The Royal Society

Application of Kohonen's self-organizing feature map algorithm to cortical maps of orientation and direction preference

Nicholas V. Swindale[1] **and Hans-Ulrich Bauer**[2]

[1]*Department of Ophthalmology, University of British Columbia, 2550 Willow Street, Vancouver, British Columbia, Canada V5Z 3N9* (swindale@unixg.ubc.ca)

[2]*Max-Planck-Institut für Strömungsforschung, PO Box 2853, 37018 Göttingen, Germany*

Cortical maps of orientation preference in cats, ferrets and monkeys contain numerous half-rotation point singularities. Experimental data have shown that direction preference also has a smooth representation in these maps, with preferences being for the most part orthogonal to the axis of preferred orientation. As a result, the orientation singularities induce an extensive set of linear fractures in the direction map. These fractures run between and connect nearby point orientation singularities. Their existence appears to pose a puzzle for theories that postulate that cortical maps maximize continuity of representation, because the fractures could be avoided if the orientation map contained full-rotation singularities. Here we show that a dimension-reduction model of cortical map formation, which implements principles of continuity and completeness, produces an arrangement of linear direction fractures connecting point orientation singularities which is similar to that observed experimentally. We analyse the behaviour of this model and suggest reasons why the model maps contain half-rotation rather than full-rotation orientation singularities.

Keywords: visual cortex; orientation preference; direction preference; cerebral cortex; visual cortex models; cortical maps

1. INTRODUCTION

Physiological studies of mammalian sensory cortical areas (Mountcastle 1957; Hubel & Wiesel 1962, 1968, 1974) have established that neurons in columns perpendicular to the pial surface respond selectively to combinations of specific stimulus features. These preferences usually, although not always, change smoothly with tangential position in the cortex. In the primary visual cortical areas of cats (Hubel & Wiesel 1962; Albus 1975; Swindale *et al*. 1987; Bonhoeffer & Grinvald 1991), ferrets (Weliky *et al*. 1996) and monkeys (Hubel & Wiesel 1968, 1974; Blasdel & Salama 1986; Blasdel 1992), preference for the orientation of a bar or an edge varies smoothly and periodically over the cortical surface, except at point singularities, each of which is surrounded by a single 180° range of orientation preferences. Preference for the direction of stimulus motion, which is usually orthogonal to orientation preference, also tends to change continuously. Single-unit recording data from cat areas 17 and 18 (Payne *et al*. 1981; Tolhurst *et al*. 1981; Swindale *et al*. 1987), and primate area MT (Albright *et al*. 1984), have suggested, and optical imaging data have confirmed (Malonek *et al*. 1994; Shmuel & Grinvald 1996; Weliky *et al*. 1996), that the direction preference map is characterized by regions of smooth change traversed by linear fractures across which direction preferences flip by 180°. These fractures, or line singularities, extend between and connect nearby orientation singularities. Their existence may reflect the geometrical necessity that, given an orthogonal relationship between direction and orientation preferences, only half of a complete 360° cycle of direction preferences can be continuously arranged around a 180° orientation singularity. As a result, line singularities in direction preference necessarily extend from, and connect together, orientation singularities (Swindale *et al*. 1987; Swindale 1996*a*).

At first sight, this arrangement poses a puzzle for theories that postulate that continuity of stimulus representation is maximized in cortical maps (see Erwin *et al*. (1995) and Swindale (1996*b*) for reviews). One might expect that a more continuous map would be obtained if the orientation map contained full rotation 360° singularities, as this would allow a continuous 360° rotation of direction preferences around each orientation singularity, and the more extensive set of linear fractures in the direction map could be avoided. We show here that this naive expectation is probably wrong, and that the experimentally observed arrangement of line and point singularities can be reproduced by a self-organizing, feature-mapping algorithm (Kohonen 1982, 1995; Obermayer *et al*. 1992*b*) that produces mappings between spaces of different dimension, subject to continuity and completeness constraints. Full rotation 360° orientation singularities are never produced by this model, although adding the direction parameter to the orientation map does produce a tendency towards their presence, inasmuch as nearby

*Author for correspondence.

Received 22 December 1997 *Accepted* 23 January 1998

singularity pairs are often of like sign. This counteracts the normal tendency in orientation maps for neighbouring singularity pairs to be of unlike sign (Obermayer & Blasdel 1997). Under some parameter regimes of the model, a distinctive manifestation of this is the occurrence of singularity 'triplets': two 180° orientation singularities of like sign close together, with a 360° direction singularity between them.

2. METHODS

Kohonen's self-organizing feature map (SOFM) algorithm can be applied to the problem of cortical map formation in the following way. The set of receptive field properties possessed by a given neuron, or cortical column, is represented as a point in a Euclidian stimulus space, the axes of which are the stimulus parameters of interest. In the present case, the dimensions of this space include two receptive field position coordinates, two components of an orientation preference vector, and two components of a direction preference vector. This representation allows for the independent coding of orientation and direction angles as well as vector magnitudes. In optical imaging experiments, the latter are usually interpreted as representing the degree of stimulus selectivity. Although a simpler representation could be adopted if direction and orientation angles were assumed to be always orthogonal, the optical data show some degree of independence between the orientation and direction maps, and this can only be captured by independent vector representations of both properties.

The model cortex is represented as a two-dimensional sheet that fills the stimulus space according to a defined set of initial conditions. The position of each cortical point in the space represents the particular set of stimulus values assigned to it. 'Stimuli'—points selected at random from a subset of points in the space—are chosen one at a time. For each, the nearest cortical point is found; this point, and surrounding ones defined by a circular neighbourhood function, are then moved closer towards the stimulus. Physiologically (Kohonen 1993), this might correspond to a process in which the region of cortex responding most strongly to a particular stimulus is first selected by a competitive mechanism likely to involve lateral inhibition. This triggers an all-or-none event, possibly involving the lateral diffusion of a trophic signal, which strengthens active synapses in the region, causing the cells to become more responsive to the activating stimulus. This has two effects: the set of points in the stimulus space that was used as stimuli gains a representation in the cortex, and, because of the cortical neighbourhood function, the resulting map tends to be locally smooth.

The learning rule can be expressed mathematically as follows: let the cortical receptive field with position j be the vector w_j. After each presentation of a stimulus vector v, cortical receptive fields change by an amount

$$\Delta \mathbf{w}_j = \epsilon h(j, j^*)(v - w_j), \tag{1}$$

where ϵ is a rate constant, j^* is the cortical point for which $\mathbf{w}_{j^*}$ is closest to v, and $h(j, j^*)$ is a positive neighbourhood function that decreases to zero with increasing cortical distance, r, between points j and j^*.

Our implementation of the algorithm was similar to that of Obermayer *et al.* (1991, 1992*b*). The cortical neighbourhood function was given by

$$h(r) = \exp(-|r|^2/2\sigma^2), \tag{2}$$

where σ is a width parameter. Stimuli and cortical receptive fields were described by a vector with components (a, b, c, d, x, y). Preferred orientation, which is π-periodic, was defined (following Swindale (1982) and others), as $\theta = 0.5 \tan^{-1}(b/a)$, while preferred direction, which is 2π-periodic, was defined as $\phi = \tan^{-1}(d/c)$. Orientation and direction magnitude values were defined as $(a^2 + b^2)^{1/2}$ and $(c^2 + d^2)^{1/2}$, respectively. These values were calculated for comparison with the magnitude values, which are obtained in optical recording experiments by vector averaging of responses across a set of stimulus orientations or directions. Receptive field position (azimuth and elevation) was given by the values of x and y.

The cortex was represented by a 128×128 grid of points with values $a_{k,l}$, $b_{k,l}$, etc., where $k, l = 1, 2, \ldots 128$. Initial values for the cortical receptive fields were calculated as follows: a roughly ordered receptive field topography was assumed to be present initially, with $x_{k,l} = kX/128 + \xi_x$ and $y_{k,l} = lY/128 + \xi_y$, where ξ_x and ξ_y are random numbers drawn from uniform distributions with means of zero and widths σ_x and σ_y, respectively, and X and Y are constants specifying the overall extent of the region of modelled retina mapped onto the cortex. Initial orientation and direction angles were uniformly randomly distributed in the intervals $[0,\pi]$ and $[0, 2\pi]$ respectively, with no initial orthogonal relationship, and with small random magnitudes drawn from zero mean Gaussian distributions with standard deviations given by σ_θ and σ_ϕ, respectively.

