ELSEVIER

Neurocomputing 11 (1996) 297–321

NEUROCOMPUTING

Learning generalisation and localisation: Competition for stimulus type and receptive field

Mike W. Oram *, Peter Földiák

School of Psychology, University of St. Andrews, St. Andrews, Fife KY16 9JU, UK

Received 17 March 1995; accepted 14 September 1995

Abstract

The evidence from neurophysiological recordings from the primate visual system suggests that sensory patterns are processed using units arranged in a hierarchical multi-layered network. Responses of these units show progressively increasing receptive field size combined with selectivity for increasing stimulus complexity at successively higher levels. It is argued that the rate of the increase in receptive field size is less than the maximum possible given the initial spread of neuronal projections that occurs during development. We show here that a competitive learning mechanism using a 'trace-Hebbian' learning rule [14] with a larger number of competing output units learns not only positional invariance for a given input feature but can also establish restricted receptive field sizes (i.e. less than the maximum size given the initial connections). Importantly the same stimulus selectivity was maintained throughout the receptive field. It is shown that this is accompanied by a relative increase in the spatial evenness of the representation of each detector type across position within the input array. The network properties were found to be robust and stable over a wide range of learning parameters. We suggest that such a competitive mechanism may help account for the reported properties of cells in the ventral stream of the primate visual system.

Keywords: Vision; Competitive learning; Invariance

1. Introduction

In this article we extend previous work [14] using a modified Hebbian learning rule. The modification allowed a network to learn invariance, e.g. positional

* Corresponding author. Email: mwo@st-andrews.ac.uk

SSDI 0925-2312(95)00099-2

invariance of orientation selectivity [14]. Here we show that networks utilising a 'trace-Hebbian' competitive learning rule can generate restricted receptive field like properties as well as show positional invariance. We review briefly below the neurobiological evidence that with increasing distance (connections) from the retina, single neurons in visual cortex show increasing levels of positional invariance while maintaining the same stimulus selectivity throughout their receptive fields. We argue that the increase in receptive field size does not occur as rapidly as the connectivity between neurons could allow, and further that there are many neurons in visually responsive cortical areas without highly specified connectivities. These conditions form the neurobiological basis for the network design. We then briefly review properties of models that are related to those of the present article.

The human brain is capable of recognising a multitude of objects under a vast number of viewing conditions (e.g. a person at different distances and orientations). Differences in shape can also be accommodated (e.g. changes in 3-D posture). Objects and object-features (e.g. colours, edges, junctions etc.) give rise to collections of highly correlated properties ([3], for review see [4,11]). It has been suggested that a goal of sensory processing may be to form a representation that explicitly represents those collections of correlated properties or 'suspicious coincidences' that have a common cause [2–4]. Evidence for single neural recordings in the macaque temporal cortex support the notion of explicit representation of complex objects (e.g. [20,30,37]). The existence of such explicit representations indicates that the nature of neuronal representations of objects is sparse (i.e. represented by activity in a relatively small proportion of a large number of units). This appears to be true not only for object class but also for object identity. Recordings from temporal cortex support this notion for the representation of facial identity [51].

A remarkable property of neurons in the temporal cortex is that not only do they show selectivity for complex features or objects, but also that they show positional invariance [19,37,45]. This means that these cells show similar stimulus selectivity across the whole receptive field (e.g. for a small vertical striped triangle anywhere in the central 20 degrees of vision). Similar position invariant stimulus selectivity exists for complex cells in the primary visual cortex (e.g. [22]). Neurophysiological recordings from the macaque monkey visual cortex also indicate that receptive field size of individual neurons increases with increasing distance from the retina (measured as number of synapses or number of cortical areas). During development the axonal spread of forward connections between visually responsive areas is considerably larger than that seen in the mature cortex [12]. Therefore the rate at which the neuronal receptive field size increases is not as fast as is theoretically possible. It is also known that in conjunction with increasing receptive field size there is an accompanying increase in the complexity of stimuli which can evoke strong responses in the neurons, for example lines and edges in V1 to 'face selective' neurons in temporal cortex [31,35].

At initial developmental stages, there are a large number of neurons without highly specified (learnt) connections. Indeed it is probably that at least 50% of cells die during the development of the visual neocortex [12]. The neurobiological

evidence therefore suggests that primate 'form processing' visual system has many neurons present which have little or no specificity of connections, particularly during development. After development there are two concurrent progressions in the properties of the neurons. These are increasing receptive field size at a rate slower than theoretically possible and increasing selectivity for stimulus complexity (see [31,35] for reviews). The selectivity of neurons is invariant across position within the receptive field. In the present article we address the issue of the slower than maximum possible increases in receptive field size combined with positional invariance for one of several stimulus types.

Mechanisms of cortical map formation have been studied extensively [e.g. 24,46]. These studies show how neighbouring neurons in the cortical sheet can develop similar response selectivities (e.g. orientation tuning). The development of localised receptive fields seen in these cortical neurons has received less attention. Models have been proposed that will generate units with increasing receptive field sizes and increasing selectivity for stimulus complexity [e.g. 17]. Models such as Fukushima's [17] use limited connectivity between pairs of layers to achieve this result. The first layer of each pair uses competitive learning to increase the complexity of the stimulus requirements but maintains small receptive fields. The second stage selectively pools together the outputs of the appropriate units to form 'complex' units with invariance properties. In the Fukushima model positional invariance was achieved by 'hard wired' pooling connections. If hard wired connections were used in the primate visual system then the connectivity observed in the mature animal should be comparable to that seen in the developing cortex. As noted above, during development the initial connections laid down by single neurons are more extensive than seen in the mature cortex [12]. Földiák [14] showed that hard wired pooling can be replaced by an adaptive pooling mechanism that can *learn* invariance. A decaying average or 'trace' measure of activity allowed pooling of one input feature detector with a second input detector that became active soon after the first. In the example, oriented lines were swept across the input array. After the learning period, the outputs showed large receptive fields covering all areas from which input was received. Within the receptive fields the same stimulus selectivity was maintained. Importantly the selectivity of a given output unit was for only one of the four possible input detector types (i.e. the units showed positional invariance and selectivity for input feature type). The observed stimulus selectivity of the output units was achieved using a 'winner-take-all' mechanism giving rise to competition between outputs. Decorrelating anti-Hebbian connections between the output units [5,13] can also be used for this purpose (Oram and Földiák, unpublished studies). Decorrelation may also be useful for establishing combinatorial encoding of stimuli and the sparse representations that appear to be found in the primate visual system [15,51].

It follows from results of studies of dipole inputs to competitive networks [39] that limited regions of the input space can be assigned to different output units. It has also been shown that output units can learn positional invariance for particular classes or types of input detector when the number of output units equals the number of input classes [14]. The question we address here is whether a network

with competitive interactions between a higher number of output units than input detector types could learn to generalise for stimulus selectivity while simultaneously generating a receptive field that is smaller than its maximum potential size. If the properties of the network were to be similar to those seen in the primate visual system, then stimulus selectivity should be constant across the receptive fields.

2. Simulations

The starting point of the simulations is the network described by Földiák [14]. In half of the present simulations the number of output nodes was increased from 4 to 8. Briefly, the network consisted of an eight by eight input grid. At each grid point there were 4 orientation detectors (horizontal, vertical and the two oblique lines) giving a total of 256 input units. All the input units were connected to all the output units (see Fig. 1), initially with random weights [0-1]. During training a line selected randomly from one of 4 orientations was swept across the input grid in one of the two directions orthogonal to its orientation. We define each sweep of a line across the entire input grid as formulating a cycle of the network simulation. Such sweeps across the input array would simulate an object moving in the environment with the observer fixating a static point or the effect of self induced motion (e.g. eye movements or head/body movements relative to static objects).

The simple Hebbian learning rule depends on instantaneous activity and therefore only associates input patterns at any one time. A trace-Hebbian rule [14] was used here where the modification to the weights was proportional not to the instantaneous pre- and post-synaptic activities, but to the pre-synaptic activity and a running average of the post-synaptic activity of output i ($\bar{y}_i$) where

$$\bar{y}_i^{(t)} = (1-\delta)\,\bar{y}_i^{(t-1)} + \delta y_i^{(t)} \quad (t = \text{time},\ \delta = \text{trace decay constant})$$

The trace activity was calculated for each output node separately and maintained between sweeps. With small values of δ the trace is slow to increase but lasts for a long time (a long trace). With large values of δ the trace level rises and falls rapidly (a short trace). Evidence of a trace in the form of neuronal activity in temporal cortical neurons has been documented, where neural responses are maintained for up to several hundred milliseconds after an effective stimulus has been removed from sight (e.g. [37]). A trace could also be implemented using biochemical signals rather than persistence of neuronal firing (see [16,32] for further discussion). With a weight change rate $= \alpha$ and using subtraction of the previous weight to give a bounding decay term, gives a learning rule of

$$\Delta w_{ij}^{(t)} = \alpha \bar{y}_i^{(t)}\left(x_j^{(t)} - w_{ij}^{(t)}\right)$$

where Δw_{ij} is the change of weight between output i and input j at time t and α is the learning rate.

A winner-take-all mechanism was used to set the activity of the output unit with the highest activity to 1 and all others to 0. We note here that it has been shown

that winner-take-all networks are robust, with many possible implementations [23]. The choice of using a simple winner-take-all mechanism does not therefore invalidate the biological validity of the network. We have run the same network using de-correlation [5,13] of output nodes and obtained similar results. The implementation of a de-correlating network may be more biologically feasible than a strict winner-take-all implementation.

