Global Motion Direction Learning in Neural Networks Using Neurons with a Physiological Spiking Mechanism

Socrates Deligeorges
V. Sundareswaran
Lucia M. Vaina
Intelligent Systems Lab, Department of Biomedical Engineering
Boston University, Boston, MA 02155 USA

In a psychophysical and computational study of perceptual learning we showed that performance in a task of direction discrimination in dynamic stochastic random dot displays improved with practice (Vaina et al., 1995). The improvement was fast, and was specific to the stimulus parameters and location in the visual field. The computational study, based on an unsupervised learning model, simulated biologically plausible synaptic weight changes that led to improved performance in the task. In the current study, we focused on obtaining a motion representation for this task, by using a supervised model. Specifically, we used feed forward networks of neurons with a physiologically realistic spiking mechanism, trained using the backpropagation rule. The spiking mechanism simulated neuronal noise and a mean-rate coding mode of transmission from cell to cell with either high or low spontaneous firing rates. Performance improvement qualitatively similar to those of human subjects was achieved. We found that (a) training these networks produced opponency between network units tuned to opposite directions of motion, and (b) the learning was slower, more stable, and more accurate for the high spontaneous firing rate. The

opponency is consistent with physiological results. The spontaneous rate provides an alternative, stable means of controlling the learning rate in backpropagation.

1 Introduction

With practice, human observers show dramatic and persistent improvement in perceptual tasks involving detection or discrimination of visual features, including complex gratings (Fiorentini and Berardi 1980), hyperacuity (McKee and Westheimer 1978; Poggio et al. 1992), stereopsis (Ramachandran and Braddick 1973; Fendick and Westheimer 1978), discrimination of texture (Karni and Sagi 1991), direction of motion (Ball and Sekuler 1987), line orientation (Vogels and Orban 1985; Shiu and Pashler, 1992), and object features (Ahissar and Hochstein 1993). In a global motion task, we found that human subjects showed rapid improvement and retention of the ability to discriminate direction (Vaina et al. 1995). In this study, we model this perceptual learning by a feed forward network using neurons with physiologically realistic spiking mechanisms.

There have been few attempts to computationally model perceptual learning. Poggio et al. (1992) pioneered this modeling, using radial basis functions to model performance improvement in a hyperacuity task. Their network learned using a supervised mechanism to alter the weights of the radial basis functions, to best approximate the input-output relationship. For the same perceptual learning task, an unsupervised model was presented by Weiss et al. (1993) who used orientation-selective basis functions. Peres and Hochstein (1994) presented a neural network model for performance improvement in an odd-element detection task (Ahissar and Hochstein 1994) based on reinforcement learning. In a recent study, we presented unsupervised models of learning to discriminate global motion direction (Sundareswaran and Vaina 1996). In that study, our goal was to model the performance improvement through Hebbian learning. In the current study, our objectives were to identify the neuronal representation which is important for the learning, and to examine the role of physiologically realistic spiking mechanisms in learning. We adopted the backpropagation method because the synaptic weight patterns generated by backpropagation may identify the representation used by the brain, even if the "brain arrives at these patterns using an algorithm more realistic than backpropagation" (Crick, 1989). In addition, learning algorithms can be used to train networks that perform a certain task, and one can then ask whether properties of the trained units in these networks resemble those of neurons in the nervous system (Zipser and Anderson 1988; Lehky and Sejnowski 1990). In this study, we were not interested in developing a more realistic learning method (as for e.g. in Mazzoni et al. 1991), but to introduce a physiological information transmission process in the forward path and study its effects on learning.

Although information transmission from layer to layer in a collection of neurons is still an unresolved issue, the concept of *rate coding* as a plausible method of transmission is increasingly gaining acceptance among researchers interested in biologically based neural networks (Softky and Koch 1993). The underlying distribution which governs the *firing rate* of the neurons is an equally ambiguous area but there is increasing evidence that the process is roughly Poisson with a non-stationary mean rate (Softky and Koch 1993; Georgopoulos 1994). An important aspect of the transmission process is the *spontaneous rate* (SR) of neuronal channels. In actual populations of neurons there is an intrinsic 'noise' to the neuron itself which is seen as spontaneous firing of the neuron without a causal input. Generally two population types are considered: those with high SR that can be viewed as more 'sensitive' and those with low SR that require greater input activity to generate the output.

In this study, we used these two types of SRs and analyzed their effects on local and long-term learning trends. To understand the internal representation useful in performing the global motion task, we examined the weight patterns generated in the backpropagation network trained to improve in this task.