Stimuli were calculated as follows: retinal position values x and y were chosen with uniform probability in the intervals $[0, X]$ and $[0, Y]$, respectively. Values of a, b, c and d were calculated by first choosing an orientation angle θ with uniform probability in the interval $[0, \pi]$, and then setting the direction angle ϕ to one of the two orthogonal directions, $\theta+\pi/2$ or $\theta-\pi/2$, with equal probability. From the resulting values of θ and ϕ the components $a = R_\theta \cos 2\theta$, $b = R_\theta \sin 2\theta$, $c = R_\phi \cos \phi$ and $d = R_\phi \sin \phi$ were calculated, where the constants R_θ and R_ϕ are the radii of the circles on which the orientation and direction stimuli lie. (The assumption that the direction of motion of a stimulus is always orthogonal to its orientation may be justified because of the aperture problem (Hildreth 1984): that is, if only a small part of a moving edge is visible, its direction of motion is ambiguous. Early stages of visual processing appear to deal with this problem by coding direction of motion in one of the two possible ways that are orthogonal to edge orientation. The direction in which an object is moving can then be determined at a later stage of processing, by vector combination of velocity signals from differently orientated edges of the same object.) Note that although the stimulus orientation and direction values were chosen in such a way as always to be orthogonal, the cortical values are not constrained in this way. The effects of relaxing the orthogonality restriction on the stimuli were not studied in detail, although it was noted that random deviations of up to $\pm 45°$ did not produce obvious changes in the model results.

Stimuli were presented to the model until the majority of cortical orientation and direction vector magnitudes were close to R_θ and R_ϕ, respectively, and the singularity density and distribution visually resembled that seen in biological maps. For the 128×128 cortical array this was about 7×10^5 presentations, i.e. roughly 42 stimuli per array point. As discussed below, the maps continue to change slowly with continued stimulus presentations. Although there is little evidence for, or against, continued slow change beyond postnatal development in real maps, it is not unrealistic to suppose that some independently

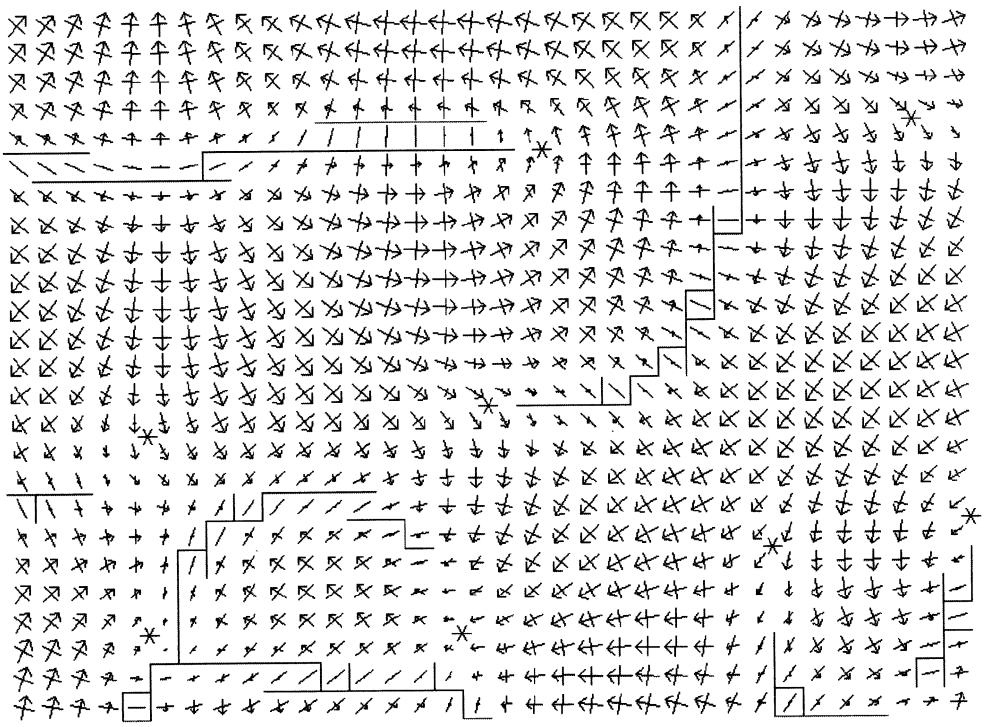

Figure 1. Combined orientation and direction preference map obtained with the Kohonen algorithm. The lengths of the lines and arrows are proportional to the magnitude of the associated orientation and direction vectors, respectively. Continuous horizontal and vertical lines indicate borders across which direction preference angles change by more than 45°. Singularities in the orientation map are indicated by asterisks. The region illustrated comes from the upper left-hand corner of the simulation results shown in figure 2*e–h*. Parameter values were $\epsilon=0.02$, $R_\theta=R_\phi=1$, $X=Y=15$, $\sigma_a=\sigma_b=\sigma_c=\sigma_d=0.1$, $\sigma_x=\sigma_y=0.5$. The cortical neighbourhood, σ, was equal to 2.5 cortical units (*ca.* 150 μm of real cortex) and was kept constant during pattern development. Periodic boundary conditions were not used. The number of stimulus presentations was about 6.9×10^5.

acting developmental program turns off, or limits, cortical plasticity at what is technically an intermediate stage in map formation.

3. RESULTS

(a) *Map properties*

Figure 1 shows a representative example of an orientation and direction preference map, obtained with $R_\theta=R_\phi=1$. Linear regions of low direction vector magnitude, across which direction preferences reverse, extend from orientation singularities, which are indicated by asterisks. Figure 2*a–d* shows angle preference and vector magnitude maps in colour-coded and grey-scale form, respectively, taken from the experimental results of Weliky *et al.* (1996) in ferret visual cortex. Figure 2*e–h* shows corresponding maps obtained with the computer model. Figure 3 shows a portion of the direction vector magnitude map in figure 2*h* in more detail, including the positions of positive and negative singularities in the orientation (red and green asterisks, respectively) and direction (red and green squares) maps.

Many features of the experimental data are reproduced by the model. In addition to an orthogonal relationship between orientation and direction preference angles (a straightforward consequence of the use of orthogonal stimuli), the results show (i) 180° singularities in the orientation map, which coincide with punctate regions of low vector magnitude in the orientation map; (ii) 180° linear fractures in the direction map, which coincide with linear regions of low direction vector magnitude and which terminate in and connect orientation singularities, although not always by the shortest route; and (iii) occasional branches in the linear fractures with an orientation singularity always present at the branch point.

Figure 3 illustrates the pairing of orientation singularities induced by adding direction preference to the set of variables represented in the map. The distances between orientation singularity pairs linked by linear regions of high rates of change of direction angle can be short or long, and pairs linked in this way can be of the same or opposite sign. Fracture trajectories between them are often curved and meandering. Pairs that are close together are often of the same sign, and form distinctive 'triplet' arrangements. A 'triplet' is defined here, somewhat loosely, as a pair of 180° orientation singularities with a 360° direction singularity in between. All three singularities have the same sign, and they are connected by a short narrow strip of low direction selectivity. This configuration has been observed in experimental maps (Shmuel & Grinvald 1996, fig. 14), although it appears that they are less common than they are in the model maps illustrated in figures 2 and 3. Factors affecting the density of triplets and their theoretical significance are discussed in more detail below.

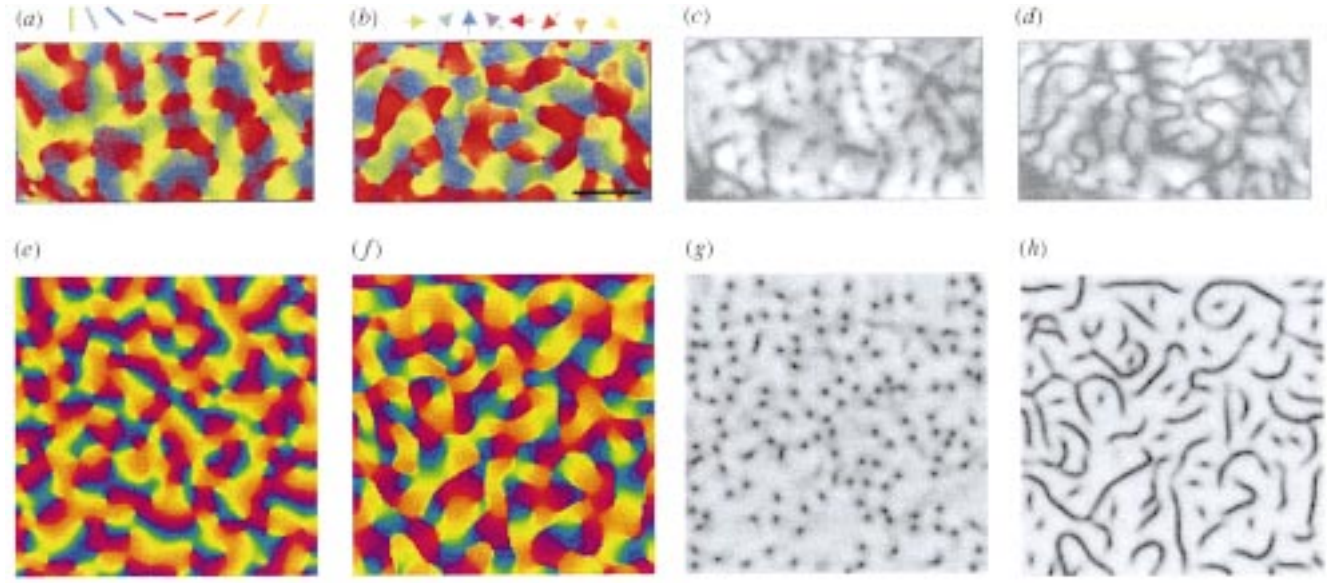

Figure 2. (*a–d*) Orientation and direction preference maps obtained by Weliky *et al*. (1996) in ferret visual cortex: (*a*) colour-coded map of orientation angle; (*b*) colour-coded map of direction angle; (*c*) grey-scale map of orientation vector magnitude, with dark spots indicating regions of low orientation preference or low responsiveness; (*d*) the corresponding map of direction preference, with dark lines indicating regions of low direction preference. The scale bar in 2*b* is 1 mm. (*e*,*f*) Model maps of (*e*) orientation angle, (*f*) direction angle, (*g*) orientation vector magnitude and (*h*) direction vector magnitude. Figures (*a–d*) are reproduced with permission from Weliky *et al*. (1996).