Fig. 1 reproduces earlier results [see 14]. In each row of the figure, the selectivity of each of the four output nodes is shown. The development of position invariance is seen running down each column. The length of the line segment at each of the input grid locations indicates the relative response of the output node to a single oriented line presented at that location. At the start, the random connections give rise to a disorganised pattern of selectivity. By cycle 100, a clear pattern is emerging, with each output unit responding to only one of the oriented lines. The output unit selectivities remains relatively stable until cycle 500. By cycle 750, the selectivity of output units 3 and 4 appear to have swapped. This transition was maintained until cycle 1000.

In a second simulation, the number of output nodes was increased to 8. The results of one run are shown in Fig. 2. With eight output nodes, a stable pattern of selectivity was reached by cycle 200. This is approximately twice as long as that taken when only 4 output units were used. The effects of having more output units (8) than detector types (4) is most clearly seen in output nodes 1,2 and 3. Output node 1 shows selectivity for an oblique line (\), but only in the lower left corner of the input grid. Node 2 shows selectivity for the same feature, but only in the upper right of the input grid, whereas node 3 shows selectivity again for the same feature, but in the central region. These results demonstrate that a simple competitive network can learn positional invariance with limited receptive field size.

Looking down the columns, it is clear that there are changes in the precise locus of maximal selectivity. For example, output node 3 shows selectivity for the central region at cycle 200, the top right at cycle 300, back to the centre at cycle 400 and 500, then shifts slightly to the lower left by cycle 750 and moves back towards the upper right by cycle 1000. However, these changes are for the receptive field of the unit, and do not involve a change in the orientation or input selectivity. Note also that changes in the 'receptive field' position in one node are accompanied by complementary changes in the receptive fields of the other nodes selective for the same feature.

A more systematic investigation into the effects of learning parameters in the network when increasing the number of output nodes was carried out. The value of the trace duration was varied ($\delta = 0.01$, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, and 0.6) as was the rate of weight change ($\alpha = 0.005$, 0.01, 0.02, 0.03 0.05). A simulation was run using each of the combinations of the δ and α with 4 and 8 output nodes, giving a total of 80 simulations. In order to investigate numerically the properties of the network an analysis was performed on the weights between input and output nodes after an initial training period.

Each of the 80 network simulations was given 2000 initial training cycles. A further 1000 cycles were then run, with the value of all the weights being stored

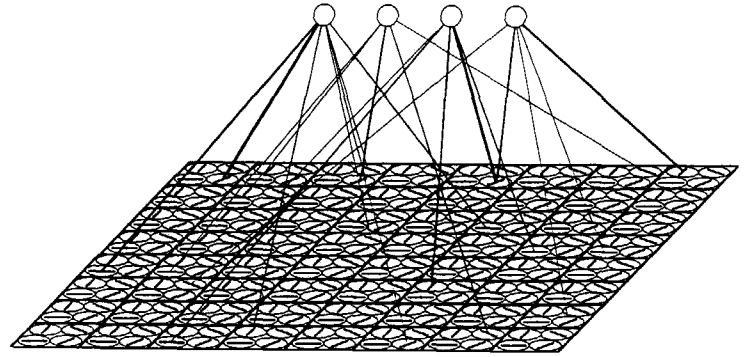

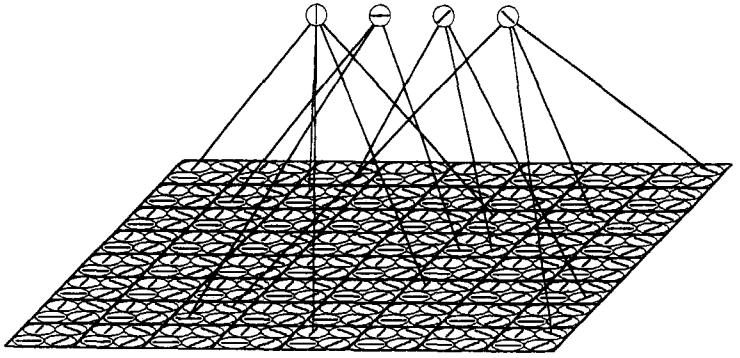

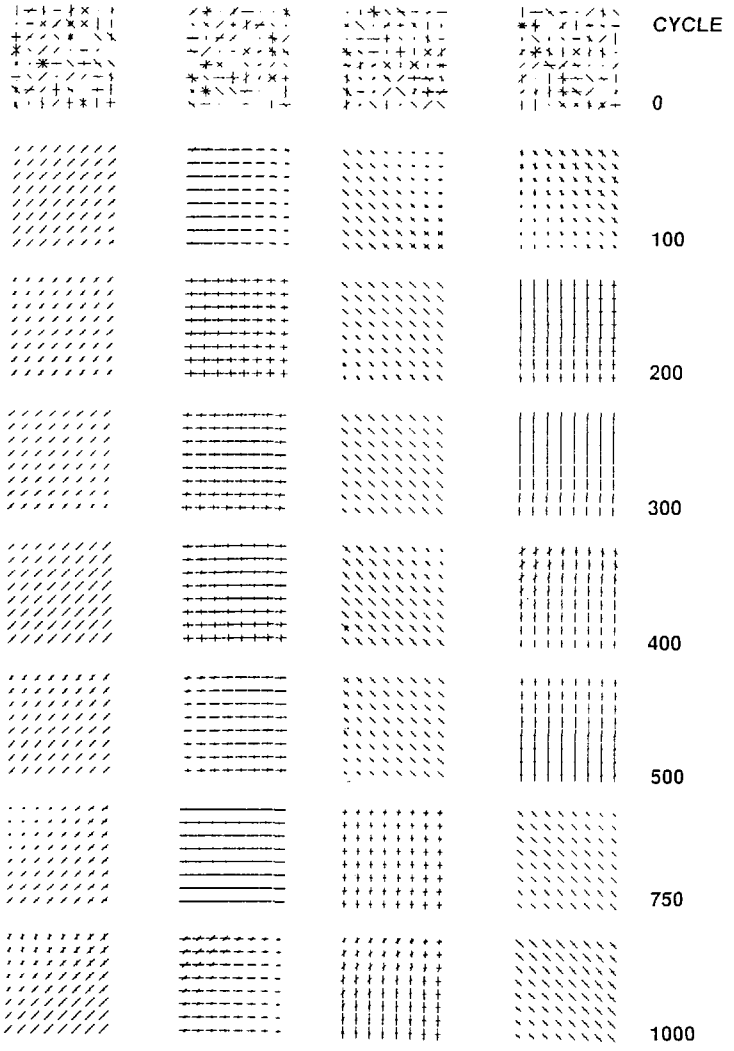
CYCLE
0
100
200
300
400
500
750
1000

every 100 cycles (sweeps). The values of the weights were then analysed to determine the effects of increasing the number of output nodes. The form of the analysis follows the method of analysis of variance (ANOVA) in order to determine the main source of variation in all of the stored weights. An estimate of the sums of squares for output nodes (SS_o) can be calculated

$$SS_o = \sum_i \left(\sum_{jkl} w_{ijkl}\right)^2 / JKL - \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL$$

where I = number of output nodes (4 or 8), J = number of detectors (4), K = number of positions (64) and L = number of sample periods (10), w_{ijkl} = weight between input detector j at position k to output i at sample l. The sums of squares for detector type, position and cycle (SS_D, SS_P, SS_C respectively) were calculated using

$$SS_D = \sum_j \left(\sum_{ikl} w_{ijkl}\right)^2 / IKL - \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL$$

$$SS_p = \sum_k \left(\sum_{ijl} w_{ijkl}\right)^2 / IJL - \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL$$

$$SS_C = \sum_l \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJK - \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL$$

SS_o is a measure of the variability of the weights from the input units onto each of the output nodes independent of the position, detector and time. A low value indicates that the weights from all detectors at all input positions was stable over time and approximately equal (i.e. each node received approximately the same total input weight). A high value indicates that one or more nodes received a noticeably different total weight input than the other nodes. Similarly a high value of SS_D would indicate that one or more of the detectors had a greater total weight than other detector types. SS_P is the analogous measure for the similarity of the weights (independent of the detector type, output node and time at which they were measured) to vary between the different positions, and SS_C is the measure of the variability over time of the total sum of the weights.

It is also possible to examine the variability of the weights due to 'interactions' of output node, detector type, input position and cycle number. For example, the tendency for different detectors to make strong connections at different input grid

Fig. 1. Development of positional invariant stimulus selectivity with 4 output nodes. *Left*: Schematic representation of the network with 4 output units. The 4 input detector types are shown in each of the 8 by 8 input array locations. The thickness of the lines connecting the input detectors to the output units represents the weight. Only a few of the weights are indicated. The upper figure illustrates the network before training, with variable strength connections from random input orientation detectors to output units. After training (lower) the connections are from a single input detector type to each output unit. The selectivity of the output unit is indicated by the oriented line in the output unit. *Right*: The stimulus selectivity of each of the 4 output nodes is indicated by the length of the oriented lines. Each column represents one output node. Each row shows the selectivity after the number of training cycles shown on the right hand side. The learning rate for this figure was $\alpha = 0.02$, trace duration $\delta = 0.2$. Note that each output node showed selectivity for only one of the orientations.

