

A Computational Model for the Development of Multiple Maps in Primary Visual Cortex

Miguel Á. Carreira-Perpiñán^{1,2}, Richard J. Lister¹ & Geoffrey J. Goodhill¹

¹Department of Neuroscience, Georgetown University Medical Center, Washington DC 20007

²Present address: Department of Computer Science, University of Toronto, Toronto, Ontario, Canada M5S 3H5

Abstract

Primary visual cortex contains multiple maps of features of the visual scene, including visual field position, orientation, direction, ocular dominance and spatial frequency. The complex relationships between these maps provide clues to the strategies the cortex uses for representing and processing information. Here we simulate the combined development of all these map systems using a computational model, the elastic net. We show that this model robustly produces combined maps of these four variables bearing a close resemblance to experimental maps. In addition we show that the experimentally observed effects of monocular deprivation and single-orientation rearing can be reproduced in this model, and we make some testable predictions. These results provide strong support for the hypothesis that cortical representations attempt to optimize a trade-off between coverage and continuity.

Introduction

The receptive fields of neurons in the primary visual cortex of cats, monkeys and ferrets are selective for several different features of the visual input. These include visual field position (VF), ocular dominance (OD), orientation (OR), direction of motion (DR), and spatial frequency (SF) (Hubel and Wiesel, 1977; Bonhoeffer and Grinvald, 1991; Shmuel and Grinvald, 1996; Weliky et al., 1996; Shoham et al., 1997; Hübener et al., 1997; Kim et al., 1999). The preferred stimulus for each variable tends to change smoothly as one moves across the surface of the cortex, with occasional discontinuities. This can be thought of as a “dimension-reducing mapping”, whereby a space of features containing multiple dimensions (VF, OD, OR, DR, SF, etc) is compressed onto an essentially two-dimensional sheet of neurons. The result is a set of coexisting maps of visual features with a rich set of geometrical relationships. For instance, OR pinwheels tend to lie in the center of OD columns, and maps tend to intersect at right angles (Bartfeld and Grinvald, 1992; Obermayer and Blasdel, 1993; Hübener et al., 1997; Kim et al., 1999). It is generally believed that the source of these relationships is the mechanisms by which visual cortical maps are generated during development (e.g. Wolf and Geisel, 1998). Thus, the structure of these maps provides an important constraint on theories for how maps form.

Several computational models have made important contributions with regard to understanding the structure of OD and OR maps (reviewed in Erwin et al., 1995; Swindale, 1996; Carreira-Perpiñán and Goodhill, 2002). In particular, “low dimensional” or “feature space” models (Durbin and Mitchison, 1990; Goodhill and Willshaw, 1990; Obermayer et al., 1992; Goodhill et al., 1997; Swindale and Bauer, 1998; Goodhill and Cimponeriu, 2000; Carreira-Perpiñán and Goodhill, 2004b), generate maps which closely reproduce the structure of OR and OD maps (Erwin et al., 1995; Swindale, 1996). However, the application of these models has been limited up to now to only one or two features (OD and/or OR) in addition to VF. This naturally raises three questions: can such models still successfully reproduce map structure when multiple features are considered; what quantitative predictions do these models make in this case; and what light does the behavior of models in this case shed on the biological mechanisms underlying map formation. In this paper we address these questions by simulating the combined development of VF, OD, OR, DR and SF using the elastic net model (Durbin and Willshaw, 1987). This is a feature space model which works by minimizing an objective function that explicitly trades off coverage versus continuity constraints (Swindale, 1996). In addition, we show that in this multiple map case the elastic net can also reproduce recent experimental results regarding monocular deprivation (Crair et al., 1997) and single orientation rearing (Sengpiel et al., 1999). Taken together, our results extend the range of phenomena to which computational modeling of map formation in visual cortex has been applied. Most importantly, they provide strong additional support for the hypothesis that the computational principles underlying algorithms such as the elastic net capture the essential biological principles underlying map formation in the real cortex.

Materials and Methods

The elastic net algorithm. The elastic net (Durbin and Willshaw, 1987; Durbin and Mitchison, 1990; Carreira-Perpiñán and Goodhill, 2002) produces maps that minimize a tradeoff $E = \alpha C + \frac{\beta}{2} R$ between coverage C of the stimulus space and continuity R of the cortical representation. The *coverage* term is defined as follows:

$$C(\mathbf{y}_1, \dots, \mathbf{y}_M; K) = -K \sum_{n=1}^N \log \sum_{m=1}^M e^{-\frac{1}{2} \left\| \frac{\mathbf{x}_n - \mathbf{y}_m}{K} \right\|^2} \quad (1)$$

where \mathbf{x}_n represents a stimulus space vector, \mathbf{y}_m an elastic net centroid (stimulus preference of neuron m), and K a receptive field size ($\|\cdot\|$ represents the Euclidean norm, or length, of a vector). Note we use the term “receptive field” to refer to all features jointly (VF, OD, OR, DR, SF), and not only to the visual field. The *continuity* term was originally defined as

$$R(\mathbf{y}_1, \dots, \mathbf{y}_M) = \sum_m \|\mathbf{y}_{m+1} - \mathbf{y}_m\|^2 \quad (2)$$

and that is the definition used here. In Carreira-Perpiñán and Goodhill (2004b) we consider the effect on OR and OD maps of generalized continuity terms, and show that definitions other than (2) result in specific distortions of the geometric relationships of the maps (see also Carreira-Perpiñán and Goodhill, 2004a). The positive ratio α/β controls the relative strength of the continuity versus the coverage terms. Biologically plausible maps arise for a range of values of α/β . The net consists of a square lattice with M centroids, representing a square array of cortical neurons (the depth dimension of the cortex is not modeled).

The stimulus space. The stimulus space consists of feature points evenly spaced along orthogonal feature dimensions. For visual field (VF_x, VF_y) the stimulus values (training set) are a grid of $N_x \times N_y$ points in the rectangle $[0, 1] \times [0, 1]$. Orientation preference and selectivity (OR, ORr) are represented by adding dimensions of $N_{OR} \times 1$ points in $[-\frac{\pi}{2}, \frac{\pi}{2}] \times \{r_{OR}\}$, where OR is conventionally represented by two variables in polar coordinates, with N_{OR} values uniformly arranged on a ring of radius r_{OR} . Direction preference and selectivity (DR, DRr) are represented by adding a dimension of points $\{\text{OR} - \frac{\pi}{2}, \text{OR} + \frac{\pi}{2}\}$ on a ring of radius r_{DR} , as in Swindale and Bauer (1998). The ocular dominance (OD) dimension has N_{OD} values in $[-l_{OD}, l_{OD}]$. Spatial frequency (SF) is represented by N_{SF} values in $[-l_{SF}, l_{SF}]$. Consistent with several experimental reports we take $N_{SF} = 2$, representing “high” and “low” frequencies (Shoham et al., 1997; Hübener et al., 1997; Kim et al., 1999). (It should be noted that others have suggested the representation of SF may be more continuous (Silverman et al., 1989; Everson et al., 1998; Issa et al., 2000)).

Training regime. We trained nets with the following configuration. Training set: $N_x = N_y = 20$, $N_{OR} = 6$, $N_{DR} = 2$, $N_{OD} = 2$, $N_{SF} = 2$ (a total of 19200 training points); $l_{OD} = 0.06$ ($l_{OD} = 0.08$ for strabismus), $l_{SF} = 0.06$, $r_{OR} = 0.08$, and $r_{DR} = 0.08$. Elastic net: $M = 128 \times 128 = 16384$ centroids; $\alpha = 1$, $\beta = 10$; nonperiodic boundary conditions. The net was started from a random initial starting configuration with some global topographic bias; where results have been averaged over multiple simulations this corresponds to taking different random initial starting configurations for the net. As is common in the elastic net, the minimization of the energy is interleaved with decreasing (annealing) K . We annealed K from a starting point of 0.2 with a rate of 0.9925 to the point at which the maps have just arisen ($K \approx 0.03$). An efficient minimization method based on Cholesky factorization was used (Carreira-Perpiñán and Goodhill, 2004b). The simulations were performed using custom software written in Matlab.

Monocular deprivation and single-orientation rearing. For the default case (no deprivation) the coverage term α was 1 for all stimulus points \mathbf{x}_n . However, in order to implement deprivation of some stimulus (e.g. OD), α was redefined as a vector with N components, where the value of the component α_n represents the relative strength with which the stimulus point \mathbf{x}_n is represented in the input. Thus, for monocular deprivation we took $\alpha_n = \text{dep}_{OD} \in (0, 1)$ for each \mathbf{x}_n matching the deprived eye. In order to implement enhancement of OR for single orientation rearing, we took $\alpha_n = \text{dep}_{OR} > 1$ for each \mathbf{x}_n matching the enhanced angle of OR. In order to capture the experimental observation that SF maps consist of patches of low frequency in a broader expanse of high frequency (Shoham et al., 1997), one value of SF was “deprived” with $\text{dep}_{SF} = 0.5$ for all simulations.

Column width. The power of the discrete Fourier transform of the angle maps was roughly isotropic and concentrated around a ring. We summarized it by the mean wavelength (mean column width), averaged over all directions.

Crossing angles. We disregarded the points lying 5 pixels or less from the boundaries of the net (a thin inner stripe framing the maps in e.g. Fig. 1) to eliminate border effects. At each pixel, we computed the angle between the gradient vectors of pairs of dimensions (see below) and mapped it to $[0^\circ, 90^\circ]$, thus obtaining the angle between contours for the dimensions. Each such angle counted with a weight proportional to the product of the gradient moduli of the dimensions, e.g. $\|\nabla \text{OR}\| \times \|\nabla \text{OD}\|$ for OR and OD, so that pixels lying in areas of either constant OR or constant OD (i.e., away from borders of OR or contours of OD) were effectively removed

from the histogram. This is because, in an area where either map is nearly constant, the gradient vector is negligible in modulus and its direction basically arbitrary; unlike along borders, where the gradient vector is large and its direction well defined. We refer to the graphs we plot as histograms, even though in a strict sense they are not.

Pinwheel distributions. An OR singularity, or pinwheel, is defined positive if the orientation angle increases in a clockwise direction around the pinwheel and negative if anticlockwise. Pinwheels were automatically located in the OR maps as follows. First, the winding number of each pixel in the OR map was estimated by summing the increments of OR angle (in $[-\frac{\pi}{2}, \frac{\pi}{2}]$) along a closed path (a square of radius 1 pixel centered in the pixel considered) in a clockwise direction and dividing the result by 2π ; this results in 0 for non-pinwheel points and $+\frac{1}{2}$ ($-\frac{1}{2}$) for pixels at or adjacent to a positive (negative) pinwheel, respectively. The exact pinwheel location was obtained by grouping clusters of nonzero winding number and computing their centers of mass. To quantify the layout of pinwheels with respect to OD or SF columns, we computed, for every pinwheel, the distribution of distances of a given pinwheel to its closest OD or SF border. We prefer this to calculating distances to OD centers since the “center of an OD/SF column” is generally not well defined, except for columns which are translations of each other, and becomes difficult to use with columns of changing width, forks, islands or other complex shapes. In contrast, the OD/SF borders are well defined in all these cases. To compute the distances between a pinwheel and its closest OD or SF border, we represented the OD/SF borders by a finely spaced collection of points (the contours generated by Matlab for the OD/SF map); the said distance was then given by the point in this collection closest to the given pinwheel. We considered that a pinwheel lay on an OD/SF border if the distance between the two was 1 pixel or less.

