target, or a unimodal auditory target, respectively, occurs in the receptive field of a deep SC neuron. At any point in time, one type of target signal, at most, may appear, or the signal may be a nontarget, denoted T_{V-A-}. The prior probabilities of these events, in obvious notation, are π_{++}, π_{+-}, π_{-+}, and π_{--} and have to add up to one.

The probability of a correct response that is being maximized in the BR model, then, is

$$\begin{aligned} P(\text{C}) = \ & P(\text{Yes} \mid T_{V+A+} \text{ or } T_{V+A-} \text{ or } T_{V-A+}) \\ & P(T_{V+A+} \text{ or } T_{V+A-} \text{ or } T_{V-A+}) \\ & + [1 - P(\text{Yes} \mid T_{V-A-})] \, P(T_{V-A-}), \end{aligned} \qquad (5)$$

generalizing Equation 1. Note that for this maximization, there is no requirement to distinguish between different types of targets, but only between targets and nontargets. For the likelihood distributions, we introduce the notation

$$f^{++}(v, a) = P(V = v, A = a \mid T_{V+A+}),$$
$$f^{+-}(v, a) = P(V = v, A = a \mid T_{V+A-}),$$

and so forth. The max-P(C) rule for a multisensory SC neuron (the bimodal decision rule), then, is

If $P(T_{V+A+} \text{ or } T_{V+A-} \text{ or } T_{V-A+} \mid V = v, A = a)$
$> P(T_{V-A-} \mid V = v, A = a)$,

then decide "Yes"; otherwise, decide "No."

By simple algebra and using the above notation, the inequality is equivalent to

$$\begin{aligned} & f_{++}(v, a) \, \pi_{++} + f_{+-}(v, a) \, \pi_{+-} + f_{-+}(v, a) \, \pi_{-+} \\ & > f_{--}(v, a) \, \pi_{--} \end{aligned}$$

or, written as a BR,

$$\frac{\pi_{++}}{\pi_{--}} \frac{f_{++}(v,a)}{f_{--}(v,a)} + \frac{\pi_{+-}}{\pi_{--}} \frac{f_{+-}(v,a)}{f_{--}(v,a)} + \frac{\pi_{-+}}{\pi_{--}} \frac{f_{-+}(v,a)}{f_{--}(v,a)} > 1. \; (6)$$

Extending the unimodal case, this BR, considered as a function of the random variables V and A, is again a random variable—$BR(V,A)$, say—where dependence on the prior probabilities is suppressed to keep the notation simple. The probability of a hit, given a bimodal target, for example, can be computed as

$$P[BR(V,A) > 1 \mid T_{V+A+}],$$

whereas the probability of a false alarm would be

$$P[BR(V,A) > 1 \mid T_{V-A-}].$$

Similarly, the max-P(C) rule for a modality-specific visual SC neuron (unimodal decision rule) is

If $P(T_{V+A+} \text{ or } T_{V+A-} \mid V = v)$
$> P(T_{V-A+} \text{ or } T_{V-A-} \mid V = v)$,

then decide "Yes"; otherwise, decide "No,"

since the auditory sensory input, $A = a$, is lacking. The corresponding BR is

$$\frac{\pi_{++} g_{++}(v) + \pi_{+-} g_{+-}(v)}{\pi_{-+} g_{-+}(v) + \pi_{--} g_{--}(v)} > 1, \qquad (7)$$

where $g_{++}(v), g_{+-}(v)$, and so forth refer to the likelihood distributions $P(V = v \mid T_{V+A+})$, $P(V = v \mid T_{V+A-})$, and so forth. For a modality-specific visual SC neuron, the probability of a hit, given a bimodal target, and of a false alarm, for example, are computed analogously from

$$P[BR(V) > 1 \mid T_{V+A+}]$$

and

$$P[BR(V) > 1 \mid T_{V-A+} \text{ or } T_{V-A-}],$$

respectively, where $BR(V)$ is the BR corresponding to Equation 7. Actual computation of these probabilities, however, requires specification of the likelihood distributions. We will consider only the Poisson case here.

The Poisson Bayes' Ratio Model

Unimodal stimulation case. The most common distribution to account for the random number of spikes oc-

curring in a fixed unit time interval is the *Poisson* distribution with a mean of $\lambda_+ > 0$,

$$P(V = v \mid T_{V+}) = \frac{\exp(-\lambda_+)\lambda_+^v}{v!}, \qquad (8)$$

and it is often seen as a reasonable first approximation (Tuckwell, 1989). For $P(V = v \mid T_{V-})$, we will assume a Poisson distribution as well, with a mean λ_- and $\lambda_+ > \lambda_-$. The hit rate in Equation 4 then equals

$$P(\text{Yes} \mid T_{V+}) = P\left[\exp(\lambda_- - \lambda_+)\left(\frac{\lambda_+}{\lambda_-}\right)^V > \frac{P(T_{V-})}{P(T_{V+})}\Bigg| T_{V+}\right]$$

$$= P(V > c \mid T_{V+}),$$

where

$$c = \frac{\ln\left[\dfrac{P(T_{V-})}{P(T_{V+})}\right] + \lambda_+ - \lambda_-}{\ln\left(\dfrac{\lambda_+}{\lambda_-}\right)}.$$

Thus, in this simple case, the hit rate can be computed directly from the Poisson distribution.

