INSTITUTE OF PHYSICS PUBLISHING NETWORK: COMPUTATION IN NEURAL SYSTEMS

Network: Comput. Neural Syst. **15** (2004) 217–242 PII: S0954-898X(04)78528-0

How different feature spaces may be represented in cortical maps

N V Swindale

Department of Ophthalmology and Visual Sciences, University of British Columbia, 2550 Willow St., Vancouver, BC, V5Z 3N9, Canada

E-mail: swindale@interchange.ubc.ca

Received 29 March 2004, accepted for publication 26 August 2004
Published 21 September 2004
Online at stacks.iop.org/Network/15/217
doi:10.1088/0954-898X/15/4/001

Abstract

This paper explores how different high-dimensional feature spaces might be represented in cortical maps, subject to continuity and completeness constraints. Spaces explored included products of circular variables (such as orientation), products of linear dimensions encoding scalar features (such as spatial frequency) and products of binary features. Maps were generated using the Kohonen algorithm, with uniform or non-uniform stimulus distributions. A 2D retina was always assumed to be present. Simulations were run with and without annealing. For uniform input distributions, coverage uniformity (Swindale 1991 *Biol. Cybern.* **65** 415–24) was used to measure how well the map was able to represent the feature space. For non-uniform distributions a weighted measure of coverage uniformity was calculated. Good coverage could be achieved for up to five or six cyclic variables but was substantially worse for a similar number of uniformly distributed scalar features. For annealed maps of multi-dimensional stimuli with Gaussian distributions, the distribution of receptive field centres and the distribution of total activity evoked on the cortex matched the stimulus distribution well. For annealed maps of non-uniformly distributed binary features there was an approximately linear relationship between the area of a map devoted to a specific feature and the probability of occurrence of the feature during development. Deviations from uniform retinotopy often led to improved coverage.

1. Introduction

The superimposed maps of visual space, orientation, ocular dominance (Hubel and Wiesel 1977) and spatial frequency (Hübener *et al* 1997) that are found in primary visual cortex have been characterized as dimension-reducing mappings, i.e. as projections from a subset of points

 © 2004 IOP Publishing Ltd Printed in the UK

in a high-dimensional feature space onto a 2D surface (Durbin and Mitchison 1990). Models that implement dimension reduction subject to completeness and continuity constraints, such as the elastic net algorithm (Durbin and Willshaw 1987) or the self-organizing feature map algorithm (Kohonen 1982, 1993, Obermayer *et al* 1990, 1992), have been remarkably successful in reproducing the properties of real visual cortex maps. For the case of visual space, orientation and ocular dominance, the relevant high-dimensional space includes two dimensions for retinal position, two more dimensions for orientation (represented as points on a circle) and a fifth dimension for ocular dominance (represented as two points at the ends of a line). Models that implement dimension-reducing projections from this five-dimensional space onto a 2D surface have reproduced many of the experimentally observed features of real maps, including the periodic stripy appearance of the ocular dominance map, a periodic orientation map including half-rotation singularities, a tendency for iso-orientation domains to run at right angles to the boundaries of the ocular dominance stripes and a weak tendency for orientation singularities to lie in the centres of the ocular dominance stripes (Bartfeld and Grinvald 1992, Obermayer and Blasdel 1993, Hübener *et al* 1997).

In addition to periodic orientation and ocular dominance maps, there is a periodic map of spatial frequency preference in primary visual cortex of the cat (Shoham *et al* 1997, Hübener *et al* 1997, Everson *et al* 1998, Issa *et al* 2000). Colour may additionally be mapped in primary visual cortex of the macaque (Landisman and Ts'o 2002). In the cat, it has been shown that the maps of spatial frequency, orientation and ocular dominance are related in a way that optimizes the uniform allocation, or coverage, of different combinations of the three variables across visual space (Swindale *et al* 2000). This optimization is predicted to occur by dimension reduction models.

Other cortical areas may, like primary visual cortex, contain superimposed maps of different stimulus features. In area V2 both colour and luminance, a three-dimensional space, may be mapped within the thin dark cytochrome stripes (Ts'o *et al* 2001, Xiao *et al* 2003); area V4 may contain superimposed maps of orientation, size and colour (Ghose and Ts'o 1997) while area MT contains maps of both stereoscopic disparity (DeAngelis and Newsome 1999) and direction of motion (Albright 1984, Diogo *et al* 2003). Representation of complex objects by combinations of columns for simpler features has been described in macaque inferotemporal cortex (Tsunoda *et al* 2001). Superimposed maps of different auditory stimulus features have been found in auditory cortex in cats (Schreiner 1995, 1998, Mendelson *et al* 1997) and evidence for superimposed maps of frequency and periodicity pitch has been found in auditory cortex in humans (Langner *et al* 1997). Maps of musical key relationships, which inhabit a toroidal stimulus space, have been proposed to exist in human prefrontal cortex (Janata *et al* 2002).

If it turns out that cortical areas generally contain mappings from high-dimensional spaces populated by various distributions of features, and are subject to the same types of continuity and completeness constraint that appear to hold in primary visual cortex, then it may be of interest to explore the impact of dimensionality and stimulus distribution from a wider variety of stimulus spaces than has previously been used in modelling studies. A previous paper (Swindale 2000) explored the limitations on the number of binary features that might be simultaneously mapped into a single cortical area and showed how coverage changed as increasing numbers of features were permuted across a model retina. The results suggested that up to seven or eight binary feature dimensions might be represented with reasonably uniform coverage within a single cortical area. In this study a binary feature space was chosen for analysis because of its simplicity: however other types of feature space are often represented in cortical maps. Cortical maps represent cyclic (as in the case of orientation) or scalar (as in the case of spatial frequency) variables but it is not known how coverage

uniformity constraints might limit the number of these variables that might simultaneously be present in a single map.

Another issue not yet examined is the likelihood that cortical maps represent spaces in which stimuli have a non-uniform distribution. For example not all possible permutations of binary sub-features might occur in a feature space, and scalar stimulus parameters are more likely to have a Gaussian than a flat distribution. This paper therefore extends the results in Swindale (2000) to non-binary feature spaces and to non-uniform stimulus distributions, with the same goal of exploring the limits placed by coverage uniformity on the numbers of stimulus dimensions that might be simultaneously mapped. Spaces explored were (a) combinations of many cyclic variables (i.e. products of circles), (b) Gaussian and uniform distributions of scalar stimuli along a line and (c) non-uniform distributions of binary features. A two-dimensional retina was always assumed to be present. Questions asked include: (1) what is the impact of dimensionality and receptive field size for coverage of scalar and cyclic feature variables?, (2) how does the representation of cyclic and non-cyclic variables compare? and (3) how well does the area of a map devoted to a stimulus correlate with the probability of its occurrence during development?

2. Methods

2.1. Simulation details

Maps were generated using Kohonen's self-organizing feature map algorithm (Kohonen 1982) applied in the following way (Obermayer *et al* 1992): a Euclidian stimulus space is defined, consisting of two retinal dimensions (corresponding to the horizontal and vertical components of receptive field position in visual space) plus N additional feature dimensions, referred to individually as sub-feature dimensions. The model cortex is represented as a two-dimensional grid of points, indexed by i and j, lying within the stimulus space. The position of each cortical point in the space defines the set of stimulus values assigned to it i.e. the centre of its receptive field along each of the dimensions. This position is given by $\mathbf{w}_{i,j} = (x, y, a_1, a_2, \dots)$ where x and y are the position in the retina of the centre of the receptive field of the point (i, j), and a_1, a_2 etc are likewise the centres of the receptive fields along each of the additional feature dimensions. 'Stimuli' (or equivalently, features) defined by the vector $\mathbf{v} = (x_s, y_s, b_1, b_2, \dots)$, are chosen one at a time, according to defined probabilities, from the space. For each stimulus, the nearest, i.e. most responsive, cortical point is found, and this point, and surrounding ones defined by a circular neighbourhood function, are then moved closer towards the stimulus. After each presentation of a stimulus vector $\mathbf{v}$, cortical receptive fields change by an amount

$$\Delta\mathbf{w}_{i,j} = \varepsilon h(r)(\mathbf{v} - \mathbf{w}_{i,j}) \tag{1}$$

where ε is a rate constant, r is the distance in the cortex between the point nearest the stimulus and point (i, j), and $h(r)$ is a neighbourhood function given by

$$h(r) = \exp(-r^2/2\kappa^2) \tag{2}$$

where κ is a width parameter. Physiologically this might correspond to Hebbian strengthening of active inputs that were connected to, or near, a cell or group of cells in the cortex that was responding to a stimulus, making it more responsive to that particular stimulus in future. Two types of simulation were carried out: one in which the width of the neighbourhood function, κ, remained constant, and another in which κ was slowly reduced to a lower value as map development progressed, a process referred to as annealing. Details of how annealing was done are given below.

As in Swindale (2000) the cortex was represented by an $M \times M$ grid of points with values $x_{i,j}$, $y_{i,j}$, $a_{1i,j}$ etc where $i, j = 0, 1, \ldots, M-1$. A roughly ordered retinal topography was assumed to be present initially, with $x_{i,j} = iX/(M-1) + \xi_x$ and $y_{i,j} = jY/(M-1) + \xi_y$ where ξ_x and ξ_y were random numbers drawn from a Gaussian distribution with a mean of zero and a standard deviation $= \sigma_x$. X and Y give the overall (fixed) extent of the region of retina mapped onto the cortex. In order to remove edge effects (namely a tendency for the cortex to shrink back from the edges of the retina causing loss of coverage in those regions) periodic boundary conditions were applied to the calculation of distances in the retina and cortex. Initial values of the non-retinal cortical receptive field centres, $a_{1,i,j}$, $a_{2,i,j}$ etc were also Gaussian random numbers with a mean of zero and standard deviation equal to σ_a.

Stimuli were calculated as follows: retinal position values x_s and y_s were chosen with uniform probability in the intervals $[0, X]$ and $[0, Y]$, respectively. For maps of multiple angular variables, the values of b_1, b_2 etc, and corresponding cortical values a_1, a_2 were grouped into N pairs. Each pair was used to define an orientation given by $\theta_k = 0.5 \tan^{-1}(a_{2k}/a_{2k-1})$, $k = 1, \ldots, N$. Angular (or orientation) stimuli always lay on the unit circle i.e. $\left(b_{2k}^2 + b_{2k-1}^2\right)^{1/2} = 1 \,\forall\, k$. Cortical values a were not constrained in any way, except by the learning rule (1). For uniform scalar feature distributions, values of b_n were randomly and independently distributed with uniform probability between limits $-d/2$ and $+d/2$ along each of the N sub-feature dimensions. For Gaussian distributions, stimulus values were drawn from a multivariate distribution with mean equal to 0 and standard deviation equal to σ_d along each of the N sub-feature axes. Because the resulting distribution is symmetric around the origin, it was often convenient to describe it in terms of the radial distance from the origin, $r = \left({b_1}^2 + {b_2}^2, \ldots, {b_N}^2\right)^{1/2}$. This distribution has the form $p(r) \propto r^{N-1}\, \mathrm{e}\left(-r^2 \big/ 2\sigma_d^2\right)$.