Results

Multiple maps

In a feature space model such as the elastic net, feature dimensions are represented as spatial dimensions in a Euclidean space. Thus there are for instance two dimensions for VF, and one for OD. Since OR is periodic it is usually represented by a ring in two dimensions, making 5 dimensions. To this we added two dimensions representing DR (also periodic), and one representing SF, making 8 dimensions in total. The precise distribution of feature points in each of these dimensions is described in Materials and Methods.

Fig. 1 shows an initial simulation of VF and OR, and the effect of adding dimensions representing DR, OD and SF respectively in subsequent simulations. The top row corresponds to OR simulated just with VF, with subsequent rows showing the addition of DR, OD and SF respectively. It is apparent that adding dimensions does not change the qualitative character of maps, but does tend to reduce their wavelengths in the model (quantified in Fig. 2). Looking at the last row of Fig. 1, it is apparent that the general structure of all four of the OR, DR, OD and SF maps have a good correspondence, at least qualitatively, with experimental data (Weliky et al., 1996; Shmuel and Grinvald, 1996; Shoham et al., 1997; Hübener et al., 1997; Kim et al., 1999). (Simulations started with different random seeds gave maps with similar global characteristics.) Also in Fig. 1 is inset a joint contour map for all 4 columnar systems, to show the extent to which they intersect at right angles, and the relation of OR pinwheels to OD and SF columns (these properties are quantitatively analyzed in Fig. 3). OR, OD, and SF all show strong orthogonality to each other, as observed experimentally (Hübener et al., 1997).

Fig. 2A–C looks at some specific relationships between the DR and OR maps, and between OR and SF, when all maps are simulated together. It can be seen that the DR map has fractures (lines of low DR modulus, which correlate with where DR angle reverses direction), and that these fractures connect OR pinwheels (Shmuel and Grinvald, 1996; Swindale and Bauer, 1998). Away from DR fractures, contours of OR and DR run parallel. DR fractures on the other hand run orthogonal to contours of OR, since they connect pinwheels. The OR and SF maps tend to intersect orthogonally (Hübener et al., 1997). Fig. 2D shows the mean wavelength for each dimension as the number of dimensions is increased. It is apparent that adding dimensions reduces the mean

wavelength for all previous maps. This effect is less pronounced if the map added has strong “deprivation”, as seen with decreased gradient for addition of SF ($\text{dep}_{\text{SF}} = 0.5$ in this case; with $\text{dep}_{\text{SF}} \lesssim 0.3$ there is no wavelength reduction).

Fig. 3 shows a quantitative analysis of the effect of adding feature dimensions on the distribution of inter-pinwheel distances, distances of pinwheels to OD and SF borders, and intersection angles between columns. As in Fig. 1 successive rows show the cases of OR simulated alone (i.e., with VF but without DR, OD and SF); OR and DR together; OR, DR and OD; and finally all four feature dimensions (OR, DR, OD, SF). The first column shows histograms of distance between pairs of pinwheels that are nearest neighbors. As in experimental data (Müller et al., 2000) singularities repel each other compared to a random distribution, though the repulsion weakens as more maps are added. For OR/VF simulated alone the percentage of nearest neighbor pinwheels which are of the same sign is $18.7 \pm 1.3\%$, which matches well with experimentally determined values of $19.7 \pm 1.5\%$ (ferret) and $21.4 \pm 2.0\%$ (cat) (Müller et al., 2000). However, in the simulations this rises to $30.9 \pm 1.1\%$ when all features are simulated.

The second column of Fig. 3 shows histograms of OR pinwheel distance to OD border and SF border (the OD case is shown for the simulations with and without SF). Distances are normalized by mean wavelength of the OD/SF map respectively, so that a normalized distance of 0.25 roughly corresponds to the “center” of the OD/SF columns (we prefer to use distance of pinwheels to column borders rather than column “centers” since borders have a less ambiguous definition than “centers”; see Materials and Methods). Pinwheels tend to lie away from the borders of OD and SF, and the histogram of pinwheel to border distance reliably has a strong peak at slightly less than 0.2 of mean wavelength. Adding SF reduces this distance slightly for OD, i.e. moves OR pinwheels slightly closer to OD borders. It is hard to make a precise quantitative comparison with experimental data due to ambiguities in the definition of “distance to column center”, but at least qualitatively there is a good match (Bartfeld and Grinvald, 1992; Obermayer and Blasdel, 1993; Hübener et al., 1997).

The third column of Fig. 3 shows intersection angles for map pairs. This shows the tendency to orthogonality for all combinations except OR and DR, where the intersection angles are low. It is clear that the shape of the intersection angles histogram for each pair of maps is very robust to the presence of additional feature dimensions. In order to make a direct comparison with numerical values quoted experimentally, table 1 shows mean intersection angles (scalar average) for the simulations. It can be seen that adding features causes maps to intersect at very slightly shallower mean angles. The mean angle for OD/OR reported by Hübener et al. (1997) was 51.7 ± 0.8 degrees, which appears statistically indistinguishable from our value of 52.5 ± 0.2 when all maps are simulated together. Our value for SF/OR is 52.0 ± 0.2 ; that found by Hübener et al. (1997) was 49.7 ± 0.8 degrees, which they did not report as significantly different from their OD/OR value. It should be noted though that such means are only a very crude way of characterizing these angle distributions, and that 90 degrees is the modal angle of intersection.

Another characteristic of visual cortical maps is the rate of change of features with respect to other features, particularly the relation between VF and OR. That is, what is the change in the VF preference versus OR preference between neighboring cortical locations? Fig. 4 shows this relationship, and those for VF/OD and OR/OD, as maps are added. In general few consistent relationships are apparent. A specific change in OR can be associated with any change in VF. Large changes in VF tend to be associated with small changes in OR when OR and VF are simulated alone, but this trend weakens considerably when DR and SF are added. These results fit better with recent reports of a lack of consistent relationship between VF and OR gradients (Hetherington and Swindale, 1999; Buzás et al., 2003; White et al., 2001; Bosking et al., 2002) than with the finding of Das and Gilbert (1997) of a strong positive correlation. Although early results with the elastic net (Durbin and Mitchison, 1990) were interpreted as implying a strong negative correlation (Das and Gilbert, 1997) this had not actually been quantified until now, and our results show a rather more subtle picture.

Strabismus

Strabismus has previously been modeled in the elastic net by increasing the spacing of points along the OD dimension (Goodhill and Löwel, 1995). However, the effect on OR, DR and SF maps of strabismus in the model was not investigated. We therefore repeated the simulations described above for the strabismic case

(see Methods). Although as previously reported OD columns become wider (Goodhill and Löwel, 1995), no significant differences were found in the interrelationships between maps reported above for the non-strabismic case (data not shown). This is consistent with the experimental finding that strabismus does not greatly alter the relationships between OR and OD maps (Löwel et al., 1998).

Monocular deprivation and single-orientation rearing

In addition to the normal development simulated above, we were also interested to examine abnormal development in response to visual deprivation. In particular we focused on the experimentally well-characterized phenomenon of monocular deprivation (MD) (reviewed in Hubel and Wiesel, 1977; Katz and Shatz, 1996), and the more recently discovered phenomenon that single-orientation rearing (SOR) leads to an over-representation of that orientation in the OR map (Sengpiel et al., 1999). In each case, deprivation was modeled by changing the strength of the coverage term in the elastic net energy function relative to the continuity term (see Materials and Methods). In effect, we decreased the influence of feature points representing inputs in the deprived eye for the cortex relative to the open eye for MD, and increased the influence of feature points representing the “single” orientation for the cortex relative to other orientations for SOR.

Goodhill and Willshaw (1994) modeled MD in the elastic net using the same method of reduced influence for a feature space consisting only of VF and OD. However, they examined only constant MD existing throughout development. Fig. 5 shows a more complete analysis of the effect of MD with varying length time windows and start times for the deprivation (here OR was simulated in addition to OD—for analysis of the relationship between OR and OD see Fig. 6). It can be seen that, analogously to experimental data (reviewed in Hubel and Wiesel, 1977; Katz and Shatz, 1996), OD deprivation has significant effects only in a “critical period” located shortly after the moment the OD map starts to arise but before it is fully developed, with OD deprivation either before or after having little effect.

Figs. 6 and 7 examine the effect of MD on the structure of OR, DR and SF maps. Although the immediate qualitative impression is that MD has little effect on the structure of other maps, there are subtle changes. Perhaps the most noticeable effect is seen in the map of OR+OD contours: the islands of cortex still dominated by the deprived eye tend to colocalize with pinwheels (Fig. 7B), as seen experimentally (Crair et al., 1997). Besides, column widths for the OR, DR and SF maps increase slightly as the deprivation ratio increases (Fig. 7C). This can be explained as follows in the coverage-continuity framework (see Materials and Methods). We know that increasing the value of the ratio β/α (within a range for which striped maps arise) increases the width of the columns for all maps (Carreira-Perpiñán and Goodhill, 2004b). In MD, we decrease the value of α (more precisely, the value of α_n for the deprived eye stimuli) in eq. (1), which has the effect of increasing β/α and so increasing the width of the columns for all maps. We also found a broadening of the histogram of distances from pinwheels to OD and SF borders for the largest deprivation ratio, as shown in Fig. 7E–F. For increasing deprivation, pinwheels move even closer to the center of ocular dominance columns, and the peak becomes wider, indicating a larger dispersion. This is because when the deprivation is high, the deprived-eye OD domains shrink and the distance to pinwheels inside them decreases; but the normal-eye domains expand and the distance to pinwheels inside them increases.

Fig. 8 shows a related influence for the structure of visual input on OR map structure via single orientation rearing (Sengpiel et al., 1999). As observed biologically, the greater the “over-representation” of one orientation in the input, the greater the area of cortex that is dominated by that orientation in the final map. The OD, DR and SF maps are not shown: as for MD, they were largely unaffected.

Discussion

Computational theories of visual map development have been proposed at several levels of abstraction. While those towards the more realistic end of the spectrum, such as “high dimensional” models (e.g. Miller et al., 1989; Goodhill, 1993), are easier to interpret biologically, they have not been very successful at reproducing the geometrical relationships described above. More successful in this regard have been self-organizing

“feature space” algorithms such as the elastic net and Kohonen algorithm, which were first applied to the development of visual cortical maps about 15 years ago (Goodhill and Willshaw, 1990; Durbin and Mitchison, 1990; Obermayer et al., 1990; Obermayer et al., 1992). These early applications considered only the development of one columnar system, either OR or OD, in addition to VF. Subsequently the joint development of two columnar systems in addition to VF was considered (OR and OD: Erwin et al., 1995; Wolf and Geisel, 1998; Goodhill and Cimponeriu, 2000; OR and DR: Swindale and Bauer, 1998), and several detailed properties of biological maps were reproduced in these studies. Here we have presented the first application of such an algorithm to the joint development of four columnar systems in addition to VF. It was not obvious a priori that such algorithms would still produce biologically relevant results in such a relatively high-dimensional space. However, we have shown that the elastic net can do a remarkably good job of reproducing a wide range of experimentally observed phenomena in this case. We have also shown that it can reproduce recently-observed phenomena concerning MD and SOR. The main experimental facts the model reproduces even when multiple maps are simulated together are now summarized.

