

Learning Cross-Modal Spatial Transformations through Spike Timing-Dependent Plasticity

Andrew P. Davison and Yves Frégnac

Unité de Neurosciences Intégratives et Computationnelles, Centre National de la Recherche Scientifique, 91198 Gif sur Yvette, France

A common problem in tasks involving the integration of spatial information from multiple senses, or in sensorimotor coordination, is that different modalities represent space in different frames of reference. Coordinate transformations between different reference frames are therefore required. One way to achieve this relies on the encoding of spatial information with population codes. The set of network responses to stimuli in different locations (tuning curves) constitutes a set of basis functions that can be combined linearly through weighted synaptic connections to approximate nonlinear transformations of the input variables. The question then arises: how is the appropriate synaptic connectivity obtained? Here we show that a network of spiking neurons can learn the coordinate transformation from one frame of reference to another, with connectivity that develops continuously in an unsupervised manner, based only on the correlations available in the environment and with a biologically realistic plasticity mechanism (spike timing-dependent plasticity).

Key words: spatial transformations; population coding; correlation-based learning; spike timing-dependent plasticity; proprioception; vision

Introduction

To integrate information from different senses and to perform sensory-guided motor tasks, the brain must perform transformations between different frames of reference. In general, these transformations must be learned and remain plastic, to some degree, in adult life.

A biologically plausible approach to performing such computations is to use basis functions (Salinas and Abbott, 1995; Pouget and Sejnowski, 1997; Pouget and Snyder, 2000), a general method for approximating nonlinear functions. In a neuronal context, the population response patterns for different stimulus locations serve as a set of basis functions. The desired transformation function can be approximated by taking a weighted sum of basis functions, i.e., by connecting the population representing the input variable to a population representing the transformed variable, with appropriate synaptic weights. Basis function networks have been used to model both sensorimotor transformations (Salinas and Abbott, 1995; Pouget and Sejnowski, 1997) and cross-modal integration (Deneve et al., 2001) with predictions that are in good agreement with experimental results.

For basis function networks to be a viable candidate for the actual mechanism of spatial transformations in biological nervous systems, we must show that the appropriate pattern of weights can be learned based only on interactions with the envi-

ronment and using biologically realistic learning rules. These criteria are not met by existing models of basis function networks, for which synaptic weights are set by hand (Pouget and Sejnowski, 1994; Deneve et al., 2001; van Rossum and Renart, 2004), learned with a supervised learning rule (Pouget and Sejnowski, 1997) or, with a Hebbian-like rule, learned during a separate training phase, in which the outputs are not driven by the inputs but are set to the desired values (Salinas and Abbott, 1995).

We hypothesized that spike timing-dependent synaptic plasticity (STDP) (for review, see Dan and Poo, 2004) could form the substrate for learning spatial transformations in a basis function network. In a spiking network model with population coding of stimulus location and with STDP, the interaction of input correlations and lateral connections imposes structure on the connectivity pattern so as to form cortical maps (Song and Abbott, 2001), i.e., the simple topographic mapping $x \rightarrow x$. It occurred to us that if a different structure could be imposed, then more complex transformations, $x \rightarrow f(x)$, could be learned.

Because potentiation–depression is favored by correlated–uncorrelated inputs, respectively (Song et al., 2000), the appropriate pattern of weights may be achieved by imposing an appropriate pattern of correlations. For learning cross-modal or sensorimotor spatial transformations, an obvious source for these correlations is the input from the other modality or measurement of the motor outcome. An output population having a preexisting, simple topographic mapping from one modality may learn to perform a more complex transformation from a different modality, because connections from input neurons selective for x to those output neurons being activated by a “training” input $f(x)$ (through the preexisting topographic connections) will be potentiated, whereas other connections will be depressed. Ultimately, an input x will elicit an output $f(x)$, even in the absence of the training input.

Received June 20, 2005; revised April 7, 2006; accepted April 11, 2006.

This work was supported by the European Union under the Life-Like Perception Programme, project reference IST-2001-34712 (SenseMaker) and the Bio-inspired Intelligent Information Systems program, project reference IST-2004-15879 (FACETS). We thank A. Destexhe and M. Rudolph for helpful comments on this manuscript.

Correspondence should be addressed to Andrew P. Davison, Unité de Neurosciences Intégratives et Computationnelles, Centre National de la Recherche Scientifique, 91198 Gif sur Yvette, France. E-mail: davison@iaf.cnrs-gif.fr.

DOI:10.1523/JNEUROSCI.5263-05.2006

Copyright © 2006 Society for Neuroscience 0270-6474/06/265604-12\$15.00/0

Here we show that such a system is indeed able to learn coordinate transformations without an external supervisor. We explore the conditions for successful learning and make specific predictions, for example, of interactions between the temporal scales of STDP and of the learning environment and of the transformation accuracy as a function of stimulus location.

Materials and Methods

Model description. For functions of one variable, the network consists of three populations, each consisting of a one-dimensional array of cells, of size N (see Fig. 1). The “input population” makes all-to-all excitatory connections to the “output population.” The strength of these connections is modifiable by STDP, with a maximum synaptic weight of w_{\max} . Connections from the “training population” to the output population are topographic and local, in the sense that neurons in the middle of the presynaptic population project only to neurons in the middle of the postsynaptic population, etc.:

$$w_{ij}^{Tr} = \begin{cases} \rho_{Tr} w_{\max} & \text{if } |d_{ij}| \leq \sigma_{Tr} \\ 0 & \text{if } |d_{ij}| > \sigma_{Tr} \end{cases}, \quad (1)$$

where w_{ij}^{Tr} is the connection weight between presynaptic neuron i and postsynaptic neuron j (see text after Eq. 4), d_{ij} is the shortest distance between i and j given the periodic boundary conditions ($0 \leq i, j < N$), σ_{Tr} is the range of connections, and ρ_{Tr} is the ratio of training–output to input–output maximum weights. (All parameter values are given in supplemental Table 1, available at www.jneurosci.org as supplemental material). In the standard model, the strength of these connections is fixed. We also consider the case in which these connections are modified by STDP. In this case, we also add local excitatory and global (all-to-all) inhibitory connections within the output population. Specification of excitatory connections follows Equation 1, but with ρ_{Tr} and σ_{Tr} replaced with ρ_{Rc} and σ_{Rc} , respectively. The strength of the global inhibitory connections is given by $\rho_{Inh} w_{\max}$. We use periodic boundary conditions, i.e., a neuron in the training population at position N will make connections to the neuron at position 1 as well as that at position $N-1$. It is possible to learn multiple functions of a given input variable simultaneously: for each function, one training and one output population is required. For functions of two variables, the input population is a two-dimensional array of size $N \times N$. The synaptic time delays are zero for the connections from the input and training populations to the output population and for inhibitory connections within the output population, and 2 ms for excitatory connections within the output population. Each cell in the output population receives excitatory, “background” Poisson input at 1000 Hz through a synapse with weight w_b . This single high-frequency input is equivalent to the combined inputs of a large number of lower-frequency afferents.

Because the input and training populations do not themselves receive any input, they are modeled simply as spike sources. The output population cells are modeled as integrate-and-fire neurons. The membrane potential of neuron j , V_j , obeys the following:

$$\tau_m \frac{dV_j}{dt} = -V_j(t) + R_m I_j^{syn}, \quad (2)$$

where τ_m is the membrane time constant, R_m is the membrane resistance and

$$I_j^{syn} = I_j^b(t) + \sum_i (I_{ij}^{In}(t) + I_{ij}^{Tr}(t) + I_{ij}^{Rc}(t) - I_{ij}^{Inh}(t)), \quad (3)$$

where I_j^b is the background noise synaptic current and $I_{ij}^k(t)$ ($k = \{In, Tr, Rc, Inh\}$) are the synaptic currents from the input, training, recurrent excitatory, and recurrent inhibitory connections. Note that the recurrent connections are only used for the subset of simulations in which the training–output connections are not fixed but are subject to plasticity (see Fig. 9), in which case recurrent connections are necessary for correct weight-pattern formation. The cell fires an action potential when V_j reaches a threshold value of 1 and is then reset to 0. There is no refractory period. Excitatory synaptic inputs are modeled as a step change in the

synaptic current followed by exponential decay, i.e., the excitatory currents $I_{ij}^k(t)$ ($k = \{In, Tr, Rc\}$) obey the following:

$$\tau_{exc} \frac{dI_{ij}^k}{dt} = -I_{ij}^k(t). \quad (4)$$

$I_{ij}^k(t)$ is incremented by an amount $w_{ij}^k I_0^{exc}$ after each spike, where w_{ij}^k is the synaptic weight and the value of I_0^{exc} is calculated such that a single presynaptic spike with $w_{ij}^k = 1$ will bring the membrane potential from rest (zero) precisely to spike threshold (1). This normalization means that the precise value of R_m is unimportant. The background noise current evolves similarly. Inhibitory synaptic currents follow an α -function-like time course. The inhibitory current $I_{ij}^{Inh}(t)$ obeys the following:

$$\tau_{inh} \frac{dI_{ij}^{inh}}{dt} = -I_{ij}^{inh}(t) + \xi_{ij}(t), \quad (5)$$

$$\tau_\xi \frac{d\xi_{ij}}{dt} = -\xi_{ij}(t). \quad (6)$$

ξ_{ij} is incremented by an amount $w_{inh} I_0^{inh}$ by every spike. The value of I_0^{inh} is calculated such that a single presynaptic spike with $w_{inh} = 1$ will bring the membrane potential from rest (zero) precisely to a membrane potential of -1 .

Use of a direct change in synaptic current rather than a change in conductance for the synaptic inputs is somewhat less realistic but has the advantage that the times of threshold crossings can be obtained by successive approximations that converge to the true value (Hines and Carnevale, 2004) rather than by integrating the differential equations. The former method is much faster than the latter. To test that this simplification did not appreciably affect our conclusions, we repeated some simulations with conductance-based (excitatory) synapses [the model described by Song and Abbott (2001)]; we found no qualitative difference in the formation of the weight pattern. The default value of w_{\max} used in our simulations was chosen to match approximately the peak conductance parameter g_{\max} used by Song and Abbott (2001). From the resting membrane potential, their default value of $g_{\max} = 0.02$ (in units of the leakage conductance) produces a depolarization of amplitude 0.233 mV, or 1.2% of the difference between rest and firing threshold. This amplitude would be somewhat smaller near threshold, because the driving force would be reduced from 74 to 54 mV. In our model, the default value of $w_{\max} = 0.02$ produces a depolarization of amplitude of which is 2% of the difference between rest and threshold. Thus, the current-based and conductance-based models are comparable, although not identical.

