Edinburgh Research Explorer

Mechanisms for Stable, Robust, and Adaptive Development of Orientation Maps in the Primary Visual Cortex

OPEN ACCESS

Systems/Circuits

Mechanisms for Stable, Robust, and Adaptive Development of Orientation Maps in the Primary Visual Cortex

Jean-Luc R. Stevens,[1] Judith S. Law,[1] Ján Antolík,[1,2] and James A. Bednar[1]
[1]Institute for Adaptive and Neural Computation, University of Edinburgh, Edinburgh EH8 9AB, United Kingdom and [2]Unité de Neuroscience, Information et Complexité, Centre National de la Recherche Scientifique, 91198 Gif sur Yvette, France

Development of orientation maps in ferret and cat primary visual cortex (V1) has been shown to be stable, in that the earliest measurable maps are similar in form to the eventual adult map, robust, in that similar maps develop in both dark rearing and in a variety of normal visual environments, and yet adaptive, in that the final map pattern reflects the statistics of the specific visual environment. How can these three properties be reconciled? Using mechanistic models of the development of neural connectivity in V1, we show for the first time that realistic stable, robust, and adaptive map development can be achieved by including two low-level mechanisms originally motivated from single-neuron results. Specifically, contrast-gain control in the retinal ganglion cells and the lateral geniculate nucleus reduces variation in the presynaptic drive due to differences in input patterns, while homeostatic plasticity of V1 neuron excitability reduces the postsynaptic variability in firing rates. Together these two mechanisms, thought to be applicable across sensory systems in general, lead to biological maps that develop stably and robustly, yet adapt to the visual environment. The modeling results suggest that topographic map stability is a natural outcome of low-level processes of adaptation and normalization. The resulting model is more realistic, simpler, and far more robust, and is thus a good starting point for future studies of cortical map development.

Introduction

Orientation-selective neurons in carnivoran and primate primary visual cortex (V1) form systematic topographic maps organized by preferred retinal location and orientation (Blasdel, 1992a,b; Kaschube et al., 2010). Optical imaging studies in ferrets indicate that these orientation maps develop in a stable way, with a high similarity in the spatial layout between the weakly selective initial maps and the final highly selective configuration (Chapman et al., 1996; Chapman and Bonhoeffer, 1998; Gödecke et al., 1997). This stability may be important for allowing regions downstream from V1 to develop, and is particularly remarkable given that V1 neurons undergo massive morphological changes over this time (for review, see White and Fitzpatrick, 2007; Huberman et al., 2008). V1 map development is also quite robust against differences in inputs, with similar final map patterns developing until nearly 3 weeks of age in cats, regardless of whether the eyes are open (Crair et al., 1998). Yet map development is not simply decoupled from vision, because the development of selectivity corresponds to a critical period during which visual experience is required for neurons to achieve fully mature levels of orientation tuning (Crair et al., 1998; Chapman and Gödecke, 2000; White et al., 2001; White and Fitzpatrick, 2007). Moreover, neurons and maps in animals reared in abnormal visual conditions during the critical period come to reflect the statistics of the abnormal visual input [e.g., lid suture, (Wiesel and Hubel, 1963; Blakemore and Van Sluyters, 1975; White et al., 2001), striped environments (Blakemore and Cooper, 1970; Sengpiel et al., 1999), or oriented blurring via goggle rearing (Tanaka et al., 2006, 2009)].

Together, these results indicate that orientation map and selectivity development is remarkably stable and robust against variability in inputs over time or across animals, yet adaptive to the statistics of available visual input. Although numerous computational models of activity-driven map development have been proposed (for review, see Swindale, 1996; Goodhill, 2007), no map model has yet been shown to develop stably in response to patterned inputs of different types and strengths. Understanding how maps can develop stably is both of direct scientific interest, and may provide a criterion that allows inherently unstable models to be rejected.

In this paper, we begin with a simple mechanistic model including lateral connectivity that captures the essential features of previous incremental Hebbian self-organizing network models for map development (von der Malsburg, 1973; Kohonen, 1982; Obermayer et al., 1990; Sirosh and Miikkulainen, 1994; Barrow et al., 1996; Burger and Lang, 1999, 2001; Bednar and Miikkulainen, 2004; Miikkulainen et al., 2005), while omitting mechanisms that destroy stability. Although this model develops maps, we show that it is still relatively unstable, and is not robust to changing inputs, e.g., varying input contrast. We then demonstrate that
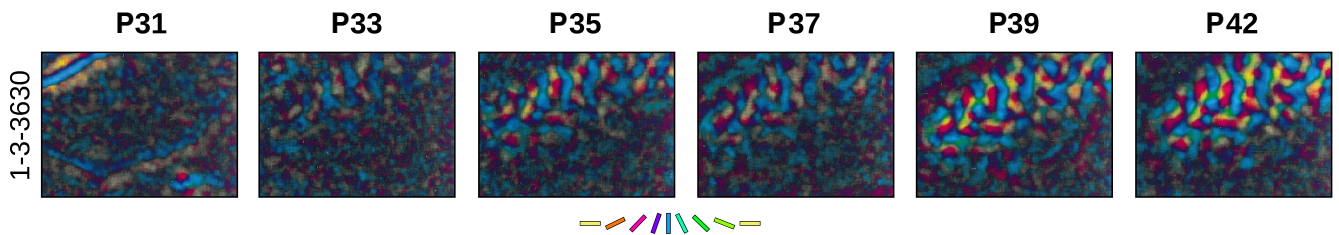
Figure 1. Stable development of orientation maps in ferret striate cortex. Orientation maps recorded in the primary visual cortex of one ferret (animal 1-3-3630) at the postnatal ages indicated are shown. In these polar maps, pixel color indicates orientation preference, and pixel brightness indicates the strength of orientation tuning. The selectivity of each map is normalized independently, making blood vessels visible at P31 (e.g., the blue streak in the top left corner), but not at P42, once orientation selectivity has developed. As the maps mature, iso-orientation domains become visible as colored regions that become more strongly responsive over time, without changing the overall map pattern. Data are reproduced from the study by Chapman et al. (1996).
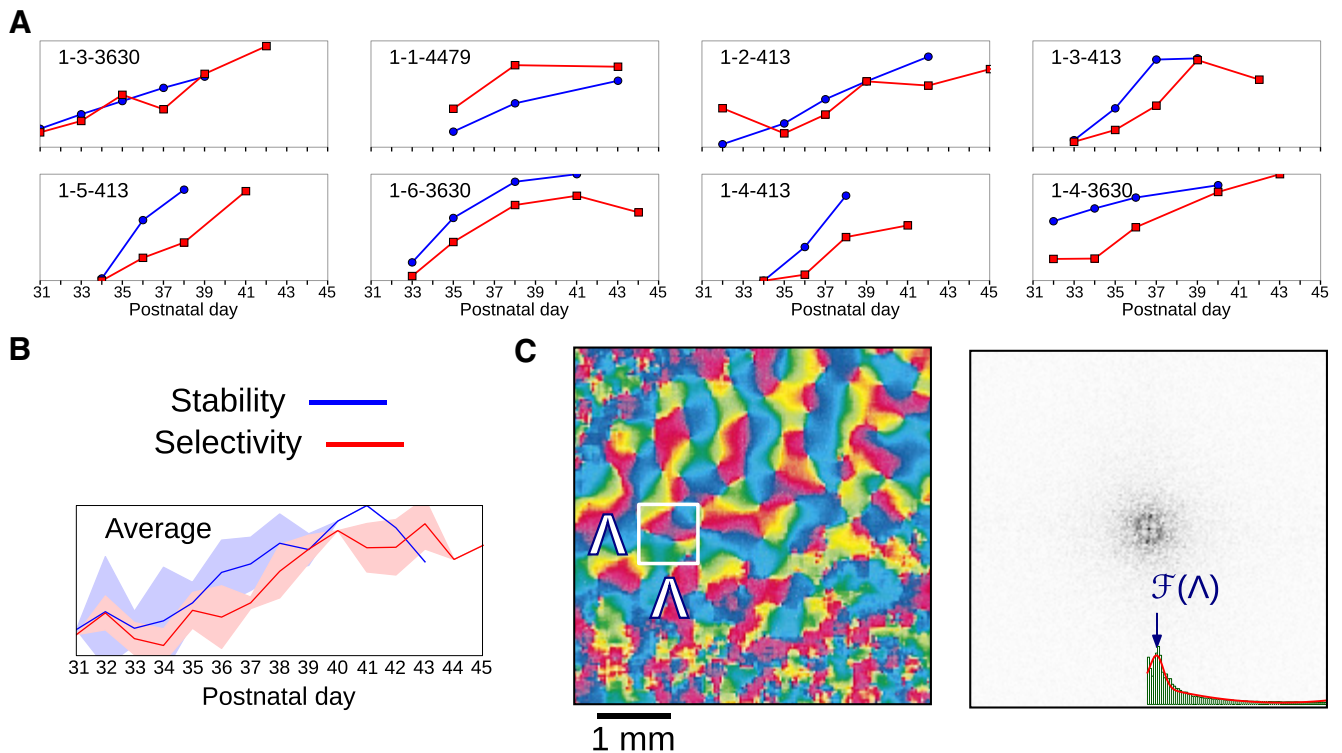