<Figure 1 about here>

2. Methods

Sim

good point

The stimuli were random dot kinematograms with 100 dots, in which 75% of the dots were displaced in random directions, and the remaining 25% were displaced in a single (coherent) direction. The stimulus is shown in Figure 1, for three different levels of coherence (15%, 25%, and 75% of the dots moving in the coherent direction). In a task to discriminate between two opposite directions in stimuli with a low proportion of coherently moving dots, observers' performance improved rapidly (200-400 trials), stabilized quickly, and was retained for days and even months (Vaina et al. 1995). Newsome and Paré (1988) first showed that this is a global motion task since to discriminate motion direction, the observer must spatially integrate information from the whole field.

To model the performance improvement, we used a standard feed forward network consisting of a four unit input layer, a two unit hidden layer, and a single unit output layer. The activity of a unit was computed as the weighted combination of the activities of units in the previous layer, transformed by a sigmoidal activation function. The input to the network were the measurements of a set of four directionally-tuned units with Gaussian tuning curves; each unit was tuned to one of the four cardinal directions of motion, i.e.

right, up, left, and down. Each directionally-tuned unit integrated motion information from a large area (the "receptive field") of the input images, by summing responses of local units with the same directional tuning. This global motion information was used to regulate the activity of the neural network's first layer. The output of the network signaled one of the two possible directions of global motion. The performance of the network was assessed by comparing the network output with the known direction of global motion of the input. The error in the output was used to modify the weights of the network using the backpropagation rule with momentum:

$$W(i,j) = \mathrm{W(i,j)} + \eta * \mathrm{Error} * \mathrm{Input} + \mu * [\mathrm{W(i,j)} - \mathrm{W_o(i,j)}], \qquad (1)$$

where:

W(i,j) : Weight value from unit i to unit j

η : Learning rate

Error : Network output error

Input : Value being passed from unit i to unit j

μ : Momentum value

$W_o(i,j)$: The previous weight value

<Figure 2 about here>

In contrast to standard feed forward networks, the transmission from the input units to the hidden units was performed by mean-rate coding of input unit activity (see Figure 2). In our model, both temporal and firing rate encoding were unified into a simple spike generation scheme. We used a simple approximation to a transmission model inspired by

Softky and Koch (1993) and Georgopoulos et al.(1994a, 1994b). When choosing a temporal "window" size for each stimulus presentation, both the computational time constraints and the number of spikes needed to transmit information through a variable mean firing rate were considered. The window size used was a maximum of 25 spikes per stimulus field presentation (for a detailed discussion of window size choice, see Heller et al., 1995). The spike arrival time was a Poisson random process. The spiking process was simulated using a Gaussian distribution of random numbers with a firing threshold regulated by the activity of the directionally tuned units (a similar threshold adjustment based on noise level has been used by McAuley and Stampfli, 1994). If the random number generated was greater than the threshold, the cell produced a spike (see Figure 2). The initial threshold was determined by the cell type chosen, either a high or low SR. Low SR threshold level was 1.65 or approximately 1 to 2 spikes per time window; high SR threshold level was 1.25 or approximately 3 to 4 spikes per time window. The activity of the neuron was measured by the number of spikes in the temporal window.

3. Simulations

Consistent with the psychophysical experiment (Vaina et al. 1995), here too we used a 2AFC (two-alternative forced-choice, e.g., left vs right) global motion direction discrimination task. A simulation run consisted of 1000 trials, divided into 40 blocks of 25 trials each. In each trial, motion of the 100 dots was simulated, with 15%, 25%, or 75% coherence. At the beginning of each run the network weights were randomized to small values following roughly a Gaussian distribution with zero mean and 0.2 unit standard deviation. Learning was supervised and implemented by standard delta rule back-propagation with momentum which exploits weight gradients to alter connections between units (1). The weight values were checked at the end of each run so that they never became smaller than +/- .0001 to ensure that a zero weight value, which would cause a permanent cessation of learning, does not occur.

From (1) it is clear that the order of the input values must be comparable to that of the error values to obtain quick and stable learning. To achieve this, the activities of the directionally-tuned units were normalized to a sum value of one. This was necessary due to the small input values (on the order of 10^{-12} at 25% signal) relative to the output error values (on the order of 1). The normalized input was used to regulate the transmission mechanism (detailed below) from input to hidden units.

We examined network performance using three measures: the final weight pattern, the percent correct by block, and the cumulative percent correct. After each run, we inspected the final weights to determine if there was a consistent pattern across runs. After each block, percent correct by block-the percentage of correct decisions in a block-was used to study the local learning trend. To study the long-term learning trend, after each trial we computed the cumulative percent correct, defined as the total number of correct decisions divided by the total number of trials.