Bimodal stimulation case. For the multisensory SC neuron, we have to specify the bivariate Poisson distribution for visual–auditory input (V, A) for each target–nontarget condition. Although many dependent bivariate Poisson models exist (e.g., Kocherlakota & Kocherlakota, 1992), we will confine discussion here to the independent Poisson case. Thus,

$$P(V = v, A = a \mid T_{V+A+}) \equiv f_{++}(v, a)$$

$$= \frac{\exp(-\lambda_+)\lambda_+^v}{v!} \frac{\exp(-\mu_+)\mu_+^a}{a!}, \quad (9)$$

and

$$P(V = v, A = a \mid T_{V+A-}) \equiv f_{+-}(v, a)$$

$$= \frac{\exp(-\lambda_+)\lambda_+^v}{v!} \frac{\exp(-\mu_-)\mu_-^a}{a!}, \quad (10)$$

and so forth, where μ_+ and μ_- are the auditory target and nontarget means, $\lambda_+ > \lambda_-$, and $\mu_+ > \mu_-$. This *conditional stochastic independence* assumption implies that

the visibility of a target indicates nothing about its audibility and vice versa.

For modality-specific visual SC neurons we assume that

$$\left.\begin{array}{l} P(V = v \mid T_{V+A+}) \equiv g_{++}(v) \\ P(V = v \mid T_{V+A-}) \equiv g_{+-}(v) \end{array}\right\} = \frac{\exp(-\lambda_+)\lambda_+^v}{v!} \quad (11)$$

and

$$\left.\begin{array}{l} P(V = v \mid T_{V-A+}) \equiv g_{-+}(v) \\ P(V = v \mid T_{V-A-}) \equiv g_{--}(v) \end{array}\right\} = \frac{\exp(-\lambda_-)\lambda_-^v}{v!}. \quad (12)$$

The likelihood functions for modality-specific auditory SC neurons are defined analogously. Note that these definitions contain two more implicit assumptions. First, the fact that the mean value parameters (the λs and μs) are the same for multisensory and modality-specific SC neurons reflects the premise that under unimodal stimulation, multisensory neurons receive the same afferent input as the modality-specific neurons. Note that it is not assumed that the mean values are the same throughout the entire SC. There is a large variability among real neurons in their response characteristics, and the specific parameter value for λ or μ is ascribed to a subpopulation of SC neurons tuned to the same level of afferent input (see also the Discussion section). Second, it is assumed that for the modality-specific visual SC neurons, the likelihood is the same whether or not an auditory target is present and vice versa. This *context independence* assumption seems very innocuous because the modality-specific visual SC neuron has no auditory input channel, by definition.

It is now straightforward to write down the probability for a "Yes" response under the various target–nontarget conditions by inserting the likelihood functions into the BRs. For a multisensory neuron, the hit probability for a bimodal target, for example, is shown at the bottom of the page in Equation 13, where the bivariate random variable (V, A) has distribution $f_{++}(v, a)$ from Equation 9. Similarly, for a modality-specific visual SC neuron, the hit probability for a bimodal target, for example, is shown at the bottom of the page in Equation 14, where random variable V has distribution $g_{++}(v)$ from Equation 11. A derivation of the explicit distribution of the BR seems difficult in the bimodal case, however, so we obtained various hit and false alarm rates through numerical simulation.

$$P[BR(V, A) > 1 \mid T_{V+A+}] =$$

$$P\left[\frac{\pi_{++}}{\pi_{--}}\exp(\lambda_- - \lambda_+)\left(\frac{\lambda_+}{\lambda_-}\right)^V \exp(\mu_- - \mu_+)\left(\frac{\mu_+}{\mu_-}\right)^A + \frac{\pi_{+-}}{\pi_{--}}\exp(\lambda_- - \lambda_+)\left(\frac{\lambda_+}{\lambda_-}\right)^V + \frac{\pi_{-+}}{\pi_{--}}\exp(\mu_- - \mu_+)\left(\frac{\mu_+}{\mu_-}\right)^A > 1 \mid T_{V+A+}\right] \quad (13)$$

$$P[BR(V) > 1 \mid T_{V+A+}] = P\left[\frac{\pi_{++} + \pi_{+-}}{\pi_{-+} + \pi_{--}}\exp(\lambda_- - \lambda_+)\left(\frac{\lambda_+}{\lambda_-}\right)^V > 1 \mid T_{V+A+}\right] \quad (14)$$