Figure 2. Development of orientation maps measured by chronic optical imaging in ferrets. **A**, Recorded selectivity (red square markers) and stability values (blue round markers) for all eight ferrets, as a function of postnatal day [replotted data from the study by Chapman et al. (1996), their Figs. 4, 7]. Stability is quantified by the average difference in preference of each orientation map with the final map, as defined by Equation 1. A common selectivity and stability scale is used to allow comparison between ferrets. All values are in arbitrary normalized units, using the lowest recorded value as the zero reference point and the highest recorded value as the maximum (see Materials and Methods). **B**, The mean selectivity and stability across ferrets as a function of postnatal day, with the $\pm 95\%$ confidence intervals for each day indicated by the shaded area around the mean line. Selectivity and stability increase steadily and simultaneously over development so that once neurons are selective, they are organized into a map with the same form as the final measured map. These results show that selectivity does not precede increasing stability, indicating stable map development. **C**, As a reference for later modeling work, the final orientation preference map (without selectivity) for ferret 1-3-360 on day P42 (Fig. 1, rightmost plot) is shown. The map is organized into regularly repeating hypercolumns in all directions, as seen by the ringness in the Fourier power spectrum, plotted with the highest amplitude component in black. The center value is the DC component, and the midpoint of each edge represents half of the highest possible spatial frequency in the cardinal directions (i.e., the Nyquist frequency). This Fourier spectrum is used to calculate the average periodicity of the map, which is then plotted as a hypercolumn area Λ^2, covering one period in the cardinal directions (white boxed area). Figure 3 illustrates how to use these calculations to determine whether a model map resembles this animal data.

extending the model with two well-established low-level properties, neuron threshold adaptation and feedforward gain control, independently improves robustness and stability, and that when used together, the final GCAL (gain control, adaptation, laterally connected) model achieves biological maps with robust and stable development, while reflecting the statistics of visual input. The results suggest that the well-established mechanisms of gain control and single-neuron homeostatic adaptation are important basic principles underlying map development, not just for individual adult or developing neurons.

Materials and Methods

The following sections first present methods used for quantifying the emergence of selectivity in biological maps, the stability of these maps over time, and the degree to which model maps resemble biological maps. We then describe the set of models used in this paper.

Measuring stability and selectivity. The stability of orientation map development is illustrated in Figure 1 with data from chronic optical imaging in ferrets, reprinted from Chapman et al. (1996). Stability and selectivity measures for maps in eight ferrets (of either sex; T. Bonhoeffer, personal communication) are shown in Figure 2 and described below.

Qualitatively, map development is considered stable if the map pattern remains constant once neurons have become highly selective; unstable development would be characterized by maps that reorganize substantially after orientation-selective patches first emerge (Chapman et al., 1996). In electrophysiological measurements from ferret V1, average selectivity (orientation tuning) is low from about postnatal day 23 (P23; when visual responses can first be measured) until P28–P34, reaching adult levels by P42–P49 (Chapman and Stryker, 1993). To examine stability of maps using optical imaging, Chapman et al. (1996) thus focused on the period from P31 to P45. This period also corresponds to the critical period for visual experience—disrupting activity before P50 eliminates selectivity even with normal visual experience after this age (Chapman and Gödecke, 2000). For this modeling study, we will also focus on this period of normal orientation selectivity emergence; stability after the critical period can be modeled trivially by disabling or greatly reducing plasticity, and will not be investigated further here.

The red lines with square markers in Figure 2A show the average selectivity of the optical imaging signal from P31–P45 for all eight ferrets (Chapman et al., 1996); one additional data point from ferret 1-1-4479 at P55 is not shown. Although the selectivity values are on an arbitrary scale due to the measurement technique, the values for each postnatal age correspond closely to those found electrophysiologically (Chapman et al., 1996), and thus cover the range from highly unselective to adult-like selectivity.

To assess stability quantitatively over this same time period, Chapman et al. (1996) computed the correlation of orientation preference at each developmental age with the organized preference map observed in the final recording for that animal. For convenience, we have linearly rescaled the "orientation similarity index" values from that study to vary from 0.0 (uncorrelated preference) to 1.0 (identical preference), which we compute as a stability index (SI):

$$\mathrm{SI} = 1 - \frac{4}{n\pi} \sum_i | (F_i - O_i) \bmod \left(\frac{\pi}{2}\right) |, \qquad (1)$$

where i iterates over all n data points (pixels in the imaging frame), F_i is the final orientation preference value for the neural unit corresponding to pixel i, and O_i is the corresponding preference value at an earlier age. The blue lines with round markers in Figure 2A show the results of this calculation for eight ferrets from Chapman et al. (1996) measured at different postnatal days. Note that even identical underlying maps would give values below 1.0 on this measure, if there is any noise or variation in measuring the map, and so the values shown have been normalized by the highest and lowest stability values across all ferrets studied.

Given this data, stable development is defined as a map having high SI values whenever the average selectivity value is high. When the selectivity is very low, presumably before the neurons have developed at all, the map is arbitrary and dominated by measurement noise, but if development is stable, the SI value should increase as soon as the selectivity value increases. Unstable development would be visible as a selectivity value that increases well before stability does, representing neurons that have achieved selectivity before assuming the final preference value observed in the last map measured. As can be seen in Figure 2A, the SI value and the average selectivity are highly correlated in every ferret shown, with both values typically increasing over time rather than selectivity predating stability (though ferret 1-1-4479 could be seen as relatively unstable). An overall correlation between selectivity and stability is a hallmark of stable map development, and is visible as a general trend in the average data from all ferrets, plotted in Figure 2B.

Assessing the organization of orientation maps. Recording the selectivity and stability index is sufficient for evaluating development in experimental recordings, where all maps are necessarily biologically realistic. In simulated models of map development, there is no guarantee that simulated maps have an orientation structure matching those observed in real animals. An unrealistic map may be both stable and selective, and so we require that a good model of map formation also results in realistic maps, similar to those found in animals. Ideally, map quality would be assessed with an automated metric that reports how close the maps are to animal data, e.g., on a scale from 0 to 1.0 as for stability.

There have been several previously published measures that could be used to quantify the degree of map organization. Biological maps tend to be smooth, which can be measured by the local homogeneity index (LHI; Nauhaus et al., 2008) or the earlier "local input orientation selectivity index" (Schummers et al., 2004), which both convey the local homogeneity in orientation preference at a particular cortical location. Biological maps typically have non-uniform orientation histograms (Coppola et al., 1998; Müller et al., 2000; Tanaka et al., 2009), which have been used to characterize model maps as well (Bednar and Miikkulainen, 2004). Biological maps tend to have a ring-shaped Fourier power spectrum (Blasdel, 1992a,b; Erwin et al., 1995; Fig. 2C), due to the regular repetition of orientation patches across the cortical surface. Orientation maps also contain pinwheels that may be identified, classified by polarity (whether orientation preference increases clockwise or counterclockwise) and analyzed; e.g., the average distance between nearest-neighbor pinwheels tends to be greater between those of matching polarity than those of opposite polarity (Müller et al., 2000).