1. **Orientation and ocular dominance:** The model reproduces the tendency of OR and OD columns to intersect at steep angles (Obermayer and Blasdel, 1993), and for pinwheels to lie far from OD borders (Bartfeld and Grinvald, 1992; Obermayer and Blasdel, 1993; Hübener et al., 1997).
2. **Direction:** In the model the DR map has fractures rather than pinwheels and pinwheels tend to be connected by fractures (Weliky et al., 1996; Shmuel and Grinvald, 1996).
3. **Spatial frequency:** The model reproduces the tendency of OR and OD columns to intersect SF columns at steep angles, and of OR pinwheels to lie far from SF borders (Hübener et al., 1997).
4. **Monocular deprivation:** MD during a critical period of development produces a shrinkage of OD domains from the deprived eye, with a magnitude proportional to the strength and duration of the deprivation. Pinwheels tend to colocalize with deprived eye patches (Crair et al., 1997).
5. **Single orientation rearing:** SOR produces an expansion of OR domains for the overrepresented orientation (Sengpiel et al., 1999).

While it was already known that some of these map properties could be produced by feature-space models (Erwin et al., 1995; Swindale, 1996; Swindale and Bauer, 1998), we have shown that these results are robust to the addition of extra feature dimensions. Further, we have shown that the relationships between maps produced by the model are generally robust to perturbations in input activity during development including strabismus, MD, and SOR. In particular, the model predicts that these perturbations have little effect on the structure of SF and DR maps. This constitutes a prediction of the model, which could relatively straightforwardly be tested by imaging all four of the maps OR, DR, OD and SF after each of these perturbations. The model also predicts a slight increase in the wavelengths of the maps of OR, DR and SF after MD. From the coverage-continuity point of view, this results from the fact that the monocular deprivation effectively increases the influence of the continuity term.

The adult arrangement of maps in visual cortex has also been analyzed from a wiring optimization perspective (Koulakov and Chklovskii, 2001; Koulakov and Chklovskii, 2002; Chklovskii and Koulakov, 2004). Although this approach does not address the development of maps, and also has not so far included topography as a variable to be mapped, it can successfully reproduce some of the experimental findings we have addressed here with the elastic net. In fact recent work has shown both a formal link between wiring optimization, Mexican-hat lateral interactions and the elastic net continuity term (Carreira-Perpiñán and Goodhill, 2004b), and that subtle differences in patterns of intracortical connectivity can have a strong influence on the relationships between individual maps.

The only other study to have examined the development of multiple columnar systems in visual cortex using a feature space model is that of Swindale (2000). Here the development of up to 9 feature dimensions in addition to VF was simulated using Kohonen’s algorithm. However, the approach was more abstract than ours in that all feature dimensions were binary, and were not interpreted in terms of specific visual features. Swindale found

that the structure of each map individually, including its wavelength, was fairly robust to additional maps, but that the relationships between maps (specifically intersection angles) changed in a gradual way as more maps were added. These results differ somewhat from our own results: we found wavelengths to decrease as maps are added, yet more robustness of intersection angles as maps are added. We did however find a shift in the location of pinwheels relative to OD columns with the addition of the SF dimension (as no periodic dimensions were simulated in Swindale (2000) the question of the location of pinwheels relative to other maps did not arise in that case). The reason for the difference in results between these two studies is unclear, but it could be a manifestation of the fact that, although the elastic net and Kohonen algorithms are similar in flavor, they can behave differently in fine details.

What light do our results shed on the biological mechanisms underlying map formation? The elastic net represents a particular mathematical instantiation of the hypothesis that visual cortical maps are the result of an optimization process. In particular, the elastic net attempts to jointly optimize both coverage, the degree to which all input features are uniformly represented, and continuity, the degree to which the spatial representation of features is “smooth” in some sense (see Materials and Methods and Carreira-Perpiñán and Goodhill, 2002). It was already known that such optimization hypotheses are sufficient to generate the detailed structure of OR and OD maps under normal conditions (Erwin et al., 1995; Swindale, 1996). Our results now show that optimization principles based on the coverage-continuity trade-off are also sufficient to reproduce a large array of additional data, both when additional maps are considered, and when particular types of visual deprivation are introduced.

How might the visual cortex actually implement such optimization principles? One way the elastic net equations can be interpreted biologically is in terms of a Hebbian process: the simple gradient descent learning rule for optimizing the elastic net objective function is equivalent to an activity-dependent strengthening of connections between the input feature that is presented at each instant and the neurons that respond most strongly to that feature. In contrast, evidence is now mounting that at least some properties of visual map formation are activity-independent (e.g. Katz and Crowley, 2002). However, the elastic net equations can be interpreted in other ways, and in fact the algorithm originally grew out of an activity-independent model that relied on the matching of molecular cues (Willshaw and von der Malsburg, 1979). The fact that such optimization models work so well, and are now the only developmental models available for reproducing the wide array of map properties addressed in this paper, suggest that the optimization hypothesis is key to understanding visual cortex, independent of details of how precisely it might be implemented.

Acknowledgements

Supported by grants NIH EY12544 and NSF BITS0130822. We thank David Swartz for help with some of the simulations.

References

- Bartfeld E, Grinvald A (1992) Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. *Proc. Nat. Acad. Sci. USA* 89: 11905–11909.
- Bonhoeffer T, Grinvald A (1991) Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* 353: 429–431.
- Bosking WH, Crowley JC, Fitzpatrick D (2002) Spatial coding of position and orientation in primary visual cortex. *Nature Neuroscience* 5: 874–882.
- Buzás P, Volgushev M, Eysel UT, Kisvárdy ZF (2003) Independence of visuotopic representation and orientation map in the visual cortex of the cat. *Eur. J. Neurosci.* 18: 957–968.
- Carreira-Perpiñán MÁ, Goodhill GJ (2002) Development of columnar structures in visual cortex. In Ascoli GA, editor, *Computational Neuroanatomy: Principles and Methods*, chapter 15, pp. 337–357. Humana Press Inc., Totowa, NJ.
- Carreira-Perpiñán MÁ, Goodhill GJ (2004a) Generalized elastic nets. Submitted.
- Carreira-Perpiñán MÁ, Goodhill GJ (2004b) The influence of lateral connections on the structure of cortical maps. *J. Neurophysiology* 92: 2947–2959.
- Chklovskii DB, Koulakov AA (2004) Maps in the brain: what can we learn from them? *Ann. Rev. Neurosci.* 27: 369–392.
- Crair MC, Ruthazer ES, Gillespie DC, Stryker MP (1997) Relationship between the ocular dominance and orientation maps in visual cortex of monocularly deprived cats. *Neuron* 19: 307–318.
- Das A, Gilbert CD (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 (1987) An analogue approach to the traveling salesman problem using an elastic net method. *Nature* 326: 689–691.
- Erwin E, Obermayer K, Schulten K (1995) Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural Computation* 7: 425–468.
- Everson RM, Prashanth AK, Gabbay M, Knight BW, Sirovich L, Kaplan E (1998) Representation of spatial frequency and orientation in the visual cortex. *Proc. Nat. Acad. Sci. USA* 95: 8334–8338.
- Goodhill GJ (1993) Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics* 69: 109–118.
- Goodhill GJ, Bates KR, Montague PR (1997) Influences on the global structure of cortical maps. *Proc. Roy. Soc. London B.* 264: 649–655.
- Goodhill GJ, Cimponeriu A (2000) Analysis of the elastic net applied to the formation of ocular dominance and orientation columns. *Network: Comput. Neural Syst.* 11: 153–168.
- Goodhill GJ, Löwel S (1995) Theory meets experiment: Correlated neural activity helps determine ocular dominance column periodicity. *Trends in Neurosciences* 18: 437–439.
- Goodhill GJ, Willshaw DJ (1990) Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network: Comput. Neural Syst.* 1: 41–59.

- Goodhill GJ, Willshaw DJ (1994) Elastic net model of ocular dominance: Overall stripe pattern and monocular deprivation. *Neural Computation* 6: 615–621.
- Hetherington PA, Swindale NV (1999) Receptive field and orientation scatter studied by tetrode recordings in cat area 17. *J. Vis. Neurosci.* 16: 637–652.
- Hubel DH, Wiesel TN (1977) Functional architecture of the macaque monkey visual cortex. *Proc. Roy. Soc. London B.* 198: 1–59.
- Hübener M, Shoham D, Grinvald A, Bonhoeffer T (1997) Spatial relationships among three columnar systems in cat area 17. *J. Neurosci.* 17: 9270–9284.
- Issa NP, Trepel C, Stryker MP (2000) Spatial frequency maps in cat visual cortex. *J. Neurosci.* 20: 8504–8514.
- Katz LC, Shatz CJ (1996) Synaptic activity and the construction of cortical circuits. *Science* 274: 1133–1138.
- Katz LC, Crowley JC (2002) Development of cortical circuits: Lessons from ocular dominance columns. *Nature Reviews Neuroscience* 3: 34–42.
- Kim DS, Matsuda Y, Ohki K, Ajima A, Tanaka S (1999) Geometrical and topological relationships between multiple functional maps in cat primary visual cortex. *Neuroreport* 10: 2515–2522.
- Koulakov AA, Chklovskii DB (2001) Orientation preference patterns in mammalian visual cortex: A wire length minimization approach. *Neuron* 29: 519–527.
- Koulakov AA, Chklovskii DB (2002) Direction of motion maps in the visual cortex: a wire length minimization approach. *Neurocomputing* 44–46: 489–494.
- Löwel S, Schmidt KE, Kim DS, Wolf F, Hoffmüller F, Singer W, Bonhoeffer T (1998) The layout of orientation and ocular dominance domains in area 17 of strabismic cats. *Eur. J. Neurosci.* 10: 2629–2643.
- Miller KD, Keller JB, Stryker MP (1989) Ocular dominance column development: Analysis and simulation. *Science* 245: 605–615.
- Müller T, Stetter M, Hübener M, Sengpiel F, Bonhoeffer T, Gödecke I, Chapman B, Löwel S, Obermayer K (2000) An analysis of orientation and ocular dominance patterns in the visual cortex of cats and ferrets. *Neural Computation* 12: 2573–2595.
- Obermayer K, Blasdel GG (1993) Geometry of orientation and ocular dominance columns in monkey striate cortex. *J. Neurosci.* 13: 4114–4129.
- Obermayer K, Blasdel GG, Schulten K (1992) 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 (1990) A principle for the formation of the spatial structure of cortical feature maps. *Proc. Nat. Acad. Sci. USA* 87: 8345–8349.
- Sengpiel F, Stawinski P, Bonhoeffer T (1999) Influence of experience on orientation maps in cat visual cortex. *Nature Neuroscience* 2: 727–732.
- 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.
- Shoham D, Hübener M, Schulze S, Grinvald A, Bonhoeffer T (1997) Spatio-temporal frequency domains and their relation to cytochrome oxidase staining in cat visual cortex. *Nature* 385: 529–533.
- Silverman MS, Groszof DH, De Valois RL, Elfar SD (1989) Spatial-frequency organization in primate striate cortex. *Proc. Nat. Acad. Sci. USA* 86: 711–715.