Comparing Bimodal With Unimodal Performance

The performance of multisensory and modality-specific deep SC neurons under the Poisson BR model is completely characterized by (1) the prior probabilities and (2) the Poisson parameters. Each data point was obtained by sampling pseudorandom numbers (or pairs of numbers) from Poisson distributions and inserting the values into Equations 13 or 14, respectively. The relative frequency with which the inequality was satisfied was taken as an estimate of the corresponding probability.[3]

It should be mentioned that in the formulation of the BR decision rule used here, the costs for false alarms and correct rejections have been set equal. An extension to unequal weighting could easily be established, if suggested by the decision situation. The results most pertinent to the question raised by the title of this article come from comparing the hit and false alarm rates between modality-specific and multisensory neurons when only a unimodal target is present. In this case, modality-specific neurons can be shown to exhibit performance superior to multisensory neurons under various conditions. This *unimodal neuron superiority* depends, in particular, on the difference between the detectabilities of the unimodal (visual and auditory) targets and on the prior probability of bimodal targets. We will consider both factors in turn.

Unimodal detectability in the Poisson case is defined similarly to the familiar d' measure of signal detection theory (cf. Egan, 1975):

$$D_V(\lambda_+, \lambda_-) = \frac{\lambda_+ - \lambda_-}{\left(\lambda_+ * \lambda_-\right)^{1/4}} \qquad (15)$$

for visual targets. $D_A(\mu_+, \mu_-)$ for auditory targets is defined analogously. The first case study will illustrate that the hit rate of a multisensory neuron does not depend on the absolute values of the means but, rather, on the detectabilities of the unimodal targets. Specifically, it will suggest that unimodal neuron superiority shows up when a unimodal target occurs in the modality of lower detectability.

Case Study 1. Setting $\pi_{++} = .45$, $\pi_{--} = .5$, $\pi_{+-} = \pi_{-+} = .025$, $\lambda_+ = 9$, and $\lambda_- = 5$, we varied both μ_+ and μ_- such that $D_A(\mu_+, \mu_-)$ was constant and equal to 3.11. Note that visual detectability remained constant: $D_V(9,5) = 1.54$. For a large range of μ_+ and μ_- values, the hit rate of a multisensory neuron for a unimodal visual target, $P[BR(V,A) > 1 \mid T_{V+A-}]$, was about constant at a value of .35 (see Figure 1).

The hit rate for a modality-specific visual neuron for the same parameter values, $P[BR(V) > 1 \mid T_{V+A-}]$, was equal to .40, demonstrating some degree of unimodal neuron superiority.


Figure 1. Curve of constant auditory detectability; for parameter values, see the text. The numbers at the arrows indicate the nearly constant hit rate of a multisensory neuron for a unimodal visual target for different (μ_-, μ_+) combinations.

Table 1
Hit and False Alarm Rates of a Multisensory Superior Colliculus (SC)
Neuron ($B_{V+A-}, B_{V+A+}, B_{V-A+}, B_{V-A-}$) as a Function of
Auditory Target Intensity μ_+

μ_+	$U_{V+} = .40$			$U_{V-} = .24$	$U_{V+} = .20$			$U_{V-} = .007$
	B_{V+A-}	B_{V+A+}	B_{V-A+}	B_{V-A-}	B_{V+A-}	B_{V+A+}	B_{V-A+}	B_{V-A-}
20	.23	.99	.99	.01	.14	.99	.99	.004
19	.24	.99	.99	.01	.13	.99	.99	.004
18	.25	.99	.99	.02	.15	.99	.99	.005
17	.28	.99	.99	.02	.16	.99	.97	.007
16	.30	.99	.99	.03	.16	.98	.96	.006
15	.32	.99	.99	.04	.16	.97	.93	.007
14	.35	.99	.97	.05	.16	.95	.88	.008
13	.41	.99	.94	.07	.17	.93	.82	.009
12	.45	.98	.88	.08	.19	.89	.73	.011
11	.50	.97	.79	.09	.21	.82	.58	.012
10	.58	.95	.68	.12	.22	.72	.41	.012
9	.66	.93	.52	.15	.26	.62	.24	.011
8	.71	.89	.35	.16	.27	.49	.09	.010
7	.76	.86	.24	.18	.27	.40	.02	.008
6	.80	.85	.12	.20	.28	.33	.01	.007
5	.85	.84	.05	.24	.32	.32	.01	.009

Note—The top row contains hit and false alarm rates for a visual SC neuron. The left part of the table presents data for prior values $\pi_{++} = .45, \pi_{--} = .5$, and $\pi_{+-} = \pi_{-+} = .025$; the right part presents data for $\pi_{++} = .05, \pi_{--} = .9$, and $\pi_{+-} = \pi_{-+} = .025$. Mean values were set to $\lambda_+ = 9, \lambda_- = 5$, and $\mu_- = 5$.