Many choices of non-uniform probability distributions for binary stimuli are possible. For the results reported here, a non-uniform distribution across the 2^N possible values of n was generated by calculating a table of probabilities p_n, uniformly distributed over a fixed interval and normalized so that $\sum_{n=1}^{n=2^N} p_n = 1$. During development of a given map, stimuli were chosen at random according to these probabilities, which remained fixed. Different simulations used different probability tables.

2.2. Terminology

As in Swindale (2000) the term 'protomap' will be used to refer to the mapping of one of the N possible dimensions, i.e. the values of a_1, a_2 etc, in the case of scalar maps, or to the mapping of one of the N angles given by pairs of values in the case of maps of multiple cyclic variables. The variable mapped in a protomap will be referred to as a sub-feature. The complete set of N such maps, will be termed a 'polymap'. In cases where there is no ambiguity, the term 'map' may refer to either an individual map (i.e. a protomap) or the complete set of maps (a polymap). Table 1 lists mathematical symbols used and their definitions.

2.3. Parameter values

In all the simulations, the array size, $M = 150$, the retina was $X = Y = 12.0$ units in size, the learning rate $\varepsilon = 0.01$, initial scatter of cortical feature values was $\sigma_x = 0.1$ and $\sigma_a = 0.1$. Orientation stimuli lay on a circle of radius $= 1$, uniformly distributed scalar features lay on a line of length $d = 2.0$, scalar features with a Gaussian distribution had $\sigma_d = 0.75$. In simulations where no annealing was done, the size of the cortical neighbourhood function, κ, was kept constant at 4.0 array units and the total number of stimulus presentations was 10^6, or about 44 stimuli per array point. Where annealing was done the following two

Table 1. Description of mathematical symbols used in the text.

Symbol	Meaning	Value
$A(\mathbf{v})$	Total cortical activity evoked by stimulus $\mathbf{v}$	
a_k	Cortical feature value for the kth map	
b_k	Stimulus feature value for the kth map	± 1
ε	Learning rate	0.01
d	Width of scalar feature dimension	2
$h(\boldsymbol{r})$	Rate of learning as a function of cortical distance $\boldsymbol{r}$	
i, j	Indices to points on the 2D cortical grid	
κ	Width of cortical neighbourhood function in pixels	$\leqslant 4.0$
L	Kullback–Leibler divergence	
N	Number of stimulus dimensions, excluding the retina	2–10
M	Size of model cortex in pixels	150
σ_a	Initial standard deviation of feature values	0.1
σ_d	Width of distribution of multivariate Gaussian stimuli	0.75
σ_r	Receptive field width for retinal position	1.12
σ_s	Receptive field width for scalar stimuli	0.25 or 0.35
σ_θ	Receptive field width for angular variable, in degrees	25°
σ_x	Initial standard deviation of values in retinal map	0.1
$\mathbf{v}$	Stimulus feature vector	
$\mathbf{w}_{i,j}$	Feature vector, i.e. receptive field, mapped to cortical point (i, j)	
$x_{i,j}, y_{i,j}$	Position in retinal space of point (i, j) in degrees	
X, Y	Size of model retina in degrees	12

schedules were used (i) for angular feature maps and for uniform scalar feature maps, $\kappa = 4.0$ for the first 2×10^5 stimulus presentations, and was then reduced in size by multiplying by 0.998 every 10^3 stimuli. When it reached a lower limit of 0.5 (at a total of 12.4×10^5 stimulus presentations) κ was kept constant and the simulation terminated at a total of 2×10^6 stimulus presentations, (ii) for maps with a multivariate Gaussian stimulus distribution, and for maps with an inhomogeneous distribution of binary features, a slower annealing schedule was used: $\kappa = 4.0$ for the first 2×10^5 stimulus presentations, and was then reduced in size by multiplying by 0.999 every 10^3 stimuli. When it reached a lower limit of 0.1 (at a total of 3.9×10^6 stimulus presentations) κ was kept constant and the simulation terminated after 10^7 stimulus presentations.

The algorithm was speeded up (by about an order of magnitude) by (a) restricting updates in cortical values to a region within which the value of $h(r) > 10^{-7}$, and (b) restricting the search for the winning cortical point to a region surrounding the point predicted to be closest to the stimulus on the basis of a uniform retinotopic map. This point is $i = x_s M/X$, $j = y_s M/Y$. The limits of the region searched were $(i \pm \Delta i, j \pm \Delta j)$ (modulo M) where every 10^3 stimulus presentations Δi and Δj were set equal to 1.5 times the largest distance between (i, j) and the actual winning point taken over the preceding 10^3 stimulus presentations. Initially, Δi and $\Delta j = M/2$, i.e. the search covered the entire cortex.

As in Swindale (2000) simulation parameters were chosen so that the protomap wavelength was $\approx$28 pixel units. Assuming that this corresponds to a distance of 1 mm in the real cortex, the cortical array units in the simulations have a spacing of about 36 μm, which corresponds roughly to the size of a cortical 'minicolumn' (Mountcastle 1978). Thus the granularity of the simulations corresponds approximately with the granularity of the real cortex.

The number of feature variables was usually varied between one and eight. For each value of N, three separate polymaps were generated using identical parameters but a different initial random number seed. Procedures *ran1* and *Gauss* (Press *et al* 1994) were used to generate random numbers.

2.4. Analyses

Periodicity was measured from the power spectrum (Swindale 2000). For maps with uniform stimulus probability distributions, coverage uniformity was calculated in the following way (Swindale 2000): for a given stimulus, $\mathbf{v}$, the total activity evoked in the cortex, $A(\mathbf{v})$ was calculated as

$$A(\mathbf{v}) = \sum_{i,j \in C} f(\mathbf{v} - \mathbf{w}_{i,j}) \tag{3}$$

where f defines the receptive field of point (i, j) and the summation is taken over all the points, C, in the cortex. Coverage uniformity, c', is then defined as

$$c' = \langle |A - \langle A \rangle|^2 \rangle^{1/2} / \langle A \rangle \tag{4}$$

i.e. the ratio between the standard deviation of the values of A, divided by the mean of A, measured over a set of stimuli representative of those used to generate the map. Calculations of c' were based on measures of response to 10^4 stimuli drawn at random from the same distribution used to generate the map. Although this is a coarse sampling of stimulus space, values of A have an approximately Gaussian distribution and accordingly estimates of c' were found to change very little (<2%) with additional numbers of stimuli.

Receptive fields were always defined as the product of a Gaussian retinal receptive field, with the receptive fields (assumed to all have the same shape) for each of the N feature variables i.e.

$$f(x, y, a_1, a_2, \ldots, a_N) = \mathrm{e}^{-(x^2+y^2)/2\sigma_r^2} \prod_{n=1}^{N} g(a_N) \tag{5}$$

where σ_r gives the retinal receptive field size. For continuously distributed scalar feature variables, $g(a)$ was assumed to be Gaussian with a receptive field width given by σ_s. Likewise for maps of angular features, $g(a)$ was a Gaussian function of angle (ignoring the modulus) with a width σ_θ. For binary features, where $b \in \{-1, 1\}$, $g(a) = 1$ for $-0.5 < a < 0.5$ and was equal to 0 otherwise. Following Swindale (2000) and Swindale *et al* (2000) receptive field sizes were chosen to be comparable to those reported for retinal receptive fields, spatial frequency and orientation tuning in visual cortex. This resulted in a choice of retinal receptive field size, $\sigma_r = 1.12$ retinal units, although the impact of other values was explored. Orientation tuning was generally assumed to have a width $\sigma_s = 25^\circ$. Tuning for scalar feature dimensions was assumed to have a width $\sigma_s = 0.35$ for uniformly distributed features and a width $\sigma_s = 0.25$ for features with a Gaussian distribution.

Maps generated with non-uniform stimulus distributions generally have activity values, $A(\mathbf{v})$ that are systematically related to the probability of occurrence of the stimulus $p(\mathbf{v})$. If every stimulus has a different probability of occurrence then the concept of uniform coverage loses its meaning because the cortex would not be expected to represent stimuli that rarely or never occurred as well as frequently occurring stimuli. However, if $p(\mathbf{v})$ is uniform across some stimulus dimensions (e.g. retinal position) then a mean and a standard deviation of A can be calculated across stimulus sets, S^k, for which $p(S^k \subset \mathbf{v})$ does not vary. An average measure of coverage uniformity, weighted by the mean response magnitude, $\langle A(S^k) \rangle$, for each stimulus

set, can thus be defined as

$$c' = \frac{1}{\sum_{k=1}^{K}\langle A(S^k)\rangle}\sum_{k=1}^{K}\frac{\langle A(S^k)\rangle\langle|A(S^k)-\langle A(S^k)\rangle^2|\rangle^{1/2}}{\langle A(S^k)\rangle}$$
$$= \frac{\sum_{k=1}^{K}\langle|A(S^k)-\langle A(S^k)\rangle|^2\rangle^{1/2}}{\sum_{k=1}^{K}\langle A(S^k)\rangle} \tag{6}$$

where the summation is taken over K stimulus sets, indexed by k, within each of which it is expected that the cortical representation is uniform. Note that if stimulus probabilities do not vary over the total set, K can be taken equal to 1 and the definition of c' is the same as that given above (4) for uniform stimuli. A weighted measure was used because with an unweighted measure c' was often dominated by highly variable responses to stimuli that evoked very low mean responses in the map. It seemed desirable that the measure should instead reflect coverage of stimuli that were well, rather than poorly, represented in the map.

This measure of map organization is map-centred because it does not take into account the values of $p(S^k)$, i.e. the frequency of occurrence of the stimuli that were presented during map development. A different criterion of the fidelity of the representation of a stimulus space by a map is to compare the fractional area of the map devoted to the representation of an equi-probable class of stimuli, with the probability of occurrence of the class during map development. If we equate the fractional area of cortex devoted to each stimulus set, $p_c(S^k)$, with the probability of a response, then a measure of similarity between the distributions is given by the Kullback–Leibler divergence (Bishop 1995) approximated by

$$L = -\sum_{k=1}^{K} p(S^k)\ln\frac{p_c(S^k)}{p(S^k)} \tag{7}$$

provided the K stimulus sets encompass the entire distribution as closely as possible. The smaller the magnitude of L the closer are the two distributions in shape.

3. Results

3.1. Combinations of angular variables

Figure 1(a) shows how c' varies with the number of angular feature parameters, for non-annealed and annealed maps. As described in the methods section, each angular stimulus feature used to generate the maps was assumed to be uniformly distributed around a circle of unit radius. An angular tuning width of $\sigma_\theta = 25^\circ$ and a retinal receptive field width $\sigma_r = 1.12$ retinal units were assumed for the coverage calculations. As for binary feature maps (Swindale 2000) c' was found to increase approximately linearly on a log scale with the number of angular variables. The slopes of the relationships were 0.51 and 0.52 for the non-annealed and annealed cases, respectively. These slopes are similar to those measured for binary feature maps (Swindale 2000), where slopes of 0.47 and 0.42 were measured for the non-annealed and annealed cases, respectively. Annealing produces an approximately three-fold decrease in c'. For $N = 6$ angular variables $c' = 0.38$. This is comparable to the values measured for feature maps of eight binary variables and is also comparable to coverage values measured in real cortical maps of ocular dominance, orientation and spatial frequency ($c' \approx 0.43$–0.49; Swindale *et al* 2000).