STDP models. In most of our simulations we use the mechanism for STDP proposed by Song et al. (2000). This takes into account multiple spike interactions and not just the most recent spikes, with two functions, $P(t)$ and $M(t)$, for each synapse. $P(t)$ starts at zero and is incremented an amount A_+ by each presynaptic spike. $M(t)$ starts at zero and is decremented an amount A_- by each postsynaptic spike. Between increments–decrements, $P(t)$ and $M(t)$ decay exponentially toward zero with time constants τ_+ and τ_- , respectively. When a presynaptic spike occurs, the synaptic weight is modified according to $w \rightarrow w + w_{\max} M(t)$, where w_{\max} is the maximum possible synaptic weight (for brevity, we henceforth use w to represent any of the connection weights w_{ij}^k). Because $M(t)$ is zero or negative, this amounts to a depression of the weight. If this would take w below zero, w is set to zero. When a postsynaptic spike occurs, the synaptic weight is potentiated according to $w \rightarrow w + w_{\max} P(t)$. If this would make $w > w_{\max}$, w is set to w_{\max} . For a single pair of spikes, a presynaptic spike at time t_{pre} and a postsynaptic spike at time t_{post} with no recent preceding activity, the change in synaptic weight is given by the following (supplemental Fig. 1A, available at www.jneurosci.org as supplemental material):

$$\Delta w = \begin{cases} w_{\max} A_+ \exp[-(t_{pre} - t_{post})/\tau_+] & \text{if } t_{pre} < t_{post} \\ -w_{\max} A_- \exp[-(t_{post} - t_{pre})/\tau_-] & \text{if } t_{pre} \geq t_{post} \end{cases}. \quad (7)$$

The values of A_+ that we used, 0.01 and 0.001, were in the region of the value (0.005) determined previously (Song et al., 2000) from experimental measurements by dividing the amount of potentiation caused by mul-

multiple spike pairs by the number of pairs. The ratio $A_- \tau_- / A_+ \tau_+$ was taken as 1.06, as in Song and Abbott (2001), although the precise value of this ratio is not critical (see supplemental note and supplemental Fig. 6, available at www.jneurosci.org as supplemental material).

We also consider a model with soft bounds for the synaptic weights, rather than the hard bounds at zero and w_{\max} . In this case, $w \rightarrow w + wM(t)$ for depression and $w \rightarrow w + (w_{\max} - w)P(t)$ for potentiation, giving, for a single pair of spikes, the following:

$$\Delta w = \begin{cases} (w_{\max} - w)A_+ \exp[(t_{\text{pre}} - t_{\text{post}})/\tau_+] & \text{if } t_{\text{pre}} < t_{\text{post}} \\ -wA_- \exp[(t_{\text{post}} - t_{\text{pre}})/\tau_-] & \text{if } t_{\text{pre}} \geq t_{\text{post}} \end{cases} \quad (8)$$

In some simulations, we consider an STDP mechanism that is symmetric, i.e., whether the change is positive or negative does not depend on the sign of the interval Δt between presynaptic and postsynaptic spikes but only on its magnitude. For short intervals, we have potentiation; for long intervals, we have depression (supplemental Fig. 1B, available at www.jneurosci.org as supplemental material). This rule does not use the $P(t)$ and $M(t)$ functions, and so only interactions between nearest-neighbor spikes are considered, as follows:

$$\Delta w = w_{\max} A_{\text{symm}} \left(1 - \left(\frac{\Delta t}{\tau_a} \right)^2 \right) \exp(-|\Delta t|/\tau_b), \quad (9)$$

where $\Delta t = t_{\text{post}} - t_{\text{pre}}$, A_{symm} controls the step size for synaptic change, and τ_a and τ_b are parameters that control the shape of the weight-change function $f(\Delta t)$ (supplemental Fig. 1B, available at www.jneurosci.org as supplemental material). τ_a is the interval at which $f(\Delta t)$ crosses from positive to negative, and τ_b controls the range of intervals within which STDP occurs and is analogous to τ_+ and τ_- in the Song–Miller–Abbott model (Song et al., 2000). Such a symmetric form for STDP has been found at GABAergic synapses in hippocampal neurons (Woodin et al., 2003). The values for τ_a and τ_b were determined by approximate fitting of the weight pattern to that obtained with the symmetric rule (supplemental Fig. 2, available at www.jneurosci.org as supplemental material).

In the Song–Miller–Abbott model (Song et al., 2000), the balance of potentiation and depression is controlled by the ratio $A_+ \tau_+ / A_- \tau_-$. In the symmetric STDP model, the balance is controlled by the ratio τ_a / τ_b . If this ratio is greater than $\sqrt{2}$, the integral of $f(\Delta t)$ from zero to infinity is positive, and so potentiation will dominate.

Learning procedure. The input and training neurons generate spikes according to a Poisson process with a time-varying mean firing rate. The following procedure is identical to that of Song and Abbott (2001), but with a different spatial distribution of firing rates. An interval T is chosen from an exponential distribution with mean interval τ_{corr} . A random location θ in the input reference frame is chosen from a uniform distribution between 0 and 2π . From this, a location $f(\theta)$ in the training/desired output reference frame is calculated, where $f(\cdot)$ is the function that we wish the network to learn, and mapped onto the range 0 to 2π . The mean firing rate for a cell with a tuning curve that is centered at location Θ is then the following:

$$R(\theta', \Theta) = (R_{\max} - R_{\min}) \exp([\cos(\theta' - \Theta) - 1]/\sigma_R^2) + R_{\min}, \quad (10)$$

where

$$\theta' = \begin{cases} \theta & \text{for the input population} \\ f(\theta) & \text{for the training population} \end{cases} \quad (11)$$

Thus $R = R_{\max}$ at $\Theta = \theta'$. Note that $R(\cdot)$ is a periodic function, i.e., a neuron with a receptive field centered at one end of the input range will respond also to inputs at the opposite end of the range. Use of such periodic boundary conditions avoids the problem of edge effects, whereby locations at the edge of the input range would be encoded by a smaller population than locations in the center, although checks showed that the choice of periodic boundary conditions was not critical for the model (supplemental Fig. 7, available at www.jneurosci.org as supplemental material). The cells fire with rates determined by $R(\theta, \Theta)$ for a time T , and then another interval and a new location are chosen and the

process is repeated. Spike arrival times may be delayed by a constant latency. For our default case, the latency was chosen to be the same for the input and training populations, but the effect of non-zero differences in latency was investigated (see Results).

For the two-dimensional case, the mean firing rate for an input population neuron at location (Θ, Φ) is as follows:

$$R(\theta, \phi, \Theta, \Phi) = (R_{\max} - R_{\min}) \exp([\cos(\delta) - 1]/\sigma_R^2) + R_{\min}, \quad (12)$$

where

$$\delta = [(\theta - \Theta)^2 + (\phi - \Phi)^2]^{1/2}. \quad (13)$$

For the training population neurons, the tuning curve is the same as for the one-dimensional case.

Position estimation from population spiking activity. The question of how nervous systems can optimally read out a population code (in our case, estimate the stimulus location) is the subject of much study (Seung and Sompolinsky, 1993; Lewis and Kristan, 1998; Deneve et al., 1999). Because this is not the subject of our current investigation, we chose to use a simple, non-neural-based method as follows: each spike train was convolved with a Gaussian kernel with SD 100 ms and area 1, to give a smooth measure of mean firing rate over time. The mean location, \bar{y} , was calculated at 10 ms intervals by summing the neuron indices weighted by the firing rate of that neuron at that time and dividing by the sum of the firing rates across the population (Georgopoulos et al., 1986), as follows:

$$\bar{y} = \frac{\sum_{i=1}^N R_i \cdot i}{\sum_{i=1}^N R_i}, \quad (14)$$

where R_i is the firing rate of the neuron and N is the population size. The periodic boundary conditions introduced a little complexity into this method, because a peak could be split, one part at one end of the array and one at the other. The method was therefore modified as follows. For $\xi \leq N/2$:

$$\bar{y} = \left(\frac{\sum_{i=1}^{\xi + \frac{N}{2} - 1} R_i \cdot i + \sum_{i=\xi + \frac{N}{2}}^N R_i (i - N)}{\sum_{i=1}^{\xi + \frac{N}{2} - 1} R_i + \sum_{i=\xi + \frac{N}{2}}^N R_i} \right) / \sum_{i=1}^N R_i, \quad (15)$$

whereas for $\xi > N/2$:

$$\bar{y} = \left(\frac{\sum_{i=\xi - \frac{N}{2} + 1}^N R_i \cdot i + \sum_{i=1}^{\xi - \frac{N}{2}} R_i (i + N)}{\sum_{i=\xi - \frac{N}{2} + 1}^N R_i + \sum_{i=1}^{\xi - \frac{N}{2}} R_i} \right) / \sum_{i=1}^N R_i, \quad (16)$$

where ξ is the nearest integer to \bar{y} . To determine ξ , every i was tested as a value of ξ , and the i that gave both the smallest difference $|\xi - i|$ and the smallest variance for y was used for ξ .

Numerical methods. The model was implemented in the NEURON simulation environment in an event-based framework (Hines and Carnevale, 2004), with the CVODE integration method and an absolute tolerance parameter of 0.001. The NEURON code for the model is available from the SenseLab ModelDB database (<http://senselab.med.yale.edu/SenseLab/ModelDB>).

Results

For clarity, we consider a specific example of a coordinate transformation: calculating the position in visual, Cartesian coordinates of the hand on an arm with one or two degrees of freedom, based on the joint angle(s) obtained from proprioception. The arm is highly simplified, with two segments of equal length with free rotation in a plane about each of two joints.

