Behavioral/Systems/Cognitive

A Two-Stage Unsupervised Learning Algorithm Reproduces Multisensory Enhancement in a Neural Network Model of the Corticotectal System

Thomas J. Anastasio[1,2] **and Paul E. Patton**[2]

[1]Department of Molecular and Integrative Physiology and [2]Beckman Institute, University of Illinois at Urbana/Champaign, Urbana, Illinois 61801

Multisensory enhancement (MSE) is the augmentation of the response to sensory stimulation of one modality by stimulation of a different modality. It has been described for multisensory neurons in the deep superior colliculus (DSC) of mammals, which function to detect, and direct orienting movements toward, the sources of stimulation (targets). MSE would seem to improve the ability of DSC neurons to detect targets, but many mammalian DSC neurons are unimodal. MSE requires descending input to DSC from certain regions of parietal cortex. Paradoxically, the descending projections necessary for MSE originate from unimodal cortical neurons. MSE, and the puzzling findings associated with it, can be simulated using a model of the corticotectal system. In the model, a network of DSC units receives primary sensory input that can be augmented by modulatory cortical input. Connection weights from primary and modulatory inputs are trained in stages one (Hebb) and two (Hebb–anti-Hebb), respectively, of an unsupervised two-stage algorithm. Two-stage training causes DSC units to extract information concerning simulated targets from their inputs. It also causes the DSC to develop a mixture of unimodal and multisensory units. The percentage of DSC multisensory units is determined by the proportion of cross-modal targets and by primary input ambiguity. Multisensory DSC units develop MSE, which depends on unimodal modulatory connections. Removal of the modulatory influence greatly reduces MSE but has little effect on DSC unit responses to stimuli of a single modality. The correspondence between model and data suggests that two-stage training captures important features of self-organization in the real corticotectal system.

Key words: superior colliculus; multisensory integration; unsupervised learning; corticotectal system; neural network model; self-organization; Hebbian learning; anti-Hebbian learning

Introduction

Integration of input from multiple senses is critical to survival in a complex environment. Research in sensory neurobiology is shifting from a focus on single sensory systems to consideration of interactions between sensory systems (Stein and Meredith, 1993; Stein, 1998). The most studied loci of multisensory convergence in mammals are the deep layers of the superior colliculus (DSC). Findings on DSC neurons raise intriguing questions about multisensory integration.

The DSC functions to detect sensory targets and initiates orienting movements toward them (Robinson, 1972; Wurtz and Goldberg, 1972; Sparks and Hartwich-Young, 1989). DSC neurons are organized topographically according to the location of their receptive fields (Middlebrooks and Knudsen, 1984; Meredith and Stein, 1990; Meredith et al., 1991). Many DSC neurons receive sensory input of multiple modalities (Wallace and Stein, 1996), and the receptive fields of the same neuron for different modalities overlap (Meredith and Stein, 1986b, 1996; Kadunce et al., 2001). Multisensory DSC neurons can exhibit multisensory enhancement (MSE), which is the augmentation of the response to stimulation of one modality by stimulation of a different modality (King and Palmer, 1985; Meredith and Stein, 1986a; Wallace et al., 1996, 1998).

The DSC receives visual, auditory, and somatosensory input from a variety of subcortical and extraprimary cortical sources (Sparks and Hartwich-Young, 1989; Wallace et al., 1993). MSE depends on input from specific regions of parietal cortex. Inactivation of these regions reduces MSE, while often only minimally affecting the responses of DSC neurons to stimulation of a single modality (Wallace and Stein, 1994; Jiang et al., 2001). It would be parsimonious to suppose that the cortical signal required for MSE is multisensory. Paradoxically, the parietal projections necessary for MSE originate from unimodal neurons (Wallace et al., 1993). The question of how unimodal cortical projections could produce multisensory enhancement in the DSC remains unanswered.

Multisensory integration apparently improves the ability of DSC neurons to detect targets (Stein et al., 1988, 1989; Wilkinson et al., 1996; Jiang et al., 2002). However, not all DSC neurons are multisensory. In the cat, approximately one-half of DSC neurons are multisensory, and, in the monkey, only approximately one-quarter, despite the availability of input of multiple modalities

(Wallace and Stein, 1996). The question of why some DSC neurons are multisensory while others are not also remains open.

This paper describes a model of the corticotectal system that simulates MSE and provides possible answers to these questions. It consists of an array of DSC units receiving unimodal primary and modulatory inputs and is trained in two stages that are unsupervised, local, and neurobiologically plausible. Primary inputs are trained first using the (Hebbian) self-organizing map algorithm. Modulatory inputs are trained second, using a novel Hebb–anti-Hebb rule. Training produces a mixture of unimodal and multisensory DSC units and causes the DSC to extract information concerning simulated targets from its inputs. The trained network can simulate MSE, with unimodal modulatory inputs preferentially augmenting DSC unit responses to stimulation of multiple modalities. The correspondence between model and data suggests that the self-organization of the real corticotectal system may involve mechanisms analogous to the two-stage algorithm.

Materials and Methods

The corticotectal network model represents neurons in the DSC and the sensory inputs they receive from subcortical and cortical sources (see Introduction). Inputs to DSC units are of two types, primary and modulatory. Primary inputs provide weighted connections to DSC units, and modulatory inputs augment the values of primary weights. Primary and modulatory inputs are specific for visual, auditory, or somatosensory modalities and can be activated by targets having the corresponding sensory attributes. Primary and modulatory input weights are trained in two separate stages of unsupervised learning. Training of primary connections produces a mixture of unimodal and multisensory DSC units, whereas training of modulatory connections produces MSE. Two-stage training causes the DSC to extract information from its sensory inputs concerning targets. The two sequential stages of training are inspired by DSC development, in which multisensory neurons first appear and later become capable of MSE, with onset of MSE corresponding to onset of parietal cortical influence (Wallace and Stein, 1997, 2000, 2001). A diagram of the corticotectal model is shown in Figure 1. Pseudocode for the two-stage algorithm is given in Table 1. A list of all variables and mathematical notation used in the paper is given in Table 2.

Architecture and activation of the network

The DSC units ($n_z = 100$) are arranged in a square 10×10 grid representing a small patch in the DSC. Neurons in the DSC have large overlapping receptive fields (Middlebrooks and Knudsen, 1984; Meredith and Stein, 1986b, 1990, 1996; Meredith et al., 1991; Wallace et al., 1996; Kadunce et al., 2001). We therefore assume that the units in the model patch have overlapping receptive fields and that the entire DSC patch receives input from the same small region of external space. Primary and modulatory inputs can be driven by targets appearing in this region.

Characterization of the target. The target T is arbitrarily assumed to be present one-half of the time and absent one-half of the time. When present, a target can exhibit any combination of visual (V), auditory (A), and somatosensory (S) attributes. The target has eight states t ($t = 0, 1, 2, \ldots, 7$), corresponding to the target-absent state of no sensory attributes plus the seven possible attribute combinations. The target-absent state (V = 0, A = 0, S = 0) has probability $P(T = 0) = \frac{1}{2}$. For simplicity, all modality-specific (single modality) targets are assigned the same probability p_s, whereas all cross-modal (multiple modality) targets are assigned probability p_c, where $p_s + p_c = \frac{1}{2}$. For present targets, the three modality-specific states (V = 1, A = 0, S = 0), (V = 0, A = 1, S = 0), and (V = 0, A = 0, S = 1) have probabilities $P(T = 1) = P(T = 2) = P(T = 3) = p_s/3$, and the four cross-modal states (V = 1, A = 1, S = 0), (V = 1, A = 0, S = 1), (V = 0, A = 1, S = 1), and (V = 1, A = 1, S = 1) have probabilities $P(T = 4) = P(T = 5) = P(T = 6) = P(T = 7) = p_c/4$. The eight target-state probabilities sum to one.

Activation of primary and modulatory inputs. Primary inputs are represented by a set of $n_x = 3$ random variables X_j ($j = 1, 2, 3$). Modulatory
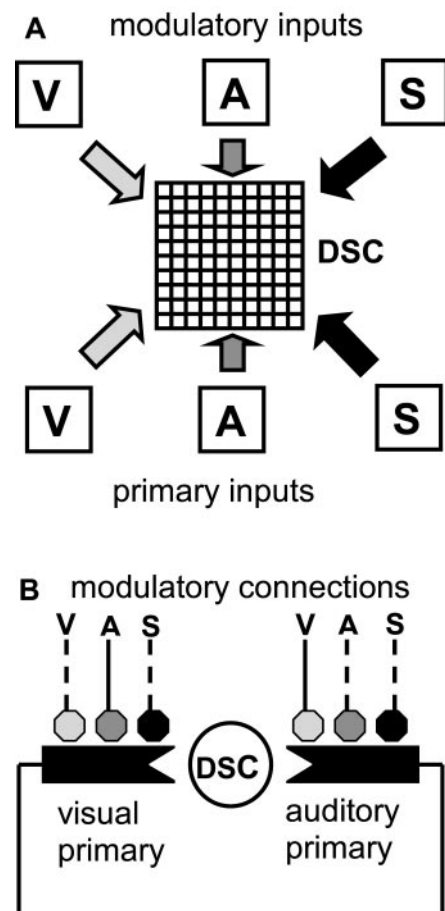
Figure 1. Schematic of the corticotectal model that produces multisensory enhancement in the DSC. *A*, The DSC is represented as a 10×10 grid of units. Primary inputs represent unimodal, excitatory projections from the visual (V), auditory (A), or somatosensory (S) systems. Modulatory inputs represent unimodal visual, auditory, or somatosensory projections from parietal cortex. Before stage-one training, each DSC unit receives primary input of all three modalities. Stage-one training causes DSC units to become specialized for specific modalities or modality combinations. As an example, a unit that receives primary input from the visual and auditory systems after stage-one training is shown in *B*. *B*, Before stage-two training, each primary connection may potentially receive modulatory input of all three modalities (solid and dashed lines), but stage-two training is restricted by the modality-matching and cross-modality constraints (see Materials and Methods). After stage-two training under these constraints, the unit shown can receive only visual and auditory modulatory input, with the primary visual connection modulated by the auditory modulatory input, and the primary auditory connection modulated by the visual modulatory input (solid lines).

inputs are represented by a set of $n_y = 3$ random variables $Y_k (k = 1, 2, 3)$. Each of the three primary or modulatory inputs is specific for one of the three sensory modalities: visual, auditory, or somatosensory. Variables x_j and y_k denote specific instances of X_j and Y_k. The variables X_j and Y_k represent whole populations of sensory neurons.

For simplicity, each discrete random variable X_j or Y_k is the sum r over a different set of $n = 20$ binary random variables. Each of the binary variables in a set is specific for the same sensory modality as the random variable X_j or Y_k that represents it. The individual binary variables take value zero or one depending only on their activation probabilities. Activation probabilities are either driven or spontaneous, depending on whether or not the target presents the modality specific to the set. For simplicity, all 60 binary variables represented by the three primary inputs X_j have the same driven and spontaneous activation probabilities of p_{x1} and p_{x0} (where $p_{x1} > p_{x0}$). Likewise, all 60 binary variables represented by the three modulatory inputs Y_k have the same driven and spontaneous activation probabilities of p_{y1} and p_{y0} (where $p_{y1} > p_{y0}$).

Characterization of primary and modulatory inputs. Because they each represent the sum of $n = 20$ binary random variables, the X_j and Y_k can

Table 1. Pseudocode for two-stage, unsupervised training of the corticotectal model of multisensory enhancement

Set unit numbers, bias, and sensitivity; set input activation and target state probabilities; set learning rates, neighborhood properties, numbers of iterations, and thresholds; initialize primary and modulatory weights.

For the required number of stage-one training iterations, do
    Determine target modality according to the target state probabilities
    Determine primary input activation using input activation probabilities
    Compute responses of DSC units to the input
    Find the DSC unit with the maximal response (winning DSC unit)
    Train primary weights of winning DSC unit and neighbors using Hebb's rule
Eliminate primary weights of any modality with values below the threshold
For the required number of stage-two training iterations, do
    Determine target modality according to the target state probabilities
    Determine primary and modulatory input activation using activation probabilities
    Modulate primary weights according to modulatory inputs and weights
    Compute responses of DSC units to primary input with modulation
    Find active primary and modulatory inputs and DSC units by thresholding
    Train modulatory weights of DSC units using the correlation–anti-correlation rule
        If a modulatory input and a DSC unit are both active, then
            increase the modulatory input weights to inactive primary inputs
            and decrease the modulatory input weights to active primary inputs
        If a modulatory input is active but a DSC unit is inactive, then
            decrease the modulatory weights to all primary inputs
End algorithm

assume any discrete value between 0 and 20. Because the individual binary variables in each sum are independent, the X_j and Y_k are binomially distributed. The general formula for the binomial distribution $b(n, p)$ is as follows (Appelbaum, 1996):

$$P(r) = b(n, p) = \frac{n!}{r!(n-r)!} p^r (1-p)^{n-r}, \qquad (1)$$

where p is the probability that any binary random variable takes value one. To use Equation 1 to describe the likelihood distributions of the primary and modulatory inputs, probability p can be replaced by the corresponding activation probability. Thus, the target drives primary input X_j when it presents the modality specific to X_j, and the driven likelihood for X_j is distributed as $b(n, p_{x1})$. Similarly, the spontaneous likelihood for primary input X_j is distributed as $b(n, p_{x0})$. The driven and spontaneous likelihoods for modulatory input Y_k are distributed as $b(n, p_{y1})$ and $b(n, p_{y0})$, respectively. All of the X_j and Y_k are distributed independently of one another given the state of the target. The binomial distribution affords a simple way to model a sensory input that represents the combined contribution of many individual inputs. The binomial approximates the Poisson distribution when n is large and p is small (Hoel et al., 1971). The Poisson distribution has been used to represent sensory inputs of different modalities in previous models of MSE (Anastasio et al., 2000).