Although these measures can each help quantify the realism of simulated maps, they do not offer a simple, unique reference value that can characterize all biological maps. A useful metric should assign biological maps to a narrow possible range of values, ideally with theoretical justification for what values should be expected in the biological system and why. Furthermore, a reliable metric should be consistently defined between species, if it is to avoid uncertainty in fitting species-dependent free parameters.

Accordingly, we developed a new map-quality metric based on the empirical observation that pinwheel count in biological orientation maps scales linearly with hypercolumn size across many different species (Kaschube et al., 2010). This linear relation gives a consistent number of pinwheels per hypercolumn area (Λ^2), implying a constant pinwheel density when averaged over sufficiently large cortical areas. This dimensionless, statistical measure of pinwheel distribution is thought to reflect a universal constant of map organization, converging to π across carnivorans, primates, cats, and tree shrews (Kaschube et al., 2010; Keil et al., 2012). This value was predicted by a theoretical model of map organization and has strong empirical evidence, with a mean pinwheel density across four species (tree shrew, galago, cat, and ferret) statistically indistinguishable from π (Kaschube et al., 2010); see data from three species in Figure 3D.

To establish the pinwheel density for any given map, the total pinwheel count must be determined. Pinwheel centers are located at the intersection of the zero contours of the real and imaginary components in the polar representation of orientation preference (Löwel et al., 1998). These contours are shown in Figure 3A as black and white lines for a sample simulated preference and selectivity map with approximately π pinwheel density. The preference-only channel of this high-quality map is shown together with corresponding pinwheel density and pinwheel count in Figure 3Ca. All the maps shown in Figure 3 are for illustration only and are derived from simulations to be described later; simulated data are necessary to demonstrate the behavior of the metric for nonbiological maps.

For simulated maps using uniform random input statistics, a single number is sufficient to characterize the Fourier plots: the radius of the isotropic ring, computed using the methods described in the study by Kaschube et al. (2010). To eliminate disruptions in hypercolumn size due to border effects at the sheet edges, only the central 1.0×1.0 area from a simulated V1 sheet of area 1.5×1.5 is analyzed for all model maps presented. The first step is to integrate the spectral power as a function of radius, as shown by the histograms in the FFT plots. The peak values of these histograms are estimated using a least-squares fit of a Gaussian curve with additional linear and quadratic terms (Kaschube et al., 2010, their Equation 7, supplementary materials). The computed fits are shown by the red curves overlaying the histograms with the fitted peaks indicated by the blue arrows, marking the estimated ring radius. This value reflects the periodicity with which hypercolumns repeat across the map and is equal to the hypercolumn area in units of Λ^2 when squared, shown for a ferret map in Figure 2C and for simulated maps in Figure 3, A and B.

An identical analysis is shown for a lower-quality simulated orientation map in Figure 3B, where the map was generated by running a simulation outside its optimal operating range. Disruptions to the smooth organization of preference result in clumps of pinwheels, explaining the higher pinwheel density seen in Figure 3Cb. These clumps of pinwheels could in principle be eliminated arbitrarily as duplicates, but because they represent genuine areas of poorly organized preferences in these
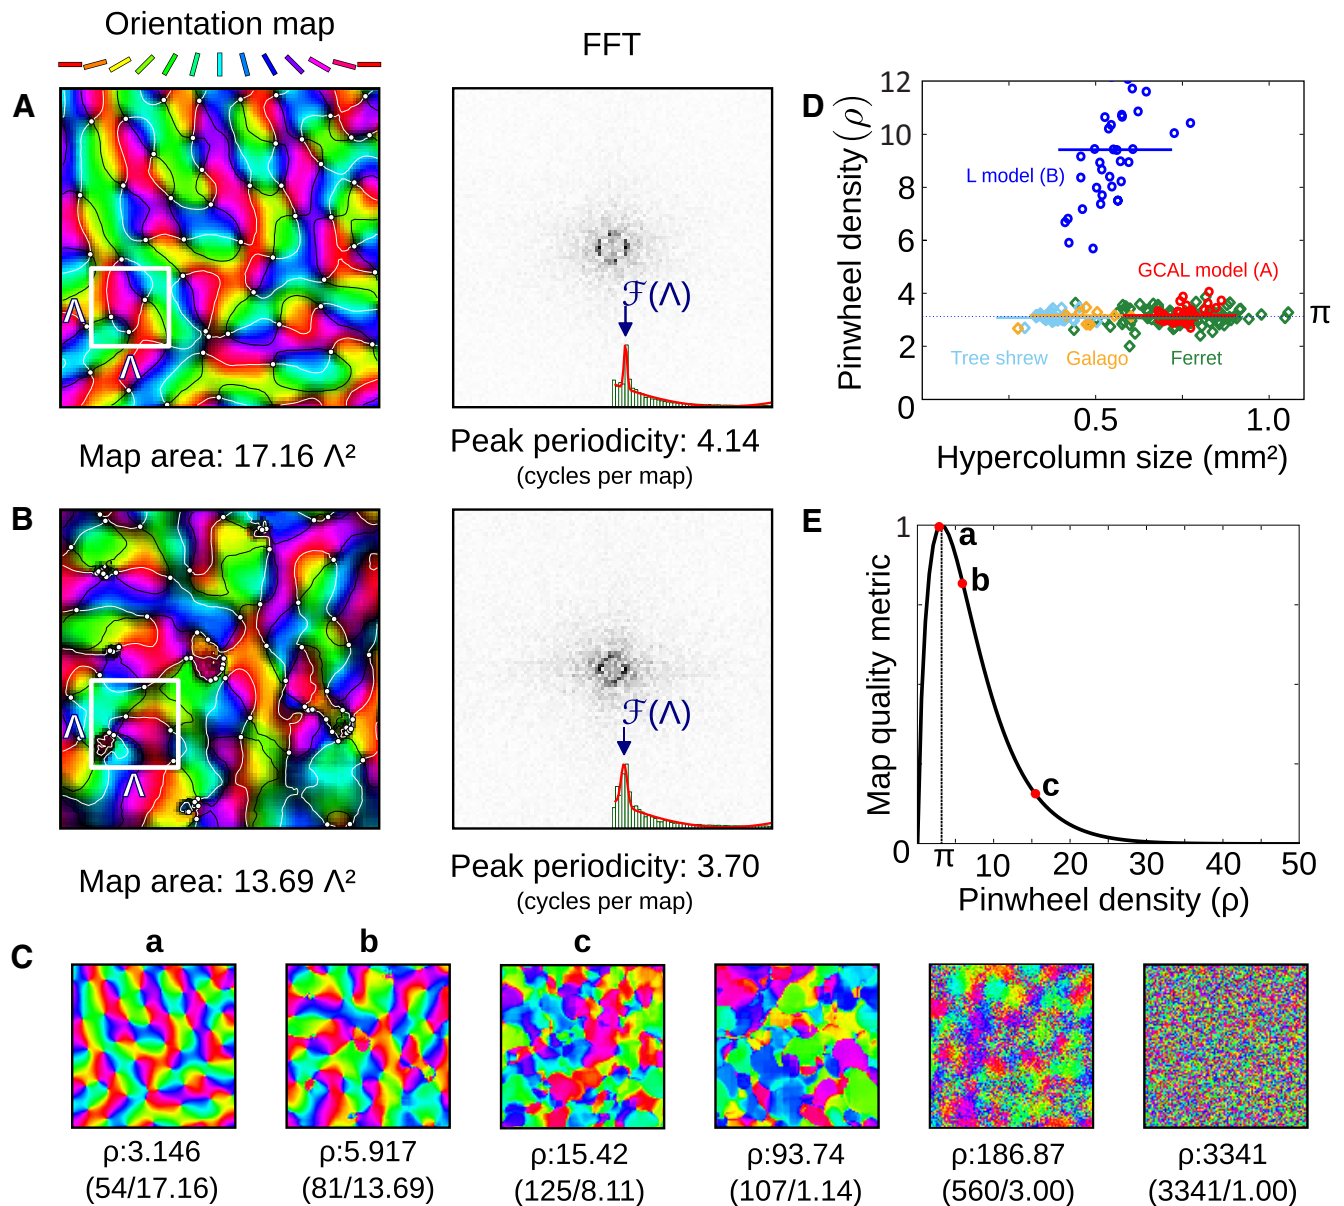
Figure 3. Evaluating map quality as the deviation from π pinwheel density (ρ). **A**, High-quality, realistic, orientation preference and selectivity map with approximately π pinwheel density (3.146), from the final model to be discussed in this paper (GCAL). Corresponding preference-only map is shown in **Ca**. Pinwheel density is defined as the average number of pinwheels (white circles) per hypercolumn area (Λ^2) indicated by the white boxed area (**A, B**). Pinwheels are identified at the intersection of the zero contours of the real and imaginary components in polar representation (white and black contours respectively). The periodicity of hypercolumns is estimated from the radius of the ring in the Fourier power spectrum (FFT) using the fitting method described in Kaschube et al. (2010). **B**, A lower-quality map generated from the first model introduced in this paper (L), which has visible discontinuities in OR preference seen in **Cb**. The greater pinwheel density (5.917) is due to a higher pinwheel count and a larger hypercolumn area due to more widely spaced orientation blobs. The histogram (FFT plot inset) indicates mean spectral power as a function of radius, the red line indicates the least-squares fit (Kaschube et al. 2010, their supporting online material, Eq. 7), and the blue arrow indicates the estimated peak spectral power radius (**A, B**). **C**, A selection of model maps ordered by pinwheel density with pinwheel count to hypercolumn area ratio, shown in parentheses. Lower-quality maps usually have higher estimated pinwheel counts and correspondingly higher pinwheel densities, with pinwheel counts so large for very poor maps as to be effectively undefined. **D**, Pinwheel density of three species (diamonds) and simulated maps (circles) as a function of hypercolumn size [data replotted from the study by Kaschube et al. (2010)]. Horizontal lines indicate median values of each cluster with the medians of high quality model maps also clustered around π. **E**, Normalized, heavily tailed gamma distribution used to transform pinwheel density into a suitable metric on a unit interval with values for maps a, b, and c.