The relationship between receptive field size and c' was studied in the non-annealed maps (figures 1(b) and (c)). As expected, decreases in receptive field size increased the values of c'. The impact of changes in angular tuning width was greater in maps with larger numbers

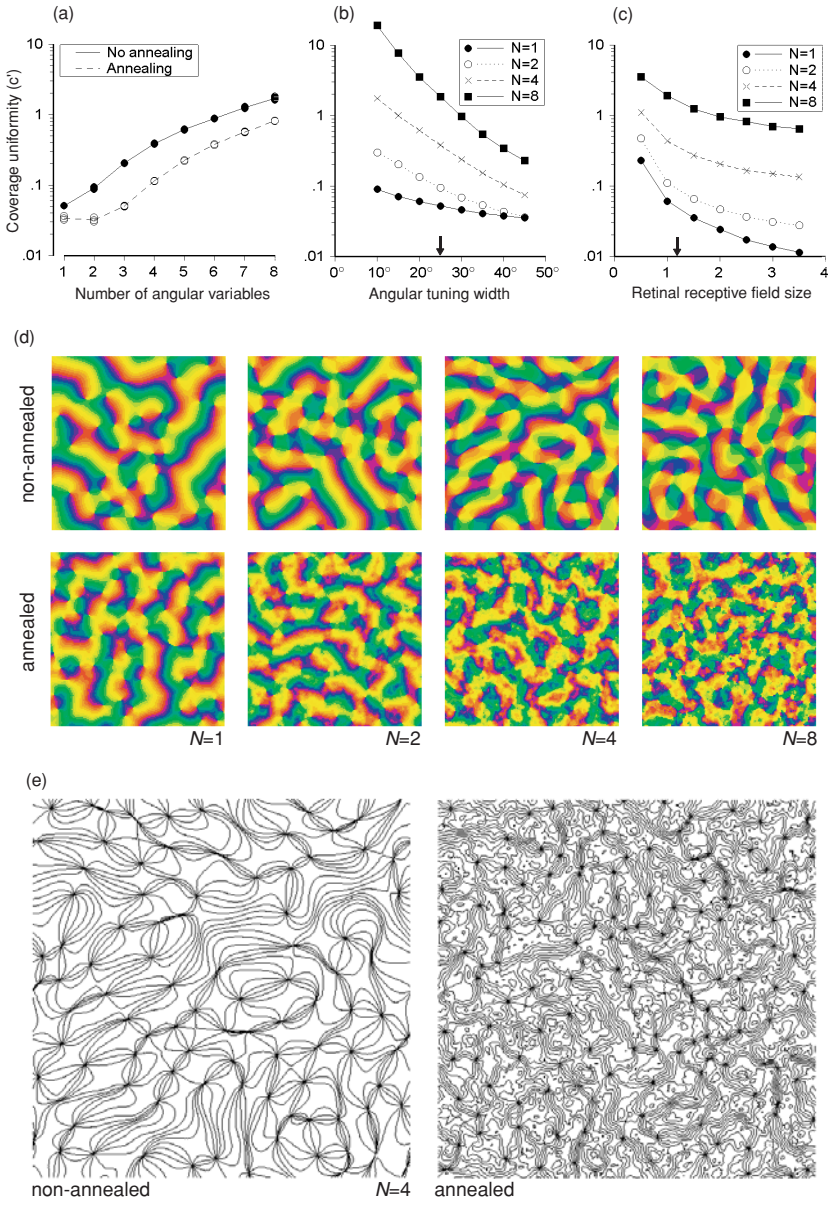

Figure 1. Results obtained from mapping different numbers (N) of angular variables. (a) The effect of N on coverage uniformity (c') for maps generated with or without annealing; lines connect the average of measures from three independently generated maps in each case; calculations were done assuming a retinal receptive field size, $\sigma_r = 1.12$ units and an angular tuning width, $\sigma_\theta = 25^\circ$, (b) the effect on c' of changing the angular tuning width, (c) the effect on c' of changing retinal receptive field size. Each point in (b) and (c) is the average of three maps; arrows indicate the parameter values used for (a). (d) protomaps obtained for different values of N. Note the increasing irregularity of map organization as N increases in the cases where annealing was used, (e) iso-angular contours,15° apart, for the non-annealed and annealed protomaps for $N = 4$.

(This figure is in colour only in the electronic version)

Table 2. Averaged protomap wavelength (array units), singularity density (number of singularities per orientation wavelength squared) and coverage uniformity, for the non-annealed and annealed maps of angular variables. Each value is the average derived from 3 map simulations.

	Wavelength		Singularity density		c'	
N	Non-annealed	Annealed	Non-annealed	Annealed	Non-annealed	Annealed
1	33.1	29.1	2.64	2.85	0.051	0.034
2	29.6	26.5	2.93	2.63	0.091	0.032
3	27.6	25.8	2.68	3.12	0.205	0.050
4	27.1	24.9	2.80	3.71	0.387	0.114
5	27.1	24.5	2.91	4.67	0.615	0.225
6	27.0	24.6	3.14	5.74	0.874	0.378
7	27.4	24.0	3.30	6.46	1.259	0.568
8	27.4	23.3	3.35	7.02	1.712	0.813

of variables (figure 1(b)). In contrast to this, changes in retinal receptive field size had effects that were roughly similar for maps of different numbers of variables (figure 1(c)).

Figure 1(d) shows protomaps for values of $N = 1, 2, 4$ and 8, using the conventional colour coding scheme for orientation. There seemed to be less long-range order present in maps with higher values of N, and iso-orientation domains seemed longer and of more uniform width for low values of N. Table 2 gives averaged singularity density and protomap wavelength for the non-annealed and annealed maps. For the non-annealed maps singularity density increased slightly with increasing N. Pattern wavelength increased slightly for values of $N < 3$. No attempt was made to correct for the impact of these small changes in wavelength on the appropriate choice of a retinal receptive field size, σ_r, which (as discussed above) is based on a protomap wavelength of 28 pixels. Although coverage values will be slightly affected the direct impact of N on coverage far exceeds that of the indirect effect due to wavelength changes.

Annealing produced a significant change in the structure of the maps, particularly for values of $N \geqslant 2$. Singularity density was increased (table 2) and the maps look more disorganized, but examination of the domain structure, using contours to represent the borders of iso-orientation domains (figure 1(e)) shows that the borders are more closely spaced in the annealed maps. In addition there are many local folds in the map as evidenced by the appearance of circular contours which are absent in the non-annealed maps. Surprisingly, the extensive local folding has little effect on the wavelengths of the maps. For example in figure 1(e) the wavelengths of the non-annealed and annealed maps are 26.7 and 26.5 pixels, respectively. In spite of this, the spacing of iso-orientation contours (and thus the local orientation gradient) is several times greater in the annealed maps.

3.2. Combinations of uniformly distributed scalar variables

These maps represent scalar parameters, each of which has a uniform distribution within fixed limits. For the sake of brevity these will be referred to as 'scalar' maps in the following text. The limits for each of the N variables were -1 and $+1$, i.e. stimuli were uniformly distributed within a hypercube with a side $d = 2$ units in length. Coverage calculations require the definition of a receptive field width for the scalar features. The choice made here was based on the relation between the range of spatial frequencies represented in area 17 of cat visual cortex and the width of spatial frequency tuning curves of individual neurons. The range on a log scale is about 3.3 octaves while individual tuning curves (full width at half height)

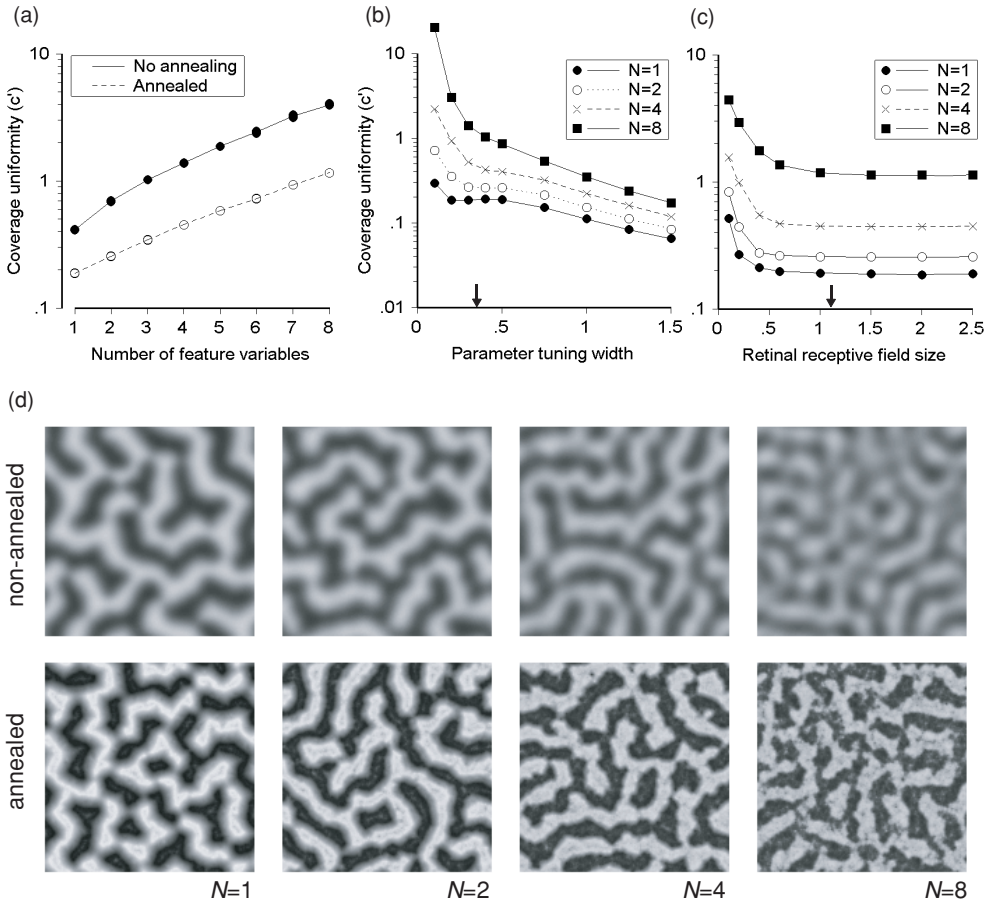

Figure 2. Results obtained from mapping different numbers (N) of scalar variables, uniformly distributed between fixed limits. (a) The effect of N on coverage uniformity (c') for maps generated with or without annealing; lines connect the average of measures from three independently generated maps; calculations were done assuming $\sigma_r = 1.12$ units and $\sigma_s = 0.35$, (b) the effect on c' of changing σ_s; (c) the effect on c' of changing retinal receptive field size. Each point in (b) and (c) is the average from three maps, the variation in values across the maps is small (<2%); arrows indicate the parameter values used for (a). (d) Protomaps for different values of N, note the increasing irregularity of map organization as N increases in the annealed cases.

are about 1.3 octaves wide (Movshon *et al* 1978). This is a ratio of about 2.5. A value of $\sigma_s = 0.35$ was chosen for tuning curve widths, since this corresponds to a full width at half height of about 0.82 and the ratio between the range of parameter values ($d = 2.0$) is similarly $2.0/0.82 \approx 2.4$. In spite of the relatively coarse coding which this implies, coverage uniformity (figure 2(a) and table 3) is substantially worse than for angular variables. Figure 2(b) shows the effect of varying σ_s on c', for a fixed retinal receptive field size. An unexpected observation is the presence of a small local minimum in c' at a value of $\sigma_s \approx 0.3$ for $N = 1$. The reason for this effect is not clear. Elsewhere, c' decreases as σ_s increases, the rate of increase being particularly large for large values of N and values of $\sigma_s < 0.5$. Figure 2(c) shows the effect of changing the retinal receptive field size, σ_r, for a fixed value of $\sigma_s = 0.35$. It was found that c' remained largely unchanged for values of $\sigma_s > 0.5$.