In the next case study, we will investigate how unimodal neuron superiority depends on the difference between unimodal target detectabilities. Note that *unimodal neuron superiority* here refers to a higher hit rate (and/or a lower false alarm rate) of a visual modality-specific neuron, as compared with a multisensory neuron, when a visual target and an auditory nontarget are presented. Interchanging the roles of the visual and the auditory modalities would lead to analogous results.

Case Study 2. Two different sets of priors are considered: the first, as before, with target probabilities equal to nontarget probabilities, and the second with rather small target probabilities (.05 for bimodal and .025 for unimodal targets). Visual detectability is kept constant at $D_V = 1.54$, with $\lambda_+ = 9$ and $\lambda_- = 5$, as before. Auditory detectability is varied from zero (last row of Table 1) to $D_A = 4.74$ (first row of Table 1) by increasing μ_+ from 5 to 20, with $\mu_- = 5$ throughout.

First, consider the case of relatively high target probability ($\pi_{++} = .45, \pi_{--} = .5, \pi_{+-} = \pi_{-+} = .025$). The hit rate for a modality-specific visual neuron, U_{V+}, equals .40, as above. Increasing auditory detectability depresses the multisensory SC neuron hit rate for visual targets, B_{V+A-} (Table 1, second column), down to a value of .23, thus exhibiting unimodal neuron superiority. For decreasing auditory detectability, unimodal superiority quickly vanishes and transforms into a clear advantage for the multisensory neuron (.85 vs. .40).

On the other hand, the multisensory SC neuron hit rate for unimodal auditory targets, B_{V-A+} (fourth column), shows the opposite behavior: It is large when B_{V+A-} is low, and vice versa. Note that, at the same time, the hit rate for the bimodal target B_{V+A+} (third column) decreases from .99 to .84 when auditory detectability diminishes. The fifth column (B_{V-A-}) contains the corresponding false alarm rates for the multisensory neuron. It starts at .24 when auditory detectability is zero and goes down to .01 for the highest auditory detectability, whereas the false alarm rate for a modality-specific visual neuron with the same (visual) parameter values equals $U_{V-} = .24$. Thus, there is a multisensory neuron superiority for the false alarm rates.

Columns 6–9 of Table 1 refer to the case of low target probabilities ($\pi_{++} = .05, \pi_{--} = .9, \pi_{+-} = \pi_{-+} = .025$). The modality-specific visual neuron hit rate is 50% lower now ($U_{V+} = .20$). The corresponding bimodal hit rate for visual targets, B_{V+A-}, is at .32 for zero auditory detectability (last row) and decreases to .14 with increasing auditory detectability. Thus, unimodal neuron superiority occurs again in line with the case of high target probabilities. Behavior of multisensory SC neuron hit rates for bimodal targets (B_{V+A+}, column 7) and for auditory targets (B_{V-A+}, column 8) parallels the previous case. On the other hand, as a consequence of the small prior target probabilities, the false alarm rate for the modality-specific visual neuron is very low (.004), and the corresponding bimodal false alarm rates do not exceed .01 for most levels of auditory detectability.[4]

The following picture emerges from this example. For the multisensory SC neuron, the hit rate for bimodal targets is high as long as visual or auditory detectability is not too low, but the exact level depends on the prior probabilities as well. Most important, its hit rate for a unimodal (say, visual) target gets severely depressed and becomes lower than the modality-specific neuron hit rate if visual and auditory unimodal detectabilities differ substantially and, simultaneously, the target happens to occur in the modality with lower detectability (the visual, in

Table 2
Hit and False Alarm Rates of Multisensory Neuron (B_{V+A-},
B_{V-A-}) and Modality-Specific Neuron (U_{V+A-}, U_{V-}) to
Unimodal Visual Target, Under Various Sets of Prior
Probabilities π_{++}, π_{+-}, π_{-+}, and π_{--}

π_{++}	π_{+-}	π_{-+}	π_{--}	B_{V+A-}	U_{V+}	B_{V-A-}	U_{V-}
.009	.0005	.0005	.99	.01	.04	.001	<.001
.001	.0045	.0045	.99	.024	.02	.001	<.001
.002	.014	.014	.97	.05	.07	.003	.001
.026	.002	.002	.97	.03	.12	.002	.003
.026	.001	.003	.97	.025	.12	.003	.002
.026	.003	.001	.97	.036	.125	.002	.002
.01	.01	.01	.97	.05	.07	.001	<.001
.03	.01	.01	.95	.08	.13	.004	.002
.001	.0245	.0245	.95	.08	.07	.003	<.001
.01	.04	.04	.91	.15	.20	.01	.006
.05	.025	.025	.90	.16	.20	.008	.007
.20	.025	.025	.75	.22	.53	.02	.07
.05	.10	.10	.75	.35	.41	.035	.031
.35	.025	.025	.60	.29	.69	.035	.13
.40	.025	.025	.55	.32	.67	.042	.13
.40	.09	.01	.50	.53	.79	.07	.24
.40	.01	.09	.50	.27	.67	.055	.14
.45	.025	.025	.50	.37	.80	.05	.24
.25	.125	.125	.50	.57	.68	.09	.13
.75	.025	.025	.20	.58	.94	.13	.56