For any given primary input X_j or modulatory input Y_k, the difference between the driven and spontaneous likelihoods can be quantified using the Kullback-Leibler divergence measures D_x and D_y, respectively, as follows (Cover and Thomas, 1991):

$$D_x = \sum_n b(n, p_{x0})\log_2[b(n, p_{x0})/b(n, p_{x1})], \qquad (2)$$

$$D_y = \sum_n b(n, p_{y0})\log_2[b(n, p_{y0})/b(n, p_{y1})]. \qquad (3)$$

D_x or D_y is used to quantify the amount of separation between the driven and spontaneous likelihoods of the primary and modulatory inputs, respectively. When D_x or D_y is small, the corresponding input is ambiguous with respect to the presence of a target. When D_x or D_y is large, the input is better able to indicate the presence of a target. A spontaneous likelihood and a series of driven likelihoods for the primary input are illustrated in Figure 2. The divergence measures associated with the spontaneous and driven activation probabilities used for the primary and modulatory inputs are enumerated in Table 3.

Table 2. List of variables and mathematical notation used in this article

|	
X_j	Discrete random variable representing primary input j ($j = 1, 2, 3$)
x_j	Specific activity level of primary input
n_x	Number of primary inputs ($= 3$)
$\mathbf{X}$	Discrete random vector representing primary input $[X_1 X_2 X_3]$
Y_k	Discrete random variable representing modulatory input k ($k = 1, 2, 3$)
y_k	Specific activity level of modulatory input
n_y	Number of modulatory inputs ($= 3$)
$\mathbf{Y}$	Discrete random vector representing modulatory input $[Y_1 Y_2 Y_3]$
Z_i	Random variable specifying activity of DSC unit i ($i = 1, 2, \ldots, 100$)
z_i	Specific activity level of DSC unit i
n_z	Number of DSC units ($= 100$)
T	Random variable representing the target
t	Target state ($t = 0, 1, \ldots, 7$)
p_s	Modality-specific target probability
p_c	Cross-modal target probability ($p_c = 1/2 - p_s$)
p_{x0}, p_{x1}	Spontaneous and driven activation probabilities for primary input X_j
p_{y0}, p_{y1}	Spontaneous and driven activation probabilities for modulatory input Y_k
p	Probability that a binary random variable takes value 1
n	Total number of binary random variables for binomial process
$b(n, p)$	Binomial distribution with parameters n and p
r	Number of binary variables taking state 1 for binomial process
u_{ij}	Unmodulated weight of primary input j onto DSC unit i
w_{ij}	Modulated weight of primary input j onto DSC unit i
v_{ijk}	Weight of the modulatory input k onto the primary input connection u_{ij}
d_{ijk}	Dummy variable for accumulating v_{ijk}
α	Learning rate for stage-one training
β	Learning rate for stage-two training
θ_u	Primary weight threshold for stage-one training
θ_x	Primary input threshold for stage-two training
θ_y	Modulatory input threshold for stage-two training
θ_z	DSC unit threshold for stage-two training
Ψ	Sum of active DSC units with activity z_i exceeding threshold θ_z
θ_I	DSC unit threshold for information gain computation
h	Index for DSC units in a neighborhood
P	Probability
D_x	Kullback-Leibler divergence measure for primary input
D_y	Kullback-Leibler divergence measure for modulatory input
I	Information in bits
H	Entropy in bits
ϕ	Tonic, bias input ($= 10$)
γ	Squashing function sensitivity ($= 1/5$)
$\leftarrow$	Update symbol
$\exp$	Exponential
Σ	Summation
$!$	Factorial

Mutual information between target and primary or modulatory inputs. The mutual information between the target and the input provides a measure of the amount of target information contained by the input. Mutual target-input information can be compared with the information content of the target alone. The information content of the event $T = t$ is defined as $I(T = t) = -\log_2[P(T = t)]$ (Cover and Thomas, 1991). The average information content of the target, equal to target entropy $H(T)$, is given as follows:

$$H(T) = \sum_t P(T)\log_2[P(T)], \qquad (4)$$

where $T = t$ is written simply as T for notational clarity. When the proportion of modality-specific to cross-modal targets is 2, target entropy $H(T)$ equals 2.32 bits (Table 3).

The DSC receives target information from its primary and modulatory inputs of all three modalities. The entire primary or modulatory input can be represented by random vectors $\mathbf{X} = [X_1, X_2, X_3]$ or $\mathbf{Y} = [Y_1, Y_2, Y_3]$. The ability of the entire primary or modulatory input to convey target information to the DSC can be quantified as mutual information $I(T; \mathbf{X})$ or $I(T; \mathbf{Y})$ as follows (Cover and Thomas, 1991):
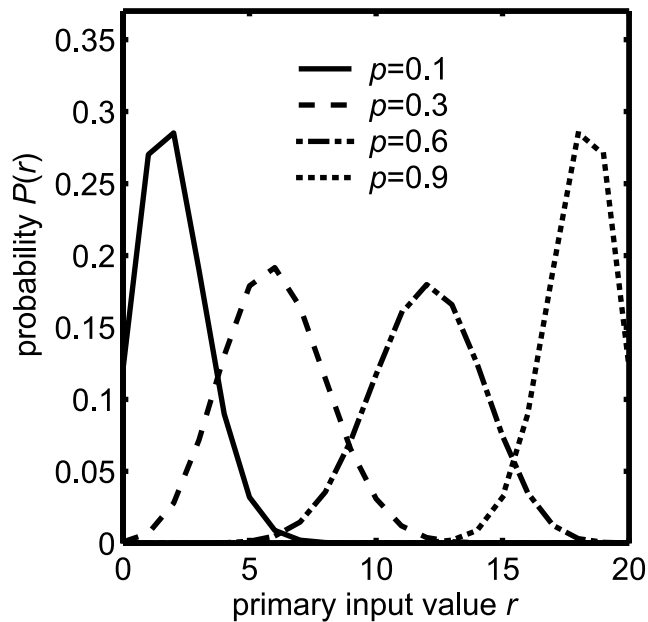
Figure 2. Input likelihoods $P(r)$ modeled as binomial distributions $b(n, p)$ (Eq. 1), where r is the number of the $n = 20$ binary variables that are active. The primary input spontaneous likelihood (solid curve) has activation probability $p = p_{x0} = 0.1$. The three primary input driven likelihoods have activation probabilities $p = p_{x1}$ of 0.3, 0.6, or 0.9 (dashed, dot-dashed, or dotted curves, respectively). For the modulatory input, the driven likelihood has activation probability $p = p_{y1} = 0.1$ (solid curve), whereas the spontaneous likelihood has activation probability $p = p_{y0} = 0$. Thus, a modulatory input of zero has probability one under spontaneous conditions.

Table 3. Information theoretic measures on target and inputs

Primary	p_{x0}	p_{x1}	D_x	$I(T; X)$
	0.1	0.3	3.36	1.36
	0.1	0.6	15.89	2.27
	0.1	0.9	50.72	2.32
Modulatory	p_{y0}	p_{y1}	D_y	$I(T; Y)$
	0	0.1	3.04	1.80
Target	Unimodal/multimodal proportion = 2/1			
	Information content [entropy, $H(T)$] = 2.32			

Relationships between spontaneous and driven activation probabilities for primary (p_{x0}, p_{x1}) or modulatory (p_{y0}, p_{y1}) inputs, Kullback-Leibler divergence measures for primary (D_x) or modulatory (D_y) inputs, and mutual information between target and primary [$I(T; X)$] or modulatory [$I(T; Y)$] inputs. The information content of the target [target entropy $H(T)$] is included for comparison.

$$I(T; X) = \sum_t \sum_X P(T, X) \log_2[P(T, X)/P(T)P(X)], \quad (5)$$

$$I(T; Y) = \sum_t \sum_Y P(T, Y) \log_2[P(T, Y)/P(T)P(Y)]. \quad (6)$$

Mutual information measures associated with the spontaneous and driven activation probabilities used for the primary and modulatory inputs are enumerated in Table 3. When the primary input spontaneous and driven activation probabilities are widely separated, as when $p_{x0} = 0.1$ and $p_{x1} = 0.9$, D_x is large and $I(T; X) = H(T)$, indicating that the input contains complete information about the target (Table 3).

Activation of model DSC units. Activities of the $n_z = 100$ DSC units are represented by random variables Z_i ($i = 1, 2, \ldots, 100$). Variables z_i denote specific instances of Z_i. The activity of a specific DSC unit z_i is computed as the weighted sum of the primary inputs to the DSC unit, passed through the sigmoidal squashing function:

$$z_i \leftarrow \frac{1}{1 + \exp\left[\gamma\left(\phi - \sum_j w_{ij}x_j\right)\right]}. \quad (7)$$

The ← symbol indicates the update that occurs with each new target presentation. The variable $\phi = 10$ is a tonic bias input. It represents inhibitory influences on the DSC from structures such as the substantia nigra (see Discussion). The sigmoid squashing function simulates the threshold and saturation properties of biological neurons. The parameter $\gamma = \frac{1}{5}$ adjusts the sensitivity of the squashing function. The variable w_{ij} represents the modulated weights of the connections to each DSC unit z_i from each primary input x_j. These weights are computed using the following formula:

$$w_{ij} \leftarrow u_{ij} + \sum_k v_{ijk}y_k. \quad (8)$$

Each u_{ij} is the unmodulated weight of the connection from primary input j to DSC unit i. Each v_{ijk} is the modulatory weight onto connection u_{ij}, from modulatory input k with activity y_k.

The learning algorithm
Each of the three primary inputs initially projects to every DSC unit. Each of the three modulatory inputs has a potential connection to every primary weight. The primary and modulatory weights are trained in two separate stages of unsupervised learning (Table 1). For both stages, training occurs only when the target is present.

Stage-one unsupervised learning. In the first stage, the primary weights u_{ij} are trained using the self-organizing map (SOM) algorithm (Willshaw and von der Malsburg, 1976; Kohonen, 1982, 1988; Haykin, 1999). First-stage training causes the model DSC to represent the primary inputs (see Results). The SOM involves selection of DSC units via competition and cooperation, followed by Hebbian modification of primary weights. The SOM is used here in its standard form.

The primary weights u_{ij} are initially set to small random values drawn from a uniform distribution in the range 0 to 0.1. The modulatory weights are fixed at $v_{ijk} = 0$ during stage-one training. At each iteration, a new target T is chosen according to target-state probabilities $P(T = t)$ for $t = 1, 2, 3, \ldots, 7$. The target state determines whether each primary input X_j will be spontaneous or driven, and values x_j are each drawn randomly from a binomial distribution with activation probability p_{x0} or p_{x1}, respectively. The DSC unit responses z_i to the primary inputs are then computed using Equation 7, where the w_{ij} are unmodulated, so that $w_{ij} = u_{ij}$. The DSC unit with the largest response is identified as the winner. The index of the winning DSC unit and of each of its 25 neighbors in the 10 × 10 grid forms subset h. The neighborhood z_h is constructed such that the winning DSC unit has activity 1, the eight nearest neighbors have activity 0.3, and the 16 neighbors once removed have activity 0.1. The neighborhood respects the boundaries of the 10 × 10 grid, so that only DSC units near the center of the grid have the full compliment of 25 neighbors. The primary weights to each DSC unit in z_h undergo the following Hebbian update:

$$u_{hj} \leftarrow u_{hj} + \alpha z_h x_j, \quad (9)$$

where α is the learning rate that is decremented after each iteration. The primary weights to each DSC unit z_h are then normalized so that $\sqrt{\sum_j (u_{hj})^2} = 1$. Training for 5000 iterations with a learning rate of 0.1 decrementing to 0.01 was found to produce stable results. The primary connections are pruned after completion of stage-one training. Any weight $u_{ij} < \theta_u$ is set to zero, where $\theta_u = 0.4$. The weights are renormalized after pruning. Primary weight pruning is necessary because the normalization procedure used does not set weights to zero.