Table 3. Averaged protomap wavelength (array units), rms map amplitude (expressed as a percentage relative to the limits of the stimulus distribution) and coverage uniformity, for the non-annealed and annealed maps of scalar variables. Each value is the average derived from 3 map simulations.

	Wavelength		rms map amplitude		c'	
N	Non-annealed	Annealed	Non-annealed	Annealed	Non-annealed	Annealed
1	30.7	28.3	36.6	54.1	0.414	0.188
2	29.1	26.0	33.0	52.9	0.694	0.256
3	27.0	25.0	29.6	50.8	1.022	0.344
4	26.2	23.9	26.9	48.9	1.379	0.451
5	25.0	23.0	24.4	47.0	1.875	0.584
6	24.3	22.5	22.6	44.9	2.422	0.727
7	24.1	22.4	21.1	42.9	3.231	0.932
8	23.9	22.1	19.9	41.1	3.975	1.161

Figure 2(d) shows different scalar protomaps. They are similar in appearance to those found in real visual cortex (e.g. ocular dominance and spatial frequency domains) and take the form of periodic blobs and branching stripes. The value of a particular cortical feature variable thus fluctuates periodically with the distance across the surface of the cortex. The rms amplitude of the fluctuations decreases as N increases (table 3). In the annealed maps the amplitude is larger and does not decrease as much. As with the angular maps the wavelength of the fluctuations decreases gradually as N increases (table 3). Annealing slightly reduces the wavelength for a given value of N. The spatial profile of the fluctuations is approximately sinusoidal for the non-annealed maps, and triangular (i.e. regions of alternating constant slope) for the annealed maps. As in the angular maps, annealing produces local folding (i.e. frequent reversals in the gradient) without disrupting the overall periodicity of the map. These reversals are visible in the case for $N = 8$ (figure 2(d)) as irregularities and apparent noise in map structure.

Constructions for a 1D cortex and simplified cyclic and non-cyclic stimulus spaces (figure 3) suggest why coverage uniformity is relatively poor for the non-cyclic space. For the cyclic case (figure 3(a)) all stimuli can be equally well represented by the map, and the space can be covered uniformly by receptive fields of any width in the orientation dimension, as long as the retinal receptive field width is equal to the repeat period of the orientation map, expressed in retinal coordinates. These behaviours are seen approximately in the 2D simulated maps, where, for $N = 1$, reducing orientation tuning width has little impact on coverage (figure 1(a)) while reducing retinal receptive field size has a relatively large effect (figure 1(b)). For non-cyclic variables (figure 3(b)) maps tend to be sinusoidal (non-annealed) or square-wave (annealed) in form and consequently many regions at the edges of the distribution are consistently poorly represented in the map. No particular receptive field shape gives good coverage. However, for fields that are sufficiently broadly tuned along the y-axis (right-hand side of the figure) retinal field sizes can be made narrow, with little impact on coverage. This behaviour is also represented in the simulated 2D maps: coverage improves with increasing y-axis tuning width up to large values, but is relatively unaffected by retinal receptive field size over a wide range.

3.3. Gaussian stimulus distributions

It may be argued that the preceding simulations are unrealistic, since, in the real world, scalar stimulus features are unlikely to be uniformly distributed between fixed limits. To address this

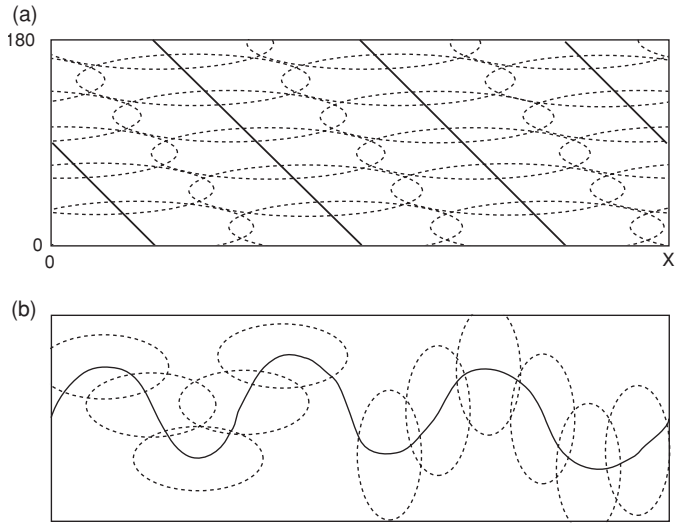

Figure 3. How receptive field size affects coverage for an idealized 1D cortex projected into a 2D stimulus space. A single retinal dimension is represented on the x-axis and a single non-retinal parameter variable is represented on the y-axis in either (a) a cyclic (orientation) stimulus space or (b) a scalar space. The continuous line shows the trajectory of the cortex through the two types of space. These are not actual maps but are representative idealizations of the behaviour of the Kohonen algorithm in the two cases. For the cyclic case (a) the space can be covered well by receptive fields of any width in the orientation dimension, as long as the retinal receptive field width is equal to the repeat period of the orientation map, expressed in retinal coordinates. For non-cyclic variables (b) maps tend to be sinusoidal and stimuli at the edges of the distribution are consistently poorly represented. No particular simple receptive field shape gives good coverage. However, for fields that are broadly tuned along the y-axis (rhs of the figure) retinal field sizes can be made narrow with little impact on coverage. This behaviour is found in the simulated 2D maps: coverage improves with increasing y-axis tuning width up to large values, but is relatively unaffected by retinal receptive field size over a wide range.

issue maps were generated with stimuli that had a Gaussian distribution, of width $\sigma_d = 0.75$, along each of the N feature dimensions. For the sake of brevity these will be referred to as 'Gaussian' maps in the following text. As before, retinal positions were uniformly distributed across a region 12×12 units in size. Figure 4(a) shows typical protomaps for different values of N in non-annealed and annealed maps. Visually, the non-annealed maps resemble those obtained for uniformly distributed stimuli; the annealed maps have an increasingly noisy appearance as N increases, although even for $N = 8$ Fourier analysis shows that the map has a band-pass spectrum similar to that observed in the non-annealed cases.

Of interest is the extent to which the distributions of (a) receptive field centres i.e. stimulus values explicitly represented in the map, and (b) the net activity evoked in the cortex, match the distribution of stimuli used to generate the map. This probability distribution was given by

$$p(b_1, b_2, \ldots, b_N)\, \mathrm{d}b_1 \mathrm{d}b_2, \ldots, \mathrm{d}b_N = \frac{1}{(2\pi)^{N/2}\sigma^N}\, \mathrm{e}^{-b_1^2/2\sigma_d^2}\, \mathrm{e}^{-b_2^2/2\sigma_d^2} \ldots \mathrm{d}b_1\, \mathrm{d}b_2, \ldots, \mathrm{d}b_N.$$

This is radially symmetric around the origin, and can be evaluated as a function of radial distance, $r = \left(b_1^2 + b_2^2 + \cdots\right)^{1/2}$, giving

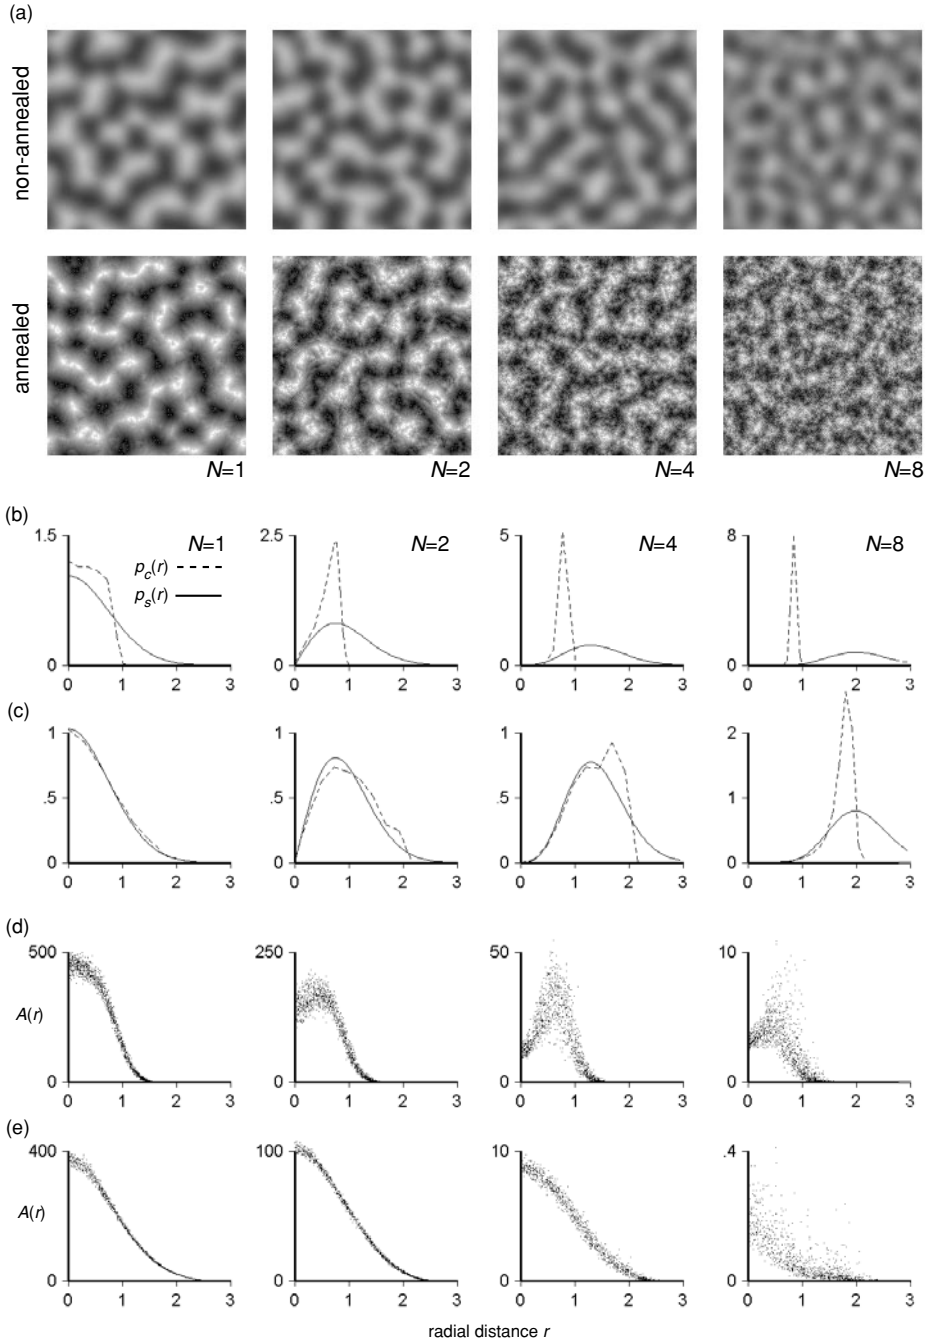

Figure 4. (a) Examples of protomaps for features with a Gaussian distribution for different values of N in non-annealed and annealed cases. Note the irregularity of the maps in the annealed case for $N = 4$ and $N = 8$. Panels (b) and (c) show, in a layout corresponding to the layout of panels in (a) the probability of stimulus occurrence as a function of radial distance, r, from the origin (solid lines) and the fractional area of cortex devoted to stimuli as a function of r (dashed lines). Each graph in panels (d) and (e) shows net responses (A, equation (3)) to 500 randomly chosen stimuli, uniformly distributed with respect to r. Each dot shows the response to one stimulus.