The learning rate η was not kept constant throughout a run, but at the end of a block it was diminished if the average error in that block decreased at least by 10% compared to the average error value at the instant of the previous alteration of the η value. This damping of the learning rate promoted a monotonically decreasing average error value.

A run was considered to have converged if the last percent correct by block was above 70%, and if the cumulative percent correct was above 80% (these thresholds were chosen empirically based on observing the output error progression patterns). It should be noted that not all runs converged; however, the non-convergent runs were infrequent (less than 5%) and no consistent cause could be identified for the non-convergence.

4. Results

4.1 Opponency: There was a consistent pattern of weights across runs. We found that the input to hidden layer weights corresponding to the two directions used in the 2AFC task had opposite signs, and had magnitudes greater than for the other directions. In Figure 3(a) we show the weight pattern corresponding to a left-right discrimination task; for example, the weight for the unit tuned to the left direction was -1.8, while the weight for the unit tuned to the right direction was 2.1, for the low SR, 15% coherence case. This weight distribution is equivalent to having an opponent connection between the units tuned to the left and right directions. This opponent connection was a general solution to the forced choice single output node neural network. The opponency is consistent with suggestions made in the literature about inhibitory connection between neurons sensitive to different directions (the waterfall illusion, Newsome et al. 1989; Qian and Andersen 1994).

The extent of opponency was measured by the magnitude of the weights of the opponent units. We observed that the extent of opponency depended on the stimulus coherence. As stimulus coherence decreased (i.e., the task became more difficult), the extent of opponency increased. Thus, the network weights seemed to compensate for the weak signal.

<Fig. 3 about here>

4.2 High and Low Spontaneous Rate Networks: The high and low SR networks had different speeds of convergence and final levels of accuracy. The low SR network converged quickly to its final state, while the high SR network converged more slowly, as seen in Figure 3(b) where the long-term learning trends are plotted. At steady state of performance, the accuracy of the decisions, as demonstrated by short-term learning trends (plotted in Figure 3(c)), was higher for the high SR networks than for the low SR networks.

We repeated the learning experiments with either one or two hidden units, and the learning curves remained qualitatively similar, indicating that the number of hidden units did not have a noticeable influence on performance.

5. Discussion

The backpropagation network learned the global motion task successfully while incorporating mean-rate coding and neuronal noise.

What properties of MT neurons may be involved in learning? Direction selectivity enables MT neurons to respond selectively to motion. By grouping responses from a set of neurons selective to the cardinal directions, we obtained a population coding of the global motion. Opponency in direction selective responses *emerged* out of training, suggesting that opponency may be a crucial property for learning to discriminate opposite directions in global motion.

Learning induced a pattern of opponent weights for opposing directional neurons for both high SR and low SR networks at all signal levels. The pattern of opponent weights (i.e., weights of opposite signs) is not a novel observation. Indeed, for a large number of classification problems, opponency is the optimal solution, and backpropagation normally arrives at this solution. A substantial body of data supports opponent connections in the motion pathway that includes psychophysical evidence (Stromeyer et al. 1984; Qian et al., 1994) and physiological evidence (Newsome et al. 1989; Qian and Andersen 1994). The waterfall illusion and the perception of counterphase gratings are attributed to opponent connections of motion mechanisms. Our backpropagation representation suggests that human subjects may exploit this opponency in successful performance of the 2AFC task of global motion discrimination. An alternative proposal for coherent motion direction computation was made by Nowlan and Sejnowski (1995), in which regions of maximal motion information are selected through a mechanism trained to recognize such regions. It is likely that a mechanism of this type operates at the V1 level, implying learning at a stage different from representation of the global motion. The exploration of this possibility is the focus of our current research.

It could be argued that the network naturally reaches an "opponent" solution because of either or both of the following: discrimination between opposite directions of motion, and the forced choice paradigm. To test whether the opponency was a result of using opposite directions in the task, we tested with two additional stimuli: left vs up, and left vs noise (0% coherence). As before, we obtained opponency for these stimuli, even though the extent of opponency was reduced. The left vs. up task resulted in opponency between the left and up nodes while for the left vs. noise case the opponency was between left and all the remaining directions. The extent of opponency was the lowest for the left vs. noise case. The possible influence of the forced choice paradigm was eliminated by using two output nodes, and allowing for four output choices. This had no effect on the opponency result.