Stage-two unsupervised learning. The modulatory weights v_{ijk} are trained in the second stage. The modulatory inputs represent neurons in parietal cortex that project to the DSC and produce MSE (see Introduction). Because the parietal inputs are unimodal but only enhance cross-modal DSC neuron responses (Wallace and Stein, 1994; Jiang et al., 2001), they are assumed to modulate inputs from other sources rather than to excite DSC neurons directly (see Results and Discussion). Electrophysiological evidence indicates that the modalities of the parietal inputs to a given DSC neuron usually match the modalities received by that neuron from other sources (Wallace et al., 1993). These findings can be used to infer two constraints on parietal-DSC connectivity: modality-matching and cross-modality. According to the modality-matching con-

straint, a DSC neuron may only receive modulatory inputs of the same modalities as those of its primary inputs. According to the cross-modality constraint, a modulatory input may only affect primary inputs of modalities different from its own (Fig. 1B). Stage two is designed to train the modulatory connections to produce MSE while respecting the constraints inferred from corticotectal neurobiology.

In the second stage, the modulatory weights v_{ijk} are trained using a novel algorithm based on correlation and anti-correlation between modulatory inputs, primary inputs, and DSC units. The modulatory weights are initially set to zero. The state of a new target is chosen at each iteration ($t = 1, 2, 3, \ldots, 7$), and the primary inputs are activated according to target state as described for stage-one training. During stage two, the modulatory inputs are also activated according to the sensory attributes of the chosen target. DSC unit responses to primary and modulatory input are determined using Equations 7 and 8.

The second-stage training algorithm requires a determination of whether each primary input, modulatory input, and DSC unit is active or inactive. A primary input x_j is active when $x_j > \theta_x$, a modulatory input is active when $y_k > \theta_y$, and a DSC unit is active when $z_i > \theta_z$, where θ_x, θ_y, and θ_z are thresholds. Primary and modulatory thresholds θ_x and θ_y are set at the integer nearest to the intersection point of the corresponding spontaneous and driven likelihood distributions (Fig. 2). Because the likelihoods are determined by the spontaneous and driven activation probabilities, each pair of activation probabilities is associated with a different threshold. Thresholds for the primary inputs are as follows: $p_{x0} = 0.1, p_{x1} = 0.3, \theta_x = 4$; $p_{x0} = 0.1, p_{x1} = 0.6, \theta_x = 6$; and $p_{x0} = 0.1, p_{x1} = 0.9, \theta_x = 10$. For the modulatory inputs, where $p_{y0} = 0$ and $p_{y1} = 0.1$, θ_y is 0. Because the probability distributions for the Z_i cannot be specified, the threshold θ_z is set empirically.

The modulatory weights v_{ijk} are trained according to the following stage-two rules. If a DSC unit and a modulatory input are both active, then increment the modulation of inactive primary inputs and decrement the modulation of active primary inputs by an amount β. If a modulatory input is active but a DSC unit is not, then decrement the modulation of all primary inputs, active and inactive, by 2β. The value of β depends primarily on the number of stage-two iterations. Increments and decrements are cumulative, but the modulatory weights themselves are constrained to be positive. Accumulation is accomplished by means of dummy variables d_{ijk} that take positive or negative values. The needed operations are summarized as follows:

$$\text{if } y_k > \theta_y \ \& \ z_i > \theta_z \ \& \ x_j \leq \theta_x \qquad \text{then } d_{ijk} \leftarrow d_{ijk} + \beta, \qquad (10)$$

$$\text{if } y_k > \theta_y \ \& \ z_i > \theta_z \ \& \ x_j > \theta_x \qquad \text{then } d_{ijk} \leftarrow d_{ijk} - \beta, \qquad (11)$$

$$\text{if } y_k > \theta_y \ \& \ z_i \leq \theta_z \qquad \text{then } d_{ijk} \leftarrow d_{ijk} - 2\beta, \qquad (12)$$

$$\text{if } y_k \leq \theta_y \qquad \text{then } d_{ijk} \leftarrow d_{ijk}. \qquad (13)$$

Although these rules train modulatory rather than direct connection weights, Rules 10 and 11 are essentially anti-Hebbian (see Discussion). Rules 10 and 11 enforce the cross-modality constraint, whereas Rule 12 enforces the modality-matching constraint. Rule 13 simply specifies that a modulatory weight is left unchanged if the associated modulatory input is inactive. The actual modulatory weights take the values of the corresponding dummy variables if they are positive and take the value zero otherwise:

$$\text{if } d_{ijk} \geq 0 \qquad \text{then } v_{ijk} \leftarrow d_{ijk}, \qquad (14)$$

$$\text{if } d_{ijk} < 0 \qquad \text{then } v_{ijk} \leftarrow 0. \qquad (15)$$

Modulatory weights can continue to grow as stage-two training proceeds, so they are bounded at an upper limit of one.

Assessing the trained model

To assess the effects of training on the model DSC, the responses of DSC units to their primary and modulatory inputs are measured, and the amount of target information extracted by DSC units from their inputs is estimated.

Quantifying MSE in the model. To simulate experiments on multisen-

sory DSC neurons, the responses of DSC units are examined as the levels of primary inputs are increased systematically from zero to $n = 20$ or held fixed at the mean of their spontaneous likelihood np_{x0}. The modulatory input levels are varied in a similar manner, but they are scaled to reflect their smaller dynamic range. The responses in the bimodal case can be used to compute percentage multisensory enhancement (%MSE) values using the following formula (Meredith and Stein, 1986a):

$$\%\text{MSE} = \frac{(\text{CM} - \text{SM}_{max})}{\text{SM}_{max}} \times 100\%, \qquad (16)$$

where CM is the cross-modal response and SM_{max} is the larger of the two modality-specific responses. By this definition, enhancement occurs whenever $\text{CM} > \text{SM}_{max}$. Enhancement is subadditive when the cross-modal response is smaller than the sum of the two modality-specific responses, and supra-additive when it is greater than the sum. The effects of the modulatory inputs can be examined by comparing %MSE values at specific levels of primary input, after various modulatory connections have been removed. This is intended to simulate the effects of experiments in which cortical input to the DSC from specific regions is selectively inactivated (Wallace and Stein, 1994; Jiang et al., 2001).

Mutual information between target and DSC units. DSC unit responses vary sigmoidally between zero and one. Even if the $n_z = 100$ DSC units are binarized by thresholding, they would still have 2^{100} different states potentially available to convey target information. Complete characterization of the mutual information between DSC units and the target is impossible because it would require determination of the joint probability of each of these DSC states and each target state. Instead, the information gain attributable to training is measured between the target and the number Ψ of DSC units whose responses z_i exceed threshold $\theta_I = 0.3$:

$$\Psi = \sum_i \left\{ \begin{array}{l} 1 \text{ if } z_i > \theta_I \\ 0 \text{ otherwise} \end{array} \right\}. \qquad (17)$$

The joint probability between the target and the number of suprathreshold DSC unit responses $P(T, \Psi)$ is estimated computationally by presenting many targets of various states, computing Ψ for each of them, and binning the Ψ values. The joint probability distribution is estimated from the resulting histogram by scaling. The marginal distributions $P(T)$ and $P(\Psi)$ are computed from the estimated joint distribution. The estimated probabilities are used to calculate an estimate of the target information gained by the DSC network:

$$I(T; \Psi) = \sum_t \sum_\Psi P(T, \Psi) \log_2[P(T, \Psi)/P(T)P(\Psi)]. \qquad (18)$$

This information gain measure is a crude estimate of the true mutual information between the target and the DSC network. However, it is adequate for the purpose of showing that the two-stage, unsupervised learning algorithm does train the model DSC to extract information concerning the target from its inputs.

Results

The two-stage learning algorithm is used to train the corticotectal model in an unsupervised manner as simulated targets are presented to it. Pseudocode for the training algorithm is given in Table 1, and a schematic of the model is shown in Figure 1. DSC units receive two types of random inputs, primary and modulatory. Primary inputs make direct, excitatory connections onto DSC units, whereas modulatory inputs can augment the primary weights (Eq. 8). Primary and modulatory weights are trained during stage one and stage two of the algorithm, respectively. Both stages of training are associated with interesting emergent properties. A mixture of unimodal and multisensory DSC units arises spontaneously from stage-one training. Multisensory enhancement arises spontaneously from stage-two training, and removal of modulatory connections has a much greater effect on

DSC unit responses to cross-modal (combined modality) stimuli than on responses to modality-specific (single modality) stimuli.

### Simulating modality specialization in the DSC

The simulations in this section use stage one of the two-stage algorithm to show how the proportion of modality-specific to cross-modal targets, and the information content of the inputs, could influence the percentage of multisensory neurons in the DSC. In the model, the modality selectivity of an individual DSC unit depends on the weights of its primary input connections. A visual–auditory DSC unit, for example, would receive nonzero weights from the visual and auditory primary inputs and zero weight from the somatosensory primary input (Fig. 1*B*). The modality selectivity of individual DSC units in the model is determined entirely by stage one of the two-stage algorithm, because stage two respects the modality selectivity established by stage one.

Stage one is based on the SOM algorithm (see Materials and Methods). The SOM is a neurobiologically plausible and well established computational tool for modeling the activity-dependent refinement of sensory maps in the nervous system. This algorithm naturally produces a mixture of unimodal and multisensory DSC units. DSC units of similar modality selectivity are colocalized by the SOM, and it is possible that such an arrangement is superimposed on the broader spatial map in the DSC. The overall organization of the DSC is beyond the scope of this study. The focus here is on the percentage of multisensory DSC units produced by the SOM in a small patch of the DSC. The proportion of unimodal to multisensory DSC units depends on several factors, including the primary weight threshold θ_u, the proportion of modality-specific to cross-modal targets, and the information content of the primary inputs.

The SOM in stage one causes the primary weight vectors to represent the primary input vectors by distributing the weight vectors approximately evenly among the input vectors. There are 100 primary weight vectors $[u_{i1}, u_{i2}, u_{i3}]$, one vector for each of the 100 DSC units with activities $z_i (i = 1, 2, \ldots, 100)$. The primary input vectors X are simply the values $[x_1, x_2, x_3]$ that are chosen randomly, on each trial, for the visual, auditory, and somatosensory primary inputs X_1, X_2, and X_3. Because the primary weights determine the modality selectivity of DSC units, factors that influence the distribution of primary input vectors can affect the modality selectivity of DSC units established by stage one.

Changing the information content of the primary inputs can affect the distribution of primary input vectors, even if the proportion of modality-specific to cross-modal targets is held constant. Primary input vectors used to train the model during stage one and the resulting primary weight vectors are illustrated in Figure 3. For *A–C* in Figure 3, the proportion of modality-specific to cross-modal targets is two to one ($p_s = 0.34$ and $p_c = 0.17$). The spontaneous activation probability p_{x0} equals 0.1 in *A–C*, and the information content of the primary inputs is altered by changing only the driven activation probability. The driven activation probability p_{x1} increases from 0.3 (Fig. 3*A*) to 0.6 (Fig. 3*B*) to 0.9 (Fig. 3*C*). Primary input information content increases and ambiguity decreases as the driven activation probability increases (Table 3). This affects the clustering of the primary input vectors.