During random arm movement, the network receives input from both sensory systems and learns the correlations between these signals at different points in space. After sufficient time, the

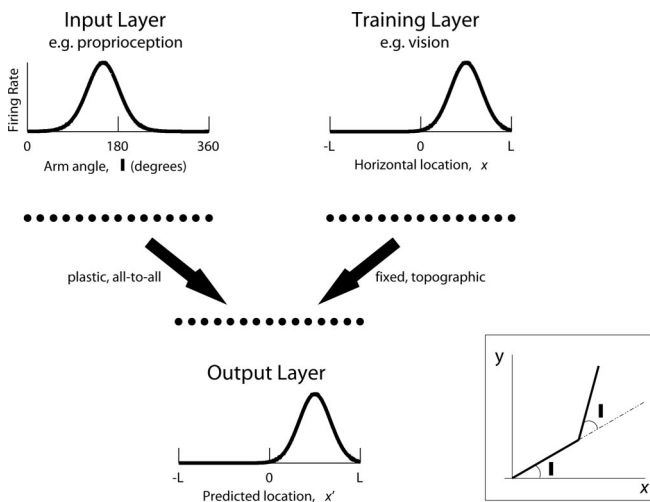


Figure 1. Network model for approximating functions of one variable. Spatial locations are represented with a population code. As an example, we consider moving an arm with 1 df and calculating the position of the end of the arm, in a vision-centered frame of reference, based on the angle at the joint. Connections are initially all-to-all from the input to output populations, and the strength of the connections is subject to modification by STDP. Connections from the training population to the output population are topographic and local in the sense that neurons in the middle of the presynaptic population project only to neurons in the middle of the postsynaptic population, etc. The strength of these connections is fixed. The inset shows an arm with 2 df. The one-dimensional case corresponds to $\phi = 0$.

visual input is no longer necessary, and the network can calculate the position of the hand in visual coordinates based only on the proprioceptive information. The network is shown in Figure 1. It consists of one population of cells for each input modality and one for the network output. Each population represents a spatial variable or location using a population code. The pattern of connections from the visual population to the output population is fixed and is assumed to have been developed at an earlier stage. The strengths of connections from the proprioceptive population to the output population are initially random and subject to change according to STDP. Full details are given in Materials and Methods.

Although we focus on a particular example, the network is capable, in principle, of learning any single-valued function of one variable. For generality, therefore, we label the three populations as the training (here vision), input (here proprioception) and output populations.

Learning spatial transformations

We consider first a one-dimensional coordinate transformation, which corresponds to an arm with one joint fixed. The function to be learned is $f(\theta) = \sin(\theta)$. After <1000 s, an “S”-shaped band becomes visible in the matrix of synaptic weights (Fig. 2A). As the training proceeds, connections within this band (henceforth described as the S-band) are potentiated, whereas those outside are depressed. The process converges by ~20,000 s (Fig. 2B). The final distribution of weights is strongly bimodal, in accord with the findings of Song et al. (2000). The mean firing rate across the population changes little over the course of the learning: from 7.0 Hz during the first 10 s to 8.4 Hz after 50,000 s. Despite the small change in mean firing rate across the population and a reduction in the total weight of all synapses onto a given cell during the course of training, there is a large increase in the incidence of small interspike intervals (ISIs) (Fig. 2C). This is in accord with

previous findings that STDP tends to reduce spike latency (Song et al., 2000; Guyonneau et al., 2005).

To test whether this S-shaped pattern of weights does indeed cause the network to perform the transformation $\sin(\theta)$ (i.e., enable the system to predict the location of the hand based only on proprioceptive information, with no visual input), the training input was removed (“eyes closed”), the STDP mechanism was inactivated, and then θ was swept over the entire range. The output population does indeed produce a good approximation to $\sin(\theta)$ (Fig. 2D), although it tends to underestimate near the extremes ($\sin(\theta) = \pm 1$). The reasons for this underestimation were investigated further (see below, Influence of population tuning curve shape and Fig. 6).

The mechanism for the formation of the bands of potentiated weights is the potentiation of correlated inputs shown by Song et al. (2000). Only weights that are activated simultaneously by both the input population and the training population are potentiated. Because of competition between synapses (Song et al., 2000), connections outside the bands are depressed. The center of each band is defined by those connections that are activated simultaneously by the peak of both input and training tuning curves. The distribution of postsynaptic–presynaptic pair intervals is plotted in Figure 3 as a function of the distance from the center of the S-band. Close to the mid-band line, we see an excess of small intervals between postsynaptic and presynaptic spikes and an excess of positive-over-negative intervals (Fig. 3, left panel). In the region of near-zero weight connections, we see the reverse effect: a deficit of short intervals in general and of short positive intervals in particular (Fig. 3, right panel). It is the location-dependent excess–deficit of positive intervals that leads to the formation of the weight patterns with the asymmetric STDP rule; however, the even greater relative excess–deficit of short intervals in general suggests that a symmetric STDP rule would give the same pattern with a shorter learning time. This proves to be the case, although the outcome is strongly dependent on the size of the time window for plasticity, because this controls the balance of potentiation–depression (supplemental Fig. 2, available at www.jneurosci.org as supplemental material).

Figure 4 illustrates some of the other functions that can be learned by the network. For linear functions, the function approximation is equally accurate for all input values, although the accuracy decreases as the number of input values that give the same output value increases. For nonlinear functions, the accuracy is poor near the extremes of the output range, particularly for trigonometric functions. The higher the order of the trigonometric function, the poorer the accuracy. This is the same phenomenon as the underestimation noted two paragraphs previously.

The network performs equally well for coordinate transformations in two dimensions (Fig. 5), although the time required for learning to converge is increased.

Influence of population tuning curve shape

We hypothesized that the poorer accuracy of trigonometric function approximation near the extremes of the function range, noted above, is attributable to the width of the tuning curves. Our hypothesis was that for functions that have multiple input values giving a single output value, the input values compete with one another, because each is correlated with the training input. The degree of competition is larger the larger the width of the tuning curve. The winning input is the one that has the most neighbors, i.e., the one in the widest part of the band at a given output neuron index. This can be seen in Figure 6A. For high values of σ_R , the peaks of the S-band are “stunted” (i.e., the peak amplitude

is smaller and the peak is flatter) compared with the small- σ_R cases. The effect of this is that the network response at the extremes of the range has a bias toward the center (Fig. 6B). The effect is also larger the closer together are the weight bands; hence the deterioration of approximation accuracy in the series of functions $\sin(n\theta)$ as n increases, seen in Figure 4. For linear functions, competition also exists, but the width of the weight band is constant across the function range, and so the effect on function accuracy is not systematic, although it may increase the noise in the approximation. This effect is in accordance with standard results from traditional radial basis function networks: too large a basis function width leads to over-smoothed output functions and too narrow a width leads to under-smoothed output functions (Bishop, 1995).

Thus, in general, the smaller the width of the population tuning curve, the better the function approximation for nonlinear, many-to-one functions such as the trigonometric functions. The narrower the tuning curve, however, the higher the input firing rates required, and there is clearly a lower limit at the point where the tuning curve is so narrow that it activates only a single cell. Close to this limit, the approximation becomes noisy (see the graphs for $\sigma_R = 0.01$ in Fig. 6A,B and the plot of approximation error against σ_R in Fig. 6C).

Turning from the width to the height of the tuning curve, the network is robust to changes in the peak input firing rate (we tested $15 \text{ Hz} \leq R_{\max} \leq 240 \text{ Hz}$) when the asymmetric STDP rule is used, but it is much less robust with the symmetric rule (supplemental Fig. 3, available at www.jneurosci.org as supplemental material).

The lower limit for R_{\max} is hard to determine because of the low rate of learning when there are few spikes. At 5 Hz, a faint S-band very slowly becomes visible, but the off-band synapses do not become depressed, so the noise in the transformation would be considerable.

The learning procedure is also robust to differences between the peak firing rates of the input and training populations. For a ratio of $R_{\max}^{\text{Training}}/R_{\max}^{\text{Input}}$ in the range 0.6 to 1.4, there is almost no discernible change in the weight pattern, and the pattern remains identifiable, although degraded, for ratios as small as 0.1: a peak training firing rate of only 6 Hz for the default input R_{\max} of 60 Hz. With no activity in the training population, the preferred input location for a given output neuron is randomly located, as shown in Song and Abbott (2001), their Figure 2.

The effect of having a non-location-specific baseline input firing rate, i.e., non-zero R_{\min} , is similar to the effect of increasing the background activity in the output population and is discussed below (see below, Influence of network activity and connectivity).