Table 4. Values of the Kullback–Leibler divergence, and weighted coverage uniformity for Gaussian input stimulus distributions of different dimensionality, N. Values are the mean of three different simulations.

	K–L divergence (L)		Weighted coverage	
N	Non-annealed	Annealed	Non-annealed	Annealed
1	−0.189	0.023	0.073	0.034
2	−0.664	−0.033	0.125	0.029
3	−1.257	−0.075	0.208	0.040
4	−2.011	−0.145	0.258	0.074
5	−2.719	−0.248	0.300	0.145
6	−3.414	−0.395	0.341	0.238
7	−4.186	−0.586	0.363	0.358
8	−4.945	−0.806	0.389	0.549

$$p(r)\,\mathrm{d}r = \frac{r^{N-1}\,\mathrm{e}^{(-r^2/2\sigma_d^2)}}{\int_0^\infty r^{N-1}\,\mathrm{e}^{(-r^2/2\sigma_d^2)}\,\mathrm{d}r}\,\mathrm{d}r. \tag{8}$$

Note that although the density per unit volume is highest at the origin, and falls off in a Gaussian manner with increasing distance from the origin, when the probability density is expressed as a radial function i.e. as the probability of a point falling within a thin spherical shell of radius r and width dr, the distribution for $N > 1$ has a maximum at a distance $r_m = \sqrt{N-1}\sigma_d$ from the origin, because the fractional volume of the shell increases as r^{N-1}. Hence the (possibly) counterintuitive result that for larger values of N the majority of values in the distribution, although very different from each other, occur within a narrow range of distances from the origin.

Figures 4(b)–(e) show distributions of stimuli, cortical receptive field centres and the net activity $A(\mathbf{v})$ (equation (3)) evoked in the cortex by specific stimuli, as a function of r, for different values of N in non-annealed and annealed maps. Figures 4(b) and (c) show the stimulus distribution used to generate the map (solid lines) and the resulting distribution of cortical receptive field centre values (dashed lines). With the exceptions of $N = 1$ and $N = 2$ in the annealed maps, the receptive field centre distributions differ systematically from the corresponding stimulus distribution. Stimuli far from the origin are typically underrepresented and field centres are narrowly distributed around a peak that is differently located from the peak in the distribution of stimulus values. Table 4 gives the Kullback–Leibler divergence, L, (a measure of the similarity of the distributions) as a function of N for the non-annealed and annealed cases.

Figures 4(d) and (e) show the distributions of cortical activity, with each dot showing the net response to a single stimulus at a distance r from the origin. This activity was calculated assuming a retinal receptive field size, $\sigma_r = 1.12$ (as above) and a stimulus receptive field width, $\sigma_s = 0.25$. Stimuli were uniformly distributed across the retina, and values of b_1, b_2 etc were uniformly distributed over the surface of a hypersphere of radius r, with values of r uniformly distributed between $0 < r < 4\sigma_d$ (=3.0). In the annealed maps (figure 4(e)) net responses varied in proportion to the probability of stimulus occurrence (expressed per unit volume) during map development, i.e. response magnitudes fell off in a Gaussian fashion with r. The agreement was less good in the non-annealed maps (figure 4(d)): for values of $N > 2$ stimuli close to the origin became increasingly underrepresented relative to those further out.

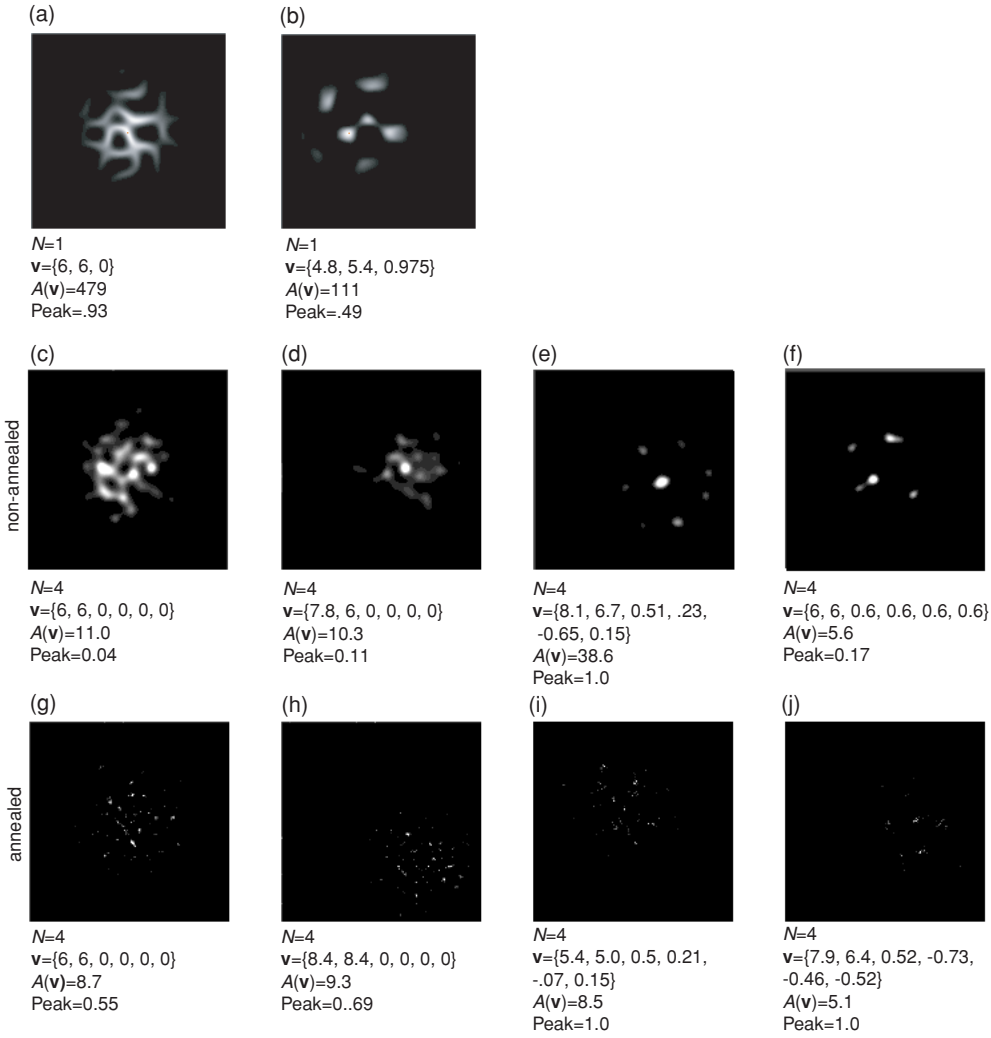

Figure 5. Spatial patterns of response to specific stimuli in maps generated with Gaussian stimulus distributions (as shown in figure 4). The value of the stimulus vector **v**, the net response A (equation (3)) and the peak response are shown below each panel. Panels (a) and (b) are from a map with $N = 1$ and show the response to a stimulus at the origin (a) and at a distance of 1.13 units from the origin (b). Panels (c)–(*f*) are from a non-annealed map, and (g)–(j) are from an annealed map, in both cases $N = 4$. Responses were calculated assuming a retinal receptive field size, $\sigma_r = 1.12$ and a stimulus receptive field width, $\sigma_s = 0.25$. In order to show the response patterns more clearly, brightness values have been normalized to the peak in each panel and a nonlinear brightness scale has been used that emphasizes small responses. Assuming a scaling of about 28 pixels to 1 mm (see methods) spots in the non-annealed maps have a diameter of $\sim$100–200 μm and in the annealed maps of $\sim$50–100 μm.

Figures 5(a), (b) show examples of spatial patterns of activity produced in a non-annealed map for which $N = 1$ (the same map shown in the top-left panel of figure 4(a)). Receptive field width values used to calculate the activity were the same as for the analysis in the previous paragraph. Figure 5(a) shows the response to a stimulus vector $\mathbf{v} = \{6.0, 6.0, 0\}$ i.e. for a retinal location = 6.0, 6.0 (i.e. in the centre of the retina, which is 12×12 units in size) and a

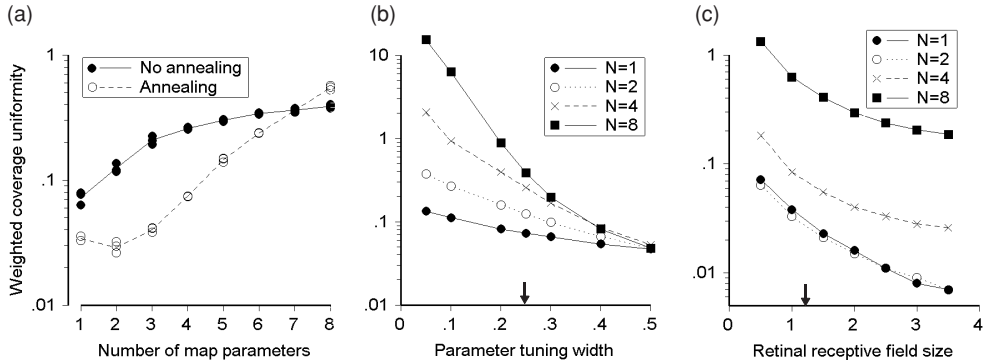

Figure 6. Values of weighted coverage (c', equation (6)) for maps of Gaussian features. (a) shows how c' varies with N for annealed and non-annealed maps, assuming a retinal receptive field size $\sigma_r = 1.12$ units and a feature tuning width, $\sigma_s = 0.25$, (b) the effect on c' of varying σ_s, keeping $\sigma_r = 1.12$ and (c) the effect on c' of varying σ_r, keeping $\sigma_s = 0.25$, for different values of N. Arrows in (b) and (c) indicate the parameter values used for (a).

stimulus at the origin, i.e. $b_1 = 0.0$. Activated regions surround and include the zero-crossings of the periodic pattern shown in figure 4(a), and are thus stripe like with a spacing about half that of the protomap periodicity. For stimuli further away from the origin, e.g. $\mathbf{v} = \{6.0, 6.0, 0.975\}$, as shown in figure 5(b), the response pattern is blob like, with the blobs corresponding to the peaks of the pattern shown in figure 4(a). This results in a blob spacing equal to that of the protomap periodicity.