The primary weight vectors, which are normalized as part of stage-one training, are plotted as plus signs in Figure 3. For comparison, the primary input vectors are also normalized before they are plotted as circles. In Figure 3*A*, in which the primary inputs are the most ambiguous and have the lowest information
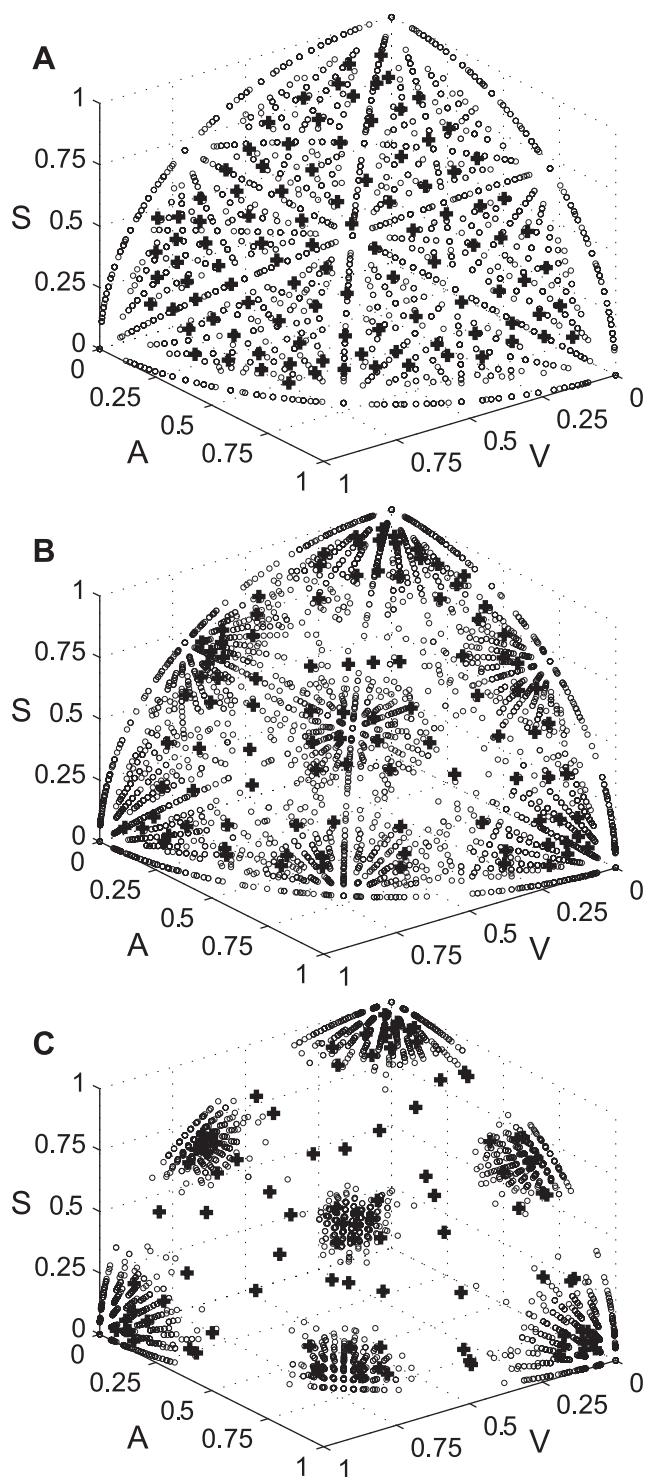
Figure 3. Stage-one training causes primary weight vectors to cluster with primary input vectors. The distinctiveness of the clusters depends on primary input ambiguity. In *A–C*, there are twice as many modality-specific as cross-modal targets, and the spontaneous primary input activation probability p_{x0} equals 0.1. The primary input becomes less ambiguous as the driven activation probability p_{x1} is increased from 0.3 (*A*) to 0.6 (*B*) to 0.9 (*C*) (Table 3). Clusters of primary input vectors (circles) become progressively more distinct. This causes more primary weight vectors (plus signs) to adopt a distinctly unimodal pattern. V, Visual; A, auditory; S, somatosensory.

content, the input vectors are evenly distributed. Predominantly unimodal inputs are located in the corners, in which primary input of one of the three modalities is near one, whereas the other two are near zero. They are rare in Figure 3*A*. The primary weight

vectors are scattered approximately evenly among the input vectors, and almost all are located in multisensory regions. As the primary inputs are made less ambiguous (Fig. 3B,C) and their information content goes up, the input vectors form distinct clusters. Clusters of unimodal primary inputs are pushed out farther into the corners. DSC units with primary weight vectors drawn into these clusters would have predominantly unimodal response characteristics. Figure 3 illustrates that stage-one training produces a greater percentage of predominantly unimodal DSC units when primary input information content is high. These results suggest that the percentage of unimodal DSC neurons in a given species could, at least in part, reflect the information content of the inputs it receives during the formation and refinement of its sensory maps.

It is clear from Figure 3 that primary weight vectors fall into clusters that are predominantly unimodal, bimodal, or trimodal. Although the primary weights from some inputs may be very small, none are zero. The DSC units could all be considered trimodal, because their weights are nonzero from all three primary inputs. Designating all DSC units as trimodal, however, would obscure the fact that primary weights are distributed throughout the input space in predominantly unimodal, bimodal, and trimodal regions. To alleviate this problem, small weights are pruned after stage-one training. Pruning is accomplished by setting to zero all primary weights u_{ij} that are less than the primary weight threshold θ_u. Pruning corresponds to a process of activity-dependent synapse elimination, such as that described for the formation of retinotopic maps in the superior colliculus (and optic tectum) and in other processes (Katz and Shatz, 1996; Lichtman et al., 1999). As a result of the removal of weak primary weights, many DSC units become unimodal or bimodal. The effects of pruning vary depending on both θ_u and modality-specific target probability p_s (where $p_c = \frac{1}{2} - p_s$). The effect of changes in these two variables on the percentage of multisensory (bimodal and trimodal) DSC units produced by stage-one training is shown in Figure 4. Multisensory DSC units are those that have nonzero weights from two or three primary inputs after pruning.

Figure 4 shows that the percentage of multisensory DSC units decreases as θ_u increases. This happens simply because more primary weights are eliminated as the threshold increases. The DSC is 100% multisensory for low values of θ_u, regardless of modality-specific target probability. However, the percentage of multisensory DSC units falls faster with increases in θ_u as modality-specific target probability increases (and as cross-modal target probability decreases). This result confirms the expectation that stage one will establish more multisensory connections when training involves a greater number of cross-modal targets.

The data in Figure 4 are generated using primary inputs of intermediate ambiguity ($p_{x0} = 0.1$ and $p_{x1} = 0.6$). The results are qualitatively similar when the primary inputs are made more ambiguous by decreasing the driven activation probability p_{x1} to 0.3 or made less ambiguous by increasing it to 0.9 (data not shown). The main difference is that the decrease in the percentage of multisensory DSC units as θ_u increases, at all levels of modality-specific target probability, is somewhat slower with more ambiguous primary input and faster with less ambiguous primary input.

The results suggest that the percentage of multisensory DSC neurons in a particular species may depend on several factors, including the proportion of cross-modal targets it encounters in its particular environmental niche. Cats, which hunt at night,
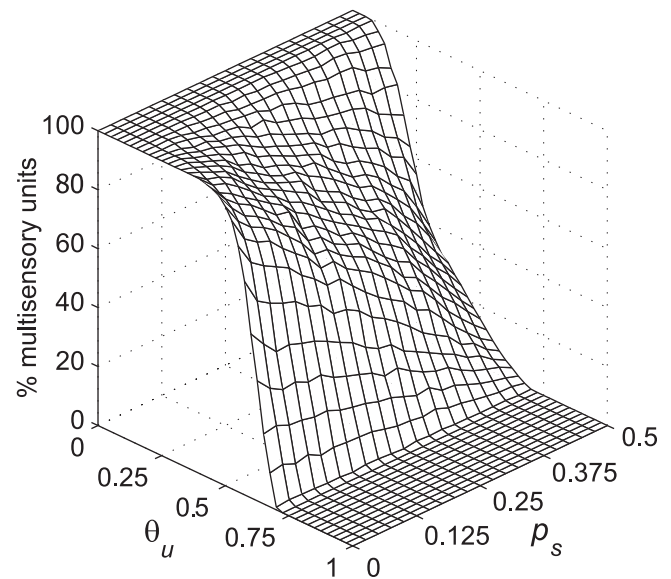
Figure 4. The percentage of multisensory DSC units resulting from stage-one training is plotted as a function of primary weight threshold θ_u and the probability of a modality-specific target p_s (where the probability of a cross-modal target p_c equals $\frac{1}{2} - p_s$). The spontaneous and driven primary input activation probabilities p_{x0} and p_{x1} are 0.1 and 0.6, respectively. The modality-specific target probability p_s is increased from 0 to 0.5 in steps of 0.025, with 10 networks trained for each p_s value. The primary weight threshold θ_u is increased from 0 to 1 in steps of 0.025. Each of the 10 networks are thresholded at each θ_u value. The percentage of multisensory units shown is the mean for the 10 networks. The percentage of multisensory DSC units falls as θ_u increases. The fall is more rapid when modality-specific targets are more probable (and cross-modal targets are less probable). Any desired percentage of multisensory DSC units can be obtained through appropriate choice of θ_u and p_s.

may encounter more cross-modal targets than monkeys, which forage during the day. This likely difference in the proportion of cross-modal targets between cats and monkeys may explain why cats have a higher percentage of multisensory DSC neurons than monkeys (Wallace and Stein, 1996; Wallace et al., 1996, 1998).

It is also possible that sensory systems are noisier in cats than in monkeys. As such, sensory input to the DSC would be more ambiguous, and carry less target information, in cats than in monkeys. The model suggests that such a difference, if present, could contribute to the difference in the percentage of multisensory DSC neurons between cats and monkeys. In the model, any relative proportion of unimodal to multisensory DSC units can be obtained by appropriate choice of primary weight threshold, primary input activation probabilities, and proportion of modality-specific to cross-modal targets. In the brain, self-organization of the corticotectal network probably involves activity-independent, genetically prespecified molecular mechanisms, as well as the activity-dependent processes modeled here. Presumably, the percentage of multisensory DSC neurons produced by these combined processes confers behavioral advantage to a species, considering such factors as its environmental niche and the properties of its sensory systems.

Simulating the parietal projection to the DSC

The corticotectal circuitry that gives rise to MSE in the model is consequent on model architecture and on training during stage two of the two-stage algorithm. Stage two is based on a novel correlation–anti-correlation rule. Experimental observations on MSE guided the design of the model and of the correlation–anti-correlation rule.

The parietal neurons that produce MSE in the DSC are them-
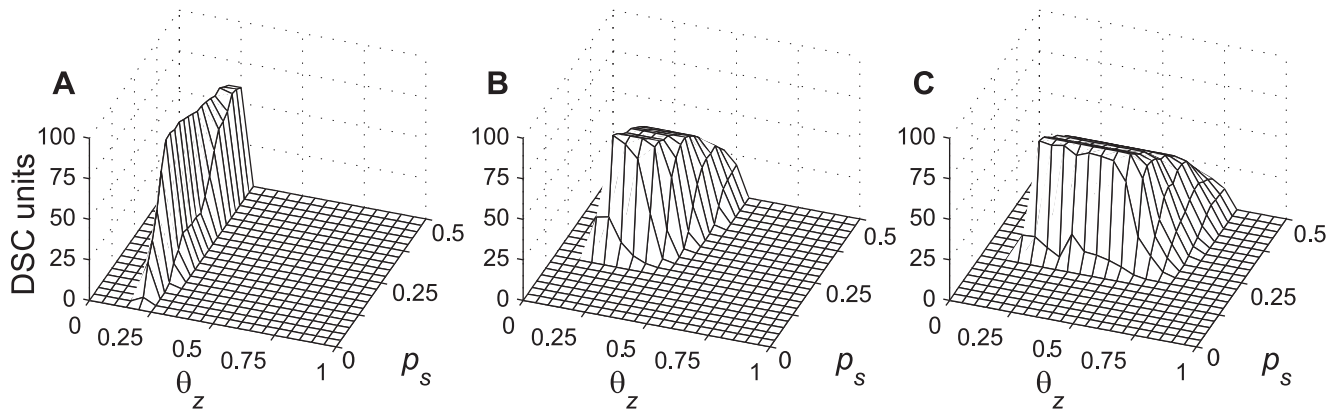
Figure 5. The abundance of correct modulatory weights resulting from stage-two training depends sensitively on DSC unit threshold θ_z and on modality-specific target probability p_s. In A–C, the spontaneous primary input activation probability p_{x0} equals 0.1. The primary input becomes less ambiguous as driven activation probability p_{x1} is increased from 0.3 (A) to 0.6 (B) to 0.9 (C) (Table 3). The primary input threshold θ_x is increased from 4 (A) to 6 (B) to 10 (C). In A–C, the modulatory input spontaneous and driven activation probabilities p_{y0} and p_{y1} are 0 and 0.1, respectively, and the modulatory input threshold θ_y is 0. Modality-specific target probability p_s is varied from 0 to 0.5 in steps of 0.025. Ten networks receive stage-one training for 5000 iterations at each p_s value. Primary weights are pruned at $\theta_u = 0.4$. DSC unit activity threshold θ_z is varied from 0 to 1 in steps of 0.05. Each of the 10 stage-one trained networks, at each p_s value, receives stage-two training for 5000 iterations at each θ_z value. This yields 10 trained networks for each combination of p_s and θ_z. Sets of 10 containing any misdirected modulatory weights (i.e., modulatory weights not respecting the modality-matching and cross-modality constraints) are excluded. For sets of 10 containing no such errors, the mean number of DSC units receiving modulatory connections is computed. Each panel plots the number of units, in error-free networks, that receive modulatory input. For $p_{x1} = 0.3$ (A) stage-two works best when $0.2 \leq \theta_z \leq 0.3$ and $p_s \geq 0.15$, for $p_{x1} = 0.6$ (B) when $0.2 \leq \theta_z \leq 0.55$ and $p_s \geq 0.23$, and for $p_{x1} = 0.9$ (C) when $0.2 \leq \theta_z \leq 0.8$ and $p_s \geq 0.23$. The number of error-free networks is greater for unambiguous than for ambiguous primary inputs.

selves unimodal (Wallace et al., 1993). If unimodal parietal neurons directly excited DSC neurons, then inactivation of the relevant parietal neurons should substantially reduce modality-specific as well as cross-modal DSC neuron responses. For many DSC neurons, however, parietal inactivation reduces MSE with little or no effect on modality-specific responses (Wallace and Stein, 1994; Jiang et al., 2001). The model therefore postulates an indirect, modulatory mechanism whereby inputs representing unimodal parietal projections could produce enhancement of cross-modal but not modality-specific responses. In the model, primary inputs directly excite DSC units and represent inputs from a variety of subcortical and cortical structures. Modulatory inputs, representing parietal projections only, do not directly excite DSC units but act by augmenting primary inputs. In some studies, parietal inactivation was found to reduce modality-specific responses in some DSC neurons (Clemo and Stein, 1986; Meredith and Clemo, 1989; Wallace and Stein, 1994). This can easily be accounted for by postulating that parietal cortex is a source of primary, as well as modulatory, input to some DSC units.