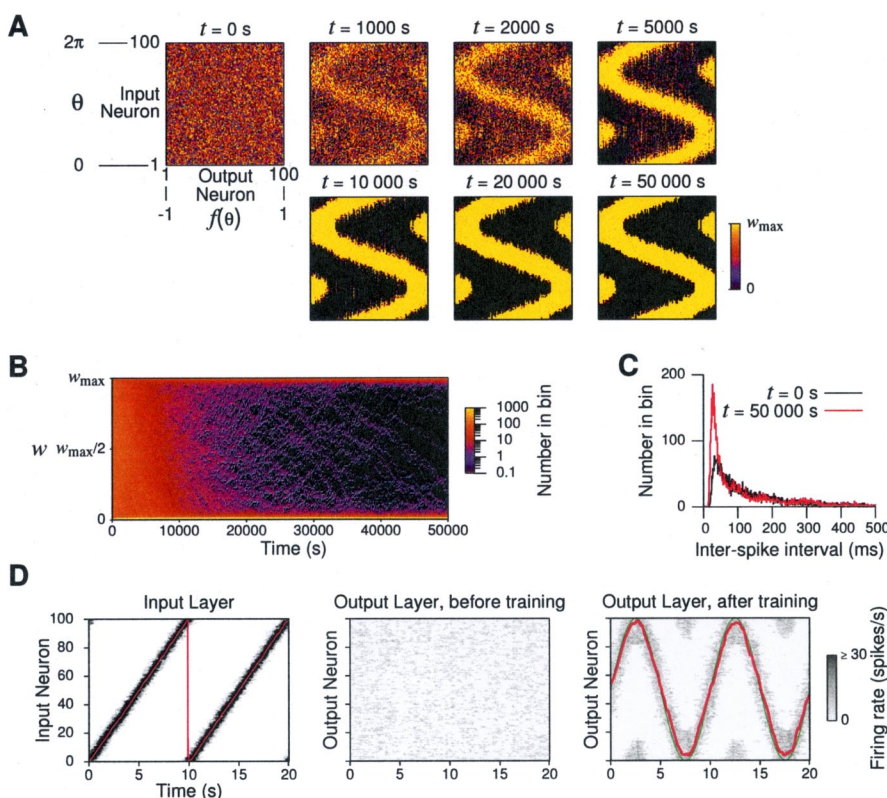


Figure 2. Learning to calculate $f(\theta) = \sin(\theta)$. **A**, Color-scale plot of the matrix of synaptic weights from input population neurons to output population neurons. Each population consists of 100 neurons. An input of θ activates the input population with a peak of activity at the neuron with index closest to $100\theta/2\pi$. The training population is activated simultaneously with a peak of activity at the neuron with array index closest to $100(\sin(\theta) + 1)/2$. Initially, the synaptic weights are distributed uniformly and randomly between 0 and w_{\max} . As training proceeds, the distribution of weights becomes bimodal, with the high-valued weights forming an S-shaped pattern. The pattern converges by $\sim 20,000$ s. Note that the patterns wrap around because of the periodic boundary conditions. **B**, Color-scale plot of the development in time of the histogram of synaptic weights, illustrating the formation of a bimodal distribution and confirming the convergence of the distribution. **C**, ISI distribution in the output population at the beginning and end of the training period. Spikes were recorded for a 10 s period. ISIs from all cells were pooled. Although the sum of the synaptic weights has been reduced during the training, there is an increase in the number of small ISIs (high instantaneous firing rate). **D**, Response of the network to an input that is swept twice across the entire input range. The grayscale plots show the firing rate (calculated as described in Materials and Methods) of each cell as a function of time. The red lines show the estimated position calculated from the population activity. The green line is the function $\sin(\theta)$. Before training, the response of the output population has no dependence on the input. After training, the position estimated from the population activity closely matches the desired function, although with a small degree of underestimation at the extremes (near $\sin(\theta) = \pm 1$). The reasons for this underestimation are discussed in Results (see Influence of population tuning curve shape).




Figure 3. Distribution of intervals $\Delta t = t_{\text{post}} - t_{\text{pre}}$ as a function of the distance d of the postsynaptic–presynaptic pair from the center of the S-band. For a presynaptic neuron with preferred input location θ_j and a postsynaptic neuron representing output location y_j ($-1 \leq y_j < 1$), $d = |y_j - \sin(\theta_j)|/2$. The gray line shows the distribution during the first 10 s of training; the black line shows the distribution between 50,000 and 50,010 s of training. Connections with small d (those that become potentiated during training) show an excess of intervals with small Δt and an excess of positive-over-negative intervals that becomes more pronounced as a result of training. Connections with large d show a deficit of intervals with small Δt that again becomes more pronounced with training, with the deficit being largest for positive Δt [although initially there is a small excess at very small (< 5 ms), positive Δt]. All spike pairs, not just nearest neighbors, were used to calculate intervals. The histograms were normalized by the mean number of intervals per bin in the range $\Delta t = \pm 1000$ ms, summed over all values of d .

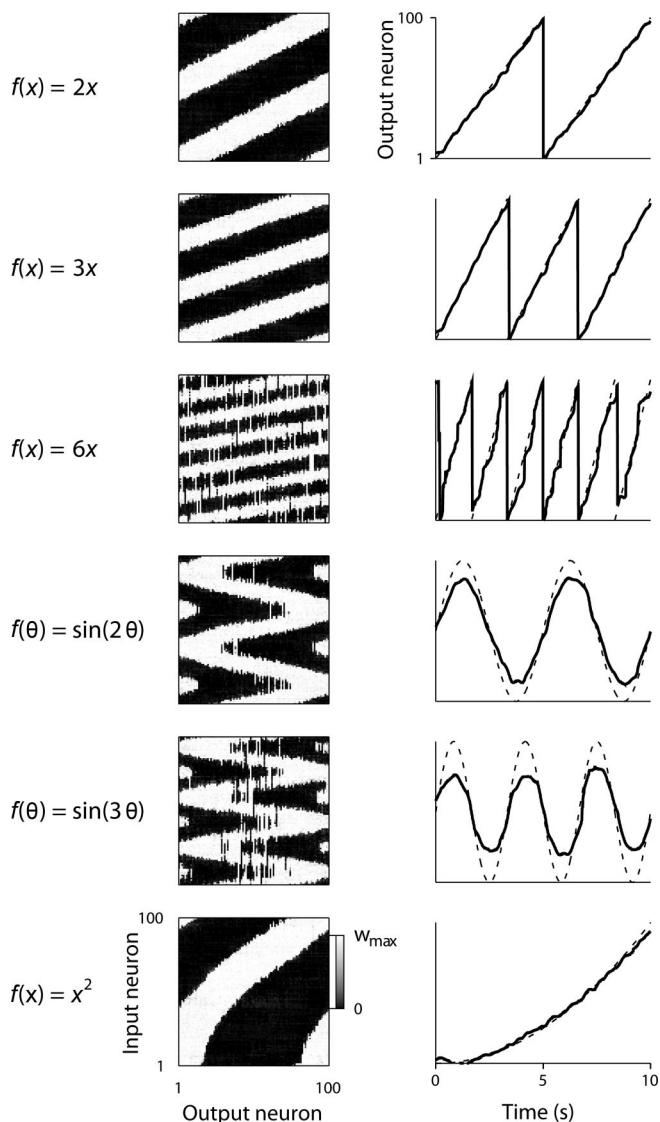


Figure 4. Learning to approximate various functions. The left column shows the converged weight patterns, and the right column shows the position estimates (solid line) compared with the actual position (dashed line) for a single sweep of the input position across its range of values. For each function, the range and domain must be mapped onto the range of cell indices (here from 1 to 100). Note that because of the periodic boundary conditions, the functions $f(x) = nx$ are actually $f(x) = nx \text{ modulo } 1$. For functions that can have the same output value for multiple inputs [such as $f(x) = nx$, $f(\theta) = \sin(n\theta)$], interference between the bands degrades the pattern as n gets larger. For nonlinear functions, for which the distance between those inputs that give the same output is not constant, this leads to a significant degradation where the input values are close together (e.g., at $\sin(n\theta) = \pm 1$).

Influence of the spatiotemporal structure of learning

The procedure that we use for generating correlated activity [taken from Song and Abbott (2001)] has a number of drawbacks as a biologically realistic training–learning procedure. First, it assumes that all points in the input space are visited with equal frequency; second, it requires dwell times (i.e., the time that an input stays in a particular location) that are approximately the same as the STDP time constants (~ 20 ms with our default parameters); and third, it ignores travel time between input locations. The second and third assumptions may be reasonable for saccades but not for arm movements, our principal example here.

To test the first point, whether all input points must be visited with equal frequency, we performed simulations in which the

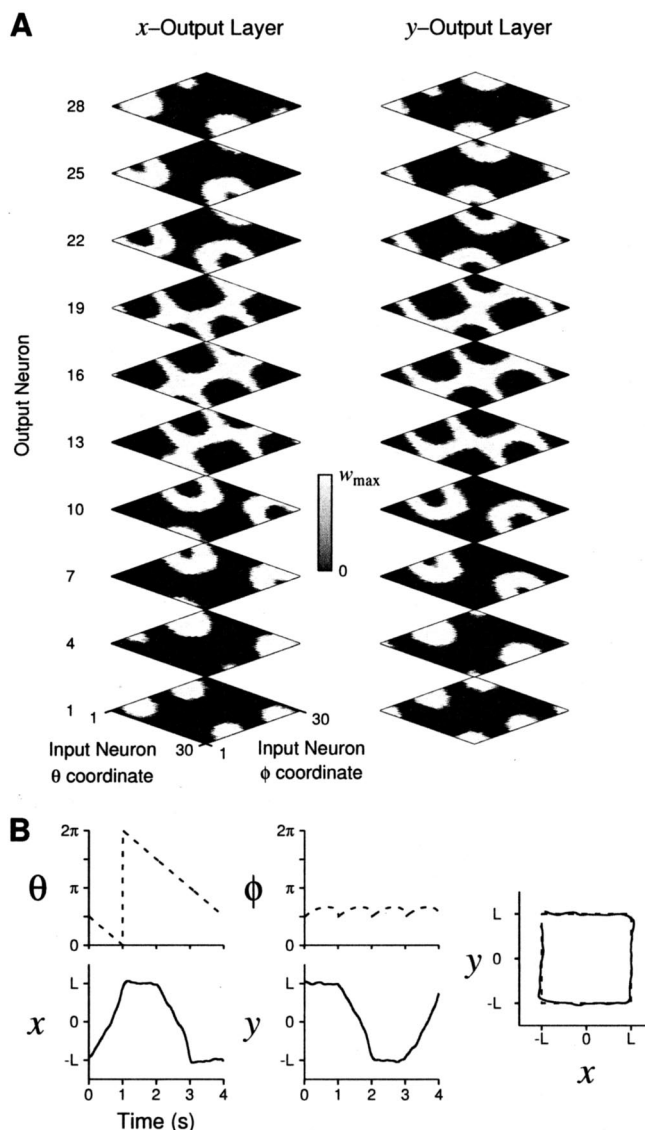


Figure 5. Generalization to two dimensions: training a network to calculate the position of the end of an arm in Cartesian (“visual”) coordinates based on joint-angle (“proprioceptive”) inputs. The geometry of the arm is shown in Figure 1, inset. The functions to be learned are $x = L \cos(\theta) + L \cos(\theta + \phi)$ and $y = L \sin(\theta) + L \sin(\theta + \phi)$. **A**, Grayscale plots showing slices through the three-dimensional weight tensor. Each slice shows the pattern of synaptic weights from the two-dimensional input population (representing θ and ϕ coordinates) to a neuron in one of the output populations [representing x (left) or y (right) coordinates]. **B**, Testing the two-dimensional network by drawing a square. The square is drawn in a clockwise direction starting in the top left corner. The four graphs on the left are plots of joint angles, θ and ϕ , and visual coordinates x and y against time. Dashed lines represent the input values, and solid lines represent the position estimated from the population activity in the output populations (see Materials and Methods). The graph on the right is a plot of y - against x -coordinate. The dashed line represents the theoretical position, and the solid line represents the estimated position.

input angle θ was chosen from a normal distribution (mean, 180° ; SD, 72°). The S-band forms first for synapses from neurons encoding locations that were most often visited and only slowly forms for rarely visited locations; however, by $\sim 40,000$ s the entire pattern has formed and stabilized, as for uniform coverage. In contrast, when the coverage is uniform except for a region that is never visited (of width 72°), a horizontal band of synapses retains their original weights, and the output population is unable to respond to inputs at that location.