Response patterns in maps for which $N > 1$ are the products of N different activity patterns, similar to those shown in figures 5(a) and (b), depending on the distances of the feature values from the origin in feature space. Representative response patterns from a non-annealed and an annealed map, in which $N = 4$ are shown in figures 5(c)–(j). Responses typically consist of several spots of activity of varying magnitude localized around the region of cortex corresponding to the retinal location of the stimulus. The spots are much smaller in annealed maps. Except for stimuli that are explicitly represented in the map (i.e. those locations in stimulus space that coincide with a point in the cortex) the peak cortical response is necessarily sub-maximal. For stimuli at the origin (figures 5(c), (d), (g) and (h)) spots have a spacing about half that of the protomap periodicity. For stimuli further from the origin (figures 5(e), (f), (i) and (j)) the spacing is wider and similar to protomap periodicity.

Weighted coverage uniformity, as defined by equation (5), was calculated and the results are shown in figure 6(a). Coverage uniformity is generally better than for the other types of map studied here, particularly for the non-annealed maps. As in the other types of map, annealing results in lower values of c'. However, because the rate of increase in c' with N falls off as N increases in the non-annealed maps, the effect of annealing is less apparent for larger values of N. For $N = 8$ annealing actually made coverage slightly worse. However, it is important to point out that the low c' values seen with the Gaussian maps, particularly with the larger values of N in the non-annealed maps are a consequence of the weighted coverage measure and the poor representation of stimulus space that occurs in the non-annealed maps for larger values of N. Thus, as shown in figure 4, most stimuli (i.e. those at a distance >1 radial unit from the origin) evoke almost no response in the cortex, and therefore make little contribution to the value of coverage uniformity. The region of stimulus space that is represented in the map lies mostly within a sphere of radius ≈ 1 and this region, although

Table 5. Analysis of maps of non-uniformly distributed binary features, generated with and without annealing. Analysis includes calculation of weighted coverage uniformity, c' and comparisons of frequency of stimulus presentation, f_n, with the fractional area of map, C_n, devoted to the stimulus. Measures of comparison include the Kullback–Leibler divergence, L, the linear correlation coefficient, r, and the slope of the correlation, m. Each value is the mean from three different simulations.

N	c'	L	R	m
Non-annealed				
4	0.163	0.139	0.842	1.48
6	0.400	0.225	0.668	1.34
8	0.709	0.559	0.401	0.977
10	1.364	0.666	0.220	0.507
Annealed				
4	0.036	0.001	0.997	0.95
6	0.079	0.005	0.989	1.08
8	0.173	0.026	0.942	1.04
10	0.383	0.079	0.929	1.37

represented with relatively uniform coverage i.e. predictable responses, represents only a small portion of the stimuli presented to the map during development. Figures 6(b) and (c) show the effects of changing receptive field size: as in angle maps increasing either stimulus receptive field size, or retinal receptive field size, always decreases c'; changes in stimulus receptive field size had large effects on coverage for large values of N, while changes in retinal receptive field size had similar effects for all values of N.

3.4. Non-uniform distributions of binary features

Swindale (2000) studied coverage uniformity for a binary feature space in which stimuli were uniformly distributed across a 2D retina and the vertices of an N-dimensional hypercube. These results are extended here for the case of a non-uniform distribution of stimuli across the vertices. It is convenient to label each stimulus (or vertex) with its corresponding N-bit binary value, n. A non-uniform distribution across the 2^N possible values of n was generated by calculating a table of probabilities p_n, uniformly distributed over an interval and then normalized so that $\sum_{n=1}^{n=2^N} p_n = 1$. (Many other ways of distributing stimuli are of course possible.) During development, stimuli were chosen at random according to these probabilities, which remained fixed during any one simulation. The distribution of stimuli across the retina was uniform. Both non-annealed and annealed simulations were done, for values of $N = 4, 6, 8$ and 10. For the non-annealed maps the total number of stimulus presentations given was 2×10^6. For the annealed maps the annealing schedule was the same as that used for Gaussian stimuli, and the total number of stimulus presentations was 10^7. Results were assessed by calculating weighted coverage (eqution (6)) across the 2^N stimulus sets, and by calculating the percentage area of cortex, C_n, devoted to a particular stimulus n, and comparing it with the frequency of presentation of the same stimulus, f_n, taken over the entire period of development. (Sufficiently large numbers of stimuli were used during development that $f_n \approx p_n$.) Comparisons were made by calculating the Kullback–Leibler divergence, L, and by calculating the correlation and slope of the relationship between f_n and C_n on a linear scale.

Figure 7 and table 5 show the results of these comparisons for the non-annealed and annealed maps. For the non-annealed maps (figures 7(a) and (b)) the correlation between f_n and C_n, was poor and got worse as N increased. In particular, for $N \geqslant 8$, many stimuli

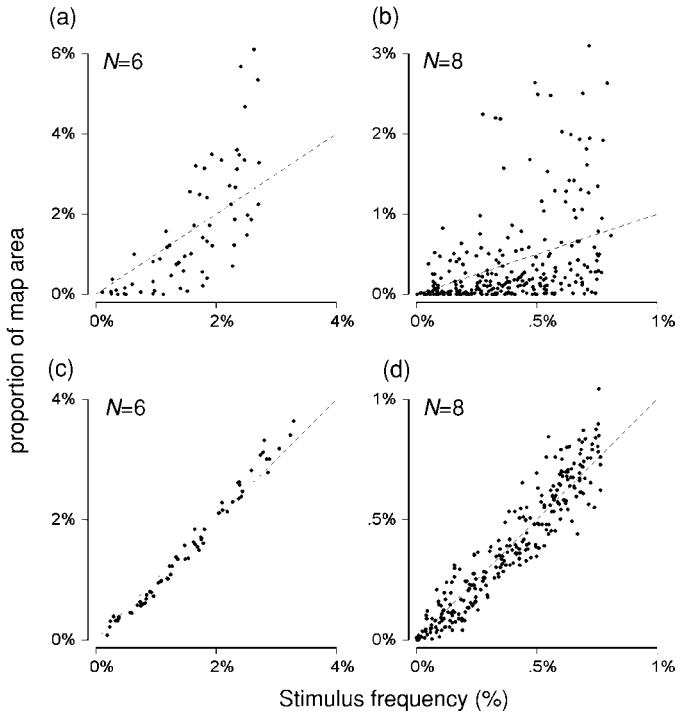

Figure 7. Graphs showing the relationship between the frequency of stimulus presentation (x-axis) and the percentage area of cortex devoted to the feature (y-axis). Each point represents a single N-bit binary feature. Because probabilities are assigned at random to individual features there is no systematic relationship between the binary value of the feature and its position along the x-axis. (a) and (b) show results from two non-annealed maps, (c) and (d) results from two annealed maps. The dashed line in all four graphs has a slope equal to 1.

failed to gain a representation in the map despite having a relatively high probability of occurrence during map formation (figure 7(b)). For $N = 4$ and $N = 6$ there was a tendency for stimuli with a lower probability of occurrence to be underrepresented and those with the highest probabilities to be over-represented (figure 7(a)). Running the simulations for longer (up to 10^7 stimuli) did not appear to change these results. For the annealed maps however (figures 7(c) and (d)), f_n and C_n were highly correlated, with a slope close to 1, and with low values of L. Weighted coverage uniformity measures are shown in table 5.

Because the values of p_n were randomly assigned across the 2^N features, the individual bit probabilities (i.e. the probability of occurrence of a particular binary subfeature) were generally within the range 0.4–0.6. As a result, individual protomaps looked similar to those obtained with uniform stimulus distributions (as in Swindale (2000)). The variability in the representation of particular features is a consequence of variations in the way in which the protomaps are overlaid, and these are not evident in individual protomaps.

3.5. Retinotopy

Figure 8 shows examples of projections of the cortex into retinotopic space for different numbers of dimensions in non-annealed and annealed maps of Gaussian features. Similar patterns of change in retinal topography were seen for the other types of feature space studied

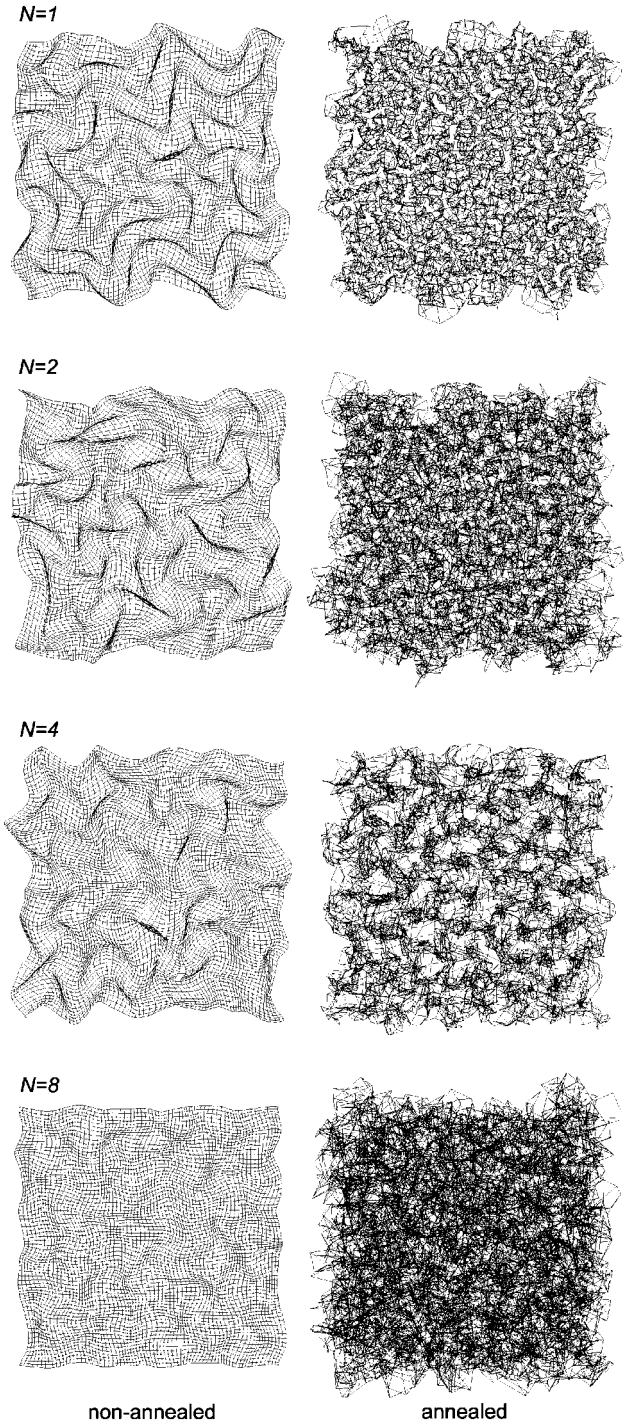

Figure 8. Examples of retinotopic maps obtained in simulations with Gaussian feature variables. Maps show the cortex projected into retinal coordinates (i.e. the position on the page represents the retinal position of each point on the cortical grid; for clarity only every second cortical point is plotted).