Constraints on learning, in addition to constraints imposed by model architecture, ensure that the performance of the trained model conforms to experimental observation. For modulatory inputs to enhance cross-modal but not modality-specific DSC unit responses, modulatory connections should be made only when a modulatory input and a primary input are of different modalities. This is the cross-modality constraint. Imposing the cross-modality constraint alone would not be sufficient to ensure that modulatory connections in the model are consistent with experimental observations on parietal projections to DSC. Under the cross-modality constraint, all DSC units could still receive every one of the three modalities, either as primary or modulatory input. All DSC units would be trimodal, which is inconsistent with observation. Maintenance of the DSC unit modality selectivities established during stage-one training requires that the modalities of the modulatory connections received by a DSC unit should match the modalities of the primary inputs received by that unit. This is the modality-matching constraint, which is

supported by orthodromic activation studies (Wallace et al., 1993). Together, the cross-modality and the modality-matching constraints ensure that a primary input connection onto a multisensory DSC unit will receive a modulatory connection only if the modality of the modulatory input is different from that of the primary input but the same as that of another primary input connection onto the DSC unit (Fig. 1B). Modulatory connections, established by the correlation–anti-correlation rule as designed, are successfully restricted by these constraints over broad ranges of model parameters.

The correlation–anti-correlation rule (see Materials and Methods) can be summarized as follows. If a DSC unit and a modulatory input are both active, then decrease the modulation of active primary inputs and increase the modulation of inactive primary inputs. If a modulatory input is active but a DSC unit is inactive, then decrease the modulation of all primary inputs. The critical parameters for stage-two training include the threshold θ_x for the primary inputs, θ_y for the modulatory inputs, and θ_z for the DSC units. These thresholds are needed for the algorithm to decide whether or not the associated model elements are active. For the primary and modulatory inputs, thresholds are set at the integer nearest the intersection points of the corresponding spontaneous and driven likelihoods (see Materials and Methods). Stage-two training depends on the spontaneous and driven activation probabilities of the primary (p_{x0}, p_{x1}) and modulatory (p_{y0}, p_{y1}) inputs, both because they determine input likelihoods and because they affect correlations among inputs and DSC units that in turn affect the behavior of the correlation–anti-correlation rule. The DSC unit threshold θ_z cannot be set on the basis of likelihoods because the likelihood distributions of DSC unit responses are not known. Stage two also depends on modality-specific target probability p_s. These factors interact in a complex way, but certain regularities in the operation of stage two can be identified.

Figure 5 plots numbers of DSC units receiving nonzero modulatory weights for those trained networks in which all modulatory connections respect the cross-modality and modality-matching constraints. The primary input spontaneous activation

Table 4. Comparison of modulatory (model) and descending (experimental) connectivity

	V	A	S	V-A	V-S	A-S	V-A-S	Total
Model results after 5000 iterations								
None	**14.80**	**13.70**	**11.90**	0.00	0.00	0.00	0.00	**40.40**
V	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
A	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
S	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V, A	0.00	0.00	0.00	**17.40**	0.00	0.00	0.00	**17.40**
V, S	0.00	0.00	0.00	0.00	**15.60**	0.00	0.00	**15.60**
A, S	0.00	0.00	0.00	0.00	0.00	**14.40**	0.00	**14.40**
V, A, S	0.00	0.00	0.00	0.00	0.00	0.00	**12.20**	**12.20**
Total	**14.80**	**13.70**	**11.90**	**17.40**	**15.60**	**14.40**	**12.20**	**100.00**
Model results after 50 iterations								
None	**14.00**	**14.40**	**13.00**	**1.70**	**2.10**	0.00	0.00	**45.20**
V	0.00	0.00	0.00	**2.70**	**1.20**	0.00	0.00	**3.90**
A	0.00	0.00	0.00	**1.80**	0.00	**3.10**	0.00	**4.90**
S	0.00	0.00	0.00	0.00	**0.20**	**1.40**	0.00	**1.60**
V, A	0.00	0.00	0.00	**9.20**	0.00	0.00	**1.00**	**10.20**
V, S	0.00	0.00	0.00	0.00	**11.50**	0.00	**1.40**	**12.90**
A, S	0.00	0.00	0.00	0.00	0.00	**10.20**	0.00	**10.20**
V, A, S	0.00	0.00	0.00	0.00	0.00	0.00	**11.10**	**11.10**
Total	**14.00**	**14.40**	**13.00**	**15.40**	**15.00**	**14.70**	**13.50**	**100.00**
Experimental results adapted from Wallace et al., 1993, their Table 1								
None	**11.76**	**1.47**	**2.94**	**2.94**	**2.94**	**0.37**	**0.37**	**22.79**
V	**18.01**	0.00	0.00	**12.50**	**1.84**	0.00	0.00	**32.35**
A	0.00	**0.37**	0.00	0.00	0.00	0.00	0.00	**0.37**
S	**0.74**	0.00	**4.04**	0.00	**5.15**	**0.74**	0.00	**10.66**
V, A	**0.74**	0.00	0.00	**5.88**	0.00	0.00	**0.37**	**6.99**
V, S	**0.37**	0.00	0.00	**0.74**	**16.18**	0.00	**2.57**	**19.85**
A, S	0.00	0.00	0.00	0.00	0.00	**0.74**	0.00	**0.74**
V, A, S	0.00	0.00	0.00	**1.84**	**0.74**	0.00	**3.68**	**6.25**
Total	**31.62**	**1.84**	**6.99**	**23.90**	**26.84**	**1.84**	**6.99**	**100.00**

V, Visual; A, auditory; S, somatosensory; V-A, visual–auditory; V-S, visual–somatosensory; A-S, auditory–somatosensory; V-A-S, trimodal.

probability p_{x0} is fixed at 0.1 in *A–C*. The primary input is made less ambiguous by increasing the driven activation probability p_{x1} from 0.3 (Fig. 5*A*) to 0.6 (Fig. 5*B*) to 0.9 (Fig. 5*C*). The primary input threshold θ_x is correspondingly increased from 4 to 6 to 10. Stage two produces large numbers of allowed modulatory connections for θ_z values ~0.2, regardless of primary input ambiguity. For the less ambiguous primary inputs ($p_{x1} = 0.6$ and $p_{x1} = 0.9$), allowed modulatory connections fail to develop when modality-specific target probability p_s is lower than ~0.2. The dependency on p_s is not as critical for the most ambiguous primary input ($p_{x1} = 0.3$). The unavoidable errors in deciding primary input activation in the ambiguous case may actually work to advantage, but only for values of θ_z ~0.2. The region over which stage two produces large numbers of allowed modulatory weights (those that respect the constraints) grows larger as the ambiguity of the primary input decreases. This is attributable to an improved ability to decide DSC unit activation in the less ambiguous networks.

In Fig. 5*A–C*, the spontaneous and driven activation probabilities for the modulatory inputs are $p_{y0} = 0$ and $p_{y1} = 0.1$, and the modulatory input threshold is $\theta_y = 0$. The ability of stage two to produce allowed modulatory connections is insensitive to the actual values of the modulatory input spontaneous and driven activation probabilities, so long as the modulatory likelihoods are well separated and decisions concerning modulatory input activation are reliable. These results demonstrate that the production of allowed modulatory connections using the correlation–anti-correlation rule depends on reliable decisions concerning input and DSC unit activation. The correlation–anti-correlation rule is robust when reliable activation decisions can be made.

The ability of the correlation–anti-correlation rule to produce modulatory connections that are consistent with experimental observations on the projection from parietal cortex to DSC is illustrated in Table 4. This table presents the modulatory connectivity produced by the model under two conditions (Table 4, top, middle) and compares it with experimental results on descending parietal connections to DSC neurons (Table 4, bottom) from an orthodromic activation study (Wallace et al., 1993). The columns of each section, labeled at the top, indicate the seven possible modality selectivities of DSC units (or neurons), as classified by the modalities of their primary inputs (or by the modalities to which the neuron responds, for the experimental data). The rows of each section, labeled at the left side, indicate the eight possible sets of unimodal modulatory inputs (or descending parietal inputs, for the experimental data). The number of units (or neurons) receiving the designated combinations of input are indicated as a percentage of the total number of DSC units in the model (or of the total number of neurons recorded, for the experimental data). For the model, DSC unit numbers are determined on the basis of 10 runs. The total percentage of units (or neurons) of each modality selectivity is indicated in the last row of each section. The total number of units (or neurons) receiving each set of modulatory (or descending) inputs is indicated in the rightmost column of each section. Wallace et al. (1993) reported the descending projections to DSC from two visual parietal structures: the lateral suprasylvian sulcus (LS) and the anterior ectosylvian visual area (AEV). To facilitate comparison with model results, data from these two structures have been grouped together as visual in Table 4, bottom.

For the model results (Table 4, top, middle), the primary input activation probabilities are $p_{x0} = 0.1$ and $p_{x1} = 0.6$, and the modulatory input activation probabilities are $p_{y0} = 0$ and $p_{y1} = 0.1$. The modality-specific target probability is set at $p_s = 0.34$ ($p_c = 0.17$). The stage-one pruning threshold is set at $\theta_u = 0.4$, and the stage-two thresholds are set at $\theta_x = 6$, $\theta_y = 0$, and $\theta_z = 0.2$. In the first network (Table 4, top), stage-two training is run for 5000 iterations and produces all of the allowed modulatory connections with no errors. In the second network (Table 4, middle), stage-two is run for only 50 iterations, and not all of the allowed modulatory connections are made. In some cases, bimodal DSC units receive a modulatory connection from only one or the other of the modulatory inputs that could connect to them or receive no modulatory connection at all. Likewise, trimodal DSC units sometimes receive modulatory connections from only two of the three modulatory inputs that could connect to them. This pattern of absent connections is consistent with experimental observation (Table 4, bottom). There is one notable difference between the modeling results of Table 4, middle, and the experimental results of Table 4, bottom. Some unimodal DSC neurons apparently receive input of the same modality from parietal cortex. As suggested above, it is possible that these inputs would be primary rather than modulatory.

In the model, a modulatory input that fails to provide an allowed modulatory connection is often associated with a weak primary input of the corresponding modality. This results because the weak primary input usually fails to activate the DSC unit (i.e., bring its activity over threshold θ_z) when it alone is activated by a modality-specific target. The consequence is that the DSC unit and the modulatory input of that modality are not consistently active together, and that modulatory input cannot establish connections to inactive primary inputs of other modalities. Bimodal DSC neurons have been studied that cannot be activated by input of one modality but show enhancement if input of that modality is presented with input of a different modality (Meredith and Stein, 1986a; Stein and Meredith, 1993). The presumption might be that the weaker modality provides the modulatory input. The model does not exclude this possibility. Instead, it predicts the existence of bimodal DSC neurons for which the stronger modality provides both strong primary and strong modulatory input, and enhancement occurs as a result of strong modulation of the weaker primary input in the event of a cross-modal stimulus. This prediction should be testable using available experimental techniques.

Information gain attributable to stage-one and stage-two training

It has been shown theoretically that the SOM algorithm not only forms maps but also causes output units to extract information from their inputs (Linsker, 1988a,b). Training with stage one (the SOM), and to a lesser extent stage two, causes the DSC to extract a substantial amount of target information from its inputs. Information gain by the DSC depends on the percentage of multisensory DSC units and actually decreases as the percentage of multisensory DSC units increases past a certain level.