- Swindale NV (1996) The development of topography in the visual cortex: A review of models. *Network: Comput. Neural Syst.* 7: 161–247.
- Swindale NV (2000) How many maps are there in visual cortex? *Cereb. Cortex* 10: 633–643.
- Swindale NV, Bauer HU (1998) Application of Kohonen's self-organizing feature map algorithm to cortical maps of orientation and direction preference. *Proc. Roy. Soc. London B.* 265: 827–838.
- Weliky M, Bosking WH, Fitzpatrick D (1996) A systematic map of direction preference in primary visual cortex. *Nature* 379: 725–728.
- White LE, Bosking WH, Fitzpatrick D (2001) Consistent mapping of orientation preference across irregular functional domains in ferret visual cortex. *J. Vis. Neurosci.* 18: 65–76.
- Willshaw DJ, von der Malsburg C (1979) A marker induction mechanism for the establishment of ordered neural mappings: its application to the retinotectal problem. *Phil. Trans. Roy. Soc. B* 287: 203–243.
- Wolf F, Geisel T (1998) Spontaneous pinwheel annihilation during visual development. *Nature* 395: 73–78.

List of Figures

- 1 The effect of adding additional feature dimensions to the elastic net. The large inset shows contours for all dimensions: OR (blue), DR (red), OD (black), and SF (green). The four rows show OR simulated alone, OR/DR, OR/DR/OD, and OR/DR/OD/SF, respectively. The same parameters were used in each case. 14
- 2 Relationships between maps when OR, DR, OD, SF are simulated together. A. OR contours (blue) plotted with DR modulus (greyscale). B. DR contours (red) plotted with OR modulus (greyscale). C. OR (blue) with SF (green) contours. D. Mean wavelength change as dimensions are added. 15
- 3 Quantitative effects of adding additional feature dimensions. Successive rows show the addition of dimensions as in Fig. 1. Column 1: distribution of pinwheel-to-closest-pinwheel distances; distances are normalized by OR wavelength ("control" means randomly distributed pinwheels). Column 2: distribution of pinwheel-to-border distances for OD and SF; distances are normalized by OD or SF wavelength. The vertical red line shows the pinwheel-border distance of a pinwheel at the center of a patch of mean wavelength. Column 3: intersection angles between all dimensions. The histograms cluster into two groups: almost uniform with bias towards parallelism (OR/DR) and bias toward orthogonality (the rest of the histograms) with almost no change as dimensions are added. All plots show mean results across several simulations; errorbars show standard error of mean ($n = 11$). Note that for the histograms of distances of OR pinwheels to OD and SF borders it is possible for a distance to exceed 0.25 of the mean wavelength in areas where columns are unusually wide relative to elsewhere in the map. The histogram for spatial frequency has a particularly thick tail because SF patches are of quite variable size. 16
- 4 Rates of change of pairs of stimulus features in the simulated maps, shown as scatterplots of the gradient moduli of visual field (VF), OD and OR. The number of points in each plot is $128 \times 128 = 16384$, the number M of points in the net. Pinwheels are marked by circles (positive) and squares (negative). 17
- 5 Simulation results for ocular dominance deprivation for a value $\text{dep}_{\text{OD}} = 0.4$ of the deprivation strength occurring only inside a window during OD development. The rows correspond to the duration k_w of the OD deprivation window, while the columns correspond to the center k_0 of that window. Both axes are in units of k , the annealing parameter of the elastic net. This is 0.2 at the start of each simulation, and is gradually reduced (see Materials and Methods), so that decreasing k means increasing developmental time. The top row thus corresponds to the longest duration of deprivation. For the columns, the earliest window ($k_0 = 0.0525$) lies entirely before the OD map starts to develop in the simulations. The latest window ($k_0 = 0.0100$) lies after the OD map has fully developed. The four intermediate values lie uniformly in this development period. The figure shows that in the elastic net, OD deprivation has a strong effect only in a critical period that starts when the OD map starts to arise (in this simulation, at $k = 0.045$) and ends when the OD map is fully developed (at $k = 0.033$). For clarity all simulations use the same starting conditions. 18
- 6 The effect of monocular deprivation on map structure for all dimensions (same starting conditions in each case). dep_{OD} is the ratio of the overall level of activity over the course of development (modeled by the elastic net parameter α) in the deprived eye relative to that in the normal eye (left-most column is control: no deprivation). The fifth row shows contours of OD (black) and OR (blue): note the tendency for the islands of cortex still dominated by the deprived eye to contain pinwheels (Crair et al., 1997). 19

- 7 Quantitative effects of monocular deprivation in the simulated maps. A. Change in proportion of cortical area covered by the deprived eye as amount of deprivation varies. Errorbars are present but very small. B. Quantitative demonstration that pinwheels tend to colocalize with deprived eye patches. The “pinwheel–OD colocalization ratio” is the ratio of (number of pinwheels lying in deprived eye patches / total number of pinwheels) to (area of deprived eye patches / total area). This is greater than one when there is a colocalization effect. C. Change in wavelengths of OR, DR and SF as MD varies. OD wavelength is not plotted. The size of left- and right-eye patches becomes quite uneven with increasing deprivation, and in this case the OD map is not well-described by the mean wavelength. D. Intersection angles between OD and OR for differing MD. E. Distribution of OR pinwheel distances from OD borders, for different values of MD. F. Distribution of OR pinwheel distances from SF borders, for different values of MD. All plots show mean results across 8 simulations; errorbars show standard error of mean. 20
- 8 The effect of single orientation rearing on OR map development, i.e., overrepresentation of one orientation in the input relative to the others (Sengpiel et al., 1999). dep_{OR} here refers to the value of α for the zero degrees orientation (cyan) compared to other orientations over the course of the simulation. Thus, $dep_{OR} = 1$ means all OR angles are equally represented in the training stimuli, while $dep_{OR} > 1$ means that zero degrees occurs more often than other angles. The histogram in B shows the proportion of cortical area dedicated to each orientation preference for different degrees of OR overrepresentation. Four histograms are shown, corresponding to the values of $dep_{OR} = 1, 1.5, 2$ and 2.5 . The histograms show that $dep_{OR} = 1$ (no overrepresentation) leads to a uniform histogram while increasing SOR increases the cortical area dedicated to zero degrees at the expense of the other OR values, in agreement with biological data. Each histogram is the result of one simulation, with little variation over different simulations (error bars are not shown to avoid cluttering the figure). 21