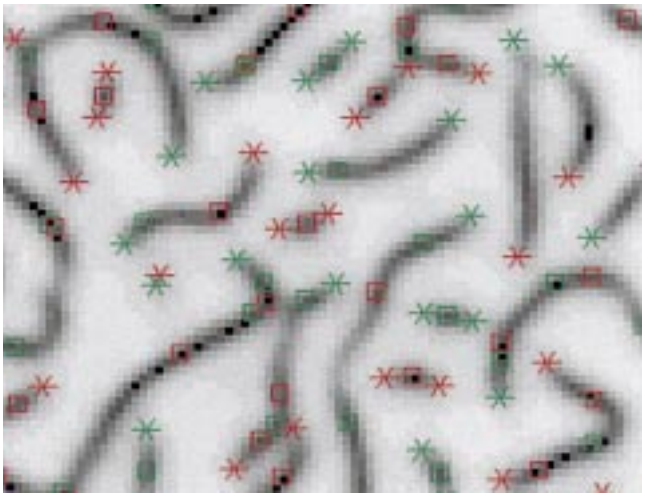

Figure 3. Portion of the map illustrated in figure 2*h*, showing grey-scale representation of direction modulus values (dark = low magnitude), together with the locations of the orientation singularities (asterisks) and direction singularities (squares). Positive singularities are shown in red and negative singularities in green. Note that two orientation singularities often straddle a direction singularity and all three have the same sign. This configuration is referred to here as a 'triplet'.

Figure 4 shows point-by-point correlations between several pairs of map variables. Orientation and direction preferences are usually orthogonal (figure 4*a*) and regions of rapid rates of change in direction preference (i.e. in the linear fractures) are confined to regions where the orientation gradient is small, and vice versa (figure 4*b*). There is a relatively weak relationship between retinal magnification factor and orientation gradient (figure 4*c*) and direction gradient values (figure 4*d*). This perhaps reflects the relatively uniform magnification within the retinotopic map (figure 5).

(b) *Developmental dynamics of the model*

Development during the initial stage of map formation required approximately 20–40 stimulus presentations per cortical array point. At the start of development, all orientation and direction vectors were small in magnitude, implying broad tuning and/or weak responses, and random in their spatial distribution. The first sign of an ordered map was the emergence of small patches with larger magnitude values; within each of these patches orientation and direction preferences were roughly constant. As development proceeded, the patches grew in size, maintaining their initial orientation and direction preferences. These aspects of the model behaviour resemble what has been observed experimentally in developing ferret visual cortex by Chapman *et al*. (1996). This stage was followed by a much longer one in which singularity density in both the orientation and direction maps declined roughly logarithmically with the total number of stimulus presentations (figure 6*a*), and the iso-orientation domains became increasingly straight and parallel. (Although the direction maps appear to contain line rather than point singularities, true line singularities are topologically unstable (S. Tanaka, personal communication), which means that they disappear if the direction values are perturbed slightly. This is not true of point singularities. In practice, it is therefore possible to identify and count point singularities in the direction maps, using the same method as for the orientation maps.) The decline in singularity density was caused by two processes: (i) annihilation caused by singularity pairs of opposite sign

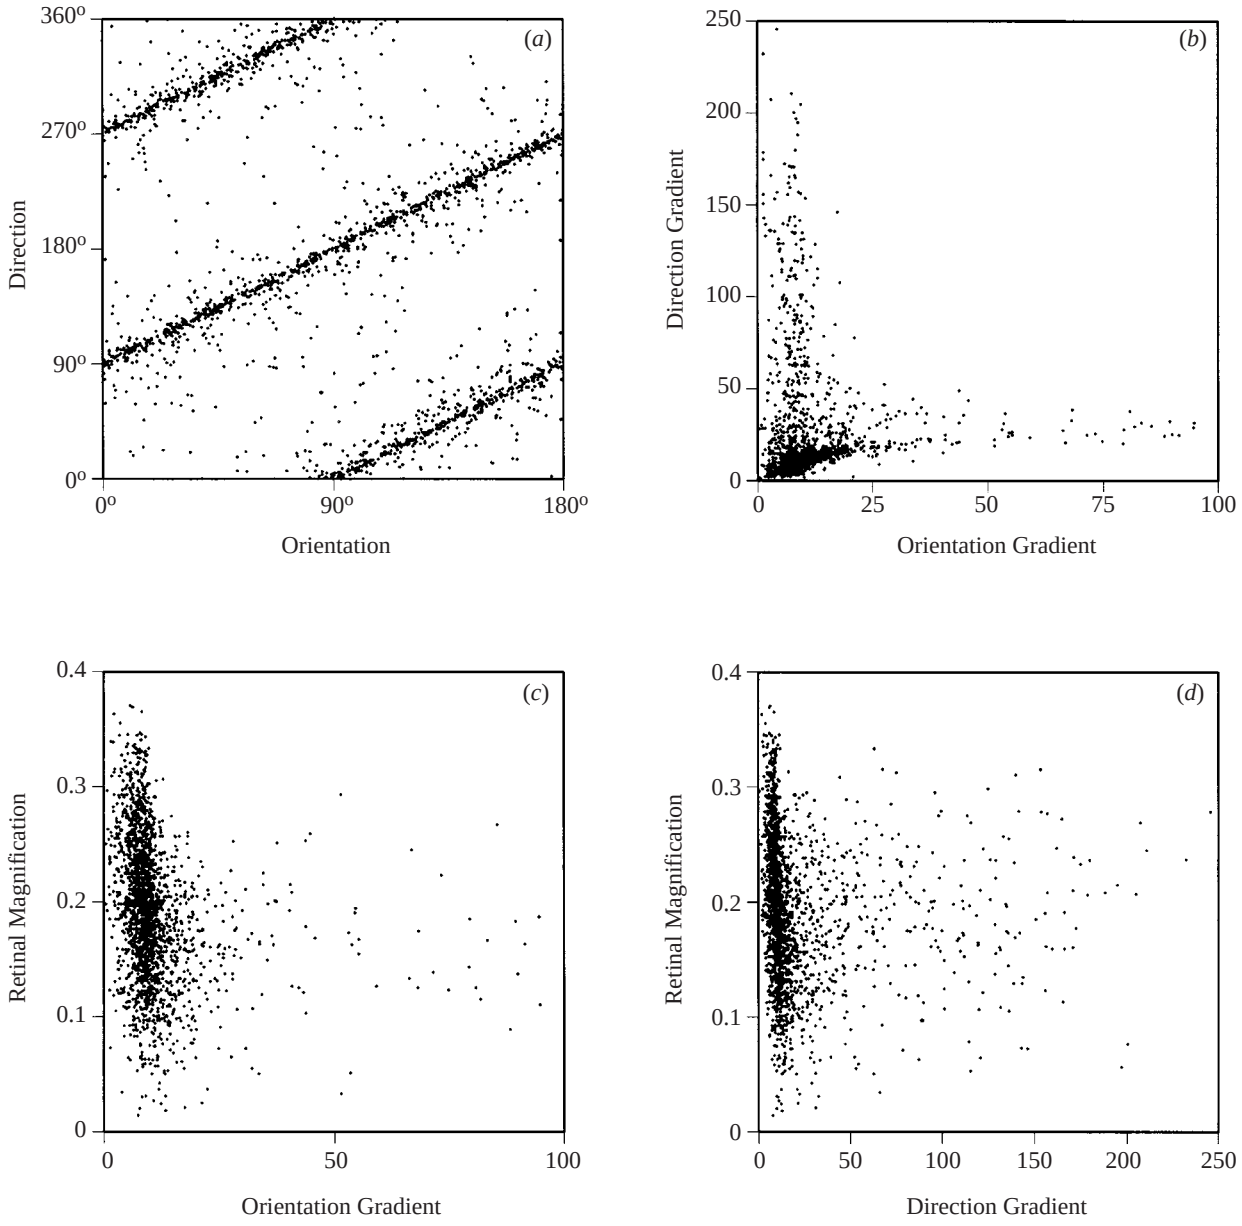

Figure 4. Scatter graphs showing correlations between (*a*) orientation and direction angles, (*b*) orientation and direction gradient values, (*c*) orientation gradient and retinal magnification factor and (*d*) direction gradient and retinal magnification factor. Points are taken from the map illustrated in figure 2. For clarity, not all of the points in the map are shown. Orientation gradient, grad($\theta_{k,l}$), was calculated as $(g(|\theta_{k+1,l} - \theta_{k,l}|_{180})^2 + g(|\theta_{k,l+1} - \theta_{k,l}|_{180})^2)^{1/2}$ where $g(\theta) = \theta$ if $\theta \leqslant 90°$; $g(\theta) = 180° - \theta$ if $\theta > 90°$. The direction gradient was calculated analogously. Retinal magnification was calculated as $((x_{k+1,l} - x_{k,l})^2 + (y_{k+1,l} - y_{k,l})^2 + (x_{k,l+1} - x_{k,l})^2 + (y_{k,l+1} - y_{k,l})^2)^{1/2}$, where x and y are the retinal elevation and azimuth, respectively, of the cortical point indexed by (k, l).