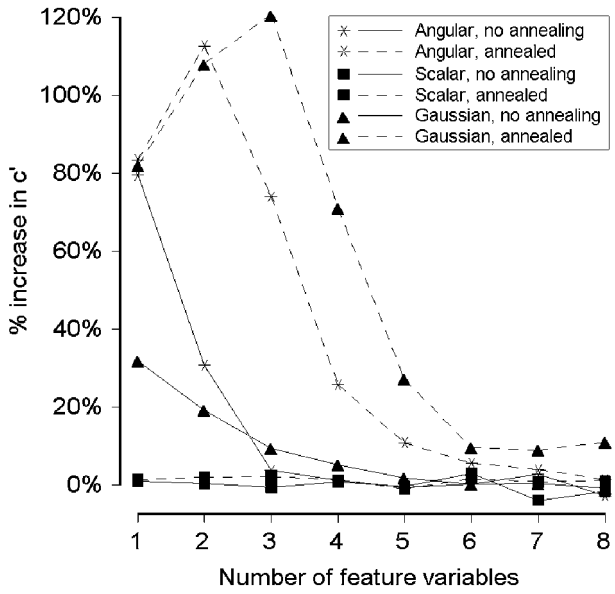

Figure 9. The percentage change in coverage caused by substituting a uniform retinotopic map for that produced by the Kohonen algorithm, for angular, scalar and Gaussian maps in non-annealed and annealed cases. Each point is the mean calculated from three maps generated with identical parameters but different random number seeds.

here. In non-annealed maps the retinotopic map was always locally smooth with periodic fluctuations in magnification factor. As N increased, the fluctuations became smaller, and magnification factor more uniform. In the annealed maps a different pattern of change was seen: deviations from uniform retinotopy were larger with extensive folding leading to local changes in the direction of retinotopy. The amount of folding and apparent local scatter of retinal positions increased as the number of feature dimensions increased. Particularly noticeable were the presence of small 'holes' in the retinal map for $N \geqslant 4$. As is characteristic of dimension reduction maps (Durbin and Mitchison 1990) regions of cortex where retinal position changed rapidly coincided with regions where other feature variables were changing slowly i.e. at the peaks and troughs of the polymap patterns for scalar and Gaussian maps, or in regions of relatively low rates of change of angle preference in the angular maps. These correlations were easy to discern for maps of one and two dimensions but became progressively less obvious as the number of feature dimensions increased.

In some cases the irregularities in retinotopy resulted in improved coverage. Figure 9 shows the change in c' caused by substituting a uniform retinotopic map for the irregular one formed during development. Large increases in c' for $N < 5$ were found in the annealed angular and Gaussian maps. Smaller increases for $N < 3$ were found for the non-annealed angular and Gaussian maps. Scalar maps, whether annealed or non-annealed, showed little or no change in c' for any value of N. The absence of any impact of retinal scatter (or systematic deviation from uniform retinotopy) in the scalar maps may be related to the negligible effect of changes in receptive field size on coverage in these maps (figure 2(c)) which in turn may be related to the geometry of the scalar space (figure 3). Since receptive fields can be made much smaller than the default value without impacting coverage, moving receptive field positions around might likewise be expected to have little effect. It is of interest that retinal maps still developed systematic folds and distortions in cases where there appeared to be little or no

effect on coverage. Non-uniform retinotopy presumably satisfies some other constraint in such cases, although it is not clear what this might be.

4. Discussion

Swindale (2000) asked 'How many maps are there in visual cortex?' and gave an answer based on an analysis of maps of different numbers of binary feature variables, spread across a 2D retina. Because c' increased (i.e. coverage uniformity degraded) by a factor of about 1.6 as each additional feature dimension was added, a fixed upper limit on the number of maps based on coverage could not be established. However, a variety of considerations led to the suggestion of an upper limit of about six to seven maps for binary features. This paper has extended the analysis of that paper to additional types of feature space. The spaces examined were combinations of angular variables, combinations of scalar variables with a uniform distribution between fixed limits, combinations of scalar variables with a Gaussian distribution and binary features with a non-uniform probability distribution. The conclusions that can be drawn about possible upper limits to map numbers are subject to a number of qualifications. One is that the use of scalar and angular variables requires the definition of a corresponding cortical receptive field size for the variable concerned, and the choice of this value needs justification. A second one is that it is unknown whether annealing, which generally results in a significant improvement in coverage, has a biological counterpart. Assuming for the moment that (a) the effects of annealing are absent, (b) the choices of receptive field size (discussed below) are justified and (c) c' has an upper limit of about 0.5, the analyses in this paper suggest upper limits of around 4 maps for angular variables, 1–2 maps for uniformly distributed scalar variables, 7 or 8 maps for scalar variables with a Gaussian distribution and 6 or 7 maps for binary variables with the particular non-uniform distribution used here. Annealing increases these limits to about 6 maps for angular variables, 5 maps for uniformly distributed scalar variables and 10 or more maps for binary variables. These conclusions of course are subject to several kinds of qualification, discussed in the following sections.

4.1. Receptive field size

Coverage uniformity is generally dependent on the choice of receptive field sizes, firstly for the retina and secondly for the particular features concerned (figures 1(b), 1(c), 2(b), 2(c), 6(b) and 6(c)). It was assumed for simplicity that field sizes are constant both within and across the different feature dimensions, although this is not necessarily the case in reality. The choice of field sizes for the retina and for orientation tuning was based on physiological values obtained in visual cortex (see Swindale *et al* 2000 for details). One source of information about the likely sizes of receptive fields in scalar and Gaussian maps comes from the spatial frequency maps in visual cortex (Shoham *et al* 1997, Issa *et al* 2000). The total range of spatial frequencies (represented by tuning curve peaks) in area 17 of the cat spans about 0.3 to 3.0 cycles/degree, a range of about 3.3 octaves, while individual neurons have, on average, bandwidths (full width at half-height) of about 1.3 octaves (Movshon *et al* 1978). Receptive fields are therefore about 2.5 times smaller than the total range of parameters represented in the map. Field widths chosen here for scalar and Gaussian maps were similarly not much smaller than the range of stimulus parameters. For the scalar maps field widths had $\sigma_s = 0.35$ which gives a full width at half height of about 0.8. Stimuli lay on a line $d = 2$ units in length giving a ratio of $2/0.8 = 2.5$ which is comparable to the estimated ratio of 2.4 (see results section) for spatial frequency tuning in cat visual cortex. For the Gaussian maps receptive fields had

widths $\sigma_s = 0.25$ and the stimulus distribution had $\sigma_d = 0.75$. The impact of receptive field width on c' was examined: for scalar maps relatively small changes in c' were observed for all values of N (figure 2(b)) while for Gaussian maps (figure 6(b)) much larger effects were seen for $N \geqslant 4$.

4.2. Annealing

Annealing, i.e. the procedure of gradually reducing the size of the cortical neighbourhood function during development, nearly always improves coverage, typically by a factor of 2 or 3 (figures 1(a), 2(a), 6(a) and table 5). The effects of annealing are evident as increased local folding of the protomaps in stimulus space, or equivalently, as increased local variability in the mapping of parameter values, or spatially as an increase in the apparent irregularity in the structure of the map. Because of the implications for improved coverage it would be interesting if experimental evidence for such local variability e.g. in the structure of the orientation map could be found. However, this evidence may not be easy to come by: firstly optical imaging probably does not have the spatial resolution required to show such fine details, secondly it might be difficult to distinguish genuine randomness in map structure from systematic local variability related to coverage. An additional possibility is that the effects of annealing might make it hard to obtain evidence for a map in the first place. For example it is difficult to discern evidence for long range periodic order in the annealed Gaussian maps for $N = 8$ (figure 4) especially if only a small region of the map is studied, as would be the case with microelectrode techniques.

The Kohonen model is necessarily vague about the interpretation of the cortical neighbourhood function and the biological interpretation of annealing is thus similarly vague. However, it is plausible that any process that leads to a decrease in the distance over which heterosynaptic facilitation is exerted—which may involve the diffusion of plasticity modulating substances—might implement effects akin to those of annealing. Early in development such a decrease might occur as a result of (a) physical expansion in the size of the brain due to growth, (b) a reduced coefficient of diffusion caused by a reduction of extracellular fluid volume or increased tortuosity of extracellular space and (c) increasing size of the molecules involved in mediating such effects. Gradual reduction in the sizes of dendritic and axonal arbors involved in mediating plasticity might have a similar effect.

4.3. Non-uniformly distributed stimuli

For maps with a non-uniform stimulus distribution, coverage uniformity defined by the weighted measure used here is not the only functional factor likely to limit the possible number of maps. As pointed out in section 3, the low values of c' obtained in the Gaussian maps, particularly for the non-annealed cases, are probably a result of the fact that the maps represented only a restricted range of the total set of the stimuli presented during development. Thus many classes of stimuli, despite being present during development, evoke little or no responses in the cortex and make little contribution to the measured value of c' which, as defined here, takes into account only the variability in coverage of stimuli that evoke relatively large cortical responses. This is the reason for using a different way of assessing the fidelity of the representation of a particular stimulus distribution, namely to compare the probability distribution of the stimuli with the corresponding net cortical response, or the fractional area of cortex devoted to the stimulus and its probability of occurrence during map development.

In the real brain it is known that the area of cortex devoted to the representation of a particular stimulus varies according to the frequency of presentation of the stimulus.

For example the area of auditory cortex devoted to a particular tone frequency increases following training in frequency discrimination at that particular frequency (Recanzone *et al* 1993); orientation maps change following continuous exposure to lines of a restricted range of orientations during development (Sengpiel *et al* 1999) and the area of somatosensory cortex devoted to the representation of parts of the body surface decreases following removal, and increases following selective stimulation, of the part concerned (Buonomano and Merzenich 1998). The evidence from these studies is not sufficiently detailed to define an exact relation between the probability distribution of the stimuli and the relative sizes of the areas of cortex devoted to them. *A priori*, however, it would seem best that the two distributions should match. The measures used here show that the agreement is generally poor (figures 4(b) and 7(a), (b)) unless annealing is used (figures 4(b) and 7(c), (d)), in which case good agreement can be obtained up to about 4–6 maps for the Gaussian case and 6–8 in the binary case.

It is of interest that mathematical analysis of the Kohonen algorithm (Ritter and Schulten 1986) shows that, for a 1D (retinal) map the local cortical magnification factor should vary as the 2/3rds power of the input stimulus density. Consequently a graph of fractional cortical area versus the probability of stimulus occurrence should have a slope of $2/3$ on a log–log plot. I confirmed this result in simulations of 1D and 2D retinal maps (not reported here) but it apparently does not apply to the case of non-uniformly distributed feature dimensions additional to the retina: in such cases approximate equality between the distributions appears to hold true (figures 4(e), 7(c) and 7(d)).