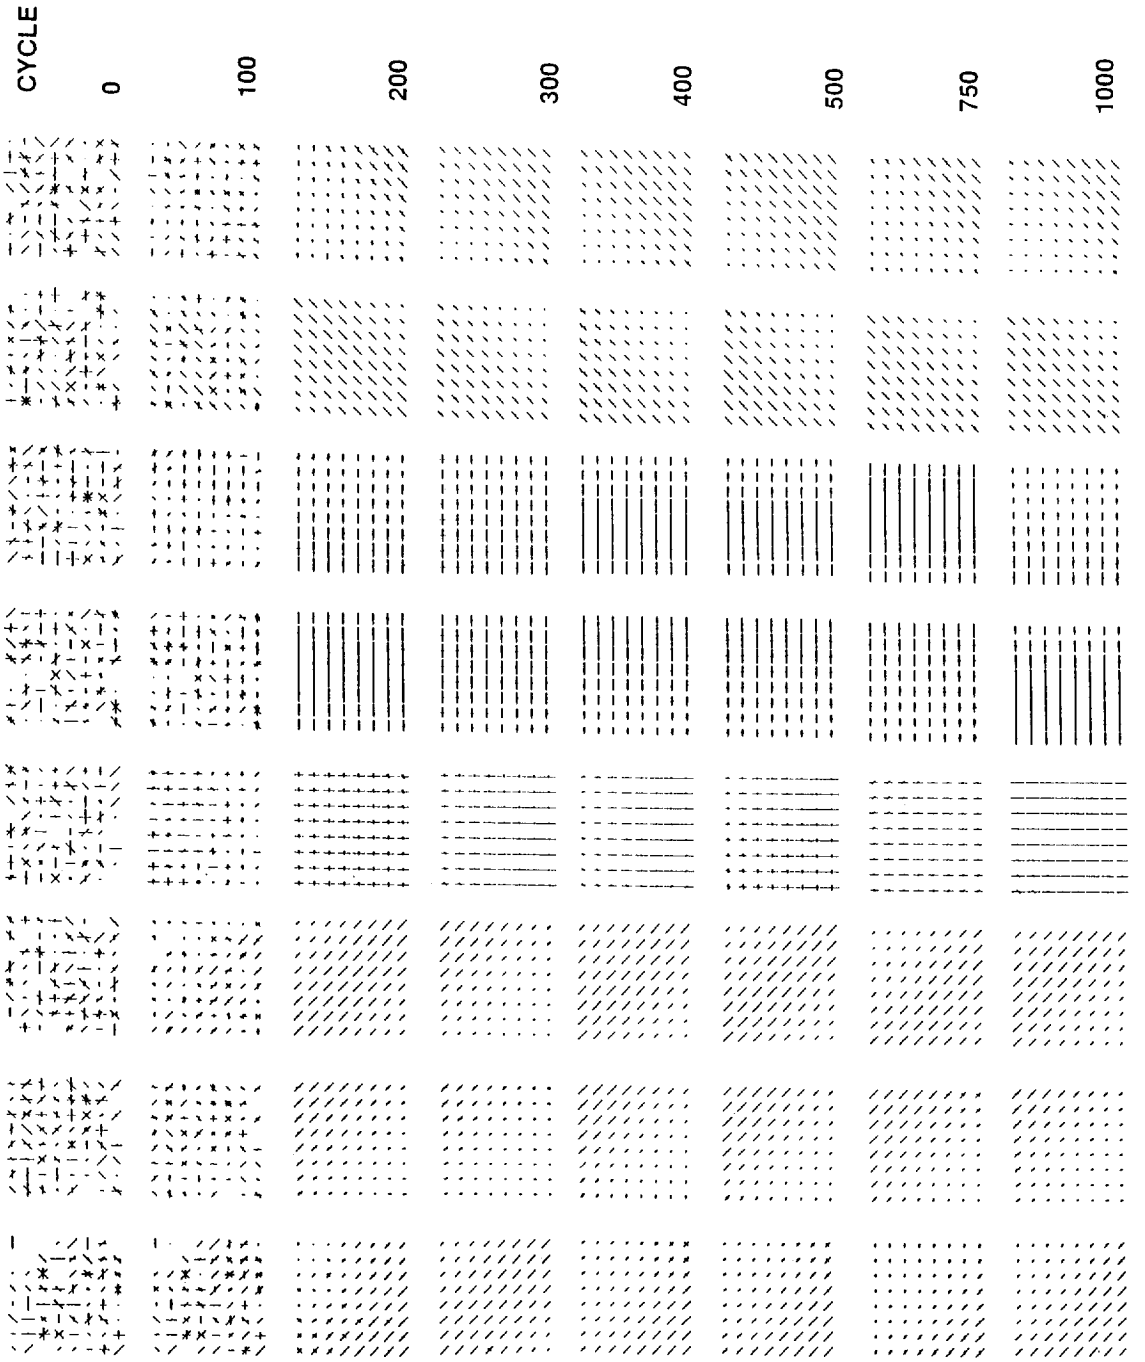
CYCLE
0
100
200
300
400
500
750
1000

Table 1
Sums of squares

	Code	
Output	O	$SS_O / SS_{Total} = 0.0018$
Detector	D	$SS_D / SS_{Total} = 0.0041$
Position	P	$SS_P / SS_{Total} = 0.0169$
Cycle	C	$SS_O / SS_{Total} = 0.00001$
Output by Detector	OD	$SS_{OD} / SS_{Total} = 0.0139$
Output by Position	OP	$SS_{OP} / SS_{Total} = 0.1737$
Output by Cycle	OC	$SS_{OC} / SS_{Total} = 0.0001$
Detector by Position	DP	$SS_{DP} / SS_{Total} = 0.1774$
Detector by Cycle	DC	$SS_{DC} / SS_{Total} = 0.0003$
Position by Cycle	PC	$SS_{PC} / SS_{Total} = 0.0030$
Output by Detector by Position	ODP	$SS_{ODP} / SS_{Total} = 0.4945$
Output by Detector by Cycle	ODC	$SS_{ODC} / SS_{Total} = 0.0005$
Output by Position by Cycle	OPC	$SS_{OPC} / SS_{Total} = 0.0231$
Detector by Position by Cycle	DPC	$SS_{DPC} / SS_{Total} = 0.0197$
Output by Detector by Position by Cycle	ODPC	$SS_{ODPC} / SS_{Total} = 0.0719$

positions is given by SS_{DP}. A high value of SS_{DP} would indicate that particular detector types in the input array are more strongly connected to the output nodes at certain positions. The terms used to calculate the tendencies for '2-way interactions' are given below:

$$SS_{OD} = \sum_{ij}\left(\sum_{kl} w_{ijkl}\right)^2 / KL - \left(SS_O + SS_D + \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL\right)$$

$$SS_{OP} = \sum_{ik}\left(\sum_{jl} w_{ijkl}\right)^2 / JL - \left(SS_O + SS_P + \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL\right)$$

$$SS_{OC} = \sum_{il}\left(\sum_{jk} w_{ijkl}\right)^2 / JK - \left(SS_O + SS_C + \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL\right)$$

$$SS_{DP} = \sum_{jk}\left(\sum_{il} w_{ijkl}\right)^2 / IL - \left(SS_D + SS_P + \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL\right)$$

$$SS_{DC} = \sum_{jl}\left(\sum_{ik} w_{ijkl}\right)^2 / IK - \left(SS_D + SS_C + \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL\right)$$

$$SS_{PC} = \sum_{kl}\left(\sum_{ij} w_{ijkl}\right)^2 / IJ - \left(SS_P + SS_C + \left(\sum_{ijkl} w_{ijkl}\right)^2 / IJKL\right)$$

Fig. 2. Development of receptive fields showing positional invariant stimulus selectivity. The stimulus selectivity of each of the 8 output nodes is indicated by the length of the oriented lines. Each column represents one output node. Each row shows the selectivity after the number of training cycles shown on the right hand side. The learning rate for this figure was $\alpha = 0.02$, the trace duration $\delta = 0.2$. Each output node showed stimulus selectivity for one of the input orientations over a restricted area of the total input array.

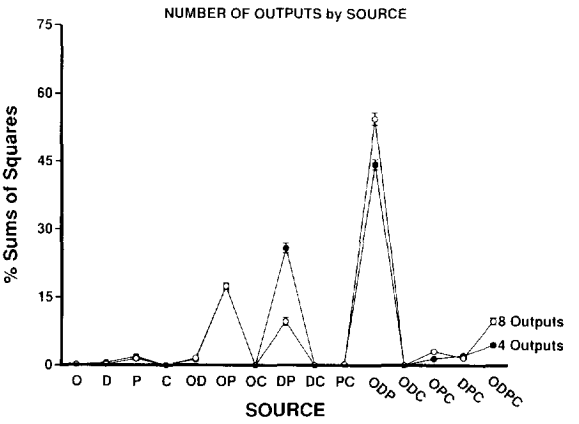

Fig. 3. The proportion of variability of the sums of squares of the weights accountable by source. The proportion (given as a percentage) of the total sums of squares of the weights are shown for the 4 and 8 output networks. The results are collapsed across the range of learning rate α and trace duration δ examined. See text for details.