colliding, and (ii) singularities disappearing off the edge of the map (this was possible because periodic boundary conditions were not used). This was shown by plotting the minimum distance between singularities over the whole map against time: when this distance jumped upwards, the number of singularities decreased by two, indicating the annihilation of a singularity pair. The minimum distance between a singularity and the edge of the map was also measured in some simulations: when this distance jumped upwards, the number of singularities decreased by one. Disappearance off the edge of the map seemed to be a less frequent process than pair annihilation. Spontaneous creation of singularity pairs was also observed, and the gradual decline in singularity density with time reflected a relatively greater rate of annihilation versus creation.

(c) *Singularity distributions*

In a spatially random distribution of singularities, the percentage of nearest-neighbour pairs that have unlike signs should average 50%. Obermayer & Blasdel (1997) measured a higher percentage of about 80% in primate visual cortex, which suggests that singularities of opposite sign tend, like charged particles, to behave as though subject to an attractive force. In the present study, we found that when direction preference was omitted from

the map (by setting $R_\phi=0$) the percentage of nearest neighbour opposite sign pairs was similarly about 78%. Adding direction to the orientation map reduced the percentage. For the map shown in figure 2, in which $R_\phi=1$, the percentage was 61.5%, and this declined roughly linearly as R_ϕ increased (figure 6*b*). This value was not stable with time, however, and it continued to decrease with increasing numbers of stimulus presentations (figure 6*c*). This is probably the result of the continuing disappearance of nearby singularity pairs of opposite sign (see Wolf & Geisel 1997) and the persistence of same-sign pairs. Thus, over time, the total number of orientation and direction singularities decreases, and orientation singularities become increasingly grouped into pairs of like sign. These pairs are typically connected by a short direction discontinuity (figure 3), forming the triplet arrangements described above.

(d) *Energetic comparison of 180° and 360° singularities*

As suggested in § 1, overall smoothness constraints might lead one naively to expect that a combined map would contain 360° (i.e. 2π) singularities, as this would allow a continuous representation of direction preference without line singularities. Instead, both the simulation results and the biological data show that 180° (i.e. π) orientation singularities prevail, and that 360° singularities rarely, if ever, occur. In the following section we discuss in more detail the possible types of singularities in coupled orientation and direction maps. We then assess these, using an analytically solvable caricature model for point singularities.

The rationale behind the caricature model is to weigh smoothness against resolution in the representation of the orientation and direction variables in the neighbourhood of a singularity. To this end we neglect the retinal dimensions and focus on the two-dimensional torus that is spanned by the (periodic) orientation and direction variables. The stimuli used in our model do not cover the whole surface of the torus, but trace out a line defined by the orthogonality constraint $\phi = \theta \pm \pi/2$. Figure 7*a* shows an unfolded projection of this stimulus space onto the flat page. The stimulus manifold winds once around the direction dimension, and twice around the orientation dimension. If we now consider closed loops of neurons in the cortex, their orientation and direction variables trace out a corresponding closed loop on the torus. Loops around singularities are particularly interesting, as they are topologically equivalent to loops in the stimulus space, which cannot smoothly be shrunk down to a point. This is the case with loops that run around the torus one or more times in one or both dimensions, with a net change in angle along the loop which is a multiple of π or 2π.

Although the stimuli are constrained to have $\phi = \theta \pm \pi/2$, the cortical receptive fields may lie anywhere on the torus. As is the case in coupled phase-orientation maps (which have the stimulus space topology of a Klein bottle, see Tanaka (1995)), it is topologically possible to have orientation-only or direction-only singularities, with an increment of $\Delta\theta=n\pi$ around the loop at $\Delta\phi=0$ or vice versa. These $(n, 0)$ or $(0, m)$ loops would correspond to horizontal and vertical lines, respectively, in figure 7*a*. Even without a detailed analysis, one can see that the deviations between vertical or horizontal lines in figure 7*a*, and the oblique lines of the stimulus manifold, is quite large. Therefore, it seems unlikely that these $(n, 0)$ and $(0, m)$ singularities will ever occur.

The $(\pm1, \pm1)$ and $(\pm2, \pm1)$ singularities are more interesting. In the former case, one possible sequence of receptive fields would have a tight constant coupling between orientation and direction variables, plus a jump of $\Delta\phi=\pi$ at one point (dashed line in figure 7*b*). Another possible sequence around a $(\pm1, \pm1)$ singularity is drawn as a dotted line in figure 7*b*. Here, orientation and direction both change smoothly, but there is no constant angle between them. In the case of the $(\pm2, \pm1)$ singularity, a loop can be realized such that all receptive fields stay on the stimulus manifold (figure 7*a*). This sequence of receptive fields is perfectly smooth in both variables, and because it matches the stimulus manifold it might be expected to predominate. However, the greater degree of smoothness must be weighed against a coarser local discretization of the stimulus space (i.e. a greater spacing of the squares in figure 7*a* compared with the triangles and crosses in figure 7*b*).

To assess the impact of these competing objectives, we assume that the receptive fields along the respective loops are adjusted according to the SOFM algorithm. The relative amplitudes of these two dimensions are scaled by a parameter s, so distances in the stimulus space are given by $d(\Delta\theta,\ \Delta\phi)=(\Delta\theta)^2+s^2(\Delta\phi)^2$. These simple distances cannot be directly transformed to the distances of the full model described above (§ 2), owing to the trigonometric functions in the distance metric. However, a comparison of distances at $(\Delta\theta=\pi/2, \Delta\phi=\pi/2)$ shows that $s\approx R_\phi/\sqrt{2}R_\theta$. The positions to which the neuron ring is mapped are assumed, for symmetry reasons, to be equidistant on the lines on the torus which correspond to the different types of singularities, i.e. the dotted, dashed and solid lines of figure 7*a*,*b*.

The relative stability of the different states of the neuron ring can be assessed using an energy formalism for SOFMs (Riesenhuber *et al*. 1996; Bauer *et al*. 1997). We calculate the value of a distortion energy for the different possible states of the map, and assume that the SOFM learning dynamics (which are closely related to the gradient of the distortion energy) drive the map towards the state with the lowest energy. Because the receptive fields are equidistant on the neuron ring, we need consider the distortion energy for only one neuron, and do not need to sum over all map neurons. Furthermore, as the explicit values of the receptive fields are part of the *ansatz*, we can simply evaluate squared differences between receptive fields and stimuli, without using the further refinements described in Riesenhuber *et al*. (1996) and Bauer *et al*. (1997). Using continuum notation, the distortion energy for a neuron j^* is

$$E_{j^*} = \int \mathrm{d}\mathbf{v}(\mathbf{v} - \mathbf{w}_{j^*})^2\ h(j(\mathbf{v}), j^*), \qquad (3)$$

where the $(\mathbf{v}-\mathbf{w}_{j^*})^2$ term evaluates the match between a particular stimulus and the receptive field $\mathbf{w}_{j^*}$, and the $h(j(\mathbf{v}), j^*)$ term weighs the distance in map space between the neuron $j(\mathbf{v})$ to which $\mathbf{v}$ is mapped, and the test neuron, j^*. In this indirect way, the continuity of the map is enforced in the learning rule and is taken into account in the energy formalism.

When evaluating the energy integral (3) for the different map states, we can pick one neuron's receptive field, calculate the squared differences between the receptive field and stimuli on the manifold, and integrate over the manifold. Because the orientation and direction variables on the manifold are tightly coupled, this amounts to only one integration. A more tricky point is how to evaluate the $h(j, j^*)$ terms. Our basic assumption of equidistant distribution of receptive fields on the neuron ring means that we can transform distances in stimulus space directly into distances on the neuron ring. However, for the $(\pm 2, \pm 1)$ singularity, we have to take into account that the neuron ring runs around the orientation variable twice. Consequently, a given distance in the orientation variable of receptive fields corresponds to only half the distance between the neurons on the ring, as compared to the ring for the $(\pm 1, \pm 1)$ singularity. This yields terms in $\exp(x^2/8\sigma^2)$, whereas for the $(\pm 1, \pm 1)$ singularity the corresponding terms are in $\exp(x^2/2\sigma^2)$. In other words, around a $(\pm 2, \pm 1)$ singularity, neurons with a given orientation distance are only half as far apart as a pair of neurons with the same orientation distance, but in the neighbourhood of a $(\pm 1, \pm 1)$ singularity.