noise-free simulations, they have been retained as indicators of poor map structure. This trend is illustrated in Figure 3C by a selection of maps of increasingly poor quality that illustrate how the pinwheel count and density tend to increase as map quality decreases.

The validity of using π pinwheel density as a reference value is demonstrated in Figure 3D. The experimental values of pinwheel density for tree shrews, galagos, and ferrets are plotted as diamonds, using data reproduced from the study by Kaschube et al. (2010). The values for simulated maps are plotted as circles for the high- and low-quality map simulations ($n = 40$ and $n = 33$ respectively). The

high-quality maps are obtained from the final model developed in this paper, while the lower-quality maps are computed using the initial, simplified model introduced in the next section. The horizontal lines through each cluster indicate the median value of the cluster points. The median is indistinguishable from the mean for all clusters except the low-quality map cluster, which has outliers of relatively high pinwheel density not visible in the plot.

It would be interesting to evaluate pinwheel density over development using the chronically recorded ferret data, illustrated in Figure 1. This measurement is not possible without access to the raw data, as prepro-

cessing may have introduced biases to the orientation structure that need to be carefully controlled for when assessing pinwheel density (Kaschube et al., 2010). The pinwheel density at the end of maturation in ferrets has been established using 82 maps recorded from adult animals, shown by a cluster of green diamonds centered around π in Figure 3D.

For the poorest-quality maps, the pinwheel density is unbounded, as illustrated by Figure 3C. In these highly disorganized maps (to the right of Fig. 3C), an unrealistically large number of locations match the automated criterion for a pinwheel; few or any of these would satisfy any subjective definition of a pinwheel. Even so, the "pinwheel count" is still a useful measure in these cases, because such high numbers reliably indicate low-quality maps.

To turn pinwheel density into a useful metric between unity (high-quality maps) and zero (low-quality maps), a heavily tailed squashing function is needed, mapping all the maps with unrealistically high pinwheel counts to zero. We chose a normalized gamma distribution, shown in Figure 3E. The labels a, b, and c show how the final metric value is computed for the three example maps shown. The gamma distribution $\Gamma(k, \theta)$ is characterized by $k = 1.8$ and $\theta = \pi/(k - 1)$. This latter constraint on θ ensures that the mode of the curve occurs at π. This kernel is then normalized by the value evaluated at π to ensure that maps with exactly π pinwheel density have a metric value of 1. The only free parameter k is set to the value $k = 1.8$ to allow detailed discrimination of high-quality maps without discarding all the information available for poor maps with a pinwheel density up to 30; other values will place more or less penalty for very poor maps with unrealistically high pinwheel densities, but will otherwise lead to similar results.

Of course, a value of unity on the pinwheel density metric (or any metric) is not sufficient by itself to guarantee that a given simulated map is indistinguishable from a biological recording. For instance, it should be possible to construct a synthetic map designed specifically to attain π pinwheel density that nonetheless appears highly unnatural. Such deliberate "gaming" of the map metric is not possible using the simulations presented in this paper, as the maps emerge gradually from a developmental process, without modeler control over the specific placement of pinwheels in the final organization. Even so, it is important to verify that simulated maps with a high metric value satisfy the types of subjective criteria used to evaluate biological results, such as those introduced by Blasdel (1992b). For each result where maps are evaluated by metric, we have thus also performed subjective evaluations of map quality, such as assessment of whether the FFT is ring shaped, and whether there are saddle points, linear zones, 180° pinwheels, fractures, and so on. We have found no examples of simulated maps that achieve a realistic pinwheel density but do not also look realistic with respect to these other properties. The automated pinwheel density metric value thus correlates well with our overall subjective assessment of map quality. We have also ensured that all the plots presented in this paper are representative of those simulations that are not shown, for the same conditions.

Model architecture. The mechanistic models we will consider are variants of the original SOM (self-organizing map) algorithm introduced by von der Malsburg (1973). In this type of model, the inputs are actual visual images or patterns of spontaneous activity, which can be directly related to the visual environment or to the imaging of retinal activity patterns. Biologically plausible properties of single neurons can be integrated into the network to explain many of the observed phenomena, as opposed to just the geometric properties of the map pattern.

Unlike correlation-based learning models that use Hebbian learning over large batches of inputs (Linsker, 1986; Miller, 1994), self-organizing map models operate using incremental Hebbian learning rules. Although the mathematical structure of the final model is often less amenable to analysis than those of linear, feedforward networks, the incremental nature of self-organizing maps make them more suitable for studying map development. Incremental learning allows gradual changes in network organization to be tracked as a stream of inputs drive the development of the network forward.

The models presented here self-organize using the same principles as the LISSOM (Laterally Interconnected Synergetically Self-Organizing Map) algorithm (Miikkulainen et al., 2005), which was inspired by earlier SOM models (Kohonen, 1982; Obermayer et al., 1990). The simplest model presented here, model L (laterally connected), may be considered a simplified version of the LISSOM model, omitting ad hoc modeler-

determined changes to the Hebbian learning rate, activation thresholds, and lateral excitatory radii over time. The fundamental operation of all these self-organizing map models has been explained in terms of dimensionality reduction, specifically a discretized approximation of the principal surface of the input (Ritter et al., 1992).

The architecture of the four models evaluated in this paper is shown in Figure 4. The models each consist of four sheets of neural units representing the input (retinal photoreceptors), the ON-center and OFF-center pathways from the photoreceptors to V1 via the retinal ganglion cells and the lateral geniculate nucleus, and V1. A sheet is a two-dimensional array of firing-rate point neurons, with activation and plasticity equations as described below. The simplest variant, model L (laterally connected) consists of four afferent connections (one from the photoreceptor sheet to each ON/OFF sheets and one from each ON/OFF sheet to V1) with lateral connectivity only in the V1 sheet.