Continuing along the same lines, there are four possible '3-way interaction' terms, given by

$$SS_{ODP} = \sum_{ijk}\left(\sum_l w_{ijkl}\right)^2/L - \left(SS_O + SS_D + SS_P + SS_{OD} + SS_{OP} + SS_{DP} + \left(\sum_{ijkl} w_{ijkl}\right)^2/IJKL\right)$$

$$SS_{ODC} = \sum_{ijl}\left(\sum_k w_{ijkl}\right)^2/K - \left(SS_O + SS_D + SS_C + SS_{OD} + SS_{OC} + SS_{DC} + \left(\sum_{ijkl} w_{ijkl}\right)^2/IJKL\right)$$

$$SS_{OPC} = \sum_{ikl}\left(\sum_j w_{ijkl}\right)^2/J - \left(SS_O + SS_P + SS_C + SS_{OP} + SS_{OC} + SS_{PC} + \left(\sum_{ijkl} w_{ijkl}\right)^2/IJKL\right)$$

$$SS_{DPC} = \sum_{jkl}\left(\sum_i w_{ijkl}\right)^2/I - \left(SS_D + SS_P + SS_C + SS_{DP} + SS_{DC} + SS_{PC} + \left(\sum_{ijkl} w_{ijkl}\right)^2/IJKL\right)$$

Fig. 4. Examples of changing trace duration and learning rate with 4 output nodes. Five examples of stimulus selectivity after 2000 training cycles are shown. The weights (of the connection) between the input (orientation type at the given locus) and each 4 output nodes is shown for each example. (A) learning rate $\alpha = 0.005$, trace duration $\delta = 0.01$. (B) learning rate $\alpha = 0.005$, trace duration $\delta = 0.6$. (C) learning rate $\alpha = 0.02$, trace duration $\delta = 0.2$. (D) learning rate $\alpha = 0.05$, trace duration $\delta = 0.01$. (E) learning rate $\alpha = 0.05$, trace duration $\delta = 0.6$.

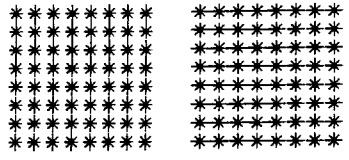

A

Trace Duration : $\delta = 0.01$ Learning Rate : $\alpha = 0.005$

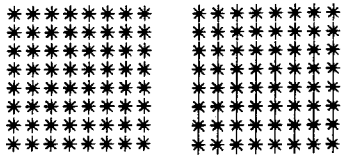

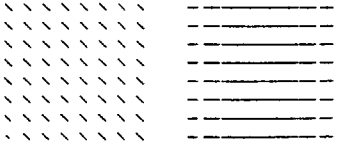

B

Trace Duration : $\delta = 0.6$ Learning Rate : $\alpha = 0.005$

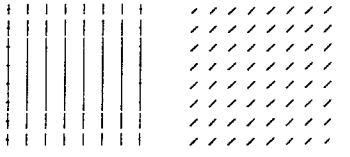

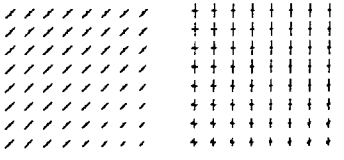

C

Trace Duration : $\delta = 0.2$ Learning Rate : $\alpha = 0.02$

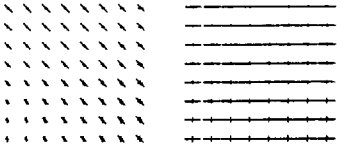

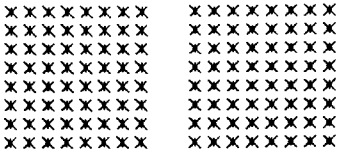

D

Trace Duration : $\delta = 0.01$ Learning Rate : $\alpha = 0.05$

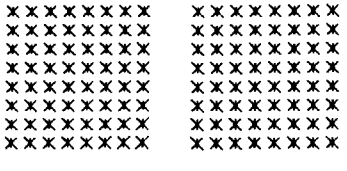

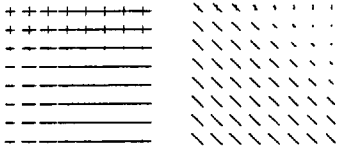

E

Trace Duration : $\delta = 0.6$ Learning Rate : $\alpha = 0.05$

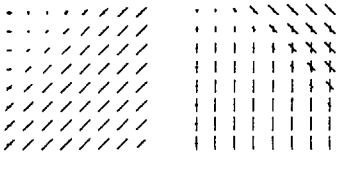

Of particular interest for this article is the value SS_{ODP}. A high value of SS_{ODP} indicates that there is high variability in the weights connecting particular detector types to the output units that depends on the position of the particular detector type within the input array. A lower value of SS_{ODP} would indicate that each output unit receives approximately the same weight of each input detector type from each position in the input array. Note that a lower value of SS_{ODP} does not imply that each detector type connects with the same strength, only that for a given input detector type the weights are approximately equivalent across all positions within the input array.

The individual weights of the input detectors onto the output nodes were used to calculate the 'total sums of squares' (SS_{Total}).

$$SS_{Total} = \sum_{ijkl}\left(w_{ijkl} - \left(\sum_{ijkl} w_{ijkl}/IJKL\right)\right)^2$$

This gives an estimate of the total variability in the weights. SS_{Total} is equal to the sum of all the other calculated SS values plus the sums of squares due to the '4-way interaction' SS_{ODPC}. To examine the relative changes in the distribution of the variability of the weights all calculated SS values were normalised using SS_{Total}. Table 1 gives the breakdown of the proportion of the sums of squares associated with each source, regardless of the number of outputs used in the network.

The low value of the output value (O) indicates that each output receives very nearly equal total weights (regardless of position, detector or time). This shows that the competitive mechanism has distributed the weights equally amongst the output nodes. Similarly, all detector types are represented equally in the network (D = 0.41%), and that all positions are also represented equally (P = 1.69%). It is worth noting that overall the network is stable after the initial training period in that the variability of the weights over time (C) is very small (0.0015% of the total variability).

It can be seen that output by position (OP) and detector by position (DP) each account for approximately 17% of the sums of squares. The high OP value reflects that the total input weights (regardless of detector type and cycle) is greater around the edges of the input grid (position) with the proviso that for any one output node this reflects increased weights from the non-preferred orientation in edge positions (e.g. Fig. 1, cycle 1000). This is probably due to the trace following a sweep of the preferred orientation allowing association of subsequent sweeps of non-preferred orientations. The effect of the trace declines over time, so this association is most strongly seen around the edges of the input grid. Similarly, the

Fig. 5. Examples of changing trace duration and learning rate with 8 output nodes. Five examples of stimulus selectivity after 2000 training cycles are shown. The weight of the connection between the input (orientation type at the given locus) and each of the 8 output nodes is shown for each example. (A) learning rate $\alpha = 0.005$, trace duration $\delta = 0.01$. (B) learning rate $\alpha = 0.005$, trace duration $\delta = 0.6$. (C) learning rate $\alpha = 0.02$, trace duration $\delta = 0.2$. (D) learning rate $\alpha = 0.05$, trace duration $\delta = 0.01$. (E) learning rate $\alpha = 0.05$, trace duration $\delta = 0.6$.