this case). Thus, the BR model predicts that, if unimodal visual and auditory detectabilities differ in a cross-modal stimulus combination, the multisensory SC neuron's responses to unimodal stimulation should also differ: High auditory detectability leads to a low visual hit rate and a high auditory hit rate, and vice versa. There is a clear unimodal neuron superiority effect for the hit rates. For the false alarm rates, the effect depends on the prior probabilities for the targets. For relatively high target probabilities (left part of Table 1), the depression of the multisensory neuron hit rate for visual targets goes along with a decrease of its false alarm rate, whereas for low prior target probabilities (right part of Table 1), the false alarm rates stay below 1% under all conditions.

Case Study 3. In this study, the effect of different sets of prior target probabilities is investigated further. The rows in Table 2 are arranged by nontarget probabilities (π_{--}, fourth column) decreasing from .99 to .2. Fixing the mean values and, thereby, visual and auditory detectability at $\lambda_+ = 9$, $\mu_+ = 14$, $\lambda_- = \mu_- = 5$, we computed the probability to detect a unimodal visual target both for a multisensory (fifth column) and for a modality-specific visual (sixth column) SC neuron and the corresponding false alarm rates (columns 7 and 8) for each combination of prior probabilities.

Clearly, the priors have a strong effect on both hit rates: The latter increase when nontarget prior probability decreases, as one would expect, and there is clear unimodal neuron superiority in the hit rate for most sets of priors. For the false alarm rates, the effects are less conspicuous. In particular, when nontarget probability is above .90, the false alarm rates stay below 1% for both types of neurons, and there is a slight tendency, if any, for smaller false alarm rates in the modality-specific neuron. This reverses into a clear advantage for the multisensory

neuron only when nontarget probability falls below 50%. Moreover, for constant nontarget prior probability π_{--}, the unimodal neuron superiority (with respect to the hit rates) tends to be larger when the prior bimodal target probability π_{++} exceeds the unimodal priors π_{+-} and π_{-+}. Finally, when the unimodal priors differ, the multisensory SC neuron detection rate is affected much more than the modality-specific SC neuron detection rate if the less likely unimodal target occurs. For example, with $\pi_{++} = .40$, the bimodal rate goes down from .53 to .27 if visual prior probability drops from .9 to .01, whereas the unimodal visual rate only drops from .40 to .34. This is in line with the observation, in the previous case study, that the performance of the multisensory neuron is affected more by differences between the unimodal conditions (priors or detectabilities).

Discussion

The question why some, but not all, deep SC neurons are multisensory is a legitimate one, given that multisensory neurons are also responsive to unimodal stimulation. One could argue that multisensory neurons are especially important when weak stimuli from several modalities co-occur, because multisensory response enhancement measured at the single-cell level is most dramatic under these conditions (*inverse effectiveness,* cf. Stein & Meredith, 1993), and that modality-specific neurons process above-threshold unimodal stimuli mostly. Note, however, that although this line of reasoning supports the existence of multisensory neurons, it does not explain why a large number of neurons in an organism remain modality specific over their lifetime. Moreover, the hypothesis that cell type (multisensory vs. modality specific) implies the tuning to specific stimulus intensities is not easy to reconcile with the observation that there is large variability in the response characteristics of the entire population of deep SC cells.

The present investigation suggests a more specific solution to the issue from a decision theoretic point of view—that is, the BR decision rule, which maximizes the percentage of correct decisions when targets are to be discriminated from nontargets. Taking for granted that deep SC neurons are, in that sense, optimally adapted to discriminate between targets (stimuli that need immediate orientation or attention) and nontargets, we were able to show that optimal performance of multisensory neurons vis-à-vis cross-modal stimuli implies, at the same time, that modality-specific neurons will outperform multisensory neurons in detecting unimodal targets. Specifically, if an auditory stimulus is better detectable than a visual, the performance of the multisensory neuron vis-à-vis the unimodal visual stimulus is severely degraded even though its performance with respect to auditory, or bimodal, stimuli can still be very high. In the terminology of decision thresholds, if the multisensory neuron weights the auditory input highly (since it is very detectable), it is difficult to exceed the decision threshold if the auditory input happens to be low. Note that this behavior of multisensory neurons is not due to some par-

ticular assumptions but simply follows from the BR decision rule. Thus, the existence of modality-specific visual neurons is quite essential and, interchanging the roles of the visual and the auditory modalities, the same holds true for modality-specific auditory neurons.