Concerning the second point, the requirement for short dwell

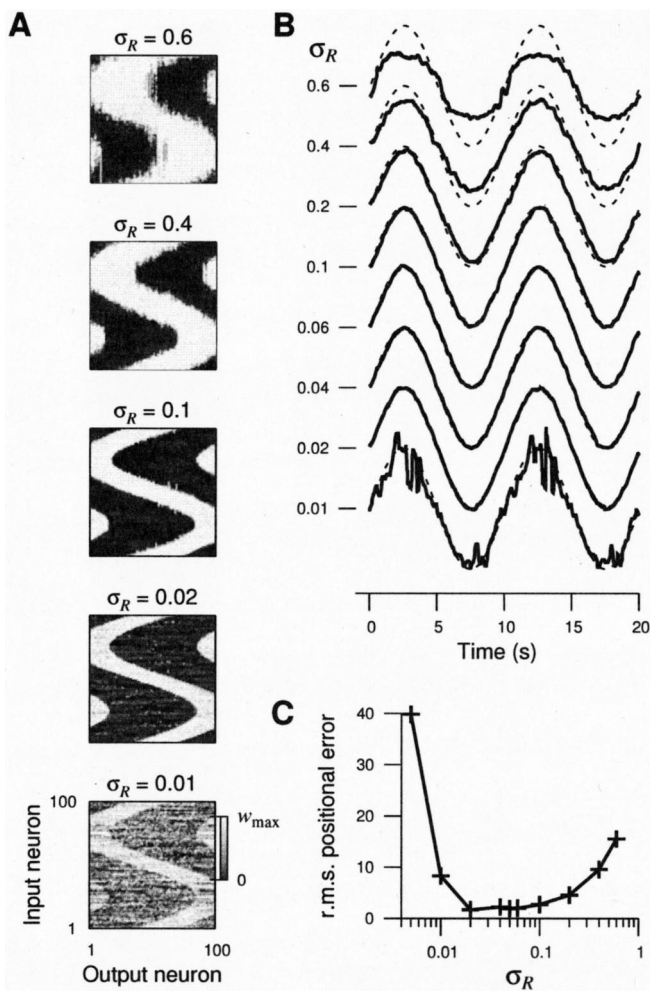


Figure 6. Increasing the width of the input tuning curve (the size of the receptive field) by increasing the parameter σ_R allows lower peak firing rates but leads to inaccuracies near the limits of $f(\theta)$. Note that when σ_R was increased, R_{\max} was decreased proportionally to maintain the total input constant. **A**, Grayscale plots show converged synaptic weight patterns for different values of σ_R . **B**, Position estimates based on the population activity (solid lines; cf. Fig. 2D) compared with the actual position (dashed lines). **C**, Root mean squared (r.m.s.) difference between the estimated and actual positions over the two cycles shown in **B** as a function of σ_R . Note that the worsening of the response when σ_R is very small is because connections distant from the S-band are inactivated too infrequently to be depressed by the STDP mechanism. This worsening could be reduced by increasing the level of background noise in the output population.

times, Song and Abbott (2001) found that there was a fall-off in the correlation-induced potentiation as the correlation time (the mean dwell-time), τ_{corr} , increased, but that this fall-off was reduced considerably by increasing the time constant of synaptic depression, τ_- . We tested how this fall-off affects the formation of weight patterns and found, as expected from the result of Song and Abbott (2001), that increasing τ_{corr} degrades the pattern severely and ultimately degrades the pattern completely; however, if a longer time constant (Feldman, 2000), on the same order as the correlation time, is used for STDP, or if a symmetric STDP rule is used, the pattern formation is preserved (Fig. 7). The symmetric STDP rule is more robust because in those regions where the input and training inputs are not strongly correlated, the deficits of both positive and negative small postsynaptic–presynaptic intervals (compare Fig. 3, right panel) contribute to the dominance of depression, whereas with the asymmetric rule, the deficit of positive intervals must outweigh the deficit of negative

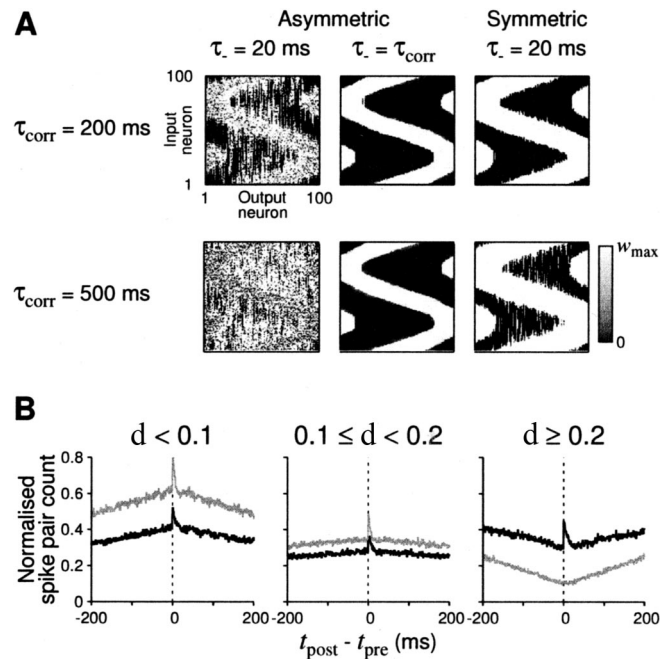


Figure 7. Learning is sensitive to the decay time constant of the input–output correlations (the correlation time, τ_{corr}), unless the time constant of the negative arm of the STDP curve (τ_-) is changed to match τ_{corr} . With a symmetric STDP rule, the sensitivity is much less. **A**, Grayscale plots show the converged weight patterns for different values of τ_{corr} and τ_- and for asymmetric and symmetric STDP rules. The ability of the network to learn the required function is degraded when the correlation time is increased from 20 to 200 ms and is lost entirely when the correlation time is increased to 500 ms. This degradation is dependent on the value of τ_- , however, and is much less pronounced for the symmetric STDP rule. **B**, Distribution of $t_{\text{post}} - t_{\text{pre}}$ with $\tau_{\text{corr}} = 500$ ms for $\tau_- = 20$ ms (black line) or $\tau_- = 500$ ms (gray line). The important difference between the two cases is the excess of intervals with small positive values for those connections with $d > 0.2$ (see Fig. 3 legend for a definition of d), when $\tau_- = 20$ ms. This excess is responsible for the large amplitude synaptic weights outside the S-band seen in **A**, left panels.

intervals, which is only the case when τ_{corr} is small. The shape of the distribution of dwell times and the individual spike train statistics do not appreciably affect the weight pattern (for details, see supplemental note, available at www.jneurosci.org as supplemental material).

To address the third point, the instantaneous movement between input locations, we trained the network with input locations that were changing smoothly and continuously rather than jumping from point to point. The input location moved alternately clockwise and counterclockwise, with a constant speed, for random time periods drawn from a uniform distribution between 100 and 1100 ms. This input pattern results in successful training, provided the speed of motion is sufficiently fast. For $\tau_- = 20$ ms, the speed must be $\geq 2^\circ/\text{ms}$; for $\tau_- = 100$ ms, a speed of $0.2^\circ/\text{ms}$ is sufficient. The latter value is similar to the maximum joint speeds seen in infants during spontaneous arm movements (Thelen et al., 1993), although more typical values are in the region of $0.05^\circ/\text{ms}$ (Bhat et al., 2005). For slower movements, regions in which the rates are correlated but low are potentiated in addition to the usual S-band of connections with correlated, high firing rates. The requirement for high speeds is related to that for short dwell times in the standard, saltatory input movement: inputs must fire together to produce potentiation but then quickly stop firing to avoid depression. With further tuning of parameters, it is probable that slower, smooth movement could give results comparable to saltatory movement; this is a subject for future study.