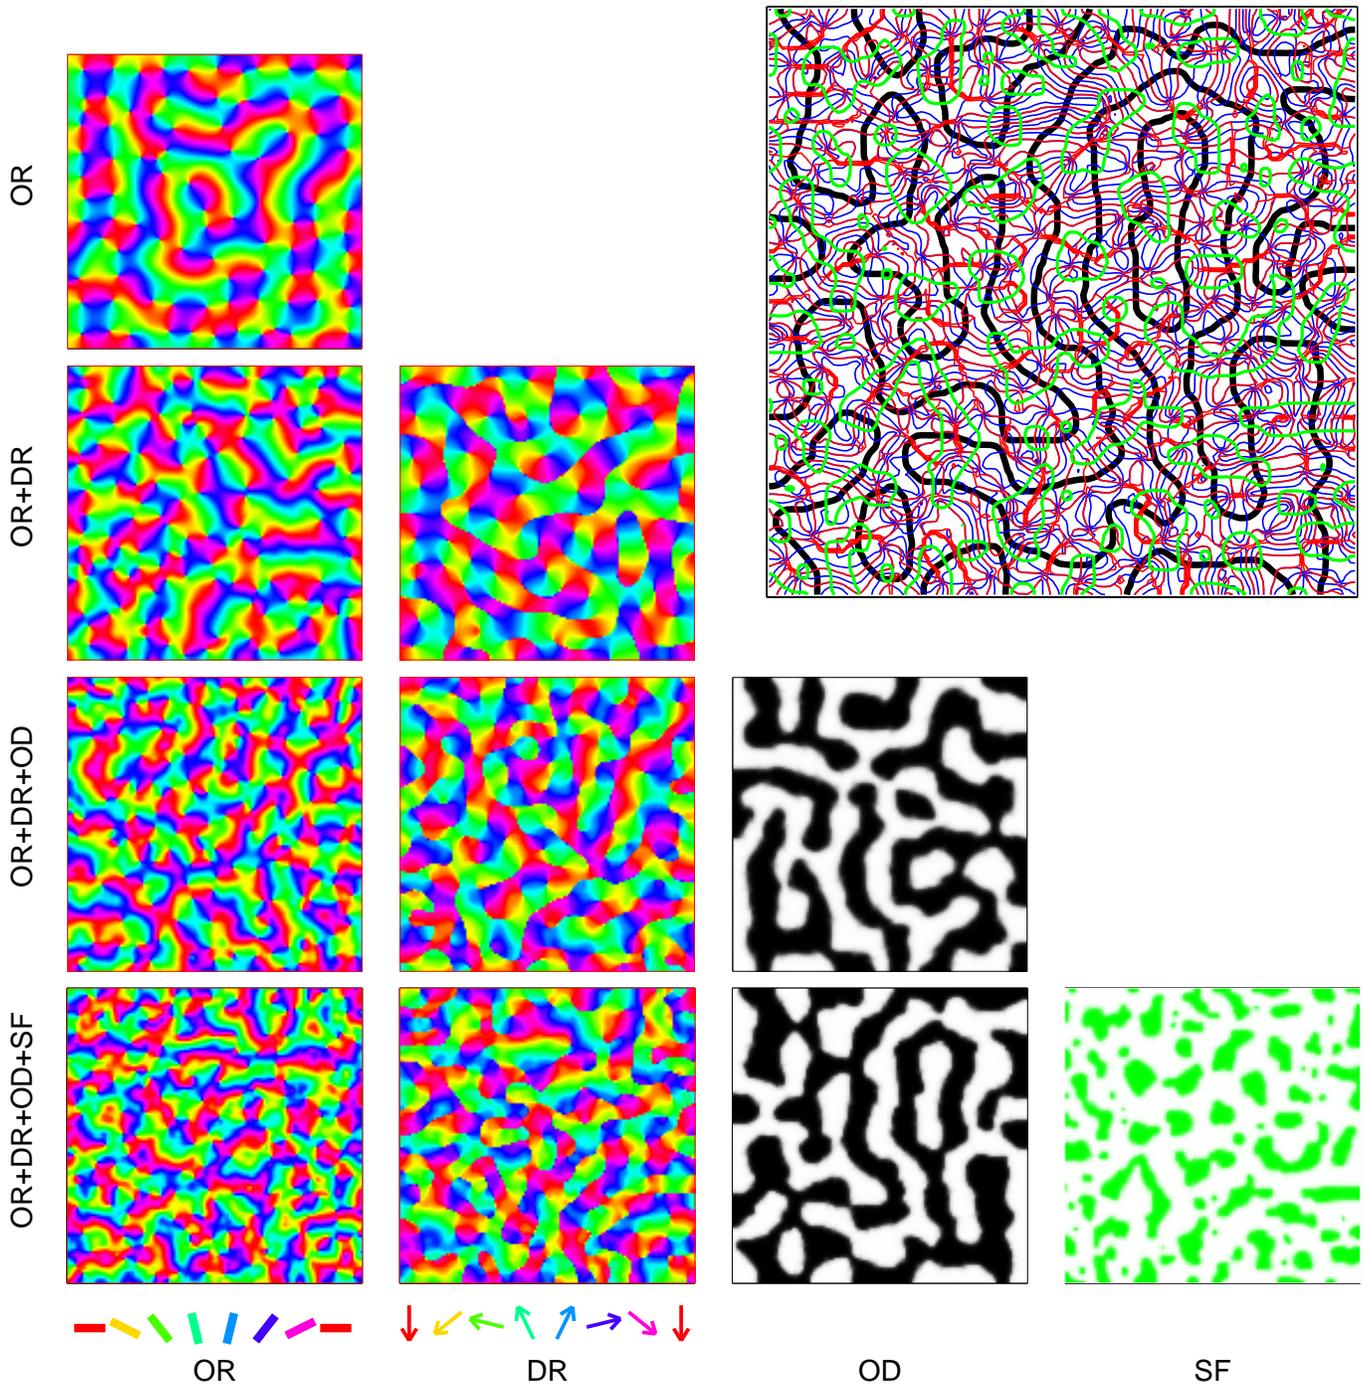
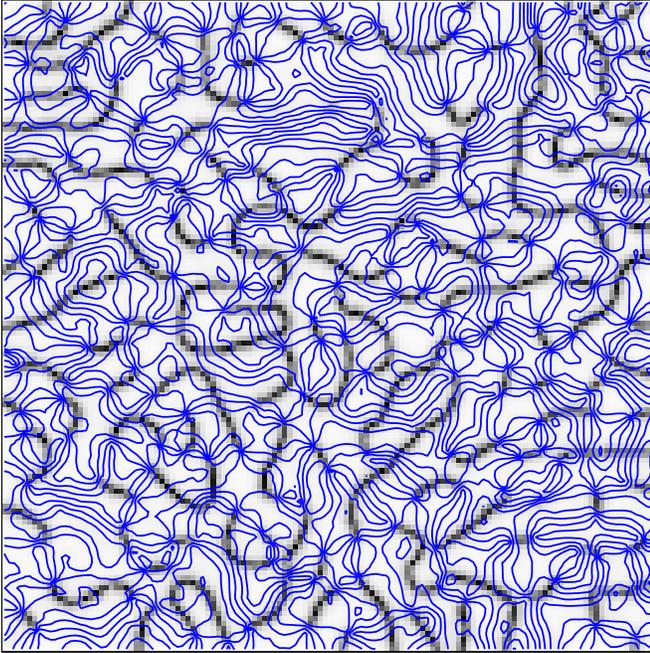
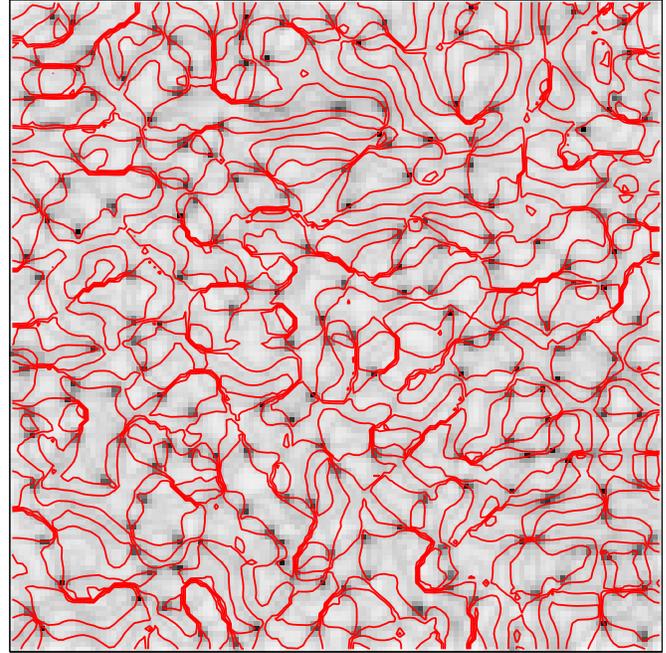


Figure 1: The effect of adding additional feature dimensions to the elastic net. The large inset shows contours for all dimensions: OR (blue), DR (red), OD (black), and SF (green). The four rows show OR simulated alone, OR/DR, OR/DR/OD, and OR/DR/OD/SF, respectively. The same parameters were used in each case.

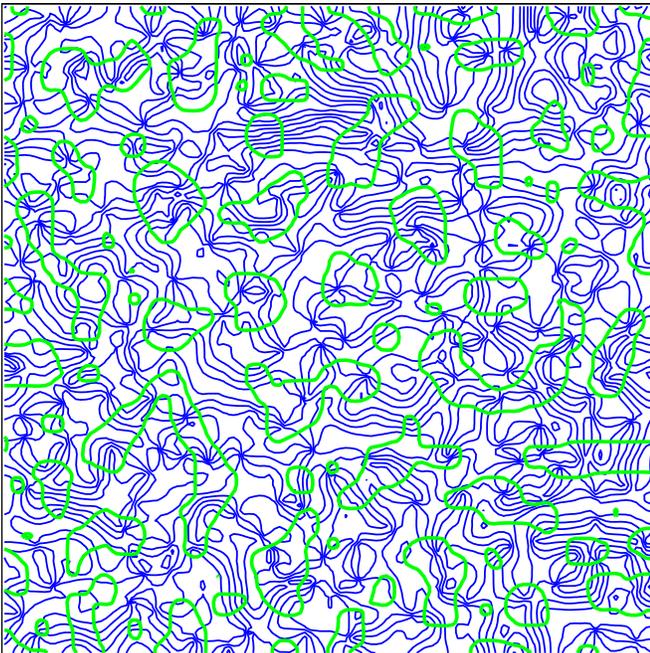
A: DR modulus + OR contours



B: OR modulus + DR contours



C: OR contours + SF contours



D: Change in wavelength

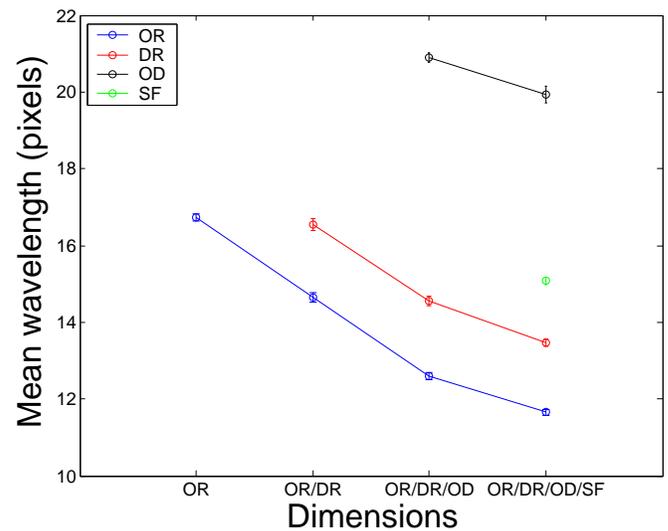


Figure 2: Relationships between maps when OR, DR, OD, SF are simulated together. A. OR contours (blue) plotted with DR modulus (greyscale). B. DR contours (red) plotted with OR modulus (greyscale). C. OR (blue) with SF (green) contours. D. Mean wavelength change as dimensions are added.