Scaling the length of the stimulus manifold to the intervals [0, 1] (orientation) and [−1, 1] (direction), we find that the distortion energy of the $(\pm 2, \pm 1)$ singularity is

$$\begin{aligned} E_{(\pm 2,\ \pm 1)} &= 2\int_0^{1/2} x^2(1+s^2)\exp(-x^2/8\sigma^2)\mathrm{d}x \\ &\quad + 2\int_{1/2}^{1}[(1-x)^2 + s^2x^2]\exp(-x^2/8\sigma^2)\mathrm{d}x \\ &= 2\{-(1+s^2)4\sigma^2\exp(-1/8\sigma^2) \\ &\quad +(1+s^2)4\sigma^3(\sqrt{2\pi})\mathrm{erf}(1/\sqrt{8}\sigma) \\ &\quad +(\sqrt{2\pi})\sigma[\mathrm{erf}(1/\sqrt{8}\sigma) - \mathrm{erf}(1/2\sqrt{8}\sigma)] \\ &\quad +8\sigma^2[\exp(-1/8\sigma^2) - \exp(-1/32\sigma^2)]\} \end{aligned} \tag{4}$$

while the minimal value for a discontinuous $(\pm 1, \pm 1)$ singularity is

$$\begin{aligned} E_{(\pm 1,\ \pm 1),\ \mathrm{discont}} &= 2\int_{-1/2}^{1/2}[x^2 + s^2(x+\tfrac{1}{2})^2]\exp(-x^2/2\sigma^2)\mathrm{d}x \\ &= 2\{-(1+s^2)\sigma^2\exp(-1/8\sigma^2) \\ &\quad +(\sqrt{2\pi})\sigma[(1+s^2)\sigma^2 + \tfrac{1}{4}s^2]\mathrm{erf}(1/\sqrt{8}\sigma)\}. \end{aligned} \tag{5}$$

Other realizations of neuron rings around a $(\pm 1, \pm 1)$ singularity, in particular the realization with smoothly varying, but uncoupled orientation and direction variables, yield analogous formulas, but with slightly larger energies than $E_{(\pm 1, \pm 1),\mathrm{discont}}$.

Numerical evaluation of equations (4) and (5) shows that when $s < s_{\mathrm{crit}} \approx 0.6$, $(\pm 1, \pm 1)$ singularities should predominate. Above s_{crit}, the smoother $(\pm 2, \pm 1)$ singularities should occur. This critical value of s depends only slightly on the map parameter σ (in contrast to other state transitions in SOFMs, where the critical parameter is proportional to σ (Obermayer *et al.* 1992*a*; Riesenhuber *et al.* 1996)). A ring with smoothly changing but uncoupled receptive field variables around a $(\pm 1, \pm 1)$ singularity systematically yields a higher energy than $E_{(\pm 1, \pm 1),\mathrm{discont}}$. Numerical simulations of SOFM learning in this caricature system also show a transition behaviour from $(\pm 1, \pm 1)$ to $(\pm 2, \pm 1)$ singularities with increasing s, roughly at s_{crit}. A more precise estimate of the transition point in the simulations is prohibited by the occurrence of distorted intermediate states near the transition. Yet, at small and large values respectively of s, the corresponding $(\pm 1, \pm 1)$ and $(\pm 2, \pm 1)$ singularities can clearly be identified.

To summarize, the above analysis of this caricature system shows that there is a regime of small direction amplitudes in which $(\pm 1, \pm 1)$ singularities are energetically more favourable than $(\pm 2, \pm 1)$ singularities, even though they offer suboptimal smoothness in the direction representation. They are preferred because they increase the resolution of the orientation variable. At large direction amplitudes, $(\pm 2, \pm 1)$ singularities are the energetically optimal solution, as would be expected if only smoothness constraints applied. Although the quantitative result of $s_{\mathrm{crit}} \approx 0.6$ should be taken with a grain of salt, considering the crudeness of the caricature model, the qualitative result of the existence of a parameter regime with discontinuous representations near singularities should not depend on technicalities of the above analysis. Indeed, the simulations of the full model showed many singularities of the $(\pm 1, \pm 1)$ type, in the limit of small R_ϕ/R_θ. It is tempting to identify the singularity triplets which occur in the full model at larger values of R_ϕ/R_θ with $(\pm 2, \pm 1)$ singularities, especially as their fraction increases with increasing s, as predicted by the above analysis. Although at present the reasons remain unclear as to why the distributed arrangement of singularities in a triplet should be preferred to a strictly local singularity, one can argue, nevertheless, that a loop enclosing a triplet of singularities is topologically equivalent to a local $(\pm 2, \pm 1)$ singularity.

4. DISCUSSION

(a) *Related models*

The results presented here suggest that cortical maps of direction preference can be added to the growing list of features of visual cortex organization which can be reproduced by relatively simple mathematical models. The most successful of these models achieve dimension-reducing mappings between Euclidian spaces of different dimension, subject to continuity and completeness constraints. For example, Durbin & Mitchison (1990) applied a model based on the elastic net algorithm (Durbin & Willshaw 1987) to a four-dimensional stimulus space similar to the six-dimensional one used here, but lacking the dimension of direction preference, and obtained realistic maps of orientation preference and retinotopy. Obermayer *et al.* (1991, 1992*b*) applied the same algorithm as that used here to a five-dimensional stimulus space, the dimensions of which were retinotopic position, orientation and eye of origin. They were able to reproduce the known structural patterns of orientation and eye dominance columns, and the experimentally observed (Blasdel 1992) orthogonal relationships between the two.

The stimulus spaces used in the present model, and others like it (Durbin & Mitchison 1990; Obermayer *et al.* 1991, 1992*b*), are highly simplified abstractions. More

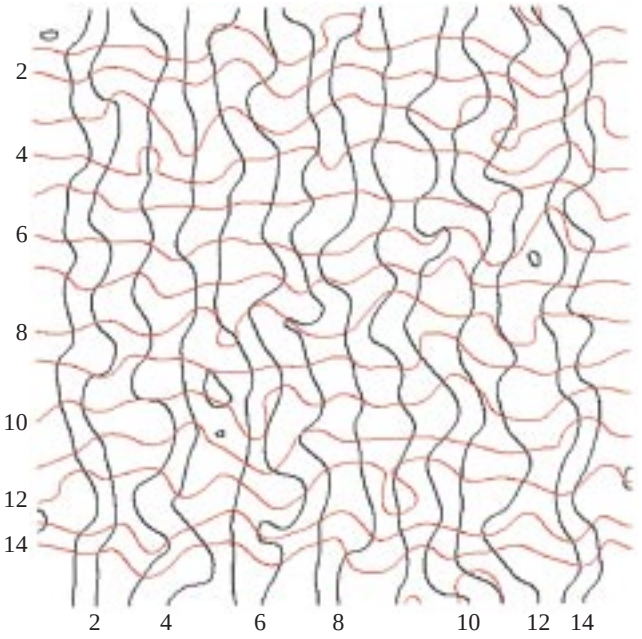

Figure 5. The map of visual field positions on the surface of the model cortex produced by the simulation illustrated in figure 2. Iso-azimuth lines (x values in stimulus space) run (notionally) vertically on the page; iso-elevation lines (y values) run horizontally. Numbers indicate the respective values of x and y in the model.

detailed versions of the models can be implemented in which the relevant stimulus dimensions are the activities of individual thalamic input neurons. In these models, stimuli are spatial patterns of activity in the input units, and cortical modifications can be directly interpreted in terms of changes in connection strengths between neurons. Models based on such 'high-dimensional' versions of the Kohonen algorithm have been applied to the problems of retinotopy, ocular dominance and orientation column formation (Obermayer *et al*. 1990; Goodhill 1993; Riesenhuber *et al*. 1998), and their behaviours have proved to be similar in most respects to those of the simpler low-dimensional ones.

Tanaka & Shinbata (1994) applied a model based on thermodynamic principles of self-organization to the organization of orientation and direction preference, and reproduced many of the structural features discussed here, including line singularities in direction preference and their association with orientation singularities: retinotopic location was not included as a parameter. Their model, like the present one, assumed the prior existence of direction and orientation selectivity, and only attempted to explain the spatial layout. However, it has been shown by mathematical analysis (Wimbauer *et al*. 1994) and by computer modelling (Miyashita *et al*. 1997) that if variation in temporal, as well as spatial, selectivity is included in the inputs to a linear, Hebbian feed-forward network, then direction selective, as well as orientation selective, receptive fields can emerge. This is possible with only random structure in the inputs to the first layer of the network. There is physiological evidence for the variation in temporal selectivity required by these models, as it has been found that ON and OFF centre receptive fields of lateral geniculate nucleus neurons can be subdivided on the basis of latency differences into 'lagged' and 'non-lagged' types (Mastronarde 1987; Saul & Humphrey 1990) and that both types of input combine to determine the receptive field properties of cortical simple cells (Saul & Humphrey 1992).