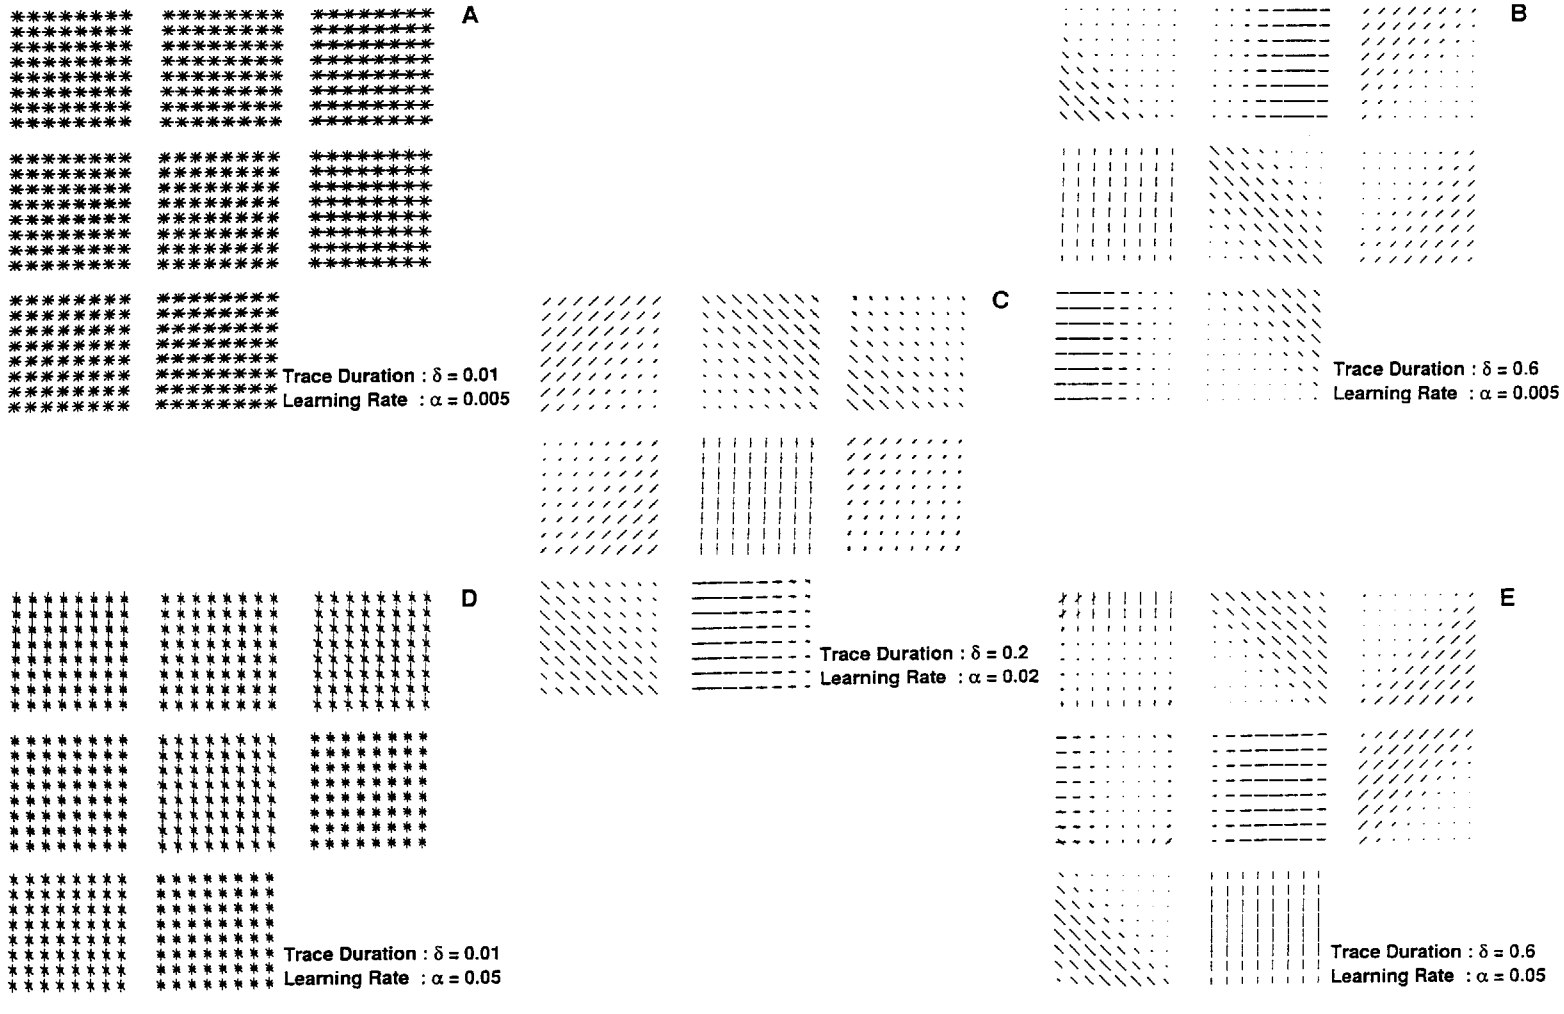
A
Trace Duration : δ = 0.01
Learning Rate : α = 0.005
B
C
Trace Duration : δ = 0.6
Learning Rate : α = 0.005
D
Trace Duration : δ = 0.2
Learning Rate : α = 0.02
E
Trace Duration : δ = 0.01
Learning Rate : α = 0.05
Trace Duration : δ = 0.6
Learning Rate : α = 0.05

detector by position (DP) sum of squares reflects the greater tendency (collapsed across output nodes and cycle) to have a greater overall connection from the edges of the input array onto the output nodes.

The percentage accounted for by output by detector by position (ODP = 49.5%) reflects even more clearly that at different input positions, the relative strengths of the detector types feed onto different output nodes with significantly different weights (i.e. output nodes are selective for different orientations and that the strength of the inputs of the particular input varies with position). This can be interpreted as variations in the degree of stimulus selectivity over position. The ODPC source accounts for a small, but non-negligible proportion of the sums of squares and reflects the slowly changing receptive field position and orientation tuning over time (e.g. Fig. 2).

Fig. 3 shows, in graph form, the proportion of the sums of squares accounted by the different sources for the 4 output node and 8 output node networks. There are four sources where a noticeable change occurs between the two types of network. These are outlined below.

Increasing the number of output nodes causes a large decrease in the proportion of the sums of squares attributable to the detector by position interaction (DP). This indicates that, regardless of which output node or the time (cycle) at which it is examined, a greater number of output nodes increases the evenness of the strength of the representation of the detector types across position.

Complementary to the increase in the evenness of the representation of detector types across position when the number of output nodes is changed from 4 to 8 is an increase in the relative variability of the weights attributable to output nodes by detector type by position (ODP). This corresponds to the development of the 'receptive field' type properties (i.e. different output nodes have different orientation selectivities at different positions across the input array). The increase is illustrated by comparing Figs. 1 and 2.

Finally, increasing the number of output nodes causes increases in the relative variability attributable changes over time (cycle) of output node by position (OPC) and output node by detector by position (ODPC). The small increase in OPC suggests that the position in the input array to which an output node is maximally responsive is more prone to change over time with 8 rather than 4 output nodes. The slightly larger increase seen in ODPC indicates that the changes in position of the centre of the receptive field are accompanied by increasing (and in some case changing) orientation selectivity. Therefore the increase in the OPC and OPDC sums of squares when the number of output nodes is increased reflects the (slowly) changing 'receptive fields' of the output nodes (see Fig. 2).

Fig. 6. Effect of changing trace duration and learning rate on the receptive field properties of the output nodes. Percentage of the total sums of squares of the network weights attributable to output by detector by position (SS_{ODP}) for the network with (A) 4 output nodes and (B) 8 output nodes. The plots show the value of SS_{ODP} for different learning rates (α) and trace duration (δ) values.

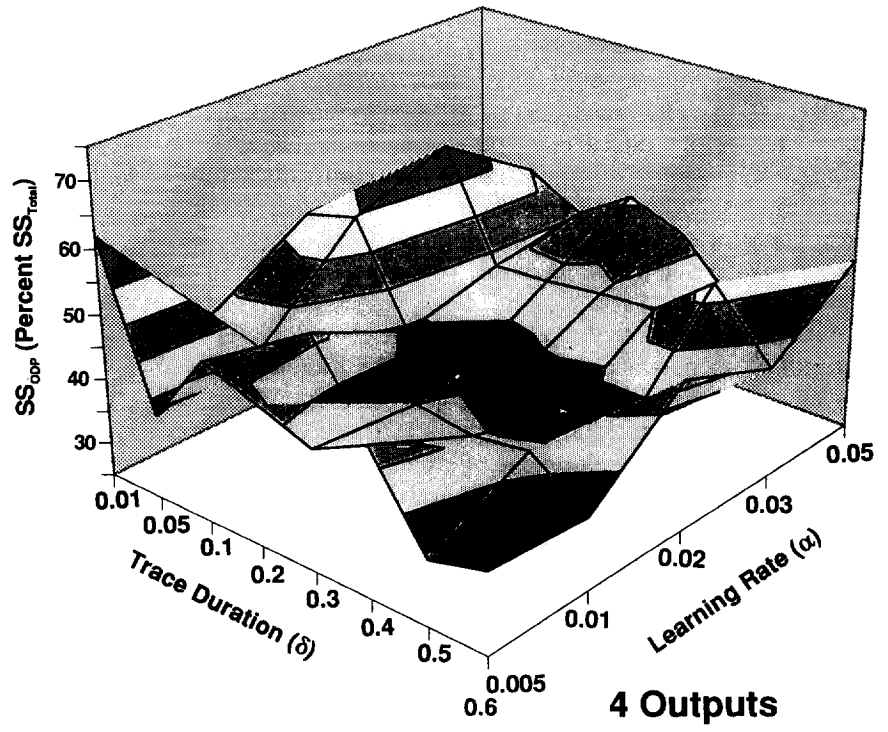
A
SS_{ODP} (Percent SS_{Total})
70
60
50
40
30
0.01
0.05
0.1
0.2
0.3
0.4
0.5
0.6
Trace Duration (δ)
0.005
0.01
0.02
0.03
0.05
Learning Rate (α)
4 Outputs