Neurophysiological implementation. Although the BR model is developed at the computational level without specifying the exact neural underpinnings, it is important to point out the neural structures that are supposed be involved in the computational process. Application of the BR rule requires the computation of the probability that the likelihood ratio is larger than the ratio of the prior probabilities of a target and a nontarget, given the afferent input (see Equation 4). For a given afferent input (a realization v of the random variable V, say), the SC neuron has to compute the likelihood ratio $L(v)$ and to determine whether or not it exceeds the ratio of priors. The former would lead to a "Yes" response and the latter to a "No" response by the neuron. In order to generate the final output—that is, a neural response proportional to the probability to respond "Yes" computed under the BR rule—the SC neuron must average its response over small unit time intervals. Note that the suggestion of conceiving of the firing rate of an SC neuron as an estimate of the (log) likelihood that a target is present has been made previously by Carpenter and Williams (1995).

An implementation of the BR model at the level of an entire cell assembly is possible. We define a *cell assembly* as a subset of multisensory and modality-specific deep SC neurons tuned to the same intensity level—that is, to certain fixed means of afferent bimodal input λ and μ. Each cell element of the assembly delivers a "Yes" or "No" vote based on the outcome of the likelihood ratio computation, as before. The final computational step of estimating the probability to respond "Yes," however, will now take place at the level of the cell assembly by evaluating the relative frequency of "Yes" responses in the cell assembly. Note that this way, the unimodal advantage is preserved at the assembly level, because an estimate of the probability of a "Yes" response, given a unimodal (visual) target, will be better in the presence of modality-specific neurons with their higher relative frequency of hits, rather than being based on the multisensory cells' hit rate only.

It is an open question at which neural site the computation of the relative frequencies would be accomplished. Several alternatives exist. Using a two-alternative forced choice task, Krauzlis and Dill (2002) have reported that build-up neurons in the rostral SC exhibit higher activity for target than for nontarget stimuli. In a similar vein, Horwitz and Newsome (2001) have reported a neural correlate of the process of discriminating the direction of motion in the activity of SC neurons. "Rather than discharge only after a target had been selected, these neurons appeared to participate in, or at least reflect, the process of accumulating sensory evidence for this process" (Stanford, 2004, p. 49). Alternatively, the FEFs may be involved as well. The FEF projects to the SC, which projects to the thalamus, which then projects back to the FEF. The transmis-

sion time of these pathways is 2–3 msec, allowing enough time for the SC and the FEF to mutually influence one another's state (see Muñoz & Schall, 2004, for references).

Finally, multisensory integration properties of most SC neurons, as well as observed orientation behavior, are mediated by influences from two parietal cortical areas, the anterior ectosylvian sulcus (AES) and the rostral aspect of the lateral suprasylvian sulcus (rLS; Jiang, Jiang, & Stein, 2002; Jiang, Wallace, Jiang, Vaughan, & Stein, 2001). When cortical activation is eliminated, multisensory SC neurons lose their enhancement facility to cross-modal combinations, and their responses parallel that for unimodal input. "Whether or not activity in SC [during decision time] is essential or simply reflects the state of computations carried out in cortex is difficult to know. . . . In principle, an experiment that combines the right behavioral task, microstimulation parameters, and microstimulation timing might distinguish the SC's respective contributions to sensorimotor decision-making and motor command formation" (Stanford, 2004, p. 49). At least as far as multisensory integration is concerned, it has been shown through ortho- and antidromic stimulation that the SC is a site of multisensory integration and not only reflects multisensory integration processes in cortical (AES) multisensory neurons (Wallace, Meredith, & Stein, 1993).

Empirical testability. The last points raise the question as to what degree the BR model may be subjected to empirical testing. It should be possible to condition an animal to an environment in which priors are set by the experimenter at the start of the recording session. When these well-defined priors are changed by modifying the environment in a controlled way, the brain should be able to adapt to these changes, and this would allow us to test predictions of the model, as is exemplified in Table 2. Similarly, it may be possible to estimate the detectability of various uni- and cross-modal stimuli by recording the afferent input. Although thousands of neurons are necessary to produce a saccade, the averaged signal from single neurons in the SC and the FEF has been shown to be sufficient to specify whether and when a saccade will occur (Muñoz & Schall, 2004, p. 74). Thus, it might also be possible, in principle, to test predictions for hit and false alarm rates, as is exemplified in Table 1.