In the literature, the information transmission process has been approached in two ways: *temporal coding* in which the exact temporal position of a neuronal spike encodes information (see for e.g., Kruglyak and Bialek 1993), and rate coding in which the average firing rate encodes information. A duplex approach to information transmission considers both temporal and rate coding aspects of the process (Bialek et al. 1991; Licklider 1951; Skaggs and Cariani 1994).

The low and high SR networks, arising out of the physiological modifications, have interesting properties. A high SR network was slower to learn but obtained a more accurate solution while a low SR network learned quickly but reaches a less accurate solution. The SR of this network construction produced results similar to changing the learning rate, namely the η in classical backpropagation. The physiological spiking process added to the network creates a more physiologically plausible means of controlling network learning rate as compared to the 'magic number' η. The SR also provides a more stable way of controlling learning using parameters in the feed-forward path, since changes in SR do not result in the instability associated with changes in η. The use of SR to control the learning rate is not restricted to the supervised learning network presented here since the spiking process is independent of the learning algorithm.

In the psychophysical experiments, we found no difference in performance whether feedback was provided or not. We modeled the performance improvement using *unsupervised* learning mechanisms (Vaina et al. 1995; Sundareswaran and Vaina 1996). The objective in that modeling was to explain the performance improvement through physiologically plausible mechanisms. In the current study, the objectives were to determine the representation useful in learning global motion direction discrimination, and to study the effects of physiological modifications on the performance of the backpropagation algorithm. We found that opponency is important for learning direction discrimination, and that spontaneous rate provides an alternative, stable means of controlling the learning rate in backpropagation.

A classical neural network with physiological modifications such as those proposed here may prove useful for more in-depth work on learning vision tasks.

Acknowledgments

We thank W. Bialek for his comments. We also thank Dan Kacher and Maya Sharon for the their help in the initial stages of the project. A preliminary version of these results were presented at ARVO, 1995. This work was supported by the ONR contract N00014-93-1-00381 and NIH 2R01EY07861-06A1.

References

Ahissar, M., and Hochstein, S. 1993. Attentional control of early perceptual learning. *Proc. Natl. Acad. Sci.* **90**, 5718-5722.

Ball, K. and Sekuler, R. 1987. Direction-specific improvement in motion discrimination, *Vision Research* **27**, 953-965.

Bialek, W., Rieke, F., van Stevenink, R. R., and Warland de Ruyter, D. 1991. Reading a neural code. *Science*.**252**, 1854-1856.

Crick, F. 1989. The recent excitement about neural networks. *Nature* **337**, 129-132.

Fendick, M. and Westheimer, G. 1983. Effects of practice and the separation of test targets on foveal and perfoveal hyperacuity. *Vision Research* **23**, 145-150.

Fiorentini, A. and Berardi, N. 1980. Perceptual learning specific for orientation and spatial frequency. *Nature* **287**, 43-44.

Georgopoulos, A.P. *et al*, 1994a. Measuring Synaptic Interactions, *Science*, **263**, 1295-1297.

Georgopoulos, A.P. *et al.*, 1994b. Cognitive Neurophysiology of the Motor Cortex. *Science,* **260**, 47-52.

Heller, J., Hertz, J.A., Troels, W.K., and Richmond, B.J. 1995. Information flow and temporal coding in primate pattern vision, *J Comp. Neurosci.* **2**, 175-193.

Heywood, M. and Noakes, P. 1993. Simple Addition to Back- Propagation Learning for Dynamic Weight Pruning, Sparse Network Extraction, and Faster Learning, *IEEE International Conference on Neural Networks,* **2**, 620-625.

Karni, A. and Sagi, D. 1991. Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proc. Natl. Acad. Aci. USA*, **88**, 4966-4970.

Kruglyak, L. and Bialek, W. 1993. Statistical Mechanics for a Network of Spiking Neurons. *Neural Comp.* **5**, 21-31.

Lehky S.T., and Sejnowski, T. 1990. Neural model of stereoacuity and depth interpolation based on a distributed representation of stereo disparity. *J. Neurosci.* **10**(7).

Licklider, J. C. R. 1951. A duplex theory of pitch perception, Experientia VII. **4**, 128-134.

Mazzoni, P., Andersen, R.A., and Jordan, M.I. 1991. A more biologically plausible learning rule for neural networks. *Proc. Natl. Acad. Sci. USA* **88**, 4433-4437.

McAuley, J.D. and Stampfli, J. 1994. Analysis of the Effects of Noise on a Model for the Neural Mechanism of Short-Term Active Memory. *Neural Comp.* **6**, 668-678.

McKee, S.P. and Westheimer, G. 1978. Improvement in vernier acuity with practice. *Perception & Psychophysics*, **24**, 258-262.