The AL (adaptation, laterally connected) model has the same set of connections as L but includes homeostatic adaptation in V1, described by the equations below. The GCL (gain control, laterally connected) model has the same V1 sheet structure as L but includes lateral, divisive inhibition in the ON/OFF sheets that implement contrast-gain control. The final GCAL model includes both the homeostatic adaptation of the AL model as well as the contrast-gain control of the GCL model. Apart from these specific differences, each model shares the same parameters and mechanisms, and can thus be compared directly against the others.

As model L has the same basic architecture as the LISSOM model, the spatial extents of connections and weight patterns are taken from and described in the study by Miikkulainen et al. (2005). Models GCL and GCAL introduce one new spatial parameter, determining the lateral extent of the contrast-gain control mechanism in the LGN layers. No parameters need to be specified with very high precision for our conclusions to hold; in the GCAL model, two significant digits of precision are sufficient to develop qualitatively indistinguishable maps. We estimate that all parameters may be changed by ~10% without affecting the overall behavior (except for the homeostatic smoothing parameter, which may not be >1.0).

The models are implemented in the Topographica simulator, freely available at www.topographica.org. Topographica allows simulation parameters to be specified in measurement units that are independent of the level of detail used in any particular run of the simulation. To achieve this, Topographica provides multiple spatial coordinate systems, called sheet and matrix coordinates. Parameters (such as sheet dimensions) are expressed in continuous-valued sheet coordinates. In practice, of course, sheets are discretized using some finite matrix of units. Each sheet has a density, which specifies how many units (matrix elements) in the matrix correspond to a unit length in sheet coordinates. Each of these simulations used a density of 98 neural units per sheet coordinate in V1, and 24 units per sheet coordinate in the retinal photoreceptor and ON/OFF sheets. Results are independent of the density used, except at very low densities where discretization artifacts can become prominent.

In all simulations shown, the area of the V1 sheet plotted and analyzed has sheet-coordinate dimensions 1.0×1.0, thus consisting of 98×98 neural units. The underlying simulated V1 area is actually 1.5×1.5 (147×147 neural units), which is then cropped to 1.0×1.0 for analysis and display, to eliminate border effects due to partial patterns of lateral connectivity for neurons near the edge of the cortical sheet. The ON/OFF channels and photoreceptors are similarly extended for simulation to 3.0×3.0 and 3.75×3.75 in sheet coordinates, respectively, to avoid having any afferent receptive field cropped off, and to avoid edge effects in the ON and OFF channels due to gain control connections when present. In Figure 4, all sheets are shown at the same size for visibility, but the actual area mapped topographically between sheets was the central 1.5×1.5 region of each sheet that corresponds to V1.

Temporal properties. As a simplification, we have ignored the detailed temporal properties of the subcortical neural responses and of signal propagation along the various types of connections. Instead, the model ON/OFF units have a constant, sustained output, and all connections have a constant delay, independent of the physical length of that connection. One training iteration in the model represents one visual fixation (for natural images) or a snapshot of the relatively slowly changing spatial pattern of spontaneous

activity (for retinal waves); i.e., an iteration consists of a constant retinal activation, followed by processing at the ON/OFF and cortical levels.

For one iteration, assume that input is drawn on the photoreceptors at time t, and the connection delay is defined as $\delta t = 0.05$. Then, at $t + 0.05$, the ON/OFF cells compute their initial activation; at time $t + 0.10$, the ON/OFF cells incorporate lateral inhibition for gain control (if present in this model); and at $t + 0.15$, V1 begins computing. At $t + 0.20$ and every 0.05 iterations until $t + 1.0$, V1 continues to compute in a series of settling steps, a fixed number sufficient to ensure that cortical activity is no longer changing significantly (data not shown). During this period, the retinal activity is assumed to be constant, to reduce computational requirements.

Stimuli. Images are presented to the model at each iteration by activating the retinal photoreceptor units. The activation value ψ_i of unit i in the photoreceptor sheet (P) is given by the grayscale value in the chosen image at that point.

The images used here are either oriented and elongated two-dimensional Gaussian patterns, disk-shaped patterns of noisy activation, or natural image patches, which together cover a range of different input types so that we can evaluate the robustness of map development. The center coordinates and orientation of each Gaussian are chosen from a uniform random distribution and cover an area of the photoreceptor sheet that is a third wider and taller than V1, so that even V1 units near the borders will see an approximately uniform distribution of each part of the Gaussian blobs. For Gaussian patterns, the contrast is defined as the percentage of the input range 0.0 to 1.0.
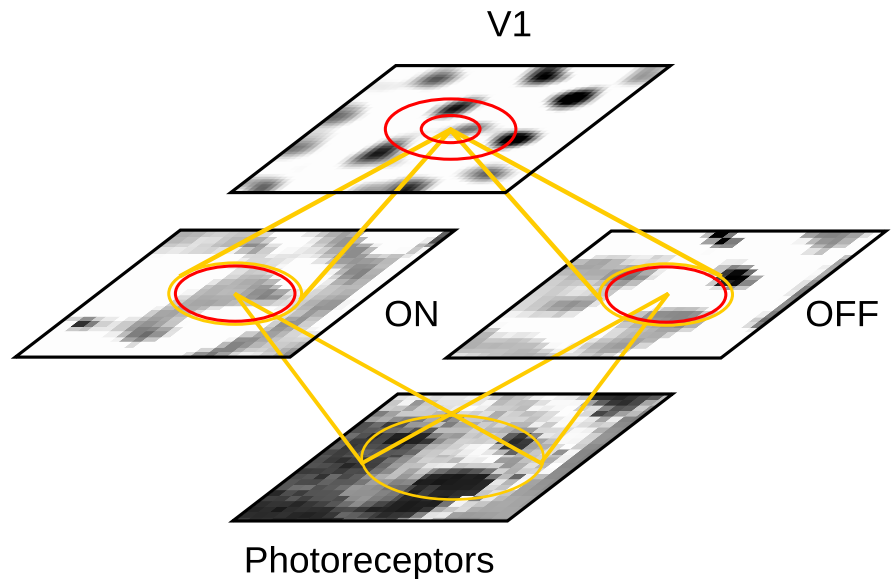

Figure 4.　General model architecture. The four models discussed in this paper consist of two-dimensional arrays of computational units representing local groups of neurons at each visual processing stage. Connections to one unit in each sheet are shown with afferent connections in yellow and lateral connections in red. The V1 region shown here covers ∼4 × 4 mm of ferret cortex, matching the cortical area represented by all the simulated orientation map plots in this paper. V1 units receive lateral excitatory (small circle) and lateral inhibitory (large circle) connections from nearby V1 units, leading to the "bubbles" of activity seen on the V1 sheet. V1 units receive afferent input from sheets representing the ON and OFF channels, representing the action of retinal ganglion cells and the lateral geniculate nucleus, which in turn receive input from the retinal photoreceptors. The difference-of-Gaussian afferent connections from the photoreceptors to the ON and OFF units form a local receptive field on the retina, and cause ON-center units to respond to light areas surrounded by dark, and OFF-center units to dark areas surrounded by light. Neighboring neurons in the ON and OFF sheets have different but overlapping receptive fields. Input to the retinal photoreceptors can be any type of patterned image, such as the small natural image patch shown here. The GCL and GCAL models have additional lateral inhibitory connections within the ON and OFF sheets, but otherwise all four models discussed share identical initial weight profiles, spatial extents, and projection strengths.

Noisy disk patterns create a circular region of activity with a radius that is a substantial portion of the size of the photoreceptor sheet to generate distinct edges; the activation is smoothed at the edges of the disk with a Gaussian blur (Bednar and Miikkulainen, 2004). Disk centers are chosen from a uniform random distribution larger than the retinal photoreceptor sheet size to allow for V1 units to see a uniform distribution of disk areas. Uniform zero-mean random noise is then added to the disk pattern.