Intriguingly, recent studies of multisensory effects under different types of lesions may offer an additional testing opportunity for the unimodal advantage prediction of the BR model even at the behavioral level. Jiang et al. (2002) showed that cryogenic blockade of the AES or the rLS disrupted the multisensory enhancement of orientation responses to visual–auditory stimuli in cats. These data were very similar to the effects of AES or rLS deactivation on individual multisensory SC neurons (Jiang et al., 2001). Under cortical deactivation, the responses of these neurons to cross-modal stimuli were indistinguishable from those of modality-specific neurons. Thus, if the effect of deactivation is functionally equivalent to an increase of the proportion of modality-specific SC neurons, the unimodal advantage should show up in the ori-

enting behavior of the cats under unimodal auditory stimulation. Unfortunately, there are two obstacles to an immediate test with the data sets currently available. First, the cats have been trained to consider only visual stimuli as targets, and second, the intensity levels of the stimuli have been at about the same (low) level. Nevertheless, the responses of the individual multisensory SC neurons under cryogenic cortical blockade would present a case in point.[5] In a similar vein, studies such as the one by Burnett, Stein, Chaponis, and Wallace (2004), where excitotoxic lesions of the SC presumably had a preferential impact on multisensory neurons, thus modifying the relative proportion of modality-specific neurons, might be used to test for a unimodal advantage in the visual–auditory orientation task.

Concluding remarks. Building a model from decision theoretic principles is certainly superior to a purely descriptive approach, and recently, there has been a strong tendency in neural modeling to appeal to Bayesian and related rules of decision making (e.g., Ernst & Banks, 2002; Glimcher, 2003). In Anastasio, Patton, and Belkacem-Boussaid (2000), it is proposed that multisensory deep SC neurons compute the posterior probability that a target is present given stochastic afferent input. In an information theoretic analysis of that model, Patton, Belkacem-Boussaid, and Anastasio (2002) showed that input of an additional modality may indeed increase target information, but only if the input received from the initial modality does not completely reduce uncertainty about the presence of a target. There are some obvious correspondences between this approach and the BR model presented here. An important difference is the explicit introduction of a signal detection theoretic perspective in the BR model, the postulate of an optimal decision rule for hits and false alarms, and the distinction between uni- and cross-modal targets that lead to specific predictions concerning unimodal neuron superiority. In a more recent development, Anastasio and Patton (2003) suggested a neural network model with a two-stage unsupervised learning algorithm that produces multisensory enhancement. Interestingly, in their information theoretic analysis of the model, the authors were able to show that target information gain is highest when the SC contains between 10% and 50% multisensory units only. Although starting from somewhat different perspectives, this coincides with the conclusion of the BR model that modality-specific SC neurons are a necessary part of an optimally adapted organism.

REFERENCES

Anastasio, T. J., & Patton, P. E. (2003). A two-stage unsupervised learning algorithm reproduces multisensory enhancement in a neural network model of the corticotectal system. *Journal of Neuroscience*, **23**, 6713-6727.

Anastasio, T. J., Patton, P. E., & Belkacem-Boussaid, K. (2000). Using Bayes' rule to model multisensory enhancement in the superior colliculus. *Neural Computation*, **12**, 1165-1187.

Bell, A. H., Corneil, B. D., Meredith, M. A., & Muñoz, D. P. (2001). The influence of stimulus properties on multisensory processing in the awake primate superior colliculus. *Canadian Journal of Experimental Psychology*, **55**, 123-132.

Burnett, L. R., Stein, B. E., Chaponis, D., & Wallace, M. T. (2004). Superior colliculus lesions preferentially disrupt multisensory orientation. *Neuroscience*, **124**, 535-547.

Carpenter, R. H., & Williams, M. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, **377**, 59-62.

Colonius, H., & Diederich, A. (2002). A maximum-likelihood approach to modeling multisensory enhancement. In T. G. Dietterich, S. Becker, & Z. Ghahramani (Eds.), *Advances in neural information processing systems 14* (pp. 181-187). Cambridge, MA: MIT Press.

Colonius, H., & Diederich, A. (2004). Multisensory interaction in saccadic reaction time: A time-window-of-integration model. *Journal of Cognitive Neuroscience*, **16**, 1000-1009.

Diederich, A., & Colonius, H. (2004). Modeling the time course of multisensory interaction in manual and saccadic responses. In G. Calvert, C. Spence, & B. E. Stein (Eds.), *Handbook of multisensory processes* (pp. 395-408). Cambridge, MA: MIT Press.

Egan, J. P. (1975). *Signal detection theory and ROC analysis*. New York: Academic Press.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, **415**, 429-433.

Frens, M. A., & Van Opstal, A. J. (1998). Visual–auditory interactions modulate saccade-related activity in monkey superior colliculus. *Brain Research Bulletin*, **46**, 211-224.

Frens, M. A., Van Opstal, A. J., & Van der Willigen, R. F. (1995). Spatial and temporal factors determine auditory–visual interactions in human saccadic eye movements. *Perception & Psychophysics*, **57**, 802-816.