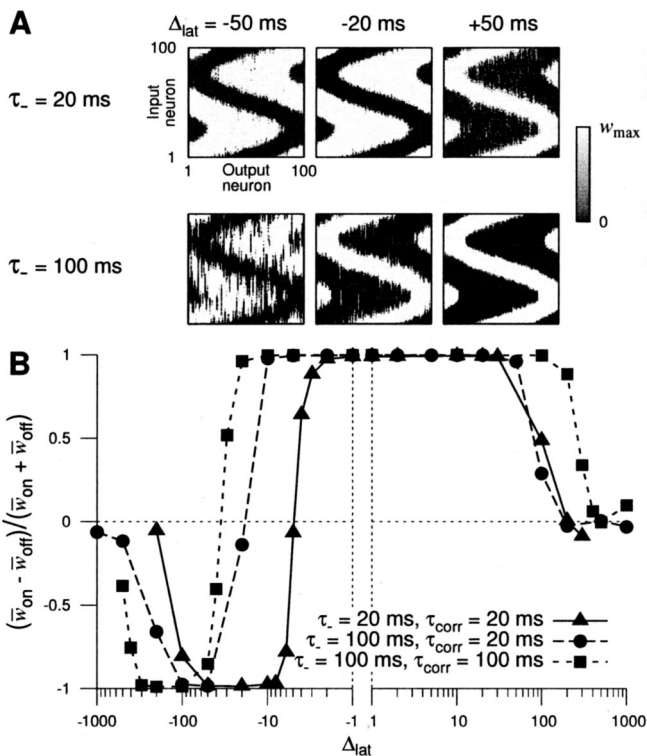


Figure 8. Learning is robust to differences in signal latency between the input and training signals, provided the difference in latency, Δ_{lat} , lies within a band whose upper limit is approximately determined by the correlation time τ_{corr} and whose lower limit is determined by both τ_{corr} and the negative STDP time constant τ_{-} (positive values of Δ_{lat} indicate that the input signal precedes the training signal). **A**, Grayscale plots show the converged weight patterns for different values of Δ_{lat} and for different values of τ_{-} and τ_{corr} ($\tau_{corr} = \tau_{-}$ in all the plots shown). When the training signal precedes the input signal by a large enough amount, the inverse pattern is learned. **B**, Graph showing the normalized difference between the mean synaptic weight, \bar{w}_{on} , within the central S-band ($d < 0.1$; see Materials and Methods for a definition of d), and the mean weight, \bar{w}_{off} , within the inverse band ($0.4 \leq d < 0.5$), as a function of Δ_{lat} . The desired weight pattern corresponds to positive values of this difference, the inverse pattern corresponds to negative values, and random connectivity corresponds to values near zero.

We would expect that the temporal structure of learning will also interact with the membrane time constant. The effect of the membrane time constant is rather complex. The optimal value, given the default parameters, is in the range 10 to 20 ms. Outside these values, the weight pattern begins to degrade, although, interestingly, after a nadir around $\tau_m = 40$ ms, the pattern recovers considerably. The root mean squared positional error (compare Fig. 6C) is 4.7% of the range at the default $\tau_m = 20$ ms, 18.0% at 40 ms, and 6.3% at 400 ms.

We have so far assumed that the input and training signals are exactly coincidental; however, the latency required for a signal to travel from peripheral receptors to the region in which coordinate transformations are performed (e.g., the parietal cortex) is likely to be different for the two modalities. For example, in the ventral intraparietal area of macaque monkeys, the mean latency of visual signals is ~ 86 ms, whereas that of tactile signals is ~ 31 ms (Avillac et al., 2005). Furthermore, in our example of a moving arm, tracking eye movements are required to follow the movement, so the visual input is likely to lag behind the proprioceptive. We therefore investigated how robust the learning is to differences in signal latency between the input and training signals (Fig. 8).

We define the difference in latency, Δ_{lat} , to be positive when

the input signal precedes (has lower latency than) the training signal and to be negative otherwise. We found that the upper limit of the window of Δ_{lat} within which the correct transformation is learned is approximately proportional to (approximately three to five times) the correlation time τ_{corr} . The above-mentioned latency difference found experimentally in macaque falls within this window. The lower limit is smaller in amplitude and depends on both τ_{corr} and the time constant for the depression part of the STDP curve, τ_{-} . When both quantities are 100 ms, the limiting value for Δ_{lat} is approximately -30 ms. As the difference in latency becomes more negative than this, there is another window in which the reverse weight pattern is learned (Fig. 8A).

The mechanism for this may be understood with reference to the histogram of postsynaptic–presynaptic intervals, Δt (Fig. 3). Initially, the postsynaptic firing time is determined mainly by spikes from the training population, because those from the input population have random weights. If the latency difference, Δ_{lat} , is positive (input earlier than training), the peak in the interval histogram is pushed toward positive Δt , into the potentiation part of the STDP curve. Therefore, learning will only begin to fail when the peak of the interval histogram exceeds the correlation time. In the converse case, in which Δ_{lat} is negative, the histogram peak is pushed toward negative Δt , because the training spike causes an action potential and then the input spike (with longer latency) arrives. With no difference in latency, there is an excess of postsynaptic–presynaptic pairs with small positive intervals. When a sufficient amount of this excess has been pushed into the depression part of the STDP curve, weights that are normally potentiated will be depressed. Because of the competitive nature of the Song–Miller–Abbott (Song et al., 2000) STDP rule, weights for which there is reduced correlation between inputs will be potentiated, leading to the observed reversal in the weight pattern.

Influence of training pathways

Learning the correct pattern of synaptic weights for the connections from the input population to the output population clearly requires a topographic pattern of connections from the training population to the output population; however, the precise strength and range of these connections are not critical (see supplemental note, available at www.jneurosci.org as supplemental material), but there is a certain minimum level of input to the output population from the training population that is required to mold the weight patterns effectively.

We now asked whether the weight pattern is stable if the training input is removed altogether but the STDP mechanism is still active. This is important because the intention is to be able to predict, for example, the location of the hand based only on proprioceptive signals without visual input. This has straightforward applications such as being in the dark or in the case of becoming blind. Must there be an upper boundary to the critical developmental period, after which weights must be frozen, or can plasticity be ongoing?

In fact, the weight pattern is very stable (supplemental Fig. 4, available at www.jneurosci.org as supplemental material), although the degree of stability depends on the peak step-size of the weight change (A_{+} and A_{-}). We conclude that differences between the plasticity rules operating during development and during adulthood are not required (nor are they precluded) and that this mechanism can operate in the adult brain and subserve slow adaptation. If a completely different training pattern is applied (consider wearing prism glasses, for example), the originally

learned weight pattern is quickly forgotten and then a new one forms.

Up to this point we have made the assumption that the connection weights from the training population to the output population (visual synapses, in our example) have been determined by an earlier phase of development and no longer undergo plasticity. We believe that this assumption is a reasonable one (see Discussion) but nevertheless examine situations in which plasticity is ongoing in the training connections. If we start from a situation in which the training connections have been prespecified, the questions are whether the preexisting pattern is stable with the addition of the new modality and whether the input to output weights develop as desired. Figure 9A indicates that this does seem to be the case: the diagonal band undergoes some changes but remains recognizably diagonal.

If both input and training connection weights start from random values, then recurrent excitatory and inhibitory connections within the output population (see Materials and Methods) and larger values of A_+ (i.e., 0.1 rather than 0.01) are required for smooth weight bands to form. Even then, the connections develop with an arbitrary phase, i.e., the network maps θ to $\sin(\theta + \psi)$, where ψ takes a random value and tends to wander as training continues (Fig. 9B). If the global inhibitory weights within the output population are too weak, then the recurrent excitatory weights tend to saturate in such a manner that ψ precesses, moving always in the same direction. On the basis of this, we conclude that it is necessary for the training pathways (visual pathways, in our example) to develop first (see Discussion) but that plasticity mechanisms can remain active even after this development has taken place.

Influence of plasticity rules

We have shown above that both asymmetric and symmetric STDP rules can produce the desired weight patterns, although the symmetric rule is more robust to the temporal structure of the sensory inputs, whereas the asymmetric rule is much less dependent on the amplitude of the population tuning curve. The precise amplitudes and time constants of potentiation and depression are not critical, provided they obey certain conditions (see supplemental note and supplemental Figs. 5 and 6, available at www.jneurosci.org as supplemental material); however, the dependence of the amplitude of weight changes on the current weight is considerably more important. Clearly, the formation of the banded pattern relies on the bimodal distribution of synaptic weights produced by the hard boundaries (no dependence of Δw on w) of the Song–Miller–Abbott (Song et al., 2000) STDP mechanism. For soft weight boundaries (linear dependence of Δw on w) that give a unimodal distribution (Gütig et al., 2003), a banded pattern also forms, but the modulation of the weight pattern by

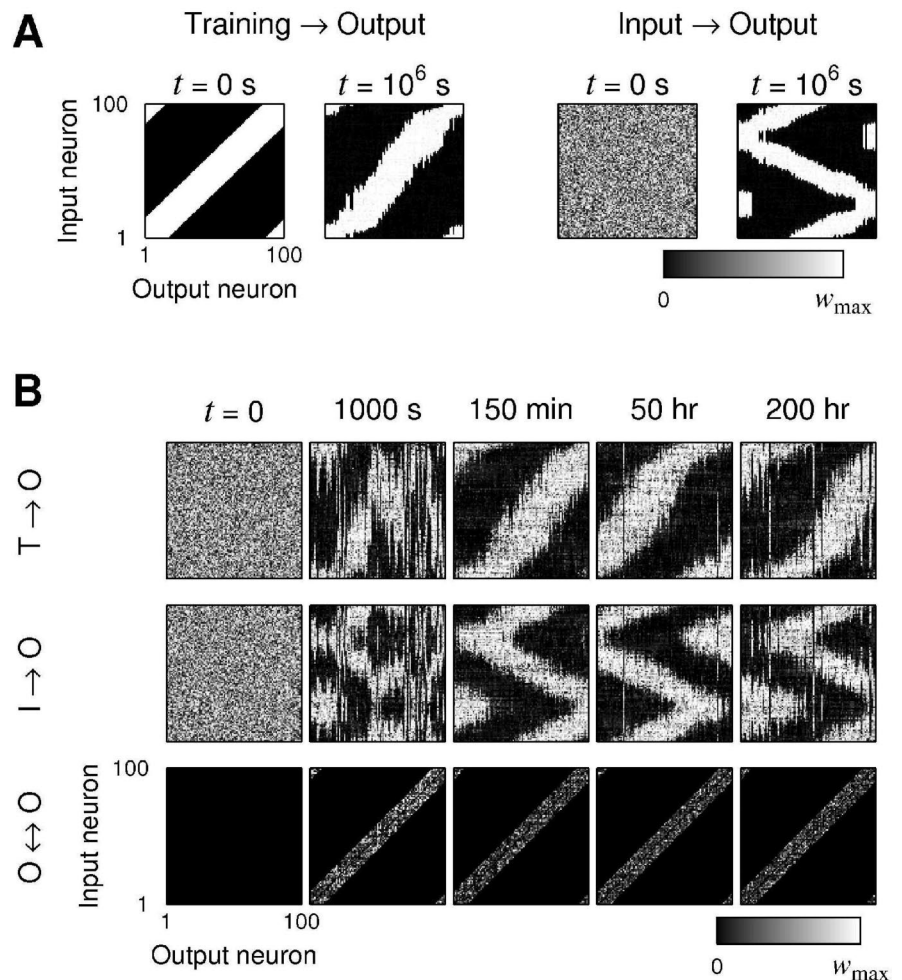


Figure 9. The effect of making training connections subject to plasticity. Grayscale plots are of the matrix of synaptic weights. Note that the model was extended with recurrent inhibitory and excitatory connections within the output population for these simulations (see Materials and Methods). **A**, A prespecified pattern of training connections remains stable, whereas the input connections develop as when training connections are fixed. Output to output connections were all weak ($w < w_{max}/10$) and are not shown. **B**, From random initial training connections, an S-band forms but its location “wanders” as the training continues. The grayscale plots show the development (left to right) of the connection matrices from training to output (top row), input to output (middle row), and recurrent output to output (bottom row) populations. Note that the amplitude of weight changes was increased for these simulations compared with the other results shown ($A_+ = 0.1$, rather than 0.01 or 0.001).

correlation is much weaker; however, the location calculated from the population activity still accurately follows the sine curve, although it is much noisier than with the bimodal distribution (Fig. 10). The combination of a unimodal distribution with pruning of weaker synapses produces patterns similar to the hard-boundary mechanism; however, with the default parameters used in our simulations, the threshold for pruning has to be a substantial fraction (~ 0.8) of the mean synaptic weight.