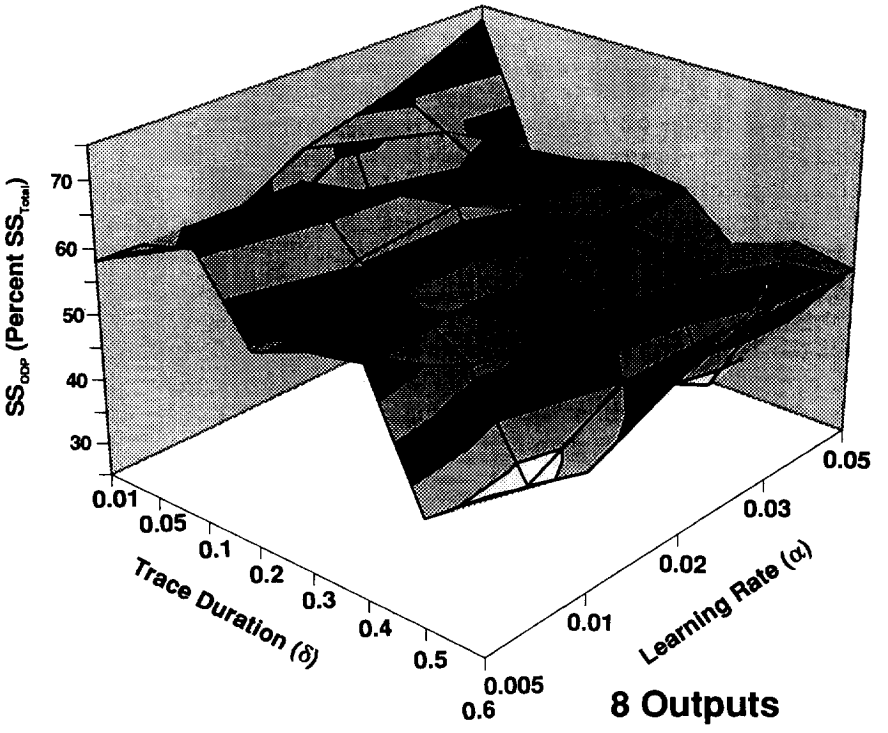
B
SS_{ODP} (Percent SS_{Total})
70
60
50
40
30
0.01
0.05
0.1
0.2
0.3
0.4
0.5
0.6
Trace Duration (δ)
0.005
0.01
0.02
0.03
0.05
Learning Rate (α)
8 Outputs

The degree of output node selectivity for particular input detector types is reflected by the relative variability attributable to the 'interaction' of these two possible sources (OD). The higher this value, then the greater the variability that is due to particular output nodes receiving strong inputs from particular detector types. With a larger number of output nodes the relative variability attributable to this source increases a very small amount (from 1.237% to 1.535%, t-test $t_{[78]} = 3.9$, $p = 0.0002$). Therefore, if anything, the degree of 'stimulus selectivity' for input orientation of the output nodes increases when increasing the number of output nodes.

It is also interesting to note that there was little change in the relative variability of the strength of weights across position onto the output nodes (OP) when the number of output nodes was increased ($t_{[78]} = 0.12$). The increase in the relative variability of selectivity by position (ODP) with increasing numbers of output nodes does therefore reflect the development of stimulus selective receptive fields, and not simply an increase in the proportion of variability due to position independent of detector type.

2.1. Effects of trace duration and learning rate

The above section considers the overall pattern of results, but does not consider the effects of changes to the rate of weight change (α) or trace duration (δ). Examples of the network with 4 output nodes after 3000 learning cycles are shown in Fig. 4. With a small value of δ (< 0.1), the duration of the trace is long, leading to associations across all positions and all detector types. When this occurs each output node receives the same input weights (Figs. 4a and d). Note that with faster learning rates (Fig. 4d) the 'selectivity' of the output nodes appears to be greater (the two oblique lines are longer than the horizontal and vertical). This reflects the most recent sweeps of the input array and the fact that the oblique lines are present in the input array for longer (15 vs. 8 positions during each cycle).

With shorter duration trace activity levels ($\delta \geq 0.1$) the output node of network receive strong inputs from only one of the four detector types (Figs. 4b,c,e). With faster learning rates (Fig. 4e) the network reflects the input which previously activated the output node and the following input sweep. For example, in the lower right of Fig. 4e, the output node is clearly activated by the inputs corresponding to vertical lines. After a vertical input sweep, an oblique starting in the top right of the input matrix occurred. With the fast learning rate, this lead to strengthening weights while the trace activity level lasted (5–6 cycles).

Examples of the network after 3000 learning cycles using various combinations of δ and learning rate with 8 output nodes are shown in Fig. 5. Selectivity of output nodes, when present, is combined with the generation of receptive fields. When δ is small (long duration trace activity) no selectivity in the output nodes is generated. δ values of 0.1 or larger generate selectivity and receptive field properties in the output nodes. The clearness of the receptive field increases with increasing values of δ (i.e. short duration trace activity) up to the largest value tried ($\delta = 0.6$).

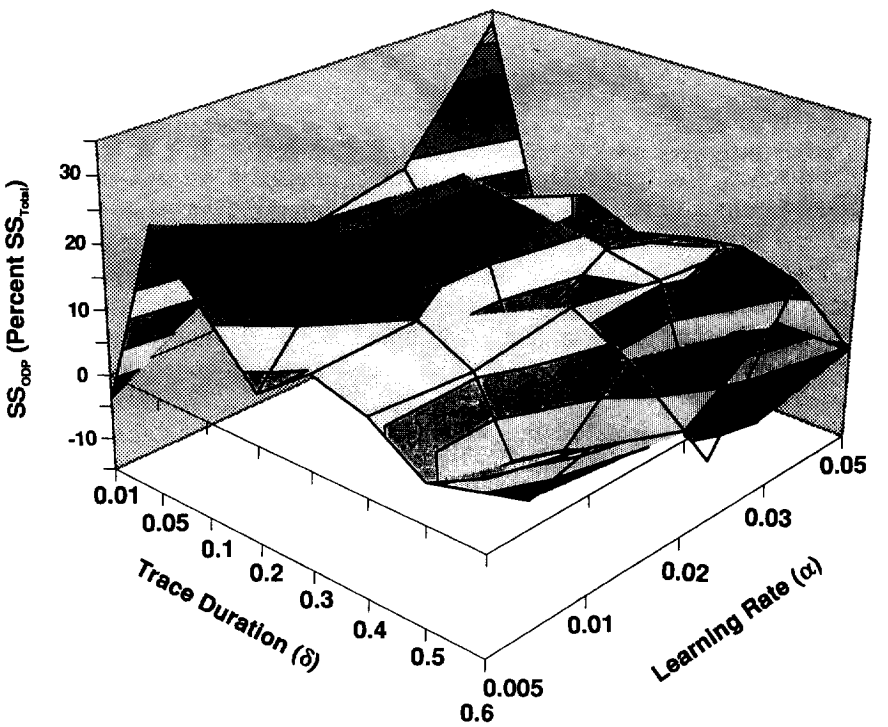

Fig. 7. The increase in the receptive field like properties when increasing the number of output nodes. The values of the percentage of the total sums of squares of the network weights attributable to output by detector by position (SSODP) for the 8 output node minus the value from the 4 output node network are shown.

To further visualise the effects of changing δ and α (the learning rate) values on the development of receptive fields, the proportion of the variability due to selectivity on position (ODP) was plotted for the network with 4 output nodes (Fig. 6a) and 8 output nodes (Fig. 6b). The two graphs show a trend for the proportion of the total variability of the weights to increase with both longer duration traces and with increasing learning rate.

The difference between Figs. 6a and b (8 outputs – 4 outputs) is shown in Fig. 7. It is clear that over most of the δ and α combinations tried a greater proportion of the total variability in the weights with 8 outputs can be attributed to 'receptive field' type properties than in the case of 4 outputs. The peak increase appears to be when the trace value was approximately 0.1. At this value, the output nodes (either 4 or 8) of the network developed selectivity for particular input types (orientations). With short duration changes ($\delta = 0.6$), only small changes in the proportion were observed. There also appears to be a small effect of learning rate, with faster learning rates reducing the development of receptive fields.

3. Discussion

3.1 Summary of results

Comparison of Figs. 1 and 2 indicates the main effects of increasing the number of available output units. These are a longer training period before stability is reached and the development of receptive fields. The competition between the output nodes is such that the only aspects that are available to compete over are the stimulus type (orientation) and the position within the input array. The correlation over time is such that, with the trace and only 4 output nodes, associations are made over position. With 8 output nodes, the competition between the output nodes forces the selectivity to compete over position, generating the limited 'receptive fields' of the output nodes. The development of receptive field properties using dipole inputs has been noted before [39]. The correlation of the orientation of the inputs means that the selectivity within the receptive field remains for one of the four possible detector types [see 14]. We show here that the network learns generalisation over space for particular stimulus types and, at the same time, learns to localise for the stimulus type by creating receptive fields covering a localised contiguous region of the input array.

The network simulated here would not show generation of selectivity for input detector type (orientation) if a rotating bar with a fixed point of rotation close to the input array were used. We stress that the input array is meant to mimic a small area of the input image (as would an 8 by 8 array of cells in V1). During normal visual experience, even a rotating bar would to all intents and purposes act as a bar sweeping across the input area, except in the small number of cases where the point of rotation was within the input area.

Analysis was performed that allowed quantitative measurements of the effects of changing the number of outputs in the network over a range of trace duration's (δ) and learning rate (α). Overall, when the number of output nodes was increased the network showed a clear tendency to learn spatial receptive field type properties while maintaining stimulus selectivity (increasing ODP value, Fig. 3). It was also shown that the representation across the input array of each detector type became more even when a larger number of outputs were used (decreasing DP value, Fig. 3). The receptive fields of the output nodes showed small changes in position over time.

Examination of Figs. 2 and 5 shows that the receptive fields in the 8 output node network did not develop evenly. The receptive fields developed for the oblique orientation selective output nodes more readily than for the vertical or horizontal selective nodes. This is due to the difference in the duration of the stimuli (vertical and horizontal passed through 8 positions whereas the oblique lines passed through 15 positions). The trace activity on which learning was based depended on time (number of positions) and consequently there was greater scope for competition between the output nodes representing the oblique inputs than there was for the vertical and horizontal inputs.