In the study of Miyashita *et al*. (1997), the spatial layout of orientation and direction preferences was similar to the experimental one, with point orientation singularities connected by line direction singularities. The similarities between the results of this model, which is derivable (Tanaka 1990) from a linear, correlation-based learning rule, and the present one, which is based on a nonlinear, competitive, learning rule, may not be coincidental, given that mathematical similarities between the two types of model can be demonstrated (Yuille *et al*. 1996). However, from a biological standpoint the models work differently, and an attempt to derive differing, experimentally testable predictions from them would be worthwhile.

(b) *Relationship between orientation gradient and retinal magnification factor*

In the dimension-reduction model of orientation and retinotopy studied by Durbin & Mitchison (1990) an inverse relationship between the orientation gradient and retinal magnification factor was found. The present model, which might have been expected to behave similarly, produced no immediately obvious relationship between these two variables (figure 4*c*). Possible reasons for this include the addition of two more dimensions to the stimulus space and the assumption of an initially relatively well-ordered retinal topography. Recently, Das & Gilbert (1997) reported a significant positive correlation between orientation gradient and retinal magnification factor in cat area 17. This correlation is strong enough that non-overlapping retinal receptive fields can be found on either side of an orientation singularity. This result cannot be explained by the present model, or by that studied by Durbin & Mitchison (1990). However, preliminary studies in the tree shrew (Bosking *et al*. 1997) and in macaque visual cortex (Campbell & Blasdel 1995) have suggested that retinal magnification factor in these species is locally uniform, and undistorted by variations in the orientation gradient map. These indications of species variability, if confirmed, may mean that a successful model, or class of models, may have to be able to produce more than one type of correlation with the retinotopic map. To do this, the effects of modifying the learning rule, using different parameter values and/or initial conditions, may need to be explored.

(c) *Sharpness of transition regions*

The present results show a loss of stimulus selectivity in the region of transition points, i.e. a loss of orientation selectivity close to singularities, and a loss of direction preference close to the direction fractures (figures 1 and 2*g*,*h*). The size of these regions of poor selectivity is determined by, and roughly equal to, the size of the cortical neighbourhood function, which was *ca*. 150 μm in the present simulations. Similar regions of poor selectivity are seen in maps obtained by optical recording experiments (figure 2*c*,*d*). However, optical recordings average the activities of neurons within regions about 150 μm wide, and would be expected to show gradual transitions even if

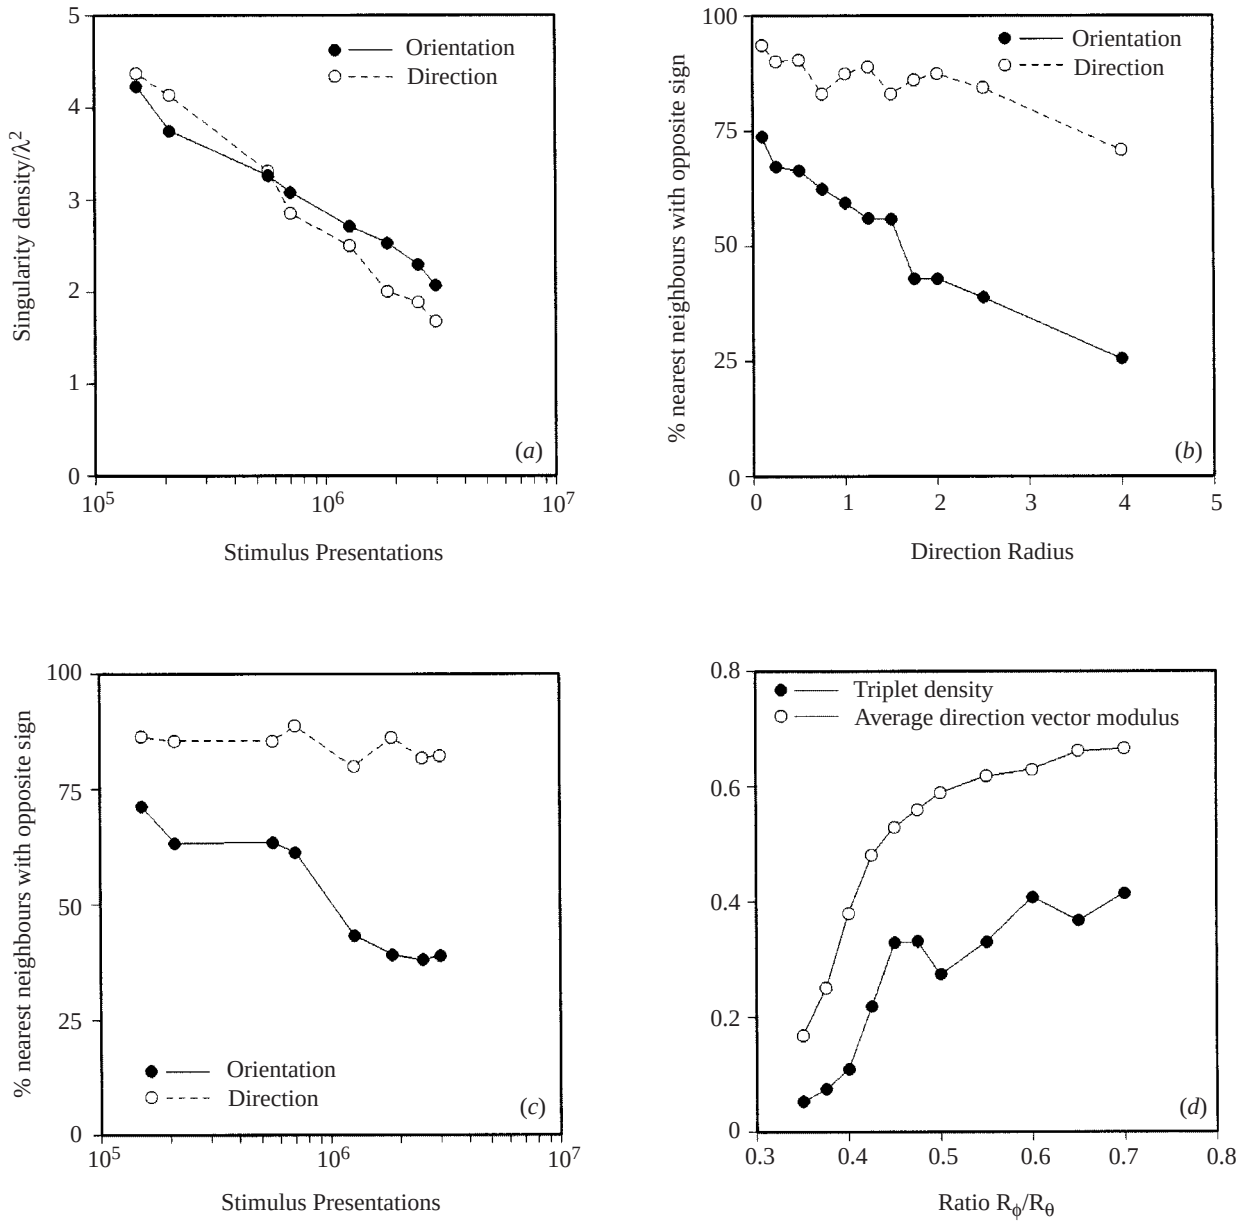

Figure 6. Measurements of orientation and direction singularity density and distribution as a function of the number of stimulus presentations and the value of direction radius. Unless otherwise stated, parameter values are the same as given for figure 2. (*a*) Change in the density of singularities, measured as the average number/λ^2, where λ is the spatial period in orientation angle, with number of stimulus presentations; (*b*) change in the percentage of nearest-neighbour singularities with opposite sign as a function of the direction radius R_ϕ; (*c*) change in the percentage of nearest-neighbour singularities with opposite sign as a function of the number of stimulus presentations; (*d*) effect of the ratio R_ϕ/R_θ on the density of triplets. In this calculation, R_ϕ and R_θ were varied reciprocally so that the function $(R_\phi \times R_\theta = 1)$ always applied, and the number of stimulus presentations was 2×10^6. Triplets were counted by an algorithm that first identified an orientation singularity and then tracked along the region of low direction modulus extending from it, until a second orientation singularity was found. A record was kept of the number of direction singularities encountered en route. A triplet was counted if (i) the distance between the two orientation singularities was less than 16 array units; (ii) only one direction singularity was counted; and (iii) all three singularities had the same sign. The proportion plotted here is the proportion of orientation singularities found in a triplet arrangement, i.e. twice the number of triplets divided by the total number of orientation singularities. The average value of the direction vector modulus is also shown, expressed as a proportion of R_ϕ.