Influence of network activity and connectivity

The network behavior is not strongly dependent on the size or connectivity of the network, provided that the level of activity to each cell is maintained approximately the same by adjusting other parameters. For example, a network of 500 neurons per population in which each input cell connects to 20% of output cells behaves similarly to a network with populations of 100 neurons in which each input cell connects to all output cells. For networks smaller than this, increasing the maximum firing rate R_{max} proportionally to the decrease in network size has a similar effect. The lower limit on network size is dependent on the num-

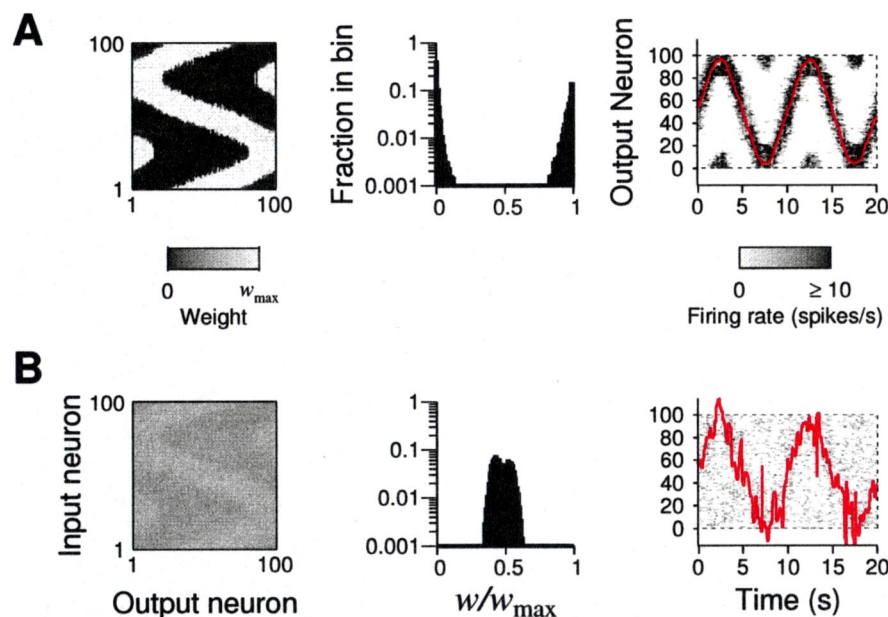


Figure 10. Comparison of “hard” and “soft” limits on synaptic weights. **A**, Hard limits (see Materials and Methods). **B**, Soft limits. Each panel has three parts, from left to right: grayscale plot showing input–output weight matrix after learning; weight histogram; firing-rate response to a sweep stimulus (grayscale map), together with population estimate of position (red line). To show more clearly the continuity of the position estimate, the estimate has been “unwrapped.” It is interesting to note that the soft-limits weight distribution is still bimodal, although the separation of the peaks is much less. Clearly, hard limits give better results but are not essential.

ber of input locations that give a similar output (compare Fig. 4), but the network performs surprisingly well with very few neurons. With $N = 20$, the network can learn to perform $f(\theta) = \sin(3\theta)$ about as well as a network of 100 neurons per population but fails on $f(\theta) = \sin(4\theta)$, whereas the network with $N = 100$ succeeds.

We chose to use periodic boundary conditions to avoid edge effects. For problems involving a body-centered coordinate system, such boundary conditions are consistent with the physical situation, because a stimulus could be at any point on a full circle around the body, although whether such conditions are implemented in the brain in the pattern of connections and how such connections could be formed are open questions. For problems involving arm movements, periodic boundary conditions are not appropriate, because the ranges of joint motion are limited. We performed simulations with nonperiodic boundary conditions in which the arrays of neurons were padded with extra neurons with receptive fields centered outside the input range. With no padding, the population estimate is inaccurate when the output location is near the edge of the array, because there is only half a tuning curve at the edge. When the padding is $0.2 N$ cells at each edge (about half the width of the S-band), there is a little nonspecific activity in the neurons at the edge of the array, but this has little effect on the population estimate, which is comparable in accuracy to the network with periodic boundary conditions (supplemental Fig. 7, available at www.jneurosci.org as supplemental material).

For an asymmetric learning rule with hard bounds on the weights, increased activity biases the distribution toward the peak at zero (Song et al., 2000). If there is insufficient activity in the network, the “off-band” weights are not depressed toward zero. If there is too much, then all weights are depressed toward zero, and the band of potentiated weights fades out (supplemental Fig. 8, available at www.jneurosci.org as supplemental material). The

effect has a complex dependence on whether the activity is presynaptic or postsynaptic. Adding background noise to each postsynaptic cell, adding a non-location-specific term (R_{\min}) to the tuning curve, increasing the height of the tuning curve (R_{\max}), or increasing the peak synaptic weight (w_{\max}) all separately have the similar effect of depressing off-band weights toward zero; however, combining both background noise and nonspecific input leads to a failure of learning. With $R_{\min} = 10$ Hz and low levels of background noise, the weight band forms as desired. If the background noise is increased, however, then off-band weights are not depressed (the opposite to what would be seen with $R_{\min} = 0$), and the weight pattern appears almost random.

Discussion

The principal conclusion of this study is that the connectivity patterns required for performing coordinate transformations between sensory frames of reference can be learned in an unsupervised manner, with the STDP learning rule, based on simultaneous observations of moving stimuli with two sensory modalities. This adds support to the hypothesis that basis function net-

works are actually used in biological neural systems [see Pouget and Sejnowski (1997) for a fuller discussion of the evidence supporting the existence of basis function networks in parietal cortex].

Correlation-based learning in spiking models

The idea of using basis function networks to perform coordinate transformations has been developed extensively (Salinas and Abbott, 1995; Pouget and Sejnowski, 1997; Deneve et al., 2001; Deneve and Pouget, 2003), but most existing models consider only the mean firing rate. The principal disadvantage of this approach is the difficulty in relating the model to underlying biological mechanisms and hence in constraining it with biological data. In contrast, an approach that models synaptic conductances and individual action potentials gives a more direct correspondence between model parameters and experimentally measurable quantities and hence allows the model to be more tightly constrained by experiments and to make more testable predictions. An example in this study is the interaction between the experimentally determined time constants of the STDP rule and the temporal structure of the motion that drives learning (Fig. 7). Another example is the prediction that position estimation should be poorer near the limits of the movement range, which is a direct consequence of the intersynapse competition inherent in the STDP rule used. One existing report has shown that a basis function network can be implemented with spiking neurons (van Rossum and Renart, 2004), but this was shown only for calculation of the sum of two variables and with synaptic weights set by hand.

Because the output of a basis function network is a linear weighted sum of the basis functions, determining the correct weights requires simply linear regression, calculated off-line or by using a supervised learning rule; however, neither one of these methods is available to biological nervous systems. A method that

uses a Hebbian-like learning rule was developed and analyzed by Salinas and Abbott (1995). As in our model, the coordinate transformation is learned based on the correlations between randomly generated inputs and independent measurements of the transformed locations. During the learning phase, however, the input and output populations are decoupled, and the desired output response is applied directly to the output population: the input neurons do not drive the output neurons. In contrast, our model has no need for a separate training phase with decoupled neurons.

The idea of learning transformations through observation of randomly generated movements [“motor babbling” (Bullock et al., 1993)] has been widely used in models of sensorimotor transformations (Kuperstein, 1988; Gaudio and Grossberg, 1991; Burnod et al., 1992; Bullock et al., 1993), with weights determined by error-correcting or generalized operant conditioning methods. Studies in kittens and infant monkeys have shown that visual feedback from limb movements is essential in the development of visually guided reaching (Hein, 1974; Held and Bauer, 1974). For humans, opposing external forces during spontaneous movements requires sight of the limb in newborns (van der Meer et al., 1995) but not in 3-month-old infants (Dibiasi and Einspieler, 2004). This also provides support for our finding that visual connections to the coordinate transformation network must develop before proprioceptive connections. This idea can also be applied to cross-modal integration as opposed to sensorimotor transformations, because the random inputs can be generated by an external mechanism rather than by the organism itself. The only requirement is for two independent measurements of the location.

Time-scale issues

The strongest influence on the learning behavior of the network is the interaction of the different time-scales in the model: of the STDP rule, of the rapidity of change in input location, and of the membrane time constant. To summarize, for saltatory movement, the mean dwell time at any one location must be on the same order as the STDP time constant for depression, τ_- (Fig. 7), although this requirement is considerably relaxed for the symmetric STDP rule. Similarly, for continuous motion, the speed must be sufficiently high that the arm moves a distance approximately equal to the receptive field size during a time τ_- . This is a strong constraint, which may possibly be eased if the two modalities do not always attend to the same stimulus, e.g., arm movements could be slower if the eyes make saccades to other regions of the visual field (during which the two modalities are not correlated) and only return periodically to observe the hand. The faster of the modalities would then be the limiting one. In general, however, we should expect to find either asymmetric STDP with $\tau_- \gg \tau_+$ or symmetric STDP in the brain regions that are responsible for coordinate transformations. If the mismatch between STDP and movement timescale is too great, further mechanisms would be required.