The development of receptive fields showing stimulus invariance is not necessarily the only change that could occur in the network when increasing the number of output units. For example it is possible that each of the 8 output nodes could have developed selectivity of different orientations within different areas of the total receptive field (i.e. not shown constancy for stimulus type over position). However this did not occur: when stimulus selectivity was established, this was always found to be constant across the whole receptive field. The simulations that were performed had very simple input patterns: there were only 4 orientations presented and they were presented sequentially. With 8 output nodes the degree of correlation remained significant between similar orientations. The 4 output node networks extracted this highly correlated feature of the input patterns. With 8 output nodes, the next significant correlation was extracted: the correlation between position. Importantly, this was extracted without disruption of the selectivity of the output nodes for particular input orientations.

With four output units, the sweeps lead to high correlation over time between the activation of successive input detectors. This correlation is higher between inputs of the same detector type since orientation of the lines was maintained during each sweep. The trace activity allows the output units to detect this correlation over time. The winner take all mechanism acts to force the output units to respond to different input orientations (see [39] for discussion). When the number of output units is increased the units are forced to sub-divide the set of stimuli. The sub-division shows contiguity over the input array. This is because the trace values are time dependent and more similar when the stimulus has moved only a small distance. This leads to neighbouring input units to be more likely to connect to the same output unit. Therefore the output unit of the network will learn restricted receptive fields showing stimulus invariance.

3.2. Effects of the trace duration (δ) and the learning rate (α)

Receptive fields were generated in the network with 8 output nodes only when the trace duration was relatively short ($\delta \geq 0.1$). The appearance of receptive fields was relatively independent of the learning rate (see Fig. 5). The relationship of the duration of the trace and the learning rate to the changes in the development of receptive fields (see Fig. 7) is predictable. With short trace duration, even with 4 output nodes, the network will associate across recently experienced input patterns only over a small area of the visual field. This was particularly evident with fast learning rates (e.g. Fig. 4e). Under these circumstances increasing the number of output nodes will lead to only small increases in the susceptibility for the formation of receptive fields. With longer duration trace activity the effect of the increased competition between the output nodes will have a greater impact. This is reflected in the downward slope from trace duration value 0.1 to 0.06 in difference of the ODP measure between the 8 and 4 output networks (Fig. 7). At values of $\delta < 0.05$ the trace duration is too long for the development of receptive fields with orientation selectivity. Thus at these values of δ there is a tendency for a decrease in the difference of the OPD measure.

There is a large range of values of both δ and α over which receptive field properties emerge when the number of outputs is increased ($0.1 \leq \delta \leq 0.6$, $0.005 \leq \alpha \leq 0.5$). The network properties described are therefore robust. This is particularly important if such a mechanism were to be adopted in a biological system since it is implausible that there would be no variation in one or other of these parameters. The ranges over which these two parameters can be varied with only a doubling of the number of output units shows that the phenomenon of developing receptive fields in a competitive learning network is robust and is therefore appropriate for use in a biological system.

3.3. Relationship to the visual system

The properties of the network described here depend on three basic assumptions. The first of these is competitive (inhibitory) interactions between output units leading to the separation of the output units selectivity. An elegant series of experiments in the primary visual cortex of the cat has shown a diffuse level of inhibition spreading over a large cortical area [8,9]. Lateral inhibitory connections have long been associated with orientation tuning in primary visual cortex (e.g. [41,42]). Although the role of these inhibitory connections in the sharpness of orientation tuning is debated (e.g. [49,50]), such connections could be used for the development of observed different neural response selectivities. During development, monocular deprivation causes a greater arborisation (spread) of geniculocortical connections than is seen in the non-deprived animals. The disruption caused by binocular deprivation is less (i.e. less spread) than seen with monocular deprived preparations. This suggests that competition between the inputs from the lateral geniculate nucleus that restricts the arborisation (see [12] for review).

Indirect evidence of lateral competitive interactions in higher visual areas is available. Some of the effects of neural modulation associated with visual attention have been linked with competition between neurons for representation of objects (see [7] for review]. Competition between units results in the formation of sparse representations. As noted in the introduction, there is evidence for sparse representations (e.g. [15,51]). It has also been argued that static upright bodies are coded in temporal cortex using just four differentially tuned cell populations [36], and that the coding of walking bodies uses just 8 cell populations [32,34]. The work of Miyashita and colleagues [27–30,39] showed that representations of novel fractal patterns in IT cortex are also represented explicitly by the responses of single units, indicating sparse representations of abstract objects. The same appears true for arbitrary 3-D wire-frame objects which assume an importance through training [25]. Interestingly the results described by Miyashita were found to show stronger associations over stimuli closely related in time rather than associations based on the visual appearance of the stimuli ([28], see also [44]). These results would be predicted by a trace based learning mechanism. Studies examining average receptive field sizes after long periods of denervation to somatosensory cortex [26] indicate no change in receptive field size [47,48], despite a reduction in the total skin area that was represented. This is suggestive that all neurons in the target

area are allocated an equal share of the input space, a property of a competitive system that depends on the number of cells in the target area.

The second assumption used in the network here is that of numerous output nodes being present. The number of neurons projecting from the monkey lateral geniculate nucleus to area V1 is approximately 1,400,000. Within area V1 there are approximately 100–200 times this number of cells (see [12] for review). With such a massive increase in cell numbers, competition between the neurons of the type described here would lead to observed overlapping small receptive field sizes. The second visual area occupies approximately the same cortical area as V1 [10]. It has been suggested that both these areas are responsible for 'edge extraction' [31,35]. The fourth visual area (V4) is approximately half the area of V1, and the final visual areas (central and anterior inferotemporal cortex and the anterior superior temporal polysensory cortex) are each approximately 1/4 the area of V1 [10]. With smaller number of cells, competitive learning of receptive field would lead to increasing receptive field size, but still maintaining coverage over the whole visual field.

As already noted there is an increase in the degree of stimulus complexity that accompanies the increase in receptive field size. Competitive learning can extract combinations of features [21,33] as well as learn positional invariance [14]. The studies here indicate that competition for visual space (receptive field) does not interfere with competition for stimulus selectivity when there is a restricted set (4) of input patterns. The lack of theoretical analysis of high order statistics of natural images means that it is not possible to estimate how many cell types would be needed to cover stimulus space in these progressively higher regions, nor is it known at what point competition for visual space would interfere or interact with competition for stimulus type.

The final assumption used in the network was that of trace activity. Recordings from temporal cortex have indicated that the response of many neurons continues for some 50–100 ms after the stimulus is removed (e.g. [37]). When performing memory tasks, many authors have noted a 'trace' activity in monkey temporal cortex neurons [e.g. 6,18] during intertrial intervals. The source of such trace activity is unclear although possibilities include intrinsic properties of the neurons themselves [16], local circuits within each area [e.g 1], or possibly inputs from the limbic system (the hippocampal complex in particular, see [35]). The trace need not be manifest in the neuron's activity: it is also possible that intrinsic biochemical cascades underlie such a mechanism (e.g. the G-protein cascades activated by increased internal calcium ion concentrations). It has also been suggested that inputs from the motion system could provide an 'association' signal [35]. Such motion inputs would aid segregation of different objects because of the common motion cues would belong to the one object [43].

One other feature that was noted was the changing receptive field position and size (e.g. Fig. 2). Dynamic changes in receptive fields have been reported in primary visual cortex (e.g. [38]). Under these circumstances input from a particular area of visual space was reversibly blocked for short periods. When permanent removal of input from a given location occurs (i.e. nerve severance) changes in the

mapping of the remaining inputs occurs in monkey somatosensory cortex (e.g. [26]). Under these circumstances there is little change in the receptive field size of the new mapped cells compared to the intact animal [47,48]. These results would both be predicted by a competitive system. Immediately following damage, cells near the cortical area corresponding to the site of the damage would receive no competition from within that area and 'encroach' into it. After a longer period the balance of competition would be restored, therefore leading to equal sized receptive fields.

3.4. Summary

The underlying assumptions of the present network have biological validity. As already noted, in the present network the property of tolerance for a large range of trace durations and learning rates is also biologically plausible. The main finding is that competition between output nodes leads to the generation of restricted receptive fields within which there is positional invariance for one of the input types. The simulations in this study concerned orientation and input positional invariance. The same principles can equally well be applied to other stimulus parameters and invariance properties. It is suggested that the primate system for form processing may constitute a series of such networks with competition between the output nodes. Such a network leads not only to sparse feature representations but also giving rise to the observed increase receptive field size with positionally invariant stimulus selectivity.

Acknowledgements

We thank the reviewers for their helpful comments. M.W. Oram is supported by a grant from the U.K. Medical Research Council (G931699IN). We thank D.I. Perrett for comments on the manuscript.

References

[1] D.J. Amit, Persistent delay activity in cortex, *Network – Computation in Neural Systems* 5 (1994) 429–436.

[2] H.B. Barlow, Single units and sensation: A neural doctrine for perceptual psychology, *Perception* 1 (1972) 371–394.

[3] H.B. Barlow, The twelfth Bartlett Memorial Lecture: The role of single neurons in the psychology of perception, *Quarterly J. Experimental Psychol.* 37A (1985) 121–145.

[4] H.B. Barlow, The neuron doctrine in perception, in: M. Gazanigga, ed., *The Cognitive Neurosciences* (Boston, MIT Press, 1994) 415–435.