the selectivities of individual neurons were maintained close to transition points. Tetrode recordings from the cat visual cortex (Maldonado & Grey 1996; Maldonado *et al.* 1997) have shown that neurons close to singularities have just as sharp orientation tuning as those further away, suggesting that transitions at point singularities can be very rapid. It is uncertain whether changes in direction preference across line reversals are similarly rapid, or whether significant numbers of non-direction selective cells are found in these regions. The present model can be made to produce sharp transitions by making the cortical neighbourhood function get smaller ('annealing') as development proceeds, although, for simplicity, we chose not to do this. Biologically, such a decrease might be

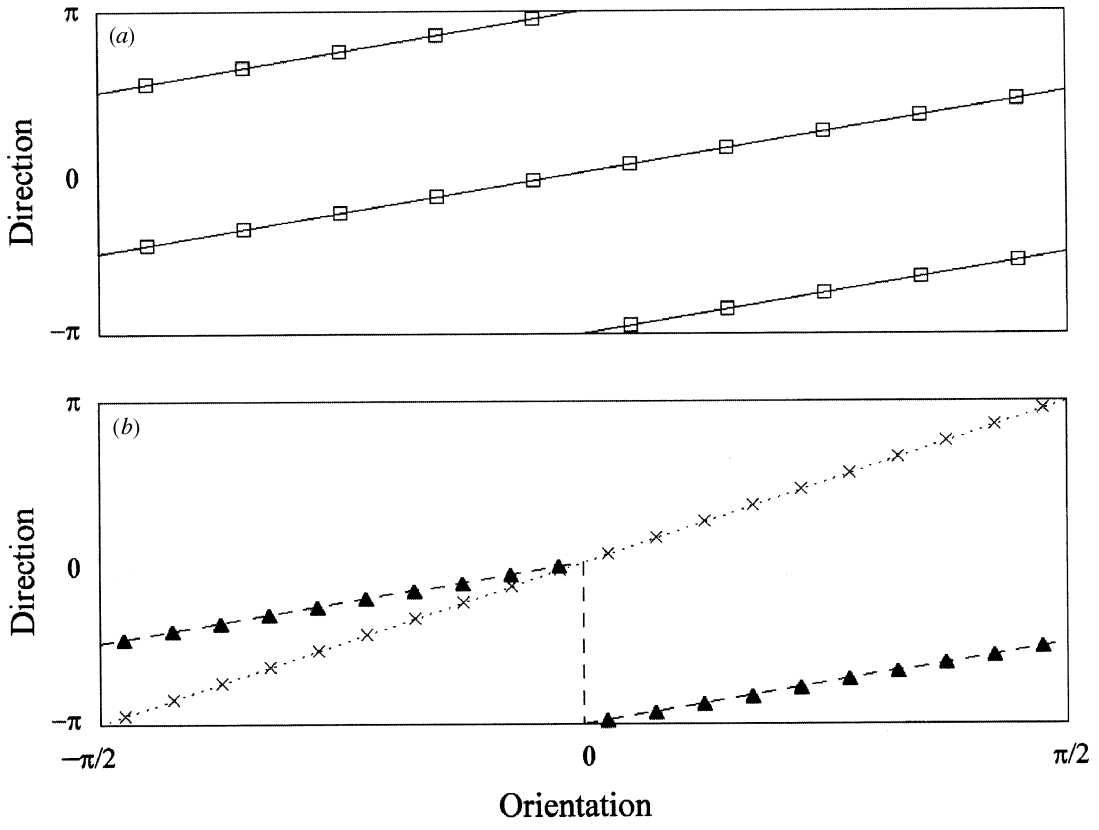

Figure 7. Diagrammatic unfolding of the stimulus space used to represent orientation and direction angles. Ignoring the dimensions of retinotopy, the space spanned by the cortex can be represented as the product of two circles, one for orientation and the other for direction, with circumferences (in angular units) of π and 2π, respectively. This can be represented by a torus, or equivalently as a rectangle with periodic boundaries (i.e. left and right edges are continuous, as are upper and lower edges). The locus of points for which direction preference is orthogonal to orientation preference (the stimulus manifold) is a series of diagonal lines on such a graph (cf. figure 4*a*). A loop on the cortical surface surrounding a singularity will map to a series of lines in this representation. (*a*) The projection of a loop around a $(\pm 2, \pm 1)$ singularity, with orientation and direction accumulating an angle of 2π along the circuit. Possible receptive field centres are shown as squares. (*b*) Two alternative types of $(\pm 1, \pm 1)$ singularity: a 'smooth' one, in which orthogonality is violated (dotted line and crosses), and a discontinuous one (dashed line and triangles), in which direction preferences flip.

attributed to physical growth of the cortex and/or a decline in the ability of small molecules to diffuse through the extracellular space.

(d) *Types and distribution of singularities*

We found that if a random stimulus sequence is used as the input to the model, it does not, in practice, converge to a fixed stable state. With continued stimulus presentations, the number of orientation and direction singularities slowly decreased, and both the iso-orientation and iso-direction domains became increasingly parallel. This suggests that the final stable state of the system (Ritter & Schulten 1988) consists of parallel iso-orientation and iso-direction stripes, with two sets of orientation domains for every complete set of direction domains. But this global arrangement was never obtained, even after very large numbers of stimulus presentations, and, although an appropriate time-scale for comparison remains uncertain, it is unlikely that it would be reached within the lifetime of any animal.

The possibility that visual cortical maps might change slowly with age has received little experimental attention. Obermayer & Blasdel (1997) found similar densities of singularities in immature and adult macaque monkeys, which suggests that large changes do not occur. A recent preliminary study (Rubin *et al.* 1997) showed that singularity density in ferrets increases between the postnatal period and puberty at about nine months of age. However, this increase was attributed to a decrease in the overall area of visual cortex, which suggests that the actual spatial pattern may have remained relatively unchanged. Without further evidence, it is always possible to assume that some factor such as a decline in the availability of neurotrophins, or an age-related decrease in Hebbian modifiability, freezes the cortical map in what is technically an intermediate or non-stable state.

The energy analysis shows that the preponderance of discontinuous $(\pm 1, \pm 1)$ singularities observed in the simulations can be understood in terms of a preference of the map formation process for good local resolution, as opposed to optimal smoothness. A balance of this kind is also achieved by the elastic net algorithm, where there is a term in the energy function (Durbin & Willshaw 1987) which is proportional to the squared differences between the receptive field parameters of neighbouring neurons. As in the framework presented here, this allows for a

trade-off between smoothness in orientation and direction variables, or in the orientation variable alone. An interesting aspect of the analysis is that the transition point depends only slightly on the width, σ, of the cortical neighbourhood function. This parameter effectively controls the degree of continuity in the resulting map. In other investigations of transitions between map states (binocular versus monocular, non-orientated versus orientated), critical values of stimulus parameters have been found to depend linearly on σ (Obermayer *et al*. 1992*a*). The qualitatively different behaviour in the present case can be attributed to the fact that a change of σ affects continuity of orientation and direction in the same way. The trade-off being made in our model when deciding between different types of singularity is between resolution, and the relative continuity of orientation and direction, which is dependent on *s*.

(e) *Predictions*

The model presented here makes a minimal set of assumptions about the underlying map formation process. These are (i) that there is a pre-existing, reasonably accurate retinal topography; (ii) that initial orientation and direction preferences are random and small in magnitude; and (iii) the map develops according to a learning rule (Kohonen 1982) that can be used to ensure that some combination of local smoothness and global completeness of map parameters is maximized. As smoothness and completeness are opposing traits that do not have unique definitions and can be traded off in different ways (Swindale 1996*b*), it is possible, in principle, that other models implementing the same set of assumptions might produce maps with different structural properties than the present one. At present, however, analytical understanding of the different models is limited and it is not always easy to predict or understand why certain models produce certain features. Quantitative and/or qualitative comparison of computer-generated maps and real ones is, therefore, a valuable way of assessing the validity of different models.

Based on the analyses that have been done so far, the following points of comparison between model and experimental data can be suggested.

1. There should be an inverse relationship between the orientation and direction angle gradient maps (figure 4*b*), i.e. regions of highest orientation gradient should occur exclusively in regions where the direction gradient is low, and vice versa. (A. Shmuel and A. Grinvald (personal communication) have obtained data showing a relationship between orientation and direction gradient values which is similar to that shown in figure 4*b*.)
2. The distribution of orientation singularities in different cortical maps may vary depending on the strength of the direction signal that is represented in the map. One simple way of measuring this distribution (Obermayer & Blasdel 1997) is to calculate the percentage of nearest-neighbour singularity pairs that have opposite signs. In maps in which direction preference is absent, or relatively weak (as seems to be the case in area 17 of the monkey, and in areas 17 and 18 of the cat), this percentage is likely to be around 80% (Obermeyer & Blasdel 1997). The percentage should be lower (e.g. 50% or less) in maps in which direction selectivity is more pronounced, for example in area MT of the primate visual cortex.
3. The distinctive triplet arrangement of two orientation singularities abutting a direction singularity, all of the same sign, should be a relatively common feature in maps in which direction preference is strong, and relatively less common otherwise.