Despite the temporal precision of the STDP rule, the model is quite robust to differences in response latency between the two modalities, particularly if the input signal precedes the training signal, as is found experimentally for tactile and visual signals in ventral intraparietal cortex in the macaque (Avillac et al., 2005). Here again, the robustness of the model is increased if τ_- is longer than τ_+ .

Biological realism of the model

This model uses integrate-and-fire neurons, which are considerably lacking in complexity compared with real neurons. We have

attempted, however, to set firing rates and excitatory postsynaptic potential sizes to biologically realistic values.

Our standard model relies on the training–output synaptic weights being fixed, and the input–output synapses being subject to plasticity (although we show that a preexisting pattern of plastic training weights is not appreciably degraded by the addition of the input connections). There are several cases in which different classes of excitatory synapses onto a single postsynaptic neuron have been shown to exhibit such differential plasticity. A classic example is the different NMDA dependence of commissural–associational and mossy fiber synaptic plasticity in hippocampal CA3 pyramidal neurons (Nicoll and Malenka, 1995). It has been shown more recently that synapses onto hippocampal CA1 interneurons that express group II metabotropic glutamate receptors (mGluRs) presynaptically exhibit mGluR-dependent long-term potentiation, whereas those that express group III mGluRs do not (Lapointe et al., 2004). To our knowledge, differential plasticity at different synapses has not been shown for postsynaptically expressed plasticity, although it has been shown that the subunit composition of AMPA receptors on hippocampal interneurons can be determined by the afferent identity (Tóth and McBain, 1998), which implies that the enzymatic machinery underlying plasticity with a postsynaptic locus might also be so regulated.

A criticism (van Rossum et al., 2000) of STDP mechanisms with hard bounds on the synaptic weights, such as that of Song et al. (2000), is that the resulting bimodal distribution of synaptic weights does not match the distributions of miniature EPSC amplitudes or AMPA receptor densities measured experimentally (cf. O’Brien et al., 1998), which are unimodal with positive skew. The validity of our model does not depend critically on a bimodal distribution of synaptic weights, because a network with soft-bounded STDP is also able to learn the desired transformation, albeit with increased output noise. We point out, however, that a bimodal distribution of weights where the maximum synaptic weight w_{\max} follows a unimodal distribution may be difficult to discriminate from a unimodal distribution of weights where w_{\max} has a single value, because synapses with zero weight [silent synapses (Malenka and Nicoll, 1997)] are likely to be difficult to detect and hence undercounted experimentally.

We have shown that the appropriate patterns of connectivity for basis function networks could be learned by using a spike timing-dependent, local learning rule. We do not make strong claims regarding in which areas of the brain such networks may exist. A strong case has been made previously that the response properties of neurons in the parietal cortex are in excellent agreement with the predictions of a basis function network model for the specific case of coordinate transformations from retinal to head-centered coordinates (Pouget and Sejnowski, 1997). The superior colliculus is another area in which there is strong evidence for sensorimotor–coordinate transformations (Jay and Sparks, 1984; Stuphorn et al., 2000).

Although we use the transformation from joint angle to visual coordinates as our principal example, the model we have presented here is not intended as a detailed model of arm movements (Baraduc et al., 2001; Joshi and Maass, 2005). In particular, proprioceptive tuning curves appear to be monotonic (Helms Tillery et al., 1996) rather than Gaussian as in our model, although two monotonic tuning curves with opposite slopes can be combined to give a peaked tuning curve (Salinas and Abbott, 1995).

Additional studies are therefore required to determine whether the principles elucidated in this study will still apply when the neuronal and circuit diversity of known cross-modal–

coordinate transformation areas is included in the model and when the model is extended to encompass the entire sensorimotor loop.

References

- Avillac M, Denève S, Olivier E, Pouget A, Duhamel JR (2005) Reference frames for representing visual and tactile locations in parietal cortex. *Nat Neurosci* 8:941–949.
- Baraduc P, Guignon E, Burnod Y (2001) Recoding arm position to learn visuomotor transformations. *Cereb Cortex* 11:906–917.
- Bhat A, Heathcock J, Galloway JC (2005) Toy-oriented changes in hand and joint kinematics during the emergence of purposeful reaching. *Infant Behav Dev* 28:445–465.
- Bishop CM (1995) *Neural networks for pattern recognition*. Oxford: Oxford UP.
- Bullock D, Grossberg S, Guenther F (1993) A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *J Cogn Neurosci* 5:408–435.
- Burnod Y, Grandguillaume P, Otto I, Ferraina S, Johnson P, Caminiti R (1992) Visuomotor transformations underlying arm movements toward visual targets: a neural network model of cerebral cortical operations. *J Neurosci* 12:1435–1453.
- Dan Y, Poo MM (2004) Spike timing dependent plasticity of neural circuits. *Neuron* 44:23–30.
- Deneve S, Pouget A (2003) Basis functions for object-centered representations. *Neuron* 37:347–359.
- Deneve S, Latham P, Pouget A (1999) Reading population codes: a neural implementation of ideal observers. *Nat Neurosci* 2:740–745.
- Deneve S, Latham P, Pouget A (2001) Efficient computation and cue integration with noisy population codes. *Nat Neurosci* 4:826–831.
- Dibiasi J, Einspieler C (2004) Load perturbation does not influence spontaneous movements in 3-month-old infants. *Early Hum Dev* 77:37–46.
- Feldman D (2000) Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27:45–56.
- Gaudiano P, Grossberg S (1991) Vector associative maps: unsupervised real-time error-based learning and control of movement trajectories. *Neural Netw* 4:147–183.
- Georgopoulos A, Schwartz A, Kettner R (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419.
- Gütig R, Aharonov R, Rotter S, Sompolinsky H (2003) Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *J Neurosci* 23:3697–3714.
- Guyonnet R, VanRullen R, Thorpe SJ (2005) Neurons tune to the earliest spikes through STDP. *Neural Comput* 17:859–880.
- Hein A (1974) Prerequisite for development of visually guided reaching in the kitten. *Brain Res* 71:259–263.
- Held R, Bauer Jr JA (1974) Development of sensorially-guided reaching in infant monkeys. *Brain Res* 71:265–271.
- Helms Tillery S, Soechting J, Ebner T (1996) Somatosensory cortical activity in relation to arm posture: nonuniform spatial tuning. *J Neurophysiol* 76:2423–2438.
- Hines M, Carnevale N (2004) Discrete event simulation in the NEURON environment. *Neurocomputing* 58–60:1117–1122.
- Jay M, Sparks D (1984) Auditory receptive fields in primate superior colliculus shift with changes in eye position. *Nature* 309:345–347.
- Joshi P, Maass W (2005) Movement generation with circuits of spiking neurons. *Neural Comput* 17:1715–1738.
- Kuperstein M (1988) Neural model of adaptive hand-eye coordination for single postures. *Science* 239:1308–1311.
- Lapointe V, Morin F, Ratté S, Croce A, Conquet F, Lacaille JC (2004) Synapse-specific mGluR1-dependent long-term potentiation in interneurons regulates mouse hippocampal inhibition. *J Physiol (Lond)* 555:125–135.
- Lewis J, Kristan W (1998) A neuronal network for computing population vectors in the leech. *Nature* 391:76–79.
- Malenka RC, Nicoll RA (1997) Silent synapses speak up. *Neuron* 19:473–476.
- Nicoll RA, Malenka RC (1995) Contrasting properties of two forms of long-term potentiation in the hippocampus. *Nature* 377:115–118.
- O'Brien R, Kamboj S, Ehlers M, Rosen K, Fischbach G, Haganir R (1998) Activity-dependent modulation of synaptic AMPA receptor accumulation. *Neuron* 21:1067–1078.
- Pouget A, Sejnowski TJ (1994) A neural model of the cortical representation of egocentric distance. *Cereb Cortex* 4:314–329.
- Pouget A, Sejnowski T (1997) Spatial transformations in the parietal cortex using basis functions. *J Cogn Neurosci* 9:222–237.
- Pouget A, Snyder L (2000) Computational approaches to sensorimotor transformations. *Nat Neurosci [Suppl]* 3:1192–1198.
- Salinas E, Abbott L (1995) Transfer of coded information from sensory to motor networks. *J Neurosci* 15:6461–6474.
- Seung H, Sompolinsky H (1993) Simple models for reading neuronal population codes. *Proc Natl Acad Sci USA* 90:10749–10753.
- Song S, Abbott LF (2001) Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32:339–350.
- Song S, Miller KD, Abbott LF (2000) Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3:919–926.
- Stuphorn V, Bauswein E, Hoffmann K (2000) Neurons in the primate superior colliculus coding for arm movements in gaze-related coordinates. *J Neurophysiol* 83:1283–1299.
- Thelen E, Corbetta D, Kamm K, Spencer J, Schneider K, Zernicke R (1993) The transition to reaching: mapping intention and intrinsic dynamics. *Child Dev* 64:1058–1098.
- Tóth K, McBain C (1998) Afferent-specific innervation of two distinct AMPA receptor subtypes on single hippocampal interneurons. *Nat Neurosci* 1:572–578.
- van der Meer A, van der Weel F, Lee D (1995) The functional significance of arm movements in neonates. *Science* 267:693–695.
- van Rossum MCW, Renart A (2004) Computation with population codes in layered networks of integrate-and-fire neurons. *Neurocomputing* 58–60:265–270.
- van Rossum MC, Bi GQ, Turrigiano GG (2000) Stable Hebbian learning from spike timing-dependent plasticity. *J Neurosci* 20:8812–8821.
- Woodin MA, Ganguly K, Poo M (2003) Coincident pre- and postsynaptic activity modifies GABAergic synapses by postsynaptic changes in Cl⁻ transporter activity. *Neuron* 39:807–820.