Natural image patterns are photoreceptor-sheet-sized patches from images of natural objects and landscapes, from a data set by Shouval et al. (1996). The pictures were taken at Lincoln Woods State Park in Rhode Island, and scanned into a 256 × 256 pixel image. No corrections were used for the optical distortions of the instruments, and there was no preprocessing of the images. The data set used to approximate vertical goggle rearing shown in Figure 11 consists of the same set of image patches, but blurred by convolution with an anisotropic Gaussian kernel. The convolution kernel is a 128 × 128 matrix, consisting of a centered, vertical Gaussian pattern with an aspect ratio of 10 ($\sigma_x = 0.25$, $\sigma_y = 0.025$ in unit sheet coordinates) and total weight of unity.

The ON/OFF sheets. At each iteration, a new retinal input is presented, and the activation of each unit in each sheet is updated. Neurons in all sheets are firing-rate point neurons, with a state characterized by an activation level. For all models, the activation level η for a unit at position j in an ON/OFF sheet O at time $t + \delta t$ is defined as follows:

$$\eta_{j,O}(t + \delta t) = f\left(\frac{\gamma_O \sum_{i \in F_{j,P}} \Psi_i(t)\omega_{ij}}{k + \gamma_S \sum_{i \in F_{j,S}} \eta_{i,O}(t)\omega_{ij,S}} \right). \quad (2)$$

The constant $\gamma_O = 14.0$ is an arbitrary multiplier for the overall strength of connections from the photoreceptor sheet to the ON/OFF sheets, chosen to give typical activations in the range 0.0 to 1.0, whereas γ_S is the strength of the feedforward contrast-gain control; ψ_i is the activation of unit i in the two-dimensional array of neurons on the photoreceptor

sheet from which ON/OFF unit j receives input (its afferent connection field $F_{j,P}$), and $\eta_{i,O}(t)$ is the activation of other ON/OFF units on the previous time step (received over the suppressive connection field $F_{j,S}$). The activation function f is a half-wave rectifying function that ensures the activation of ON/OFF units is always positive.

In the L and AL models where contrast-gain control is not applied, $k = 1$ and $\gamma_S = 0$, whereas in the GCL and GCAL models, $k = 0.11$ and $\gamma_S = 0.6$. The constant k ensures that the output is always well defined for weak inputs, and γ_S is chosen to rescale activation values so that the numerical results are comparable with and without gain control.

The weights ω_{ij} represent the fixed connection weights from photoreceptor i to the ON or OFF unit j defined with a standard difference-of-Gaussians kernel. The connection fields for ON units have a positive center and negative surround, and vice versa for OFF units. More precisely, the weight ω_{ij} from an ON-center cell at location $(0, 0)$ in the ON sheet and a photoreceptor sheet in location (x, y) on the photoreceptor sheet is given by the following:

$$\omega_{ij} = \frac{1}{Z_C} \exp\left(-\frac{x^2 + y^2}{2\sigma_C^2} \right) - \frac{1}{Z_S} \exp\left(-\frac{x^2 + y^2}{2\sigma_S^2} \right). \quad (3)$$

The width of the central Gaussian is defined by $\sigma_c = 0.037$, and $\sigma_s = 0.15$ determines the width of the surround Gaussian, where Z_c and Z_s denote the normalization constants that ensure the center and surround weights each always sum to 1.0. The weights for an OFF-center cell are the negative of the ON-center weights (i.e., surround minus center). The center of the connection field of each ON/OFF unit is mapped to the location in the photoreceptor sheet corresponding to the location of that unit in sheet coordinates, making the projection retinotopic.

The weights $\omega_{ij,S}$ in the denominator of Equation 2 specify the spatial profile of the lateral inhibition received from other ON/OFF units when contrast-gain control is active. The weights of these connections have a fixed, circular Gaussian profile so that a neuron located at $(0, 0)$ in either the ON or OFF sheet will have the following weights:

$$\omega_{ij,S} = \frac{1}{Z_S} \exp\left(-\frac{x^2 + y^2}{2\sigma_S^2}\right), \qquad (4)$$

where (x, y) is the location of the presynaptic neuron, $\sigma_S = 0.125$ determines the width of the Gaussian, and Z_S is a normalizing constant that ensures that the total of all the lateral inhibitory weights ω_{ij} to neuron j sum to 1.0. When contrast-gain control is enabled, these recurrent lateral connections are activated once per iteration, before activity is sent to the V1 sheet.

The V1 sheet. Each V1 neuron in each model receives connections from three different connection types or "projections" (p), i.e., the afferent projection from the ON/OFF sheets (both channels concatenated into one input vector; $p = A$), the recurrent lateral excitatory projection ($p = E$), and the recurrent lateral inhibitory projection ($p = I$) from other V1 neurons.

The contribution $C_{j,p}$ to the activation of unit j from each projection type ($p = A, E, I$) is calculated as follows:

$$C_{j,p}(t + \delta t) = \sum_{i \in F_{j,p}} \eta_{i,p}(t)\omega_{ij,p}, \qquad (5)$$

where $\eta_{i,p}$ is the activation of unit i taken from the set of neurons in V1 to which unit j is connected (its connection field F_j), and $w_{ij,p}$ is the connection weight from unit i in V1 to unit j in V1 for the projection p. Afferent activity ($p = A$) remains constant after the first update from the retina, but the other contributions change over 16 settling steps, depending on the activity in V1.

The contributions from all three projections to V1 (afferent and lateral) described above are combined using Equation 6 to calculate the activation of a neuron j in V1 at time t:

$$\eta_{j,V}(t) = f\left(\sum_p \gamma_p C_{jp}(t)\right). \qquad (6)$$

The projection strength scaling factors for each projection type p are $\gamma_A = 1.5$, $\gamma_E = 1.7$, and $\gamma_I = -1.4$ for all models, set to provide a balance between excitation and inhibition, and between afferent and lateral influences, to provide robust formation of activity bubbles that allows smooth maps to form. For models L and GCL, f is a half-wave rectifying function that ensures positive activation values. In the case where single-neuron automatic adaptation is included (models AL and GCAL), f has a variable threshold point (θ) dependent on the average activity of the unit as described in the next subsection, but in all cases the gain is fixed at unity.

Once all 16 settling steps are complete, the settled V1 activation pattern is deemed to be the V1 response to the presented pattern. At this point we use the V1 response to update the threshold point (θ) of V1 neurons (using the adaptation process described below) and to update the afferent and lateral inhibitory weights via Hebbian learning. V1 activity is then reset to zero, and a new pattern is presented. Note that both adaptation and learning could instead be performed at every settling step, but this would greatly decrease computational efficiency.

Adaptation. To set the threshold for activation, each neuron unit j in V1 calculates a smoothed exponential average of its settled activity patterns ($\bar{\eta}_j$):

$$\bar{\eta}_j(t) = (1 - \beta)\eta_j(t) + \beta\bar{\eta}_j(t - 1). \qquad (7)$$

The smoothing parameter ($\beta = 0.991$) determines the degree of smoothing in the calculation of the average. ($\bar{\eta}_j$) is initialized to the target average V1 unit activity (μ), which for all simulations is $\bar{\eta}_{jA}(0) = \mu = 0.024$. The threshold is updated using the following:

$$\theta(t) = \theta(t - 1) + \lambda(\bar{\eta}_j(t) - \mu), \qquad (8)$$

where $\lambda = 0.01$ is the homeostatic learning rate. The effect of this scaling mechanism is to bring the average activity of each V1 unit closer to the specified target. If the activity in a V1 unit moves away from the target during training, the threshold for activation is thus automatically raised

or lowered to bring it closer to the target. Note that an alternative rule with only a single smoothing parameter (rather than β and λ) could be formulated, but the rule as presented here makes it simple for the modeler to set a desired target activity μ.