Newsome, W.T., Britten, K.H., Movshon, J.A., and Shadlen, N.M. 1989. Single neurons and the perception of visual motion. In: Lam D.M.K., Gilbert, C.D. (eds) Neural mechanisms of visual perception: proceedings of the retinal research foundation. Portfolio Publishing, Huntington, NY. pp. 171-198.

Newsome, W. T., and Pare, E. 1988. "A Selective Impairment of Motion Perception Following Lesions of the Middle Temporal Visual Area," *J Neurosci.*, **8**, 2201-2211.

Nowlan, S. and Sejnowski, T. 1995. A selection model for motion processing in area MT of primates. *J. Neurosci.* **15**, 1195-1214.

Peres and Hochstein, S. 1994. Modeling perceptual learning with multiple interacting elements: a neural network model describing early visual perceptual learning, *J Comp. Neurosci.* **1**, 328-338.

Poggio, T., Fahle, M., and Edelman, S. 1992. Fast perceptual learning in visual hyperacuity, *Science*, **256**, 1018-1021.

Qian,N., Andersen, R.A., and Adelson, E.H. 1994. Transparant motion perception as detection of unbalanced motion signals. I. Psychophysics, *J. Neurosci.* **14**(12), 7357-7366.

Qian, N. and Andersen, R.A., 1994. Transparent motion perception as detection of unbalanced motion signals. II. Physiology, *J. Neurosci.* **14**(12), 7367-7380.

Rumelhart, D.E., Hinton, G.E., and Williams, R.J. 1988. Learning Internal Representations by Error Propagation. In *Parallel Distributed Processing*, Vol. 1, Rumelhart, D.E., and McClelland, J.L., eds, The MIT Press, Cambridge.

Shadlen, M.N. and Newsome, W.T. 1994. "Noise, Neural Codes, and Cortical Organization,"*Current Opinion in Neurobiology*, **4**, 569-579.

Softky, W.R. and Koch, C. 1993. "The Highly Irregular Firing of Cortical Cells is Inconsistent with Temporal Integration of Random EPSPs," *J Neurosci*, **13**, 334-350.

Shiu, L.-P., and Pashler, H. 1992. Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Perception and Psychophysics* **52**(5), 582-588.

Sundareswaran, V., and Vaina, L.M. 1996. Adaptive computational models of fast learning of motion direction discrimination, *Biological Cybernetics* (in press).

Vaina, L.M., Sundareswaran, V. and Harris, J.G. 1995. Learning to Ignore: Psychophysics and Computational Modeling of Fast Learning of Direction in Noisy Motion Stimuli, *Cognitive Brain Research* **2**(3), 155-163.

Vogels, R. and Orban, G.A. 1985. The effect of practice on the oblique effect in line orientation judgments. *Vision Research* **25**,1679-1687.

Weiss, Y., Edelman, S., and Fahle, M. 1993. Models of perceptual learning in vernier hyperacuity. *Neural Comp.* **5**, 695-718.

Zipser, D. and Andersen, R.A. 1988."A Back-propagation Programmed Network that Simulates Response Properties of a Subset of Posterior Parietal Neurons," *Nature*, **331**, 679-683.

Figure Legends

Figure 1. - Schematic depiction of the stimulus used for both the psychophysical and neural network experiments. The three examples represent different levels of coherence-defined as the percentage of the dots that move in the coherent (signal) direction-used in the simulations (15%, 25%, and 75%). Filled circles represent signal dots while open circles represent noise dots. The arrow for each dot shows the direction of motion of the dot.

Figure 2. - The neural network structure and physiological modifications applied to the network are shown. On the left, the randomly fluctuating cell potential is modulated by the activities of the directionally tuned units. The spontaneous rate of each neural channel, shown on top, is increased as the firing threshold is depressed with directional stimulation. At the network summation nodes the spikes in each "time window" are integrated to create the values passed to the neural network input nodes. The network consists of a 4 unit input layer, a 2 unit hidden layer, and a single unit output layer.

Figure 3. - (a) The opponency of the network weights are shown for both high and low SR networks for all coherence levels. The figure shows the magnitude of the weights for each directional unit and its polarity with respect to the other units. (b) The cumulative responses for the high and low SR networks are shown. The results have been averaged over three simulations. (c) The percent correct by block, averaged over three simulations, for the high and low SR networks.