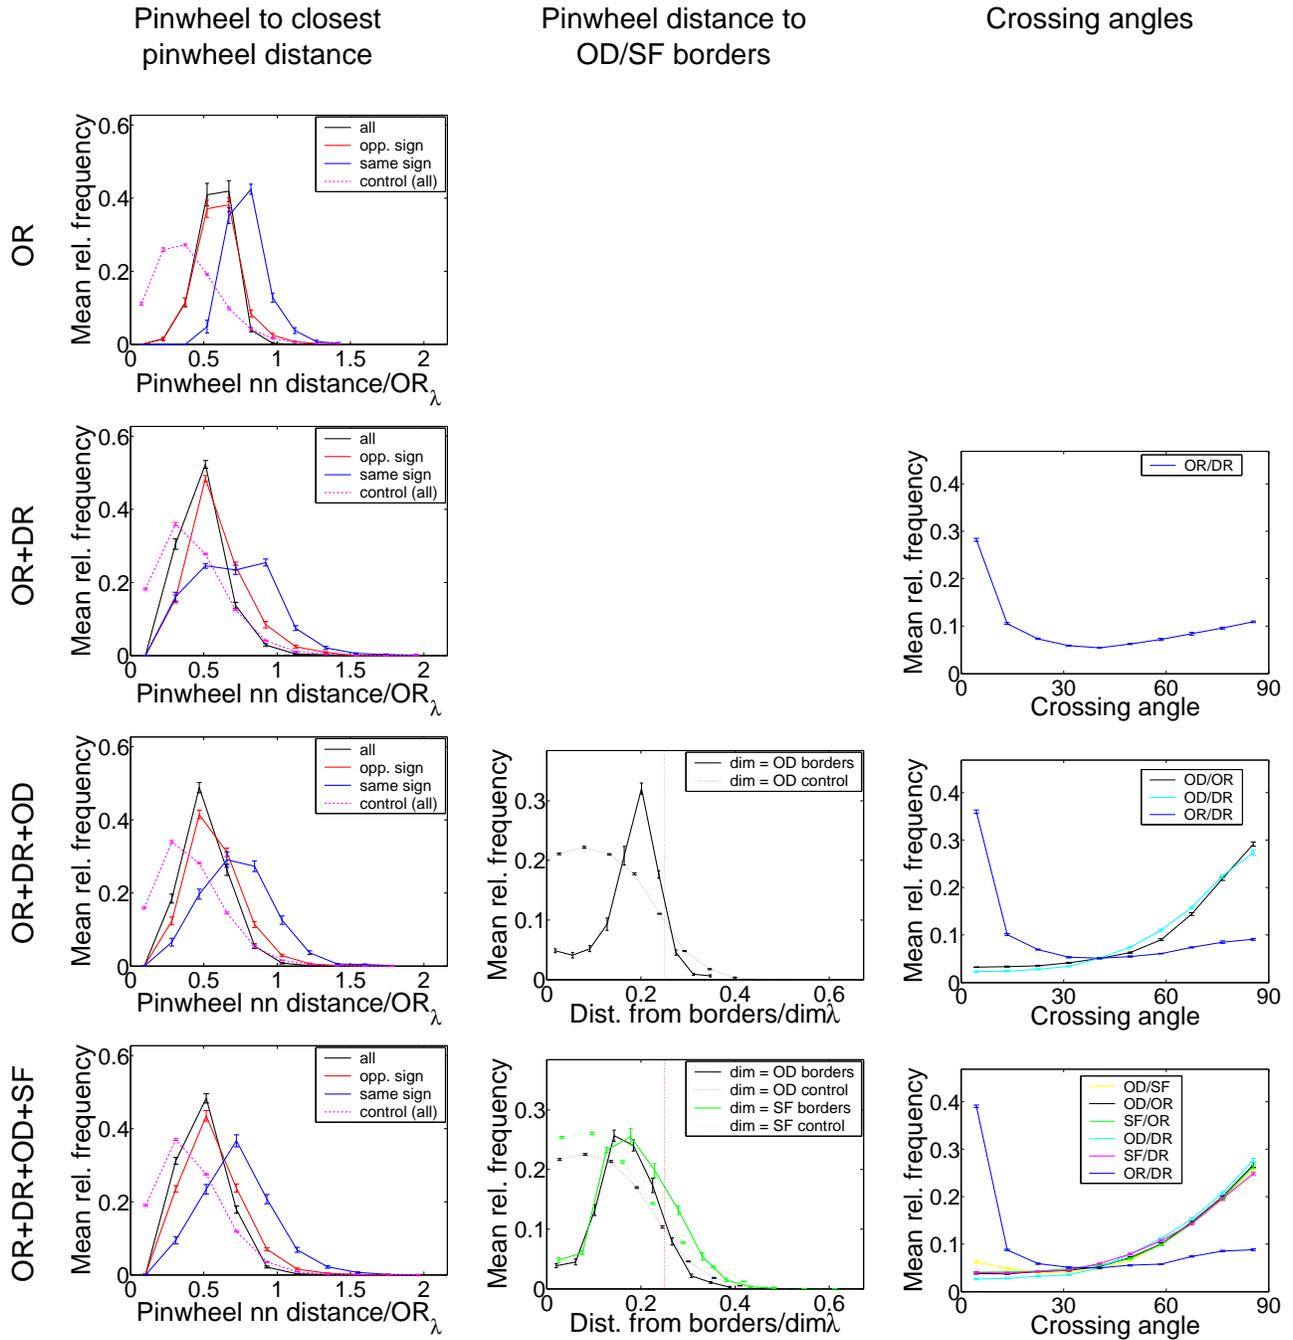


Figure 3: Quantitative effects of adding additional feature dimensions. Successive rows show the addition of dimensions as in Fig. 1. Column 1: distribution of pinwheel-to-closest-pinwheel distances; distances are normalized by OR wavelength (“control” means randomly distributed pinwheels). Column 2: distribution of pinwheel-to-border distances for OD and SF; distances are normalized by OD or SF wavelength. The vertical red line shows the pinwheel-border distance of a pinwheel at the center of a patch of mean wavelength. Column 3: intersection angles between all dimensions. The histograms cluster into two groups: almost uniform with bias towards parallelism (OR/DR) and bias toward orthogonality (the rest of the histograms) with almost no change as dimensions are added. All plots show mean results across several simulations; errorbars show standard error of mean ($n = 11$). Note that for the histograms of distances of OR pinwheels to OD and SF borders it is possible for a distance to exceed 0.25 of the mean wavelength in areas where columns are unusually wide relative to elsewhere in the map. The histogram for spatial frequency has a particularly thick tail because SF patches are of quite variable size.

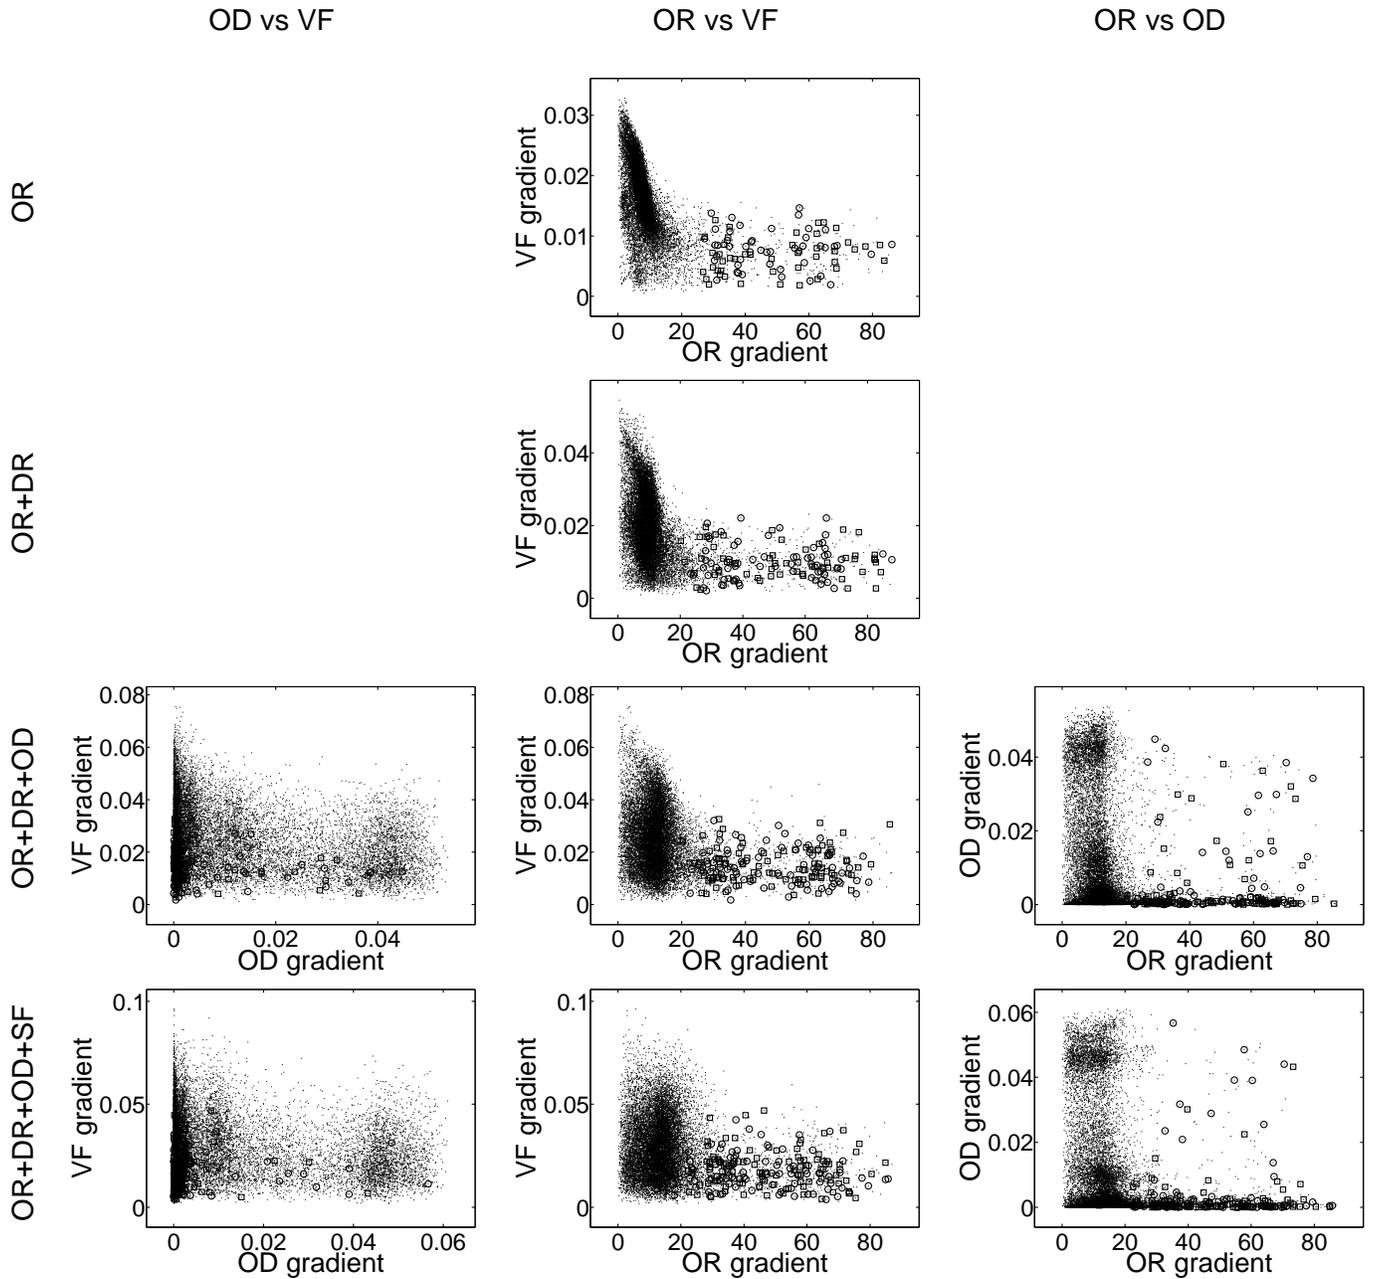


Figure 4: Rates of change of pairs of stimulus features in the simulated maps, shown as scatterplots of the gradient moduli of visual field (VF), OD and OR. The number of points in each plot is $128 \times 128 = 16384$, the number M of points in the net. Pinwheels are marked by circles (positive) and squares (negative).