The DSC response to a target is characterized simply as the number Ψ of DSC units that, on target presentation, show activity exceeding threshold θ_I (see Materials and Methods). A value of $\theta_I = 0.3$ is chosen, although the results are similar over a range of θ_I. The target information gain, or the mutual information $I(T; \Psi)$ between the target and the number of suprathreshold DSC unit activities, is computed (Eq. 18) and compared for various network configurations.
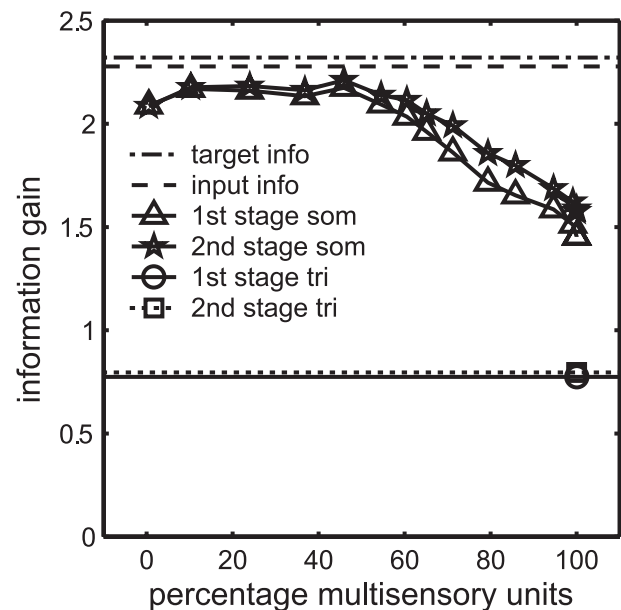
Figure 6. Two-stage training causes the DSC to extract most of the target information content of the primary inputs, especially when the DSC contains a mixture of unimodal and multisensory units. The percentage of multisensory DSC units is varied by manipulating the primary weight threshold θ_u. Ten networks receive stage-one training for 5000 iterations ($p_{x0} = 0.1$, $p_{x1} = 0.6$, $p_s = 0.34$, and $p_c = 0.17$). Each network is then pruned using θ_u varying from 0 to 1 in steps of 0.005. This produces mean percentages of multisensory DSC units over the 10 networks ranging from 0 to 100%. Each pruned network receives stage-two training for 5000 iterations ($p_{y0} = 0$, $p_{y1} = 0.1$, $\theta_x = 6$, $\theta_y = 0$, and $\theta_z = 0.2$). For each of the 10 networks associated with each θ_u value, both before and after stage-two training, the mutual information between the target and the number of suprathreshold DSC unit responses is computed (Eqs. 17 and 18; $\theta_I = 0.3$). The mean information gain after stage-one and stage-two training is plotted against the mean percentages of multisensory units. The mutual information between the target and the primary inputs (2.27 bits; dashed line; Eq. 5) is nearly as high as the information content of the target (2.32 bits; dot-dashed line; Eq. 4). Stage-one training causes the DSC to extract a large amount of target information (triangles), and stage-two (stars) causes a small increase in this amount. The increase is significant when the percentage of multisensory units is 60% or larger (t test, 0.05 significance level). The mutual information between target and DSC is nearly as large as the mutual information between target and primary inputs, but only for percentages of multisensory DSC units between ~10 and 50%. The mutual information between target and DSC decreases steadily as the percentage of multisensory DSC units increases above 50%. This decrease in mutual information between target and DSC approaches that of a uniformly trimodal DSC, with (0.80 bits; square) and without (0.77 bits; circle) modulatory connections. Variability in the DSC response after two-stage training keeps DSC information content above that of the uniformly trimodal DSC.

The corticotectal model can be used to explore the relationship between target information gain and the relative proportion of unimodal to multisensory DSC units. The model is retrained 10 times from a random initial condition. Manipulating the primary weight threshold θ_u varies the percentage of multisensory DSC units as a result of stage-one training. The threshold θ_u is increased in steps of 0.05 to produce percentages of multisensory DSC units ranging from 0 to 100%. Stage-two training follows but does not alter the modality selectivity established for DSC units during stage-one (see above). Target information gain at the DSC is computed before and after stage-two training. The results are plotted in Figure 6.

For more than one-half of the range of percentage multisensory DSC units, target information gain is almost as high as the information content of the primary inputs. For the example explored in Figure 6, the primary inputs are of intermediate ambiguity, with spontaneous and driven activation probabilities of $p_{x0} = 0.1$ and $p_{x1} = 0.6$, respectively. The mutual information

A modulatory connections intact **B** modulatory connections cut
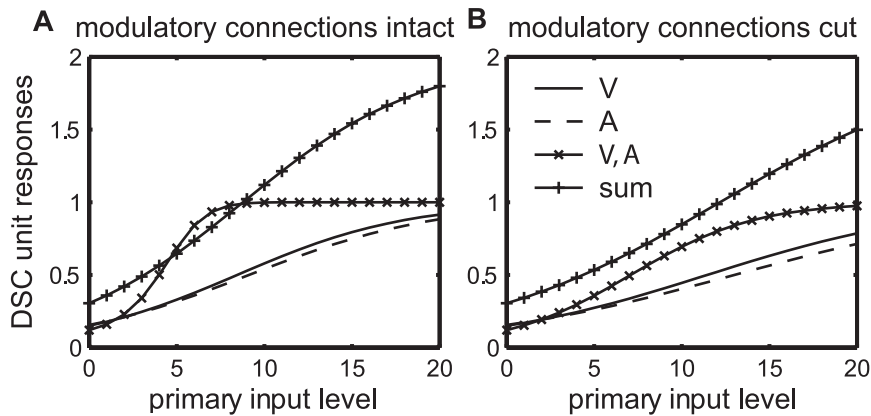
Figure 7. Responses of a bimodal, visual–auditory DSC unit over the full range of primary input levels. The responses were determined after the network containing this unit was trained using the two-stage algorithm with the following parameters: $p_s = 0.34, p_c = 0.17, p_{x0} = 0.1, p_{x1} = 0.6, \theta_u = 0.4, p_{y0} = 0, p_{y1} = 0.1, \theta_x = 6, \theta_y = 0$, and $\theta_z = 0.2$. The solid and dashed curves show the visual and auditory modality-specific responses, respectively. The curves with $\times$ symbols show the cross-modal response, and the curves with $+$ symbols show the sum of the modality-specific responses. Responses with and without modulatory connections are shown in *A* and *B*, respectively. Responses at all input levels are subadditive without modulatory connections (*B*) but can be supra-additive over a narrow range with modulatory connections (*A*). V, Visual; A, auditory; S, somatosensory.

between the primary inputs and the target is 2.27 bits (Table 3). The information content of the target is 2.32 bits. Thus, the primary input in this case contains almost complete target information. The modulatory inputs have spontaneous and driven activation probabilities of $p_{y0} = 0$ and $p_{y1} = 0.1$, respectively. The mutual information between the modulatory inputs and the target (1.74 bits) is lower than that of the primary inputs in this case. Modulatory inputs can increase the estimate of DSC information gain by producing MSE and helping DSC unit activities exceed threshold θ_I. Stage two provides a small increase in information gain that is significant when the percentage of multisensory DSC units is 60% or larger (*t* test, 0.05 significance level).

The most striking feature of the plot in Figure 6 is that target information gain at the DSC is highest for percentages of multisensory DSC units between 10 and 50% and falls steadily as the percentage of multisensory DSC units rises above 50%. Insight into this result can be obtained through comparison with a DSC network in which all of the units are trimodal and receive primary connections of identical weight of all three modalities. To make a uniformly trimodal DSC, all primary weights are set to $\sqrt{1/3}$ ($u_{ij} = \sqrt{1/3}$ for all *i* and *j*). This sets the lengths of the primary weight vectors to one, to match the lengths of the normalized primary weight vectors produced by stage one. Stage-two training can start from the uniform, trimodal primary weight configuration. The target information gain of the uniformly trimodal DSC network is only 0.77 bits without modulatory input. It increases to only 0.80 bits with modulatory input. The target information gain of the DSC network trained from a random state with the two-stage algorithm approaches this low level as the percentage of multisensory DSC units increases.

These results demonstrate that a uniformly trimodal DSC network, in which all units respond identically to all targets, is very uninformative. Target information gain in trained DSC networks with 100% multisensory units is somewhat higher, because primary weight vectors in trained networks are non-uniform, and units may vary in their activation by different targets. Still, the results clearly indicate that target information gain is highest when the DSC contains between 10 and 50% multisensory units. Networks in this range have a mixture of unimodal, bimodal, and

trimodal units and best convey information concerning the target in its various states.

The results shown in Figure 6 are representative of those obtained with different primary input ambiguities and proportions of modality-specific to cross-modal targets. As suggested above, the actual percentage of multisensory DSC neurons found in a species may reflect the combined effects of sensory input ambiguity and the proportion of cross-modal targets encountered in its environmental niche. As the results in this section suggest, the proportions of unimodal, bimodal, and trimodal DSC neurons may also reflect the needs of an organism for target information gain by the DSC.

Simulating multisensory enhancement in the DSC

MSE requires input from unimodal regions of parietal cortex. Inactivation of these regions can drastically reduce MSE but may have little effect on the modality-specific responses of DSC neurons (Wallace and Stein, 1994; Jiang et al., 2001). Stage two is designed to produce MSE at the DSC using unimodal modulatory inputs (see Materials and Methods). The result is that cross-modal responses can be significantly larger than modality-specific responses and that MSE depends on modulatory connections. MSE is examined in a network trained with primary input of intermediate ambiguity (activation probabilities are $p_{x0} = 0.1$ and $p_{x1} = 0.6$) and targets that are twice as likely to be modality-specific as cross-modal ($p_s = 0.34$ and $p_c = 0.17$). After training, the responses of DSC units to two-modality targets show MSE (Fig. 7).

DSC unit responses are determined for modality-specific or two-modality targets (target states $t = 1$ to $t = 6$; see Materials and Methods). If the target presents the modality specific to primary input X_j, then the activity of that primary input is increased from 0 to 20 in steps of 1 ($n = 20$ is the number of binary variables in the binomial processes that define the input likelihoods; see Materials and Methods). If the target does not present the modality specific to X_j, then the activity of that primary input is fixed at the mean of its spontaneous likelihood, which is 2 (i.e., 20 times the primary input spontaneous activation probability p_{x0}). If the target presents the modality specific to modulatory input Y_k, then the activity of that modulatory input is increased from 0 to 4 in steps of 0.2. The smaller range of the modulatory compared with the primary inputs is meant to reflect the five times greater dynamic range of primary compared with modulatory inputs ($p_{x0} = 0.1$ and $p_{x1} = 0.6$, whereas $p_{y0} = 0$ and $p_{y1} = 0.1$). Modulatory input Y_k takes value zero if the target does not present its specific modality, because the modulatory input spontaneous activation probability p_{y0} is zero. With the input specified, the DSC unit responses z_i are found by application of Equations 7 and 8.

DSC unit responses z_i are computed with and without modulatory connections. To find the responses without modulatory connections, the modulatory weights are simply set to zero. The results with and without modulatory connections are shown in Figure 7, *A* and *B*, respectively, for a bimodal visual–auditory DSC unit that is typical of the other multisensory DSC units in the network. Without modulatory connections (Fig. 7*B*), the cross-

modal responses (× symbols) can be larger than either of the two modality-specific responses (visual, solid line; or auditory, dashed line) but are smaller than the sum of the modality-specific responses (+ symbols) over the entire range. Thus, cross-modal responses are subadditive without modulatory connections. With modulatory connections (Fig. 7A) there is a range of input in which cross-modal responses are supra-additive.

DSC unit responses can be used to compute percentage MSE (%MSE) values according to Equation 16. For the responses of the bimodal, visual–auditory (V–A) DSC unit shown in Figure 7, maximal %MSE occurs when the visual and/or auditory primary inputs take the value of six. Responses at this level are detailed in Figure 8A–D, in which primary inputs take value six when they are driven by a target with the appropriate sensory attribute. When the appropriate target sensory attribute is absent, primary inputs are considered spontaneously active and take value two. Modulatory inputs take driven and spontaneous values of 6/5 = 1.2 or 0, respectively. Responses are shown with all modulatory connections intact (Fig. 8A) or with visual modulatory connections cut (Fig. 8B), auditory modulatory connections cut (Fig. 8C), or all modulatory connections cut (Fig. 8D).

For a two-modality target (V, A), both the visual and auditory primary and modulatory inputs are driven. For a modality-specific target (V only or A only), one primary and one modulatory input are driven while the others are spontaneous. With all modulatory connections intact (Fig. 8A), the cross-modal DSC unit response (V, A) is substantially larger than either of the two modality-specific responses (V only or A only), and %MSE equals 123%. Cutting modulatory connections reduces the amount of MSE. Cutting the visual or auditory modulatory connections alone reduces %MSE to 86 and 75%, respectively. Cutting both sets of modulatory connections reduces %MSE to 39%. Thus, MSE for DSC units in the corticotectal model depends on unimodal modulatory connections, which correspond to descending projections from neurons in unimodal regions of parietal cortex. Individual DSC units in the model can be modulated by more than one cortical region, and the reduction in cross-modal responses is greater when modulatory connections from multiple regions are interrupted. These effects, which are observed for all other multisensory DSC units in this network, are consistent with experimental findings (Wallace and Stein, 1994).