Glimcher, P. W. (2003). *Decisions, uncertainty, and the brain: The science of neuroeconomics*. Cambridge, MA: MIT Press.

Horwitz, G. D., & Newsome, W. T. (2001). Target selection for saccadic eye movements: Prelude activity in the superior colliculus during a direction-discrimination task. *Journal of Neurophysiology*, **86**, 2543-2558.

Jiang, W., Jiang, H., & Stein, B. E. (2002). Two corticotectal areas facilitate orientation behavior. *Journal of Cognitive Neuroscience*, **14**, 1240-1255.

Jiang, W., Wallace, M. T., Jiang, H., Vaughan, J. W., & Stein, B. E. (2001). Two cortical areas mediate multisensory integration in superior colliculus neurons. *Journal of Neurophysiology*, **85**, 506-522.

Kocherlakota, S., & Kocherlakota, K. (1992). *Bivariate discrete distributions*. New York: Dekker.

Krauzlis, R. J., & Dill, N. (2002). Neural correlates of target choice for pursuit and saccades in the primate superior colliculus. *Neuron*, **35**, 355-363.

Krauzlis, R. J., Liston, D., & Carello, C. D. (2004). Target selection and the superior colliculus: Goals, choices and hypotheses. *Vision Research*, **44**, 1445-1451.

Meredith, M. A., & Stein, B. E. (1986a). Spatial factors determine the activity of multisensory neurons in cat superior colliculus. *Brain Research*, **365**, 350-354.

Meredith, M. A., & Stein, B. E. (1986b). Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration. *Journal of Neurophysiology*, **56**, 640-662.

Muñoz, D. P., & Schall, J. D. (2004). Concurrent, distributed control of saccade initiation in the frontal eye field and superior colliculus. In W. C. Hall & A. Moschovakis (Eds.), *The superior colliculus: New approaches for studying sensorimotor integration* (pp. 55-82). Boca Raton, FL: CRC Press.

Patton, P., Belkacem-Boussaid, K., & Anastasio, T. J. (2002). Multimodality in the superior colliculus: An information theoretic analysis. *Cognitive Brain Research*, **14**, 10-19.

Rao, R. P. N. (2004). Bayesian computation in recurrent neural circuits. *Neural Computation*, **16**, 1-38.

Rowe, D. B. (2003). *Multivariate Bayesian statistics*. Boca Raton, FL: CRC Press.

Stanford, T. R. (2004). Signal coding in the primate superior colliculus revealed through the use of artificial signals. In W. C. Hall & A. Moschovakis (Eds.), *The superior colliculus: New approaches for studying sensorimotor integration* (pp. 35-53). Boca Raton, FL: CRC Press.

Stein, B. E., Magalhães-Castro, B., & Kruger, L. (1976). Relationship between visual and tactile representations in cat superior colliculus. *Journal of Neurophysiology*, **39**, 401-419.

STEIN, B. E., & MEREDITH, M. A. (1993). *The merging of the senses*. Cambridge, MA: MIT Press.

TUCKWELL, H. C. (1989). *Stochastic processes in the neurosciences*. Philadelphia: Society for Industrial and Applied Mathematics.

WALLACE, M. T., MEREDITH, M. A., & STEIN, B. E. (1993). Converging influences from visual, auditory, and somatosensory cortices onto output neurons of the superior colliculus. *Journal of Neurophysiology*, **69**, 1797-1809.

WALLACE, M. T., MEREDITH, M. A., & STEIN, B. E. (1998). Multisensory integration in the superior colliculus of the alert cat. *Journal of Neurophysiology*, **80**, 1006-1010.

WALLACE, M. T., & STEIN, B. E. (1996). Sensory organization of the superior colliculus in cat and monkey. *Progress in Brain Research*, **112**, 301-311.

WALLACE, M. T., WILKINSON, L. K., & STEIN, B. E. (1996). Representation and integration of multiple sensory inputs in primate superior colliculus. *Journal of Neurophysiology*, **76**, 1246-1266.

WICKENS, T. D. (2002). *Elementary signal detection theory*. New York: Oxford University Press.

NOTES

1. Such estimates of proportions of multisensory cells are always fairly rough, because it is difficult to exclude the possibility that an experimenter may simply be using the wrong type of stimuli.

2. It is a special case of the *maximum a posteriori* (MAP) rule (see, e.g., Rowe, 2003).

3. The Mathematica programming system was used for the simulations. A sample size of $N = 5,000$ yielded probability estimates sufficient for a precision of about 1%. For false alarm rates, sample size was increased to $N = 10,000$ for greater precision.

4. Given that the accuracy of our simulations is not smaller than 1%, the values observed do not differ from each other significantly.

5. Figure 3 in Jiang et al. (2001) even suggested a unimodal auditory advantage of the SC neuron under cortical deactivation.