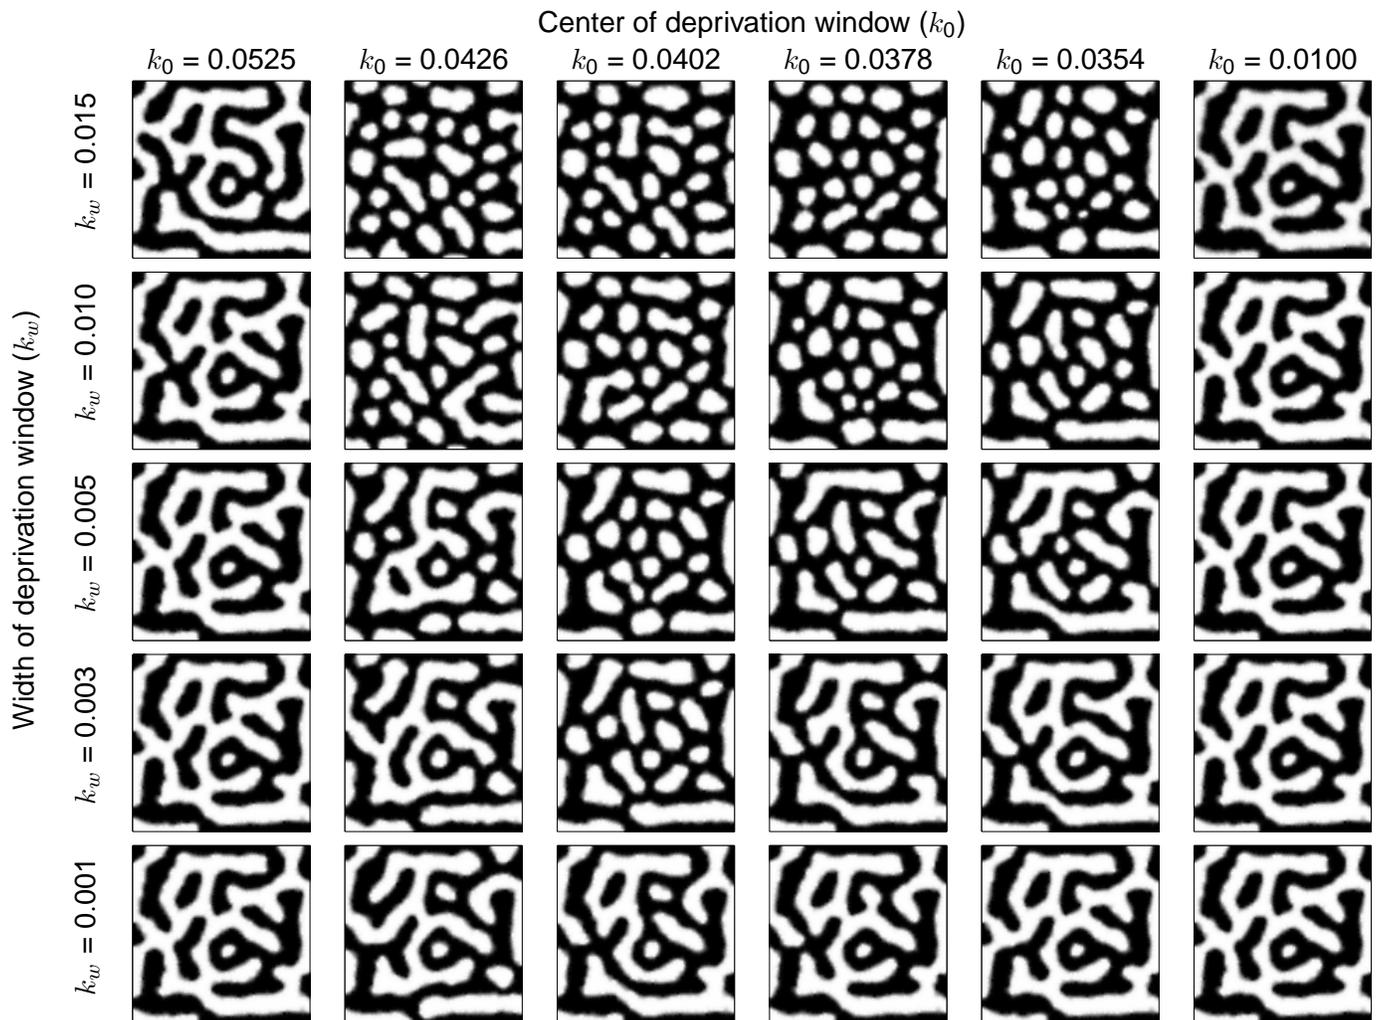


Figure 5: Simulation results for ocular dominance deprivation for a value $\text{dep}_{\text{OD}} = 0.4$ of the deprivation strength occurring only inside a window during OD development. The rows correspond to the duration k_w of the OD deprivation window, while the columns correspond to the center k_0 of that window. Both axes are in units of k , the annealing parameter of the elastic net. This is 0.2 at the start of each simulation, and is gradually reduced (see Materials and Methods), so that decreasing k means increasing developmental time. The top row thus corresponds to the longest duration of deprivation. For the columns, the earliest window ($k_0 = 0.0525$) lies entirely before the OD map starts to develop in the simulations. The latest window ($k_0 = 0.0100$) lies after the OD map has fully developed. The four intermediate values lie uniformly in this development period. The figure shows that in the elastic net, OD deprivation has a strong effect only in a critical period that starts when the OD map starts to arise (in this simulation, at $k = 0.045$) and ends when the OD map is fully developed (at $k = 0.033$). For clarity all simulations use the same starting conditions.

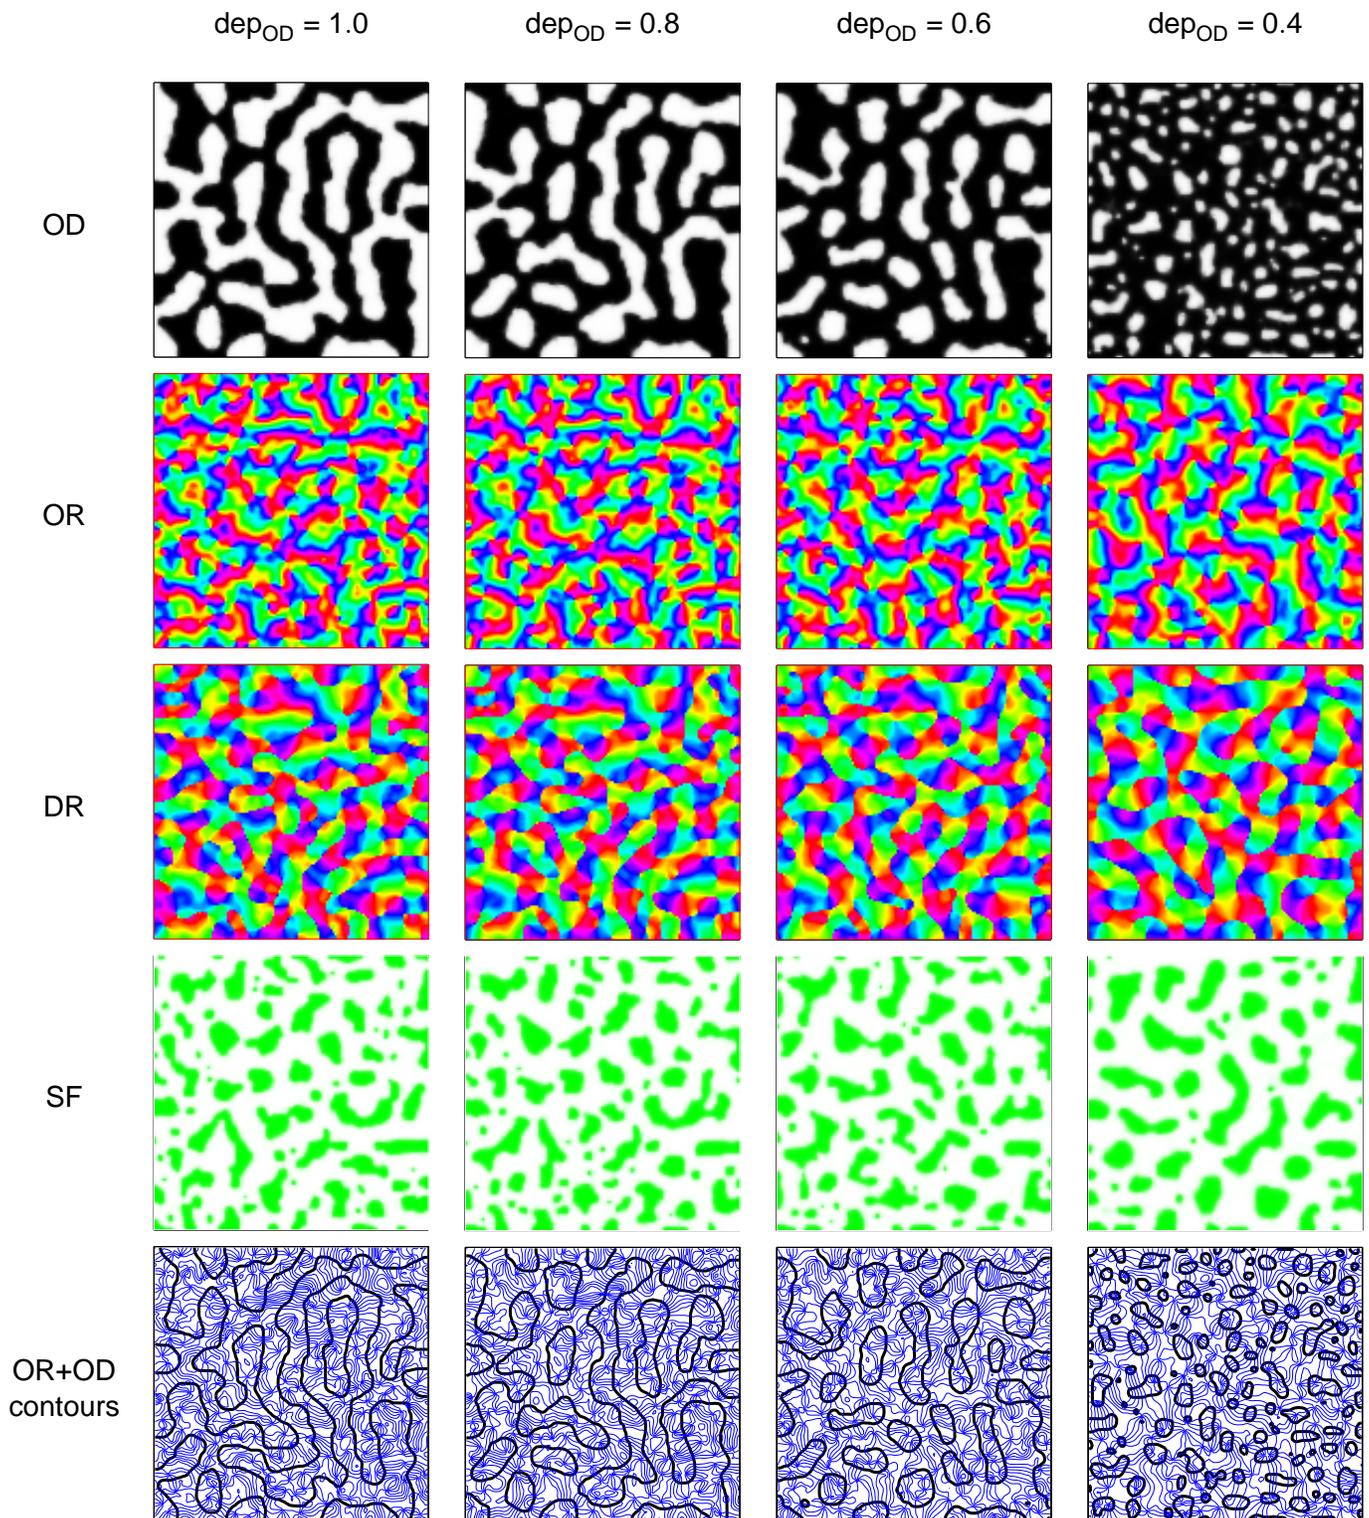


Figure 6: The effect of monocular deprivation on map structure for all dimensions (same starting conditions in each case). dep_{OD} is the ratio of the overall level of activity over the course of development (modeled by the elastic net parameter α) in the deprived eye relative to that in the normal eye (left-most column is control: no deprivation). The fifth row shows contours of OD (black) and OR (blue): note the tendency for the islands of cortex still dominated by the deprived eye to contain pinwheels (Crair et al., 1997).