Cortical cooling experiments show that inactivation of multiple regions of parietal cortex can not only reduce MSE but, in some cases, can eliminate MSE entirely (Wallace and Stein, 1994; Jiang et al., 2001). Elimination of enhancement brings cross-modal responses to the level of the largest modality-specific response. In rare cases, inactivation of regions of parietal cortex can produce negative enhancement, in which the cross-modal response is actually smaller than the largest modality-specific response (Jiang et al., 2001). Complex single-neuron models, involving multiplicative nodes and inhibitory connections, can simulate MSE at any level whether positive, negative, or zero (Patton and Anastasio, 2003). Modified versions of these complex neural elements could be used as DSC units and would allow the corticotectal model to simulate zero or negative enhancement after removal of modulatory connections. For simplicity in this initial presentation, only simple neural elements are used in the corticotectal model (Eq. 7). For that reason, the model can simulate the reduction in MSE brought about by cortical inactivation but cannot currently simulate zero or negative enhancement.
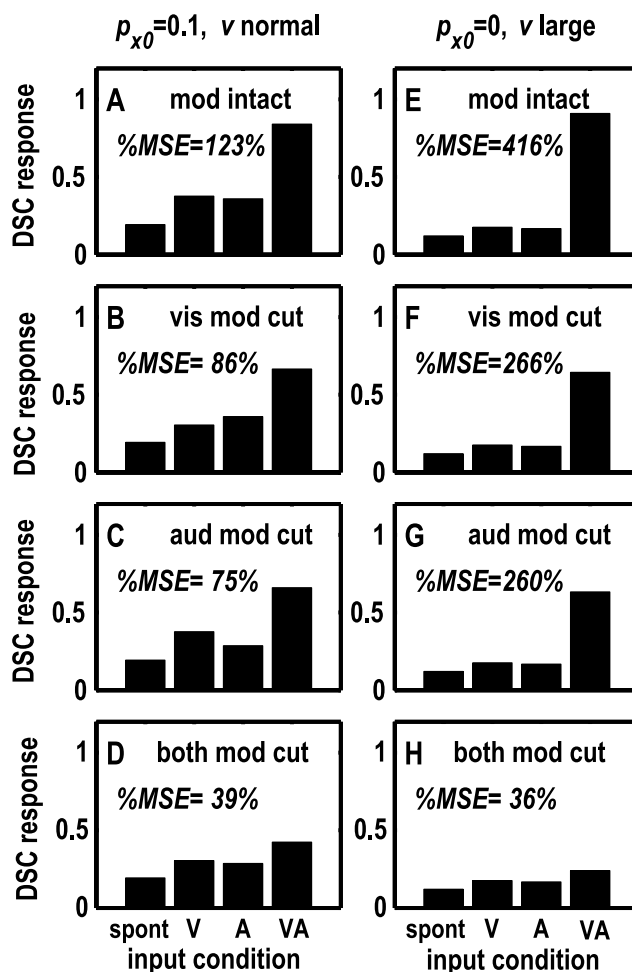
Figure 8. MSE in the corticotectal model depends on unimodal modulatory inputs. Responses of the bimodal, visual–auditory DSC unit from Figure 7 are computed, to targets of visual (V), auditory (A), both (V, A), or neither (spont) modality. Active primary inputs were assigned a value of six, because this value was found to produce maximal MSE for this DSC unit. Responses are shown for the intact model (A) and after interruption of modulatory connections of the visual (B), auditory (C), or both (D) modalities. Interruption of modulatory connections has little effect on modality-specific responses but greatly decreases cross-modal responses and reduces MSE. The reduction in MSE is greatest when modulatory connections of both modalities are interrupted (D). This result is a consequence of training with the correlation–anti-correlation rule in stage two, which produces cross-modal but not modality-specific modulatory connections. The amount of MSE in the model is affected by both the magnitude of modulatory weights and the primary input spontaneous activation probability. In E–H, the modulatory weights have been increased by seven times (v large), and the spontaneous activation probability for the primary inputs has been reduced to zero ($p_{x0} = 0$). Active primary inputs are assigned a value of three, because this value now produces maximum MSE. Responses are shown for the intact model (E) and after interruption of the modulatory connections of the visual (F), auditory (G), or both (H) modalities. The effect of the modulatory connections is qualitatively the same as before ($p_{x0} = 0.1$, v normal), but maximal percentage enhancement is higher. Also, the effect of removal of modulatory connections is greater than before for cross-modal responses and nil for modality-specific responses.

Spontaneous activity may limit multisensory enhancement in the DSC

MSE occurs whenever the cross-modal response is larger than the maximal modality-specific response (Meredith and Stein, 1986a; Stein and Meredith, 1993). While percentage MSE in excess of 1000 has been observed, most reported enhancements are considerably smaller than that (Meredith and Stein, 1986a; Wallace and Stein, 1997; Jiang et al., 2001). The model suggests that the

spontaneous activity of direct, excitatory inputs to the DSC, which would be considered primary in the model, limits the amount of MSE. Descending inputs from parietal cortex would modulate the spontaneous as well as the driven activity of primary inputs, and this could affect the ability of descending inputs to produce large cross-modal enhancements.

This effect can be seen in the simulated responses of Figure 8. Although the main effect of modulation is on cross-modal responses, small effects can be discerned on modality-specific responses. In the modality-specific case, an active modulatory input of one modality can modulate the spontaneous activity of primary inputs of a different modality. For example, cutting the visual modulatory connection (Fig. 8*B*) causes a slight reduction in the response to a unimodal visual target (V only). This is not attributable to removal of visual modulation of the visual primary input, because the cross-modality constraint already excludes the visual–modulatory to visual–primary connection. The reduction occurs because the cut visual–modulatory connection no longer modulates the spontaneous activity of the auditory primary input. As the spontaneous activity is increased toward the driven activity of the primary input, the effect of modulatory input on modality-specific responses gets bigger, and the amount of MSE gets smaller.

Very large enhancements can be produced when there is very low spontaneous primary input activity in the corticotectal model. To illustrate this, the same bimodal, visual–auditory DSC unit shown in Figure 8*A–D* is examined again in Figure 8*E–H*, but the spontaneous activation probability of the primary inputs ($p_{x0} = 0.1$) is reduced to zero ($p_{x0} = 0$). The modulatory weights, deliberately kept small during stage-two training by imposing an upper bound (v normal), are increased by seven times (v large). Now maximal enhancement is observed at a primary input level of three. As before, MSE depends on unimodal modulatory input from multiple sources, but now the magnitude of enhancement is much greater, and the effect of modulation on modality-specific responses is nil. In principal, with zero spontaneous activity, the modulatory weights and the amount of MSE could be increased without bound. Contrariwise, limitations imposed by the presence of primary input spontaneous activity may explain the typically low percentage enhancements observed for most DSC neurons (Wallace and Stein, 1997; Jiang et al., 2001).

Even without increasing the modulatory weights, simply removing the spontaneous input to the DSC unit of Figure 8 doubles its maximal %MSE (data not shown). The model predicts that MSE should be increased by factors that decrease the spontaneous activity of primary inputs. Anesthesia may be one such factor. Experiments that uncovered large enhancements were conducted on anesthetized cats (Meredith and Stein, 1986a,b, 1996; Wallace and Stein, 1997; Kadunce et al., 2001). In contrast, experiments in alert, behaving cats failed to reveal large enhancements (Populin and Yin, 2002). The model opens the possibility that the larger enhancements seen in anesthetized animals may be attributable, in part, to a reduction by anesthetic of the spontaneous rate of primary inputs. This possibility could be explored experimentally.

The spontaneous activity of the modulatory inputs could also limit the amount of MSE. Spontaneous firing of the modulatory inputs would enhance the ongoing spontaneous activity of the primary inputs and produce potentially large DSC unit activations in the absence of targets. Limits on the strength of modulatory connections would reduce the magnitude of such spurious enhancements but would also reduce appropriate enhancements. This trade-off is avoided entirely in the model by setting the

spontaneous activity of the modulatory inputs to zero. This solution is based on findings from anesthetized cats showing that AES neurons have very low spontaneous rates (1–8 Hz) (Mucke et al., 1982). Data from alert animals on the spontaneous rates of neurons in AES, and other parietal areas projecting to DSC, are currently lacking. Presumptive excitatory pyramidal neurons in other cortical areas are known to exhibit low spontaneous rates in alert cats (9.4 ± 1.7 Hz) (Steriade et al., 2001).

Discussion

The two-stage algorithm works in a local, unsupervised, and neurobiologically plausible way. Training the corticotectal model using the two-stage algorithm causes the tectal component, which represents the DSC, to extract a substantial amount of target information from its inputs. The model offers possible answers to two of the most pressing questions concerning multisensory integration in the DSC: why some but not all DSC neurons are multisensory, and how MSE exhibited by multisensory DSC neurons could be produced through descending input from unimodal parietal cortical neurons. The corticotectal model provides insight into how MSE might be produced in the actual nervous system.

Information gain and modality specialization

The DSC receives input of three sensory modalities. Despite the potential availability of trimodal input, most DSC neurons only respond to stimuli of one or two sensory modalities (cat, 43% unimodal, 45% bimodal; monkey, 73% unimodal, 21% bimodal) (Wallace and Stein, 1996). Trimodal neurons are rarely observed in the DSC (cat 9%; monkey 6%) (Wallace and Stein, 1996). A principle result of the corticotectal model is the demonstration that a DSC composed of a mixture of unimodal, bimodal, and trimodal units extracts substantially more target information from its inputs than a uniformly trimodal DSC. The model also demonstrates how such a mixture of modality selectivities could emerge automatically from an unsupervised learning process.

That process, used in stage one of the two-stage algorithm, is the SOM algorithm (Willshaw and von der Malsburg, 1976; Kohonen, 1982, 1988; Haykin, 1999). The SOM is neurobiologically plausible because it is based on a local Hebb rule. The process of selection of the winner and its neighborhood in the DSC could occur through the type of burst production that is involved in the generation of saccadic commands (Wurtz and Goldberg, 1972; Munoz and Wurtz, 1995). Lateral connectivity profiles, consisting of short-range excitation and long-range inhibition, have been identified in this structure (McIlwain, 1982; Meredith and Ramoa, 1998; Munoz and Istvan, 1998). This connectivity could mediate a winners-take-all process in the DSC.

The SOM has been widely applied in modeling map formation in the brain (Udin, 1988). Whereas findings in molecular neuroscience underscore the importance of activity-independent processes in map formation (Flanagan and Vanderhaeghen, 1998), the SOM remains an important model of the activity-dependent processes that refine those maps (Katz and Shatz, 1996; Cline, 1998; Zhang et al., 1998). Activity-dependent refinement may have as much to do with information extraction as with map formation. Linsker (1988a, b) has shown that the SOM, by essentially creating a neighborhood of specialists, causes a network to extract information from its inputs.

Ambiguous inputs carry less target information than do unambiguous inputs (Table 3). Previous theoretical work suggested, on that basis, that unimodal DSC units receive unambiguous input of one modality, but that multisensory DSC units

integrate ambiguous inputs of multiple modalities to increase the amount of target information they receive (Patton et al., 2002). That view, which considers DSC units individually rather than collectively, should be broadened in light of the results of the corticotectal model. Input ambiguity can increase the percentage of multisensory DSC units produced in the model during training, and this is consistent with the previous theory. The percentage of multisensory DSC units in the model also increases with the proportion of cross-modal targets presented during training. However, the model DSC as a whole extracts the most target information from its inputs when the percentage of multisensory DSC units falls between 10 and 50%. In the model, the tendency for ambiguous inputs and cross-modal targets to increase the percentage of multisensory DSC units would have to be balanced by the need for the model DSC as a whole to extract target information. The percentage of multisensory DSC neurons in the brain similarly may be determined by multiple factors.

Descending modulation and multisensory enhancement

The other principle result of the corticotectal model is that it reproduces findings on MSE in the DSC, which requires descending input from parietal cortex. Inactivation of parietal cortical neurons, in the AES or LS area of the cat, reduces MSE but may have little effect on the responses of DSC neurons to modality-specific stimulation (Wallace and Stein, 1994; Jiang et al., 2001). Paradoxically, the parietal projections critical for MSE originate from unimodal, not multisensory, neurons (Wallace et al., 1993). These data argue against a direct excitatory effect of descending parietal projections onto DSC neurons.

The paradox is resolved in the corticotectal model by treating the relevant descending projections from parietal cortex as modulatory. There are a variety of neural mechanisms that might mediate the proposed modulation of excitatory input. Experimental evidence suggests that NMDA-sensitive receptors may be involved in amplifying the responses of DSC neurons (Binns and Salt, 1996; Binns, 1999). Presynaptic enhancement by metabotrophic glutamate receptors (Anwyl, 1999) is another possible way in which descending parietal projections could modulate DSC neuron responses. Ultrastructural studies of somatosensory terminals in the DSC (Harting et al., 1997) suggest a possible neuroanatomical substrate for modulation. Ascending trigeminal somatosensory inputs terminate on small, presumably distal dendrites of DSC neurons, whereas descending cortical somatosensory inputs terminate on proximal dendrites. This synaptology suggests a gating role for descending projections. These data lend support to the idea that many cortical descending projections to DSC are modulatory rather than directly excitatory.