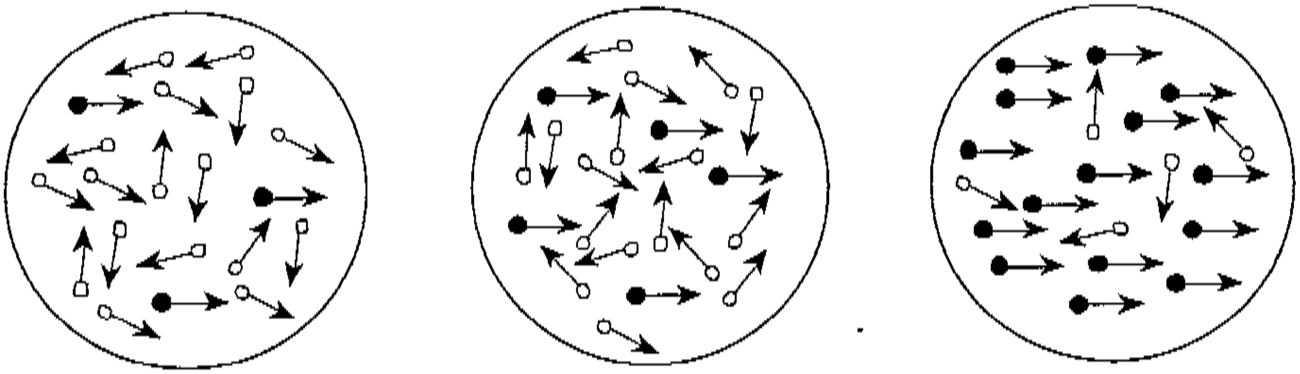

15% Coherence 25% Coherence 75% Coherence

Figure 1

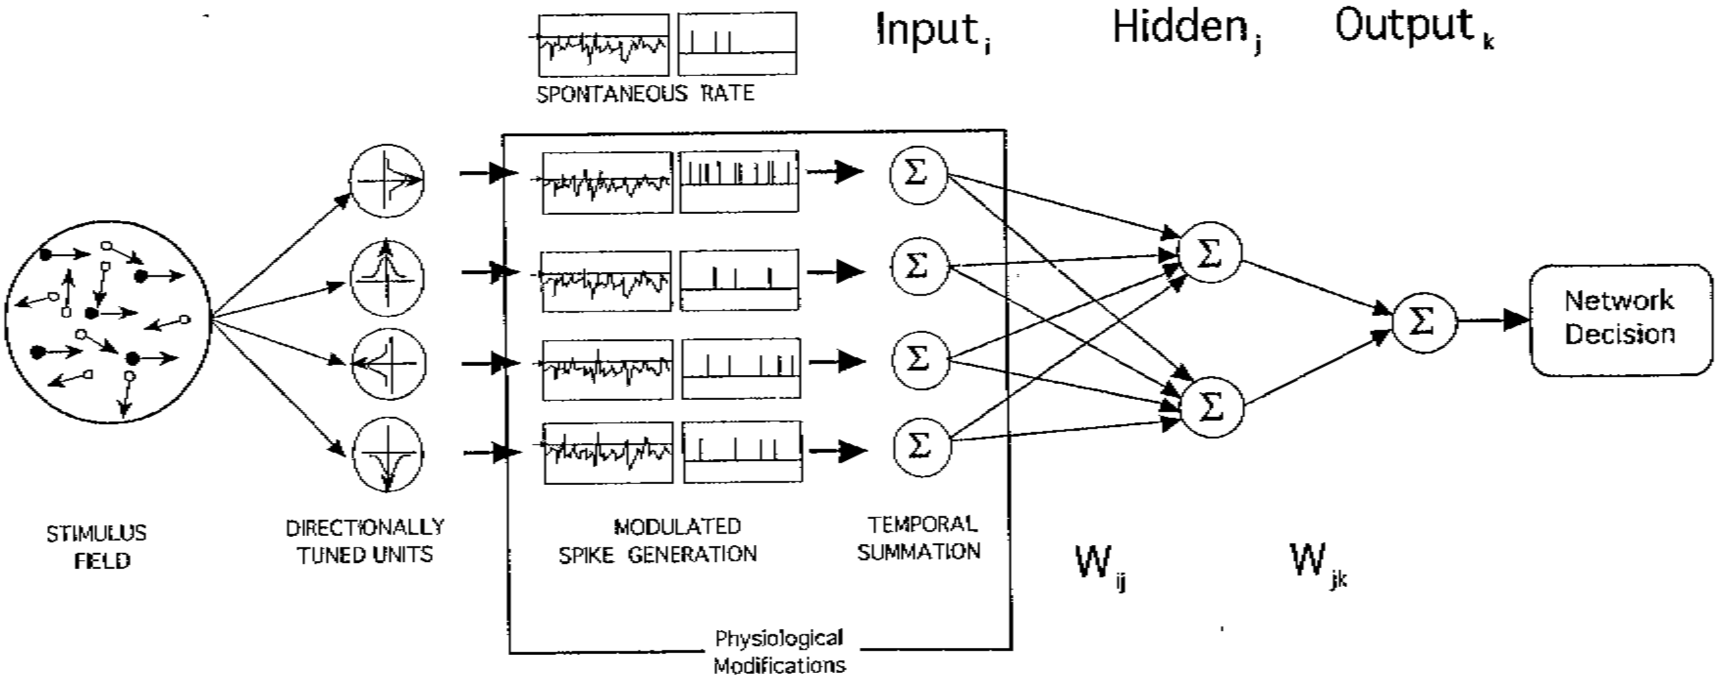
SPONTANEOUS RATE
Input$_i$
Hidden$_j$
Output$_k$
Σ
Network Decision
STIMULUS FIELD
DIRECTIONALLY TUNED UNITS
MODULATED SPIKE GENERATION
TEMPORAL SUMMATION
W_{ij}
W_{jk}
Physiological Modifications