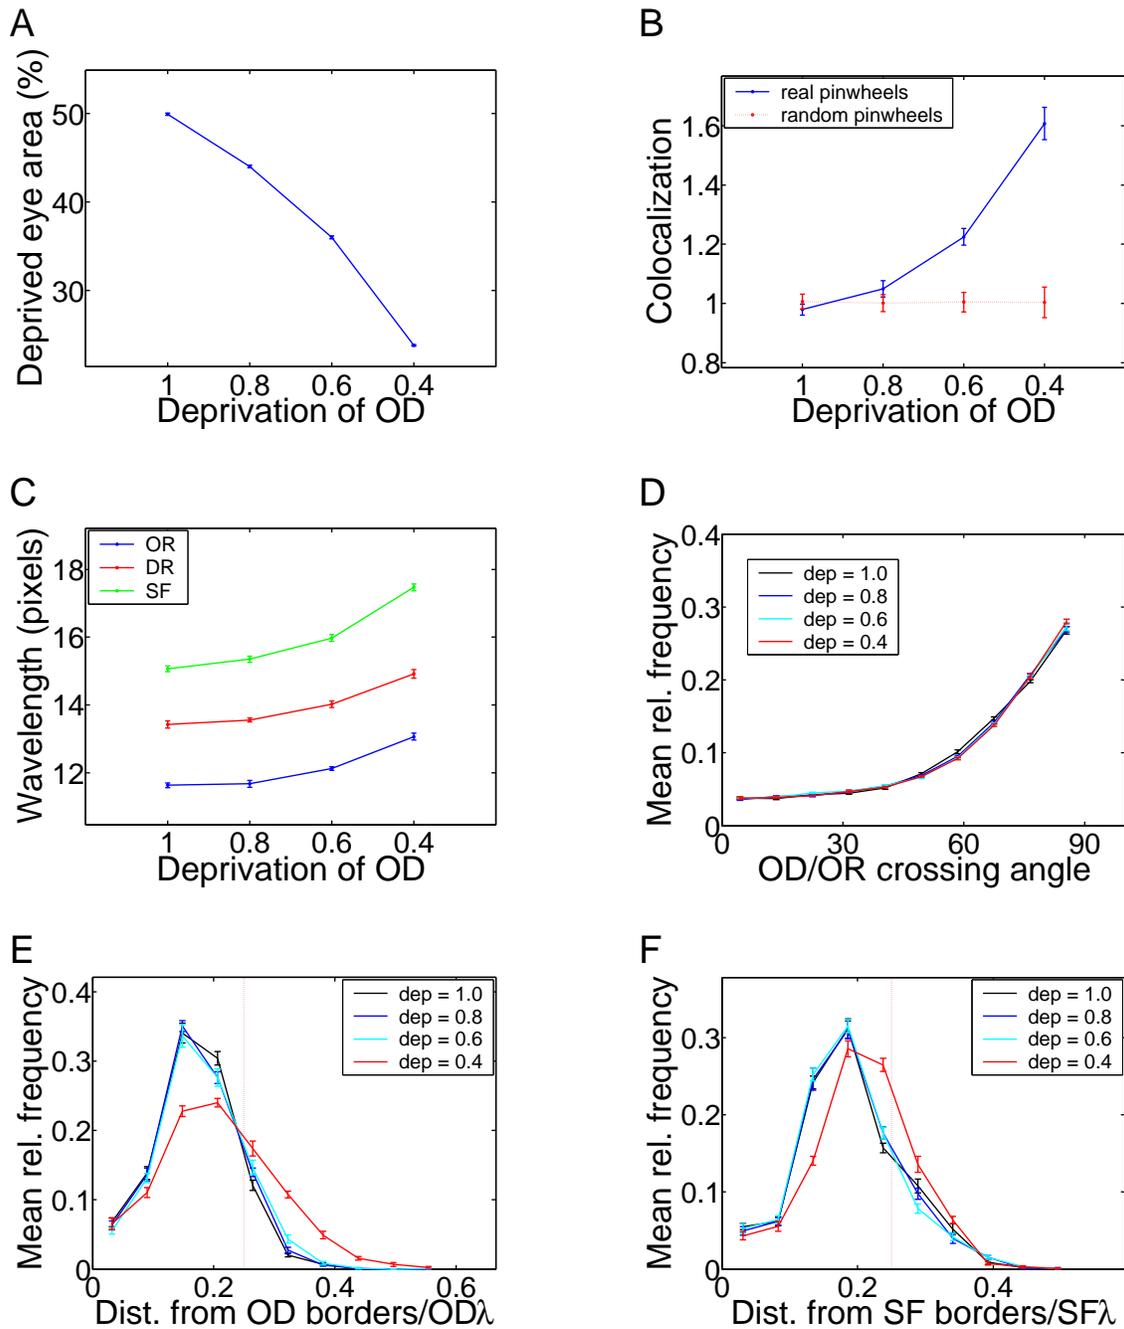


Figure 7: Quantitative effects of monocular deprivation in the simulated maps. A. Change in proportion of cortical area covered by the deprived eye as amount of deprivation varies. Errorbars are present but very small. B. Quantitative demonstration that pinwheels tend to colocalize with deprived eye patches. The “pinwheel–OD colocalization ratio” is the ratio of (number of pinwheels lying in deprived eye patches / total number of pinwheels) to (area of deprived eye patches / total area). This is greater than one when there is a colocalization effect. C. Change in wavelengths of OR, DR and SF as MD varies. OD wavelength is not plotted. The size of left- and right-eye patches becomes quite uneven with increasing deprivation, and in this case the OD map is not well-described by the mean wavelength. D. Intersection angles between OD and OR for differing MD. E. Distribution of OR pinwheel distances from OD borders, for different values of MD. F. Distribution of OR pinwheel distances from SF borders, for different values of MD. All plots show mean results across 8 simulations; errorbars show standard error of mean.

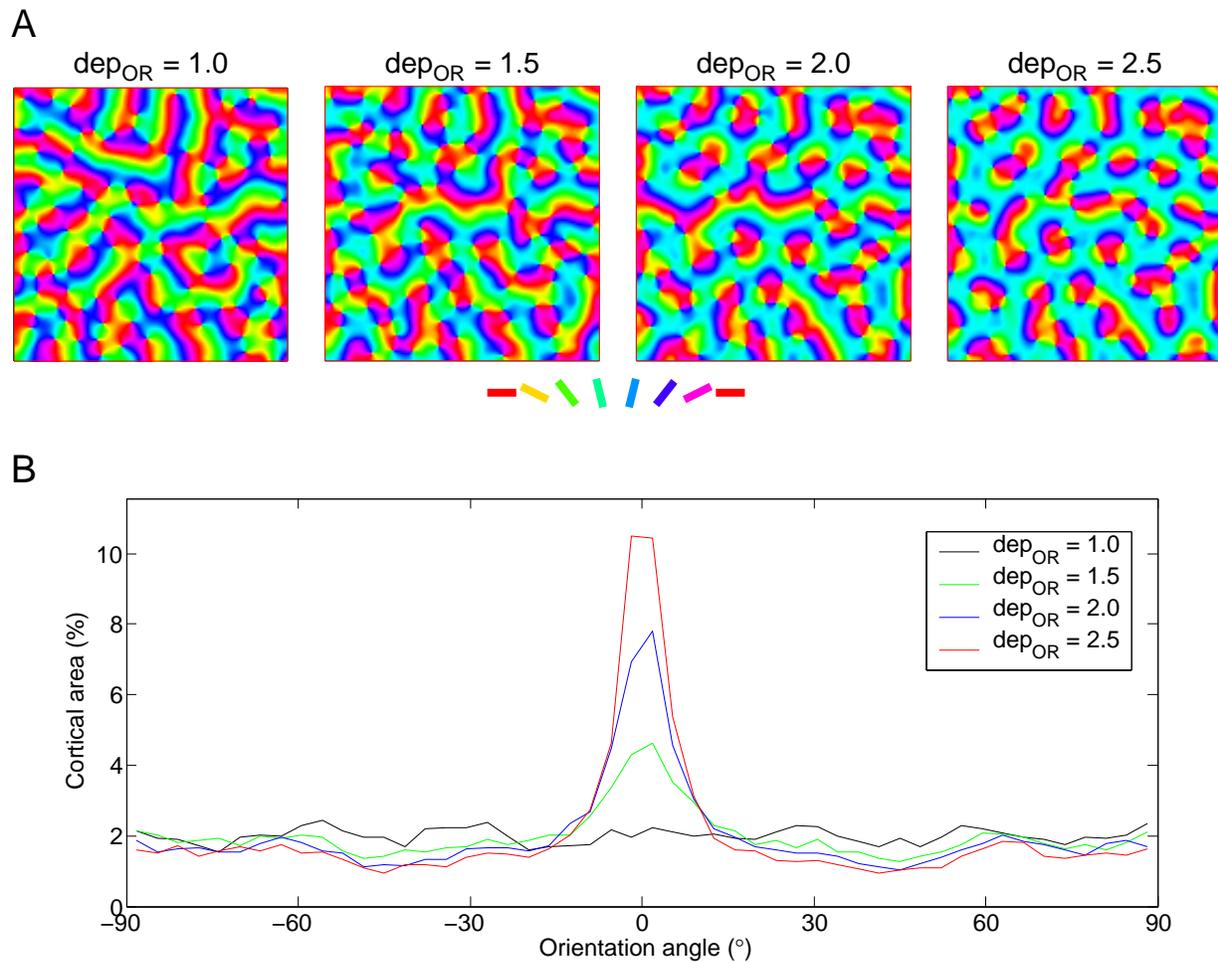


Figure 8: The effect of single orientation rearing on OR map development, i.e., overrepresentation of one orientation in the input relative to the others (Sengpiel et al., 1999). dep_{OR} here refers to the value of α for the zero degrees orientation (cyan) compared to other orientations over the course of the simulation. Thus, $\text{dep}_{\text{OR}} = 1$ means all OR angles are equally represented in the training stimuli, while $\text{dep}_{\text{OR}} > 1$ means that zero degrees occurs more often than other angles. The histogram in B shows the proportion of cortical area dedicated to each orientation preference for different degrees of OR overrepresentation. Four histograms are shown, corresponding to the values of $\text{dep}_{\text{OR}} = 1, 1.5, 2$ and 2.5 . The histograms show that $\text{dep}_{\text{OR}} = 1$ (no overrepresentation) leads to a uniform histogram while increasing SOR increases the cortical area dedicated to zero degrees at the expense of the other OR values, in agreement with biological data. Each histogram is the result of one simulation, with little variation over different simulations (error bars are not shown to avoid cluttering the figure).

List of Tables

- 1 Mean crossing angles between the maps of OD, OR, DR and SF for the model, and values reported by Hübener et al. (1997) for the cat. 23

	Model of OR+DR	Model of OD+DR+OD	Model of OR+DR+OD+SF	Experimental results of Hübener et al. (1997)
OR/DR	19.7 ± 0.2	16.9 ± 0.2	15.6 ± 0.1	
OD/OR		54.6 ± 0.2	52.5 ± 0.2	51.7 ± 0.8
OD/DR		56.7 ± 0.2	54.1 ± 0.2	
OD/SF			47.2 ± 0.3	
SF/OR			52.0 ± 0.2	49.7 ± 0.8
SF/DR			52.6 ± 0.1	

Table 1: Mean crossing angles between the maps of OD, OR, DR and SF for the model, and values reported by Hübener et al. (1997) for the cat.