Many DSC neurons are activated at short latencies by electrical stimulation of corticotectal regions of parietal cortex (Wallace et al., 1993). This observation could be taken parsimoniously as evidence for monosynaptic excitation. It could instead result from activation of a modulatory input as postulated here, given a constant subthreshold level of excitation at the primary inputs. Activation of modulatory input attributable to electrical stimulation of the cortex could augment otherwise subthreshold primary input activity, thereby activating DSC neurons at short latency.

The main features of MSE, as observed for multisensory DSC neurons, are that the cross-modal response is larger than the maximal modality-specific response (in many cases, even larger than the sum of the modality-specific responses), and the amount of enhancement is magnitude dependent, decreasing as the magnitudes of the modality-specific responses increase (Meredith and Stein, 1986a). Previous theoretical work showed that these features are

consistent with the hypothesis that DSC neurons use their sensory inputs to compute the probability that a target has appeared (Anastasio et al., 2000; Patton and Anastasio, 2003). That theory does not explain findings regarding the cortical role in MSE, but the corticotectal model presented here does. The corticotectal model also simulates MSE and the magnitude dependency of MSE (Fig. 7A), but it does not compute target probabilities.

As the output nodes of a neural network, DSC units could be trained using supervised learning to estimate target probabilities to arbitrary accuracy (Bishop, 1995). The unsupervised two-stage algorithm does not endow DSC units with that capability, and it is not clear that any unsupervised scheme could do so. It is possible that something like the two-stage algorithm sets up the basic corticotectal circuitry, but that some form of supervised learning must tune that circuitry to accurately compute target probabilities. Although the modulatory inputs that produce MSE provide little in the way of information gain (Fig. 6), they can produce augmentation of cross-modal responses that could easily cause DSC units to overestimate target probabilities. This raises the intriguing possibility that the corticotectal circuit may be tuned by inhibition. Such inhibition could arise from a number of sources, including the well studied projection to DSC from substantia nigra (Hikosaka and Wurtz, 1983, 1985a,b; Mize, 1992).

Although the two-stage algorithm does not endow DSC units with the ability to compute target probabilities, the correlation–anti-correlation rule in stage two of the algorithm is based on a probabilistic argument. If primary and modulatory inputs are consistently active together, then their coactivation does not indicate a higher target probability and descending cortical modulation should be reduced. However, if primary and modulatory inputs are not consistently active together, then their coactivation does indicate a higher target probability and descending cortical modulation should be increased. The rule also depends on the coactivation of DSC and cortical units and on the modality selectivity of DSC units established in stage one of the algorithm. The correlation–anti-correlation rule is local and neurobiologically plausible, especially given recent evidence for anti-Hebbian forms of synaptic plasticity (Linden, 1995). The correlation–anti-correlation rule and the corticotectal model provide a new view of top-down organization and processing in the corticotectal system.

References

Anastasio TJ, Patton PE, Belkacem-Boussaid K (2000) Using Bayes' rule to model multisensory enhancement in the superior colliculus. Neural Comput 12:997–1019.

Anwyl R (1999) Metabotrophic glutamate receptors: electrophysiological properties and role in plasticity. Brain Res Rev 29:83–120.

Appelbaum D (1996) Probability and information: an integrated approach, Chap 5.7, pp 81–84. Cambridge, UK: Cambridge UP.

Binns KE (1999) The synaptic pharmacology underlying sensory processing in the superior colliculus. Prog Neurobiol 59:129–159.

Binns KE, Salt TE (1996) Importance of NMDA receptors for multimodal integration in the deep layers of the cat superior colliculus. J Neurophysiol 75:920–930.

Bishop CM (1995) Neural networks for pattern recognition. Oxford: Clarendon.

Clemo HR, Stein BE (1986) Effects of cooling somatosensory cortex on response properties of tactile cells in the superior colliculus. J Neurophysiol 55:1352–1368.

Cline HT (1998) Topographic maps: developing roles for synaptic plasticity. Curr Biol 8:R836–R839.

Cover TM, Thomas JA (1991) Elements of information theory, Chap 2, pp 12–49. New York: Wiley.

Flanagan JG, Vanderhaeghen P (1998) The ephrins and EPH receptors in neural development. Annu Rev Neurosci 21:309–345.

Harting J, Feig S, Van Lieshout D (1997) Cortical somatosensory and tri-

geminal inputs to the cat superior colliculus: light and electron microscopic analyses. J Comp Neurol 388:313–326.

Haykin S (1999) Neural networks: a comprehensive foundation, Chap 9, Ed 2, pp 443–483. Upper Saddle River, NJ: Prentice-Hall.

Hikosaka O, Wurtz RH (1983) Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. J Neurophysiol 49:1285–1301.

Hikosaka O, Wurtz RH (1985a) Modification of saccadic eye movements by GABA-related substances. I. Effect of muscimol and bicuculline in monkey superior colliculus. J Neurophysiol 53:266–291.

Hikosaka O, Wurtz RH (1985b) Modification of saccadic eye movements by GABA-related substances. II. Effects of muscimol in monkey substantia nigra pars reticulata. J Neurophysiol 53:292–308.

Hoel PG, Port SC, Stone CJ (1971) Introduction to probability theory, Chap 3.4.2, pp 69–70. Boston: Houghton Mifflin.

Jiang W, Wallace MT, Jiang H, Vaughan JW, Stein BE (2001) Two cortical areas mediate multisensory integration in superior colliculus neurons. J Neurophysiol 85:506–522.

Jiang W, Jiang H, Stein BE (2002) Two corticotectal areas facilitate multisensory orientation behavior. J Cogn Neurosci 14:1240–1255.

Kadunce DC, Vaughan JW, Wallace MT, Stein BE (2001) The influence of visual and auditory receptive field organization on multisensory integration in the superior colliculus. Exp Brain Res 139:303–310.

Katz LC, Shatz CJ (1996) Synaptic activity and the construction of cortical circuits. Science 274:1133–1138.

King AJ, Palmer AR (1985) Integration of visual and auditory information in bimodal neurones in guinea-pig superior colliculus. Exp Brain Res 60:492–500.

Kohonen T (1982) Self-organized formation of topologically correct feature maps. Biol Cybern 44:59–69.

Kohonen T (1988) Self organization and associative memory, Ed 2. Berlin: Springer-Verlag.

Lichtman JW, Burden SJ, Culican SM, Wong ROL (1999) Synapse formation and elimination. In: Fundamental neuroscience (Zigmond MJ, Bloom FE, Landis SC, Roberts JL, Squire LR, eds), pp 547–580. San Diego: Academic.

Linden DJ (1995) Long-term synaptic depression. Annu Rev Neurosci 18:319–357.

Linsker R (1988a) Self-organization in a perceptual network. Computer 21:105–117.

Linsker R (1988b) Towards an organizing principle for a layered perceptual network. In: Neural information processing systems (Anderson DZ, ed), pp 485–494. New York: American Institute of Physics.

McIlwain JT (1982) Lateral spread of neural excitation during microstimulation in intermediate gray layer of cat's superior colliculus. J Neurophysiol 47:167–178.

Meredith MA, Clemo HR (1989) Auditory cortical projection from the anterior ectosylvian sulcus (field AES) to the superior colliculus in the cat: an anatomical and electrophysiological study. J Comp Neurol 289:687–707.

Meredith MA, Ramoa AS (1998) Intrinsic circuitry of the superior colliculus: pharmacophysiological identification of horizontally oriented inhibitory interneurons. J Neurophysiol 79:1593–1602.

Meredith MA, Stein BE (1986a) Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration. J Neurophysiol 56:640–662.

Meredith MA, Stein BE (1986b) Spatial factors determine the activity of multisensory neurons in cat superior colliculus. Brain Res 365:350–354.

Meredith MA, Stein BE (1990) The visuotopic component of the multisensory map in the deep laminae of the cat superior colliculus. J Neurosci 10:3727–3742.

Meredith MA, Stein BE (1996) Spatial determinants of multisensory integration in cat superior colliculus neurons. J Neurophysiol 75:1843–1857.

Meredith MA, Clemo HR, Stein BE (1991) Somatotopic component of the multisensory map in the deep superior colliculus of the cat superior colliculus. J Comp Neurol 312:353–370.

Middlebrooks JC, Knudsen EI (1984) A neural code for auditory space in the cat's superior colliculus. J Neurosci 4:2621–2634.

Mize RR (1992) The organization of GABAergic neurons in the mammalian superior colliculus. Prog Brain Res 90:219–248.

Mucke L, Norita M, Benedek G, Creutzfeldt O (1982) Physiologic and anatomic investigation of a visual cortical area situated in the ventral bank of the anterior ectosylvian sulcus of the cat. Exp Brain Res 46:1–11.

Munoz DP, Istvan PJ (1998) Lateral inhibitory interactions in the intermediate layers of the monkey superior colliculus. J Neurophysiol 79:1193–1209.

Munoz DP, Wurtz RH (1995) Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cells. J Neurophysiol 73:2313–2333.

Patton P, Anastasio T (2003) Modeling cross-modal enhancement and modality-specific suppression in multisensory neurons. Neural Comput 15:783–810.

Patton P, Belkacem-Boussaid K, Anastasio T (2002) Multimodality in the superior colliculus: an information theoretic analysis. Brain Res Cogn Brain Res 14:10–19.

Populin LC, Yin TCT (2002) Bimodal interactions in the superior colliculus of the behaving cat. J Neurosci 22:2826–2834.

Robinson DA (1972) Eye movements evoked by collicular stimulation in the alert monkey. Vision Res 12:1795–1808.

Sparks DL, Hartwich-Young R (1989) The deep layers of the superior colliculus. In: The neurobiology of saccadic eye movements (Wurtz RH, Goldberg M, eds), pp 213–255. Amsterdam: Elsevier.

Stein BE (1998) Neural mechanisms for synthesizing sensory information and producing adaptive behaviors. Exp Brain Res 123:124–135.

Stein BE, Meredith MA (1993) The merging of the senses. Cambridge, MA: MIT.

Stein BE, Huneycutt WS, Meredith MA (1988) Neurons and behavior: the same rules of multisensory integration apply. Brain Res 448:355–358.

Stein BE, Meredith MA, Huneycutt WS, McDade L (1989) Behavioral indicies of multisensory integration: orientation to visual cues is affected by auditory stimuli. J Cogn Neurosci 1:12–24.

Steriade M, Timofeev I, Grenier F (2001) Natural waking and sleep states: a view from inside neocortical neurons. J Neurophysiol 85:1969–1985.

Udin SB (1988) Formation of topographic maps. Annu Rev Neurosci 11:289–327.

Wallace MT, Stein BE (1994) Cross-modal synthesis in the midbrain depends on input from cortex. J Neurophysiol 71:429–432.

Wallace MT, Stein BE (1996) Sensory organization of the superior colliculus in cat and monkey. Prog Brain Res 112:301–311.

Wallace MT, Stein BE (1997) Development of multisensory neurons and multisensory integration in cat superior colliculus. J Neurosci 17:2429–2444.

Wallace MT, Stein BE (2000) Onset of cross-modal synthesis in neonatal superior colliculus is gated by development of cortical influences. J Neurophysiol 83:3578–3582.

Wallace MT, Stein BE (2001) Sensory and multisensory responses in the newborn monkey superior colliculus. J Neurosci 21:886–8894.

Wallace MT, Meredith MA, Stein BE (1993) Converging influences from visual, auditory, and somatosensory cortices onto output neurons of the superior colliculus. J Neurophysiol 69:1797–1809.

Wallace MT, Wilkinson LK, Stein BE (1996) Representation and integration of multiple sensory inputs in primate superior colliculus. J Neurophysiol 76:1246–1266.

Wallace MT, Meredith MA, Stein BE (1998) Multisensory integration in the superior colliculus of the alert cat. J Neurophysiol 20:1006–1010.

Wilkinson LK, Meredith MA, Stein BE (1996) The role of anterior ectosylvian cortex in cross-modality orientation and approach behavior. Exp Brain Res 112:1–10.

Willshaw DJ, von der Malsburg C (1976) How patterned neural connections can be set up by self-organization. Proc R Soc Lond B Biol Sci 194:431–445.

Wurtz RH, Goldberg ME (1972) Activity of superior colliculus in behaving monkey. III. Cells discharging before eye movements. J Neurophysiol 35:575–586.

Zhang LL, Tao HW, Holt CE, Harris WA, Poo M-M (1998) A critical window for cooperation and competition among developing retinotectal synapses. Nature 395:37–44.