Learning. Initial connection field weights are isotropic 2D Gaussians for the lateral excitatory projection and uniformly random within a Gaussian envelope for afferent and lateral inhibitory projections. Specifically, a neuron located at (i, j) will have the following weights:

$$\omega_{ij} = \frac{1}{Z_p} u \exp\left(-\frac{x^2 + y^2}{2\sigma_p^2}\right), \qquad (9)$$

where (x, y) is the sheet-coordinate location of the presynaptic neuron, $u = 1$ for the lateral excitatory projection ($p = E$), and u is a scalar value drawn from a uniform random distribution for the afferent and lateral inhibitory projections ($p = A, I$), σ_p determines the width of the Gaussian in sheet coordinates ($\sigma_A = 0.27$, $\sigma_E = 0.025$, $\sigma_I = 0.075$), and Z_p is a constant normalizing term that ensures that the total of all weights ω_{ij} to neuron j in projection p is 1.0. Weights for each projection are only defined within a specific maximum circular radius r_p ($r_A = 0.27$, $r_E = 0.1$, $r_I = 0.23$).

In the model, as images are presented to the photoreceptors, V1 afferent connection weights $\omega_{ij,A}$ from the ON/OFF sheets are adjusted once per iteration (after V1 settling is completed) using a simple Hebbian learning rule. This rule results in connections that reflect correlations between the presynaptic ON/OFF unit activities and the postsynaptic V1 response. Hebbian connection weight adjustment at each iteration is dependent on the presynaptic activity, the postsynaptic response, and the Hebbian learning rate:

$$\omega_{ij,p}(t) = \frac{\omega_{ij,p}(t - 1) + \alpha\eta_j\eta_i}{\sum_k(\omega_{kj,p}(t - 1) + \alpha\eta_j\eta_k)}, \qquad (10)$$

where for unit j, α is the Hebbian learning rate for the afferent connection field F_j. Unless it is constrained, Hebbian learning will lead to ever-increasing (and thus unstable) values of the weights (Rochester et al., 1956). In all of the models, the weights are constrained using divisive postsynaptic weight normalization (Eq. 10), which is a simple and well-understood mechanism. Afferent connection weights from ON and OFF units are normalized together in the model. We expect that a more biologically motivated homeostatic mechanism for normalization such as multiplicative synaptic scaling (Turrigiano, 1999; Turrigiano and Nelson, 2004; Sullivan and de Sa, 2006) or a sliding threshold for plasticity (Bienenstock et al., 1982) would achieve similar results, but we have not tested these.

The learning rates for the afferent projection, lateral excitatory projection, and lateral inhibitory projections are $\alpha_A = 0.1$, $\alpha_E = 0.0$, and $\alpha_I = 0.3$, respectively. The density-specific value used in the equation above is then calculated as $\alpha = \alpha_A/\tau_A$, where τ_A is the number of connections per connection field in the afferent projection. To increase computational efficiency, lateral excitatory connections do not learn during development in the simulations presented here. The effect of lateral learning has been explored in detail in previous similar models (Miikkulainen et al., 2005), and the maps generated by the GCAL model are visually indistinguishable regardless of whether or not lateral excitatory connections are plastic (data not shown).

Analysis of model maps. Model orientation maps are calculated based on the vector average method (Blasdel and Salama, 1986; Miikkulainen et al., 2005). Sine grating inputs that cover the full range of parameter values (combinations of all orientations, frequencies, and phases) are presented, and for each orientation, the peak response of the neuron is recorded. The orientation preference is calculated by constructing a vector for each orientation θ (between 0 and 180°), with the peak response as the length and 2θ as its orientation. These vectors are summed and the preferred orientation is calculated as half of the orientation of the summed vector. The selectivity is given by the magnitude of the summed vector. Average orientation selectivity values are the mean value of orientation selectivity across all units in the map and are normalized by dividing by the maximum average selectivity measured across all simulations reported in this paper.
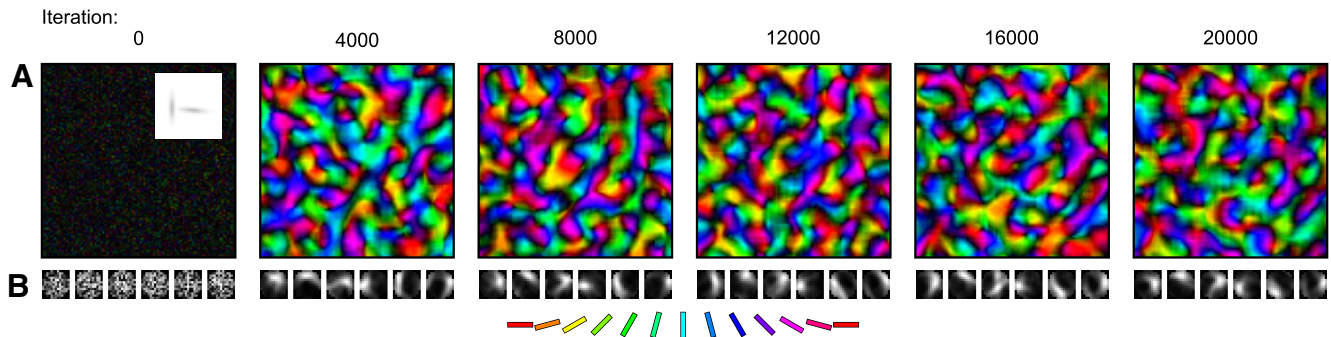
Figure 5. Model L develops maps, but is not stable. ***A***, Model development at six different iteration time points for a single simulation. Self-organization was driven by two elongated Gaussian patterns at 25% contrast per iteration, with an example inset into the plot for iteration 0. Polar orientation maps from the beginning of development to the final map at iteration 20,000 are shown. Each unit is color coded according to orientation preference, as shown by the color key. The brightness and saturation of the color indicates the strength of orientation tuning of the afferent connections, and each panel corresponds to ~4 × 4 mm of visual cortex (a 1.0 × 1.0 area in sheet coordinates). ***B***, Afferent connections from the ON sheet to V1 are shown for an arbitrary set of V1 units throughout development. Initially random connections are strengthened and weakened by Hebbian learning, forming orientation-selective receptive fields, but the map patterns change significantly over time. These results, while an improvement over existing models that have additional, biologically implausible mechanisms for reducing stability to improve their final map organization, represent a baseline for the results of the later models (AL, GCL, and GCAL). The behavior of L is analyzed further in Figure 6, and the last map shown corresponds to Figure 6E.

Orientation tuning curves are measured by presenting sine gratings with 20 different orientations, each at eight different phases and for a range of contrasts. For each contrast, the responses of each neuron unit are measured, and the maximum response at each orientation over all phases is recorded. The tuning curve is constructed from these maximum responses.

To ensure consistent map measurement across all conditions, orientation maps in Figures 3–9 are measured before lateral interactions and the V1 activation function take affect (i.e., on the afferent input activity only). Maps for models with gain control change little if these steps are included. However, for models without gain control, if activation thresholds and lateral interactions were simulated during map measurement, maps would need to be measured at many different contrasts, because orientation tuning in these models is not contrast invariant, and therefore orientation selectivity is difficult to determine. Instead, we use the orientation selectivity of the afferent connections as a measure of the map selectivity in all cases. This simplification allows fair and consistent comparison of selectivity values between the simulations of all four models. Similarly, connection field plots show the weights of connections between the ON sheet and V1 at different times during development. These weights are used as an approximation to the receptive fields that can be measured using a computationally intensive reverse correlation process that takes into account the full connectivity, which again must control for contrast in the models without gain control. The true orientation selectivity of neurons in the full recurrent network is reflected in the orientation tuning curves for the GCAL model (as shown in Fig. 12).

Results

In this section we present results from four closely related developmental models of orientation map formation, all sharing the architecture illustrated in Figure 4. As described in Materials and Methods, each model is a network of single-compartment firing-rate units, each receiving afferent and lateral connections that are modified by Hebbian learning. The models share the same four sheets of simulated neurons representing the pathway from the retina to V1, with the same dimensions, connection radii, initial weights, and all other parameter settings.