[5] H. Barlow and P. Földiák, Adaptation and decorrelation in the cortex, in: R. Durbin, C. Miall and G. Mitchison, eds., *The Computing Neuron* (Addison-Wesley, Wokingham, England, 1989) 54–72.

[6] L. Chelazzi, E.K. Miller, J. Duncan and R. Desimone, A neural basis for visual search in inferior temporal cortex, *Nature* 363 (1993) 345–347.
[7] R. Desimone and J. Duncan, Neural mechanisms of selective visual attention, *Annual Rev. Neurosci.* 18 (1995) 193–322.
[8] R.J. Douglas and K.A.C. Martin, A functional microcircuit for cat visual cortex, *J. Physiol. (London)* 440 (1991) 735–769.
[9] R.J. Douglas, K.A.C. Martin and D. Whitteridge, An intracellular analysis of the visual responses of neurons in cat visual cortex, *J. Physiol. (London)* 440 (1991) 659–696.
[10] D.J. Felleman and D.C. Van Essen, Distributed hierachical processing in the primate cerebral cortex, *Cerebral Cortex* 1 (1991) 1–47.
[11] D.J. Field, What is the goal of sensory coding? *Neural Computat.* 6 (1994) 559–601.
[12] B.L. Finlay and S.L. Pallas, Control of cell number in the developing mammalian visual system, *Progress in Neurobiol.* 32 (1989) 207–234.
[13] P. Földiák, Forming sparse representation by local anti-Hebbian learning, *Biol. Cybernet.* 64 (1990) 165–170.
[14] P. Földiák, Learning invariance from transformation sequences, *Neural Computat.* 3 (1991) 194–200.
[15] P. Földiák and M.P. Young, Sparse coding in the primate cortex, in: M.A. Arbib, ed., *The Handbook of Brain Theory and Neural Networks* (Bradford Books/MIT Press, 1995) 895–898.
[16] P. Földiák, Learning constancies for object perception in: V. Walsh and J.J. Kulikowski, eds. *Visual Constancies: Why things look as they do* (Cambridge University Press; Cambridge, 1996).
[17] K. Fukushima, Neocognition: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position, *Biol. Cybernet.* 36 (1980) 193–202.
[18] J.M. Fuster, Inferotemporal units in selective visual attention and short term memory, *J. Neurophysiol.* 64 (1990) 681–697.
[19] C.G. Gross, Representation of visual stimuli in inferior temporal cortex, *Philosophical Trans. Royal Soc. London (Series B)* 335 (1992) 3–10.
[20] C.G. Gross, C.E. Rocha-Miranda and D.B. Bender, Visual properties of neurons in inferotemporal cortex of the monkey, *J. Neurophysiol.* 35 (1972) 96–111.
[21] S. Grossberg, Adaptive pattern classification and universal recording, I. Parallel development and coding of neural feature detectors, *Biol. Cybernet.* 23 (1976) 121–134.
[22] D.H. Hubel and T.N. Wiesel, Receptive fields and functional architecture of monkey striate cortex, *J. Physiol. (London)* 195 (1968) 215–243.
[23] S. Kaski and T. Kohonen, Winner-take-all networks for physiological models of competitive learning, *Neural Networks* 7 (1994) 773–984.
[24] T. Kohonen, *Self-organization and Associative Memory* (Springer, Berlin, 1984).
[25] N.K. Logothetis, J. Pauls and T. Poggio, Shape representation in the inferior temporal cortex of monkeys, *Current Biol.* 5 (1995) 552–563.
[26] M.M. Merzenich, R.J. Nelson, M.P. Stryker, M. Cynader, A. Schopmann and J.M. Zook, Somatosensory cortical map changes following digit amputation in adult monkeys, *J. Comparative Neurol.* 224 (1984) 591–605.
[27] Y. Miyashita, Associative representation of the visual long-term memory in the neurons of the primate temporal cortex, in: E. Iwai and M. Mishkin, eds., *Vision, Memory and the Temporal Lobe*, (Elsevier Science, New York, USA, 1990) 75–87.
[28] Y. Miyashita, Neuronal correlate of visual associative long-term memory in the primate temporal cortex, *Nature* 335 (1988) 817–820.
[29] Y. Miyashita and H-S. Chang, Neuronal correlate of pictorial short-term memory in the primate temporal cortex, *Nature* 331 (1988) 68–70.
[30] Y. Miyashita, A. Date and H. Okuno, Configurational encoding of complex visual forms by single neurons of monkey temporal cortex, *Neuropsychologia* 31 (1993) 1119–1131.
[31] M.W. Oram and D.I. Perrett, Modelling visual recognition from neurobiological constraints, *Neural Networks* 7 (1994) 945–972.
[32] M.W. Oram and D.I. Perrett, Integration of form and motion in the anterior superior temporal polysensory area (STPa) of the macaque monkey, *J. Neurophysiol.* (under revision).

[33] R.C. O'Reilly and M.H. Johnson, Object recognition and sensitive periods: A computational analysis of visual imprinting, *Neural Computat.* 6 (1994) 357–389.
[34] D.I. Perrett, M.H. Harries, R. Bevan, S. Thomas, P.J. Benson, A.J. Mistlin, A.J. Chitty, A.J., J.K. Hietanen and J.E. Ortega, Frameworks of analysis for the neural representation of animate objects and actions, *J. Experimental Biol.* 146 (1989) 87–14.
[35] D.I. Perrett and M.W. Oram, The neurophysiology of shape processing, *Image and Vision Computing* 11 (1993) 317–333.
[36] D.I. Perrett, M.W. Oram, M.H. Harries, R. Bevan, J.K. Hietanen, P.J. Benson and S. Thomas, Viewer-centred and object-centred coding of heads in the macaque temporal cortex, *Experimental Brain Res.* 86 (1991) 159–173.
[37] D.I. Perrett, E.T. Rolls and W. Caan, Visual neurons responsive to faces in the monkey temporal cortex, *Experimental Brain Res.* 47 (1982) 329–342.
[38] M.W. Pettet and C.D. Gilbert, Dynamic changes in receptive field size in cat primary visual cortex, *Proc. Nat. Acad. Sci. U.S.A.* 89 (1992) 8366–8370.
[39] D.E. Rummelhart and D. Zipser, Feature discovery by competitive learning, in: D.E. Rummelhart and J. McClelland, eds., *Parallel Distributed Processing, vol. 1.* (MIT Press, Cambridge MA 1986) 151–193.
[40] K. Sakai and Y. Miyashita, Neural organization for the long-term memory of paired associates, *Nature* 354 (1991) 152–155.
[41] A.M. Sillito, The contribution of inhibitory mechanisms to the receptive field properties of neurons in the striate cortex of the cat, *J. Physiol. (London)* 250 (1975) 305–329.
[42] A.M. Sillito, The effectiveness of biculline as an antagonist of GABA and visually invoked inhibition in the cat's striate cortex, *J. Physiol. (London)* (1975) 250, 305–329.
[43] G.R. Stoner and T.D. Albright, Image segmentation cues in motion processing: Implications for modularity in vision, *J. Cognitive Neurosci.* 5 (1993) 129–149.
[44] M.P. Stryker, Neurobiology – Temporal associations, *Nature* 354 (1991) 108–109.
[45] M.J. Tovee, E.T. Rolls, and P. Azzopardi, Translation invarience in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque, *J. Neurophysiol.* 72 (1994) 1049–1060.
[46] C. von der Malsburg, Self-organization of orientation sensitive cells in the striate cortex, *Kybernetik* 14 (1973) 85–100.
[47] J.T. Wall, M.F. Huerta and J.H. Kaas, Changes in the cortical map of the hand following postnatal median nerve injury in monkeys: Modification of somatotopic aggrgates, *J. Neurosci.* 12 (1992) 3445–3455.
[48] J.T. Wall, M.F. Huerta and J.H. Kaas, Changes in the cortical map of the hand following postnatal median nerve injury in monkeys: Organization and modification of nerve dominance aggregates, *J. Neurosci.* 12 (1992) 3456–3465.
[49] F. Worgotter and G. Holt, Spatiotemporal mechanisms in receptive fields of visual cortical simple cells: A model, *J. Neurophysiol.* 65 (1991) 494–510.
[50] F. Worgotter and C. Koch, A detailed model of the primary visual pathway in the cat: Comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity, *J. Neurosci.* 11 (1991) 1959–1979.
[51] M.P. Young and S. Yamane, Sparse population coding of faces in the inferotemporal cortex, *Science* 256 (1992) 1327–1331.

Mike Oram received a BSc. in zoology at Bristol University, then an MSc in Biological computation at York University. He was awarded his PhD in neurophysiology at the University of St. Andrews. His research interest is primarily the neurophysiology underlying primate visual information processing, particularly the use of neurophysiological data to constrain the neural network approach to understanding sensory mechanisms.

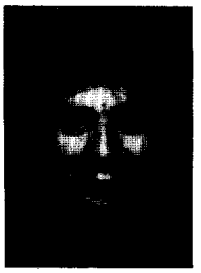

Peter Földiák received an MSc in electrical engineering at the Technical University of Budapest and a PhD at the University of Cambridge. He studied mechanisms of unsupervised learning and their role in perception and cortical function. He was research fellow at the MRC Research Centre in Brain and Behaviour, University of Oxford and is currently a lecturer at University of St Andrews.