4.4. Comparison between cyclic and non-cyclic feature spaces

The results obtained here suggest that cyclic features may have a particular advantage over uniform scalar ones when it comes to being represented in cortical maps. This is because the cortical map can wrap around without compromising continuity, and edge effects, or difficulties representing the extremes of the distribution uniformly, do not occur (figure 3). An analogy is that it is simple to cover a cylindrical surface uniformly with tape by wrapping the tape in a helical pattern but much harder to do the same for a long rectangle. Probably because of this, coverage was generally several times worse for scalar stimuli uniformly distributed along a line than for angular stimuli (although the difference became smaller as the number of maps increased—cf figures 1(a) and 2(a)). This difference would have been still larger if equivalent receptive field sizes, i.e. the same ratio between the total range of parameter values represented in each protomap and the receptive field width for the parameter concerned, had been chosen for the angular and scalar feature variables. For the orientation maps this ratio was $180/\sigma_\theta = 180/25 \approx 7$ while for the scalar maps the ratio was $d/\sigma_s = 2.0/0.35 \approx 5.7$. As discussed above these ratios are roughly the same as those found in orientation and spatial frequency maps in visual cortex. Thus the relatively coarser coding of spatial frequency compared to orientation may perhaps be related to the geometry of the spaces that these features are best represented in.

4.5. Choice of algorithm for the simulations

Other developmental models or algorithms could have been used to generate the types of mappings studied here, possibly with different results. The Kohonen algorithm was chosen here, and in a previous related work (Swindale 2000) because it is capable of producing highly realistic (at least to visual inspection) simulations of real visual cortex maps (Obermayer *et al* 1990, Obermayer and Blasdel 1993) and because it is computationally fast. This is a significant

practical advantage because computation times rise exponentially with N, and for larger numbers of dimensions it can take a day or more to generate a single map on a fast PC. Although the algorithm works in a way that is not obviously based in biology, it does have a biological interpretation, aspects of which are discussed briefly above (in sections 2.1 and 4.2) and more extensively by Kohonen (1993).

The main alternative to the use of Kohonen algorithm in the present setting is the elastic net algorithm (Durbin and Willshaw 1987, Durbin and Mitchison 1990, Goodhill and Cimponeriu 2000). Comparisons of the two algorithms (Swindale, unpublished results) have shown similar coverage values for comparable types of mapping although an extensive comparison of the two algorithms has not been carried out. However, the main goal of this paper is to establish the existence of biologically plausible solutions to mapping problems that have potential biological relevance, and to set rough upper limits on the numbers of dimensions that might be represented within particular types of cortical map. Exploring the possibility of better solutions, and making a detailed comparison of the solutions produced by different developmental algorithms will arguably be of little use unless the solutions can be tested against biological data, which at present do not exist.

4.6. Gaussian maps

Many types of complex objects may be described in terms of combinations of lower-level features, each of which is likely to have some characteristic distribution, possibly Gaussian, in the real world (Biederman 1987). For example faces can be characterized as points in a 'face space' the dimensions of which are the variables that characterize the variation between individual faces (Valentine 1991). Such variables might be the distance between the eyes, length of the nose, etc. Little is known about the numbers of dimensions likely to be used by the brain in such representations (six were used in the study by Lee *et al* 2000). Nor is it known whether or not such dimensions have an associated cortical mapping. A suggestive optical imaging study by Wang *et al* (1998) showed that face stimuli evoked a number of closely spaced small spots of activity in inferotemporal cortex (similar to the pattern shown in figure 5(b)) and that the positions of individual spots changed as the angle of view of the face changed. This suggests that angle of view is one of the parameters mapped in this area. To the author's knowledge there have been no attempts to demonstrate variation in response selectivity, or mapping of feature variations likely to code facial identity in any visual cortical area. However, the analysis done here suggests that a single cortical area might be able to represent a five–six-dimensional feature space reasonably faithfully.

It is shown here (figure 4(d), (e)) that stimuli close to the origin of a multi-dimensional space represented in multi-dimensional feature maps usually evoke larger net responses than those more distant from the origin. Given this result, it is of interest that 'average' stimuli—that is, averages of stimuli taken across restricted sets of dimensions—are also rated as more attractive by human observers. This is true for faces (Galton 1878, Langlois and Roggman 1990, Rhodes and Tremewan 1996) and for non-face stimuli (birds, fish and cars: Halberstadt and Rhodes (2003)). This suggests that, within a narrowly restricted stimulus set, the greater attractiveness of some stimuli relative to others may simply be related to the magnitude of the response evoked within the cortical area where the dimensions in question are mapped.

Acknowledgments

This research was supported by grants from the Canadian Institutes of Health Research and the National Science and Engineering Research Council of Canada.

References

Albright T D 1984 Direction and orientation selectivity of neurons in visual area MT of the macaque *J. Neurophysiol.* **52** 1106–30

Bartfeld E and Grinvald A 1992 Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular dominance columns in primate striate cortex *Proc. Natl. Acad. Sci. USA* **89** 11905–9

Biederman I 1987 Recognition-by-components: a theory of human image understanding *Psychol. Rev.* **94** 115–47

Bishop C M 1995 *Neural Networks For Pattern Recognition* (Oxford: Oxford University Press)

Buonomano D V and Merzenich M M 1998 Cortical plasticity: from synapses to maps *Annu. Rev. Neurosci.* **21** 149–86

DeAngelis G C and Newsome W T 1999 Organization of disparity-selective neurons in macaque area MT *J. Neurosci.* **19** 1398–415

Diogo A C, Soares J G, Koulakov A, Albright T D and Gattass R 2003 Electrophysiological imaging of functional architecture in the cortical middle temporal visual area of *Cebus apella* monkey *J. Neurosci.* **23** 3881–98

Durbin R and Mitchison G 1990 A dimension reduction framework for understanding cortical maps *Nature* **343** 644–7

Durbin R and Willshaw D J 1987 An analogue approach to the travelling salesman problem using an elastic net method *Nature* **326** 689–91

Everson R M, Prashanth A K, Gabbay M, Knight B W, Sirovich L and Kaplan E 1998 Representation of spatial frequency and orientation in the visual cortex *Proc. Natl. Acad. Sci. USA* **95** 8334–8

Galton F 1878 Composite portraits *Journal of the Anthropological Institute of Great Britain & Ireland* **8** 132–42

Ghose G M and Ts'o D Y 1997 Form processing modules in primate area V4 *J. Neurophysiol.* **77** 2191–6

Goodhill G J and Cimponeriu A 2000 Analysis of the elastic net model applied to the formation of ocular dominance and orientation columns *Network* **11** 153–68

Halberstadt J and Rhodes G 2003 It's not just average faces that are attractive: computer-manipulated averageness makes birds, fish, and automobiles attractive *Psychon. Bull. Rev.* **10** 149–56

Hubel D H and Wiesel T N 1977 Functional architecture of macaque monkey visual cortex *Proc. R. Soc. Lond.* B **198** 1–59

Hübener M, Shoham D, Grinvald A and Bonhoeffer T 1997 Spatial relationships among three columnar systems in cat area 17 *J. Neurosci.* **17** 9270–84

Issa N P, Trepel C and Stryker M 2000 Spatial frequency maps in cat visual cortex *J. Neurosci.* **20** 8504–14

Janata P, Birk J L, Van Horn J D, Leman M, Tillmann B and Bharucha J J 2002 The cortical topography of tonal structures underlying Western music *Science* **298** 2167–70

Kohonen T 1982 Self-organized formation of topologically correct feature maps. *Biol. Cybern.* **43** 59–69

Kohonen T 1993 Physiological interpretation of the self-organizing map algorithm *Neural Netw.* **6** 895–905

Landisman C E and Ts'o D Y 2002 Color processing in macaque striate cortex: relationships to ocular dominance, cytochrome oxidase, and orientation *J. Neurophysiol.* **87** 3126–37

Langlois J H and Roggman L A 1990 Attractive faces are only average *Psychol. Sci.* **1** 115–21

Langner G, Sams M, Heil P and Schulze H 1997 Frequency and periodicity are represented in orthogonal maps in the human auditory cortex: evidence from magnetoencephalography *J. Comput. Physiol.* **181** 665–76

Lee K, Byatt G and Rhodes G 2000 Caricature effects, distinctiveness, and identification *Psychol. Sci.* **11** 379–85

Mendelson J R, Schreiner C E and Sutter M L 1997 Functional topography of cat primary auditory cortex: response latencies *J. Comput. Physiol.* A **181** 615–33

Mountcastle V B 1978 *The Mindful Brain* ed F O Schmitt (Cambridge, MA: MIT Press) pp 7–50

Movshon J A, Thompson I D and Tolhurst D J 1978 Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex *J. Physiol. (Lond.)* **283** 101–20

Obermayer K and Blasdel G G 1993 Geometry of orientation and ocular dominance columns in monkey striate cortex *J. Neurosci.* **13** 4114–29

Obermayer K, Blasdel G G and Schulten K 1992 Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps *Phys. Rev.* A **45** 7568–89

Obermayer K, Ritter H and Schulten K 1990 A principle for the formation of the spatial structure of cortical feature maps *Proc. Natl. Acad. Sci. USA* **87** 8345–9

Press W H, Teukolsky S A, Vetterling W T and Flannery B P 1994 *Numerical Recipes: The Art of Scientific Computing* 2nd edn (Cambridge: Cambridge University Press)

Recanzone G H, Schreiner C E and Merzenich M M 1993 Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys *J. Neurosci.* **13** 87–103

Rhodes G and Tremewan T 1996 Averageness, exaggeration, and facial attractiveness *Psychol. Sci.* **7** 105–10

Ritter H and Schulten K 1986 On the stationary states of Kohonen's self-organizing sensory mapping *Biol. Cybern.* **54** 99–106

Schreiner C E 1995 Order and disorder in auditory cortical maps *Curr. Opin. Neurobiol.* **5** 489–96
Schreiner C E 1998 Spatial distribution of responses to simple and complex sounds in the primary auditory cortex *Audiol. Neurootol.* **3** 104–22
Sengpiel F, Stawinski P and Bonhoeffer T 1999 Influence of experience on orientation maps in cat visual cortex *Nat. Neurosci.* **2** 727–32
Shoham D, Hübener M, Schulze S, Grinvald A and Bonhoeffer T 1997 Spatial frequency domains and their relation to cytochrome oxidase staining in cat visual cortex *Nature* **385** 529–33
Swindale N V 1991 Coverage and the design of striate cortex *Biol. Cybern.* **65** 415–24
Swindale N V 2000 How many maps are there in visual cortex? *Cereb. Cortex* **10** 633–43
Swindale N V, Shoham D, Grinvald A, Bonhoeffer T and Hübener M 2000 Visual cortex maps are optimized for uniform coverage *Nature Neurosci.* **3** 822–6
Ts'o D Y, Roe A W and Gilbert C D 2001 A hierarchy of the functional organization for color, form and disparity in primate visual area V2 *Vis. Res.* **41** 1333–49
Tsunoda K, Yamane Y, Nishizaki M and Tanifuji M 2001 Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns *Nature Neurosci.* **4** 832–8
Valentine T 1991 A unified account of the effects of distinctiveness, inversion, and race in face recognition *Q. Exp. Psychol.* A **43** 161–204
Wang G, Tanifuji M and Tanaka K 1998 Functional architecture in monkey infero-temporal cortex revealed by *in vivo* optical imaging *Neurosci. Res.* **32** 33–46
Xiao Y, Wang Y and Felleman D J 2003 A spatially organized representation of colour in macaque cortical area V2 *Nature* **421** 535–9