The robustness of each model will initially be analyzed across a wide range of input contrasts for an artificial input pattern (randomly positioned and oriented elongated two-dimensional Gaussians). These simple patterns help make our reasoning about how each mechanism affects map development clear and intuitive and are sufficient for evaluating each model's robustness to input contrast. In turn, contrast is used as a well-defined and easily characterized proxy for a variety of changes in input pat-

terns that could affect how much V1 is activated for a given input, and could thus affect stability and robustness. Once all four models have been evaluated with respect to contrast for Gaussians, we will show that the results still hold (with no changes in model parameters) for more complex changes in the input images that are more difficult to characterize.

For each model, we evaluate the orientation selectivity, map stability, and map quality, and relate the results to data from ferrets. The first model is a simple laterally connected model of orientation map development that we will refer to as model L. Model L retains the essential features of previous models in this genre (Sirosh and Miikkulainen, 1997; Burger and Lang, 1999, 2001; Miikkulainen et al., 2005), while omitting biologically unrealistic mechanisms like lateral radius shrinking that also reduce stability, as well as eliminating arbitrary manual interventions like threshold changes or learning rate changes over development.

Subsequent models build on L, adding single-neuron homeostatic threshold adaptation (model AL), feedforward gain control (model GCL), and finally both mechanisms together (model GCAL). Justification for each mechanism is discussed as it is introduced, but the details of these two mechanisms are described in Materials and Methods. The results show that as well as reproducing stable development, the final GCAL model is robust to extreme changes in the input statistics, even during the process of development.

Model L: a simple model of map development

Results from simulated development in the L model are shown in Figure 5. Each V1 model neuron receives excitatory connections from an already retinotopically aligned set of ON/OFF neurons and excitatory and inhibitory connections from other V1 neurons (Fig. 4). At each iteration, two elongated and oriented Gaussian patterns are presented on the model retina. The afferent and lateral connections lead to the formation of isolated bubbles of activity in V1 in response to the input pattern. Initially random afferent weights between the ON/OFF units and V1 then adapt based on the presynaptic and postsynaptic activity, via a simple Hebbian learning rule with divisive normalization (Eq. 10). Neurons in each bubble thus learn their topographically corresponding input pattern, while neurons in other bubbles learn different patterns. Over the course of development (20,000 iterations),
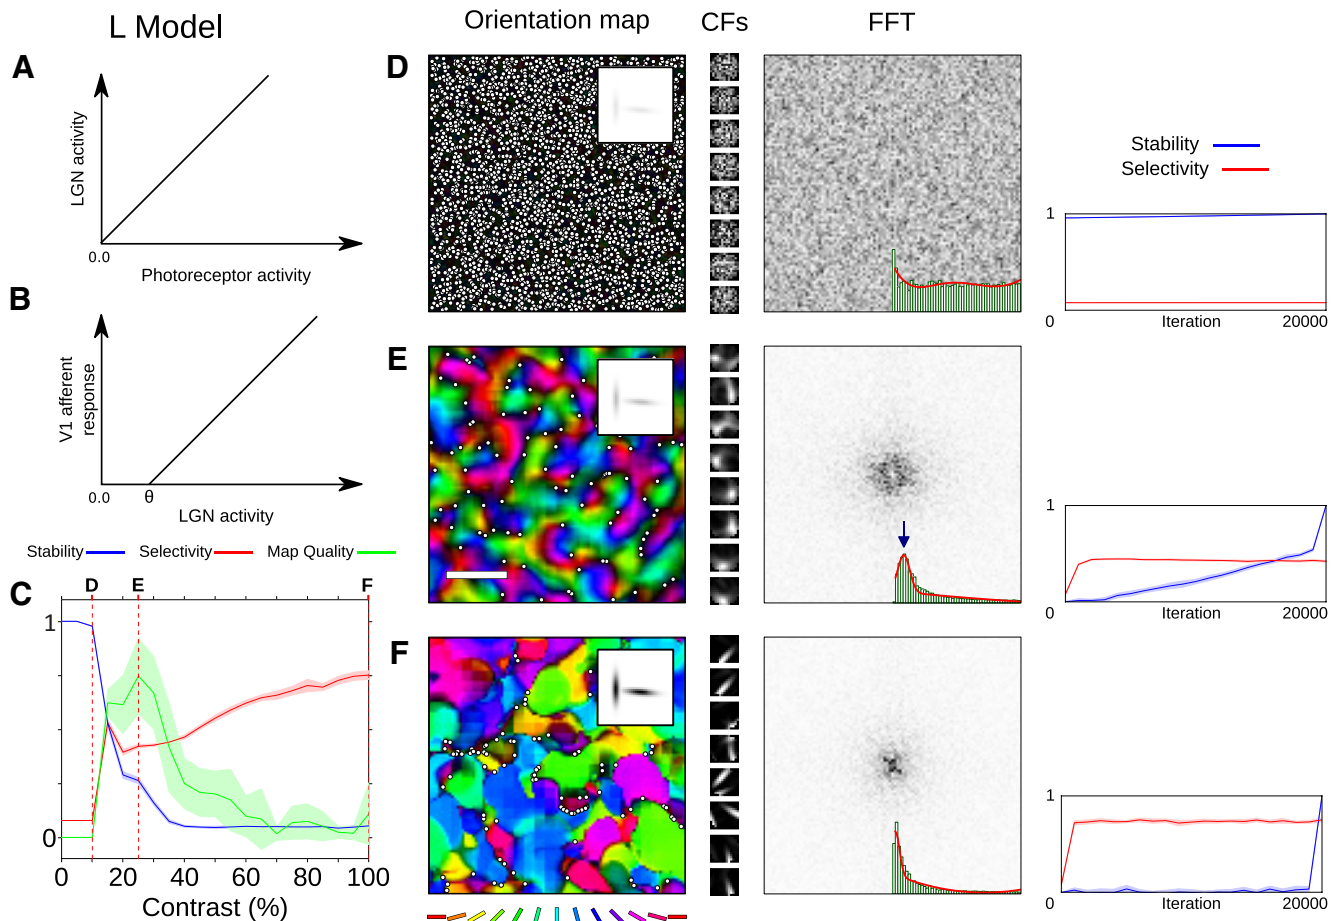
Figure 6. Model L: basic model has relatively poor stability, map quality, and robustness to contrast. *A*, *B*, The transfer functions from photoreceptor activity to LGN activity (*A*) and from LGN activity to V1 afferent response (*B*). Both transfer functions have unit slope, and V1 units have a fixed threshold of $\theta = 0.2$. *C*, The mean map stability (green), selectivity (red), and map quality (blue) as a function of contrast, across 10 simulations with randomized input sequences. Shaded areas indicate ±95% confidence intervals. For low-contrast inputs, neurons are not activated, and thus no learning occurs (point D); at slightly higher contrasts neurons have higher stability and selectivity (point E), but at the highest contrasts (point F) neurons are very unstable. Neurons have a relatively high selectivity and stability in a small range of contrasts for which the model has been tuned (contrasts 15–30%). *D–F*, Organization of model L at the end of development (iteration 20,000) at 10% (*D*), 25% (*E*), and 100% (*F*) contrast. Polar orientation preference maps with estimated pinwheel positions, along with sample inputs (inset) and afferent connection fields (CFs) from the ON sheet to V1 neurons for an arbitrary selection of model neurons (evenly spaced along the vertical midline of the map), are shown. The corresponding two-dimensional FFTs are shown, with 1D the spectral power histogram (green), function fit (red), and estimated peak position (blue arrow) used to estimate the hypercolumn distance (white scale bar). These values determine the value of the map metric, which is fairly poor even for the best L maps due to the large number of pinwheels identified. On the right, the stability and selectivity are plotted as a function of simulation time for that contrast, showing the average across all 10 random seeds with the ±95% confidence intervals. *D*, For low-contrast inputs (contrast 10%) orientation maps do not develop, yielding nominally high stability values that are meaningless because the selectivity remains low. *E*, Maps are well ordered, and development is somewhat stable, within the "tuned" range of contrasts (contrast 25%, as in Fig. 5). *F*, Orientation maps for high-contrast patterns (100% contrast) are highly disorganized, with sharp boundaries between hypercolumns and a non-ring-shaped FFT