We thank Graeme Mitchison for his comments on an earlier version of the manuscript. This work was supported by a grant from the Medical Research Council of Canada, and by the Deutsche Forschungsgemeinschaft, SFB 185, TP E6. H.U.B. also acknowledges the kind hospitality of the Institute for Theoretical Physics, Santa Barbara.

REFERENCES

Albus, K. 1975 A quantitative study of the projection area of the central and paracentral visual field in area 17 of the cat. II. The spatial organization of the orientation domain. *Expl Brain Res*. **24**, 181–202.

Albright, T. D., Desimone, R. & Gross, C. G. 1984 Columnar organization of directionally selective cells in visual area MT of the macaque. *J. Neurophysiol*. **51**, 16–31.

Bauer, H.-U., Brockmann, D. & Geisel, T. 1997 Analysis of ocular dominance pattern formation in a high-dimensional self-organizing map model. *Network* **8**, 17–33.

Blasdel, G. G. 1992 Orientation selectivity, preference, and continuity in monkey striate cortex. *J. Neurosci*. **12**, 3139–3161.

Blasdel, G. G. & Salama, G. 1986 Voltage sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* **321**, 579–585.

Bonhoeffer, T. & Grinvald, A. 1991 Orientation columns in cat are organized in pin-wheel like patterns. *Nature* **353**, 429–431.

Bosking, W. H., Crowley, J. C. & Fitzpatrick, D. 1997 Fine structure of the map of visual space in the tree shrew striate cortex revealed by optical imaging. *Soc. Neurosci. Abstr*. **23**, 1945.

Campbell, D. & Blasdel, G. G. 1995 Optical measurement of cortical magnification factors in new and old world primates. *Soc. Neurosci. Abstr*. **21**, 771.

Chapman, B., Stryker, M. P. & Bonhoeffer, T. 1996 Development of orientation preference maps in ferret primary visual cortex. *J. Neurosci*. **15**, 6443–6453.

Das, A. & Gilbert, C. D. 1997 Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature* **387**, 594–598.

Durbin, R. & Mitchison, G. 1990 A dimension reduction framework for understanding cortical maps. *Nature* **343**, 644–647.

Durbin, R. & Willshaw, D. J. 1987 An analogue approach to the travelling salesman problem using an elastic net method. *Nature* **326**, 698–691.

Erwin, E., Obermayer, K. & Schulten, K. 1995 Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. *Neural Comput*. **7**, 425–468.

Goodhill, G. J. 1993 Topography and ocular dominance: a model exploring positive correlations. *Biol. Cybern*. **69**, 109–118.

Hildreth, E. C. 1984 *The measurement of visual motion*. Cambridge, MA: MIT Press.

Hubel, D. H. & Wiesel, T. N. 1962 Receptive fields, binocular interaction, and functional architecture of cat striate cortex. *J. Physiol., Lond*. **160**, 106–154.

Hubel, D. H. & Wiesel, T. N. 1968 Receptive fields and functional architecture of monkey striate cortex. *J. Physiol., Lond*. **195**, 215–243.

Hubel, D. H. & Wiesel, T. N. 1974 Sequence regularity and geometry of orientation columns in the monkey striate cortex. *J. Comp. Neurol.* **158**, 267–294.

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 Networks* **6**, 895–905.

Kohonen, T. 1995 *Self-organizing maps*. Berlin and Heidelberg: Springer.

Maldonado, P. E., Gray, C. M. 1996 Heterogeneity in local distributions of orientation-selective neurons in the cat primary visual cortex. *Vis. Neurosci.* **13**, 509–516.

Maldonado, P. E., Gödecke, I., Gray, C. M. & Bonhoeffer, T. 1997 Orientation selectivity in pinwheel centers in cat striate cortex. *Science* **276**, 1551–1555.

Malonek, D., Tootell, R. B. & Grinvald, A. 1994 Optical imaging reveals the functional architecture of neurons processing shape and motion in owl monkey area MT. *Proc. R. Soc. Lond.* B **258**, 109–119.

Mastronarde, D. N. 1987 Two classes of single-input X-cells in cat lateral geniculate nucleus. 1. Receptive field properties and classification of cells. *J. Neurophysiol.* **57**, 357–413.

Miyashita, M., Kim, D.-S. & Tanaka, S. 1997 Cortical direction selectivity without directional experience. *Neuroreport* **8**, 1187–1191.

Mountcastle, V. B. 1957 Modality and topographic properties of single neurons of cat's somatic sensory cortex. *J. Neurophysiol.* **20**, 408–434.

Obermayer, K. & Blasdel, G. G. 1997 Singularities in primate orientation maps. *Neural Comput.* **9**, 555–575.

Obermayer, K., Ritter, H. & Schulten, K. 1990 A principle for the formation of the spatial structure of cortical feature maps. *Proc. Natn. Acad. Sci. USA* **87**, 8345–8349.

Obermayer, K., Blasdel, G. G. & Schulten, K. 1991 Neural network model for the formation and for the spatial structure of retinotopic maps, orientation- and ocular dominance columns. In *Artificial neural networks* (ed. T. Kohonen, K. Mäkisara, O. Simula and J. Kangas), pp. 505–511. Elsevier.

Obermayer, K., Blasdel, G. G. & Schulten, K. 1992*a* Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Phys. Rev.* A **45**, 7568–7589.

Obermayer, K., Ritter, H. & Schulten, K. J. 1992*b* A model for the development of the spatial structure of retinotopic maps and orientation columns. *IEICE Trans. Fundamentals* **E75-A**, 537–545.

Payne, B. R., Berman, N. & Murphy, E. H. 1981 Organization of direction preferences in cat visual cortex. *Brain Res.* **211**, 445–450.

Riesenhuber, M., Bauer, H.-U. & Geisel, T. 1996 Analyzing phase transitions in high dimensional self-organizing maps. *Biol. Cybern.* **75**, 397–407.

Riesenhuber, M., Bauer, H.-U., Brockmann, D. & Geisel, T. 1998 Breaking rotational symmetry in a self-organizing map model for orientation map development. *Neural Comput.* **10**. (In the press.)

Ritter, H. & Schulten, K. 1988 Convergence properties of Kohonen's topology preserving maps: fluctuations, stability and dimension selection. *Biol. Cybern.* **60**, 59–71.

Rubin, B. D., White, L. E. & Fitzpatrick, D. 1997 Late developmental changes in the map of orientation preference in ferret visual cortex. *Soc. Neurosci. Abstr.* **23**, 2059.

Saul, A. B. & Humphrey, A. L. 1990 Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *J. Neurophysiol.* **64**, 206–224.

Saul, A. B. & Humphrey, A. L. 1992 Evidence of input from lagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 in the cat. *J. Neurophysiol.* **68**, 1190–1208.

Shmuel, A. & Grinvald, A. 1996 Functional organization for direction of motion and its relationship to orientation maps in cat area 18. *J. Neurosci.* **16**, 6945–6964.

Swindale, N. V. 1982 A model for the formation of orientation columns. *Proc. R. Soc. Lond.* B **215**, 211–230.

Swindale, N. V. 1996*a* Looking into a Klein Bottle. *Curr. Biol.* **6**, 776–779.

Swindale, N. V. 1996*b* The development of topography in the visual cortex: a review of models. *Network* **7**, 161–247.

Swindale, N. V., Matsubara, J. A. & Cynader, M. S. 1987 Surface organization of orientation and direction selectivity in cat area 18. *J. Neurosci.* **7**, 1414–1427.

Tanaka, S. 1990 Theory of self-organization of cortical maps: mathematical framework. *Neural Networks* **3**, 625–640.

Tanaka, S. 1995 Topological analysis of point singularities in stimulus preference maps of the primary visual cortex. *Proc. R. Soc. Lond.* B **261**, 81–88.

Tanaka, S. & Shinbata, H. 1994 Mathematical model for self-organization of direction columns in the primate middle temporal area. *Biol. Cybern.* **70**, 227–234.

Tolhurst, D. J., Dean, A. F. & Thompson, I. D. 1981 Preferred direction of movement as an element in the organization of cat visual cortex. *Expl Brain Res.* **44**, 340–342.

Weliky, M., Bosking, W. H. & Fitzpatrick, D. 1996 A systematic map of direction preference in primary visual cortex. *Nature* **379**, 725–728.

Wimbauer, S., Gerstner, W. & van Hemmen, J. L. 1994 Emergence of spatio-temporal receptive fields and its application to motion detection. *Biol. Cybern.* **72**, 81–92.

Wolf, F. & Geisel, T. 1997 Activity-dependent self-organization of orientation preference predicts a transient overproduction of pinwheels during visual development. In *Computational neuroscience: trends in research* (Proceedings of CNS 96) (ed. J. Bower), pp. 879–884. New York: Plenum Press.

Yuille, A. L., Kolodny, J. A. & Lee, C. W. 1996 Dimension reduction, generalized deformable models and the development of ocularity and orientation. *Neural Networks* **9**, 309–319.

As this paper exceeds the maximum length normally permitted, the authors have agreed to contribute towards production costs.