Figure 2

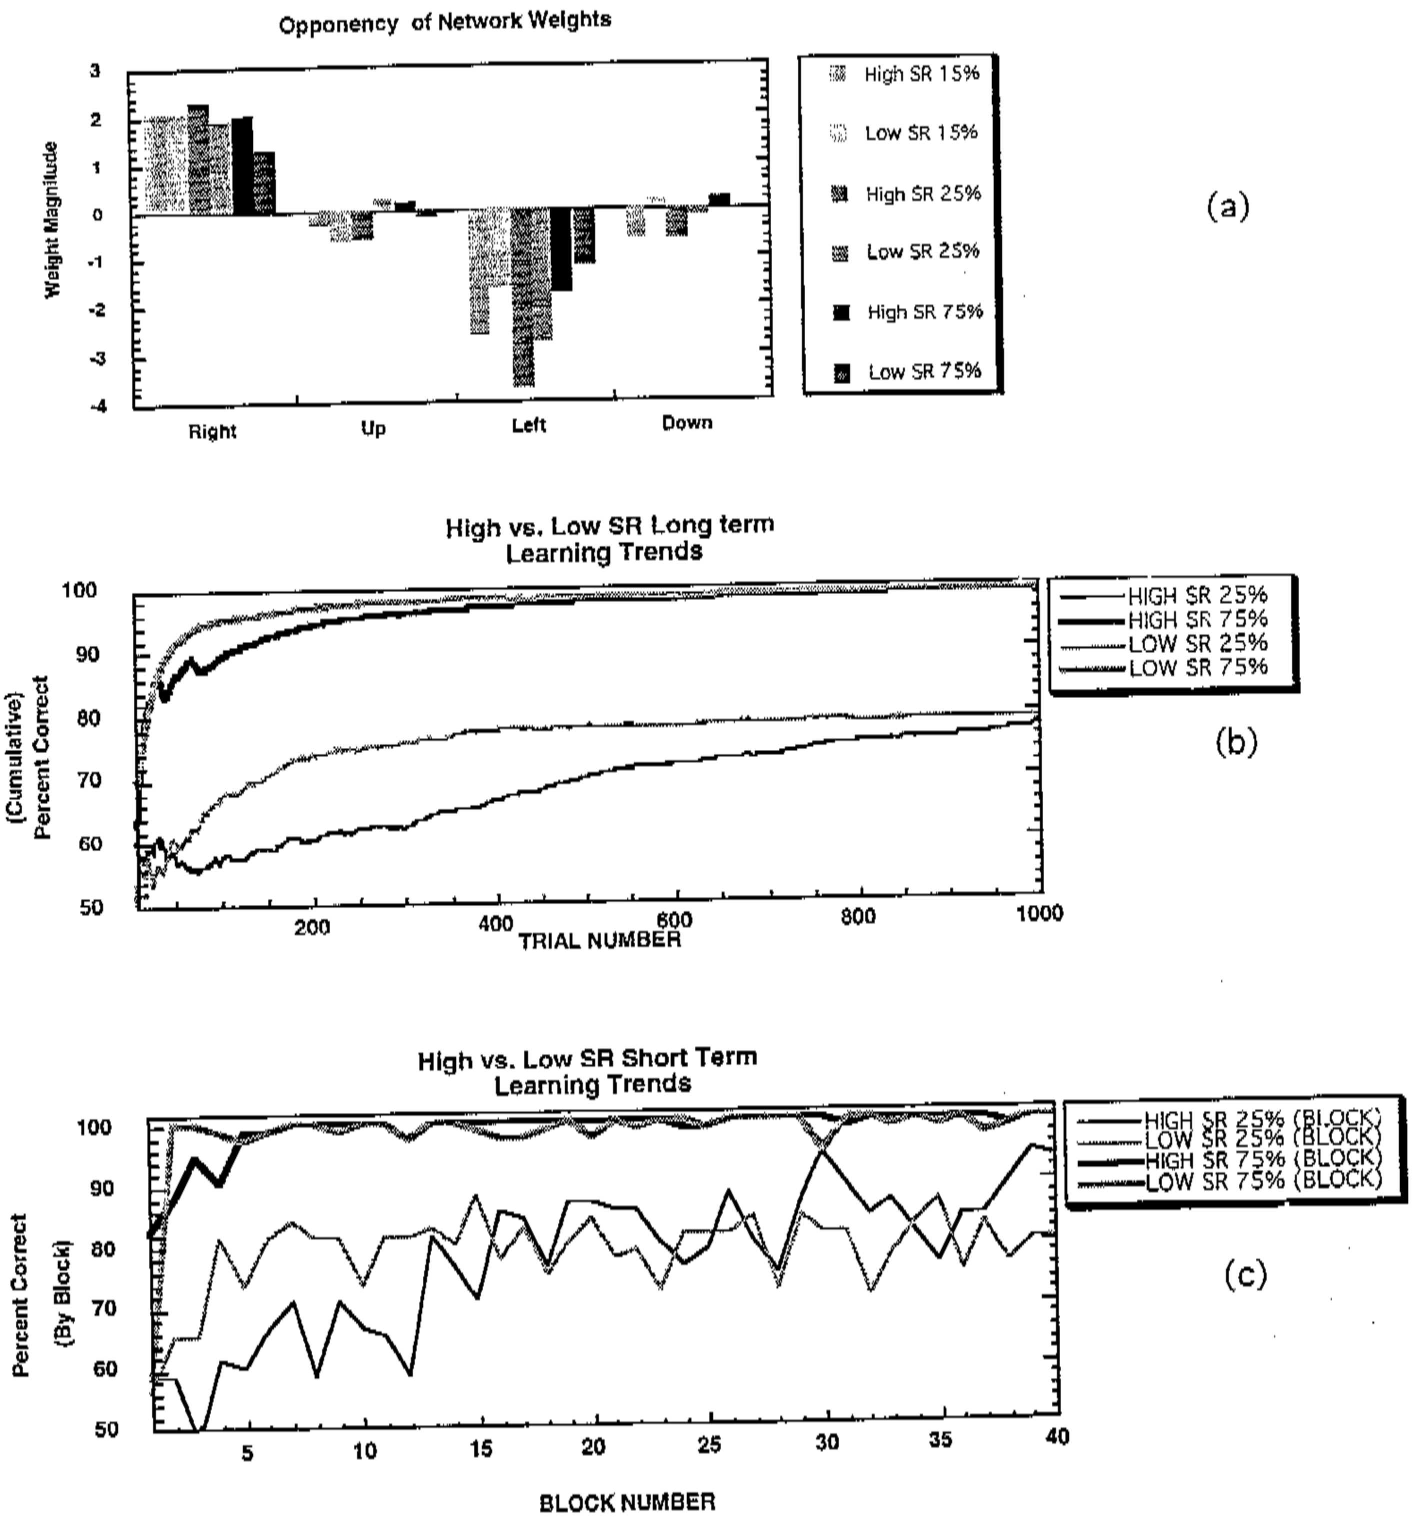
Opponency of Network Weights
Weight Magnitude
3
2
1
0
-1
-2
-3
-4
Right
Up
Left
Down
High SR 15%
Low SR 15%
High SR 25%
Low SR 25%
High SR 75%
Low SR 75%
(a)
High vs. Low SR Long term Learning Trends
(Cumulative) Percent Correct
100
90
80
70
60
50
200
400
600
800
1000
TRIAL NUMBER
HIGH SR 25%
HIGH SR 75%
LOW SR 25%
LOW SR 75%
(b)
High vs. Low SR Short Term Learning Trends
Percent Correct (By Block)
100
90
80
70
60
50
5
10
15
20
25
30
35
40
BLOCK NUMBER
HIGH SR 25% (BLOCK)
LOW SR 25% (BLOCK)
HIGH SR 75% (BLOCK)
LOW SR 75% (BLOCK)
(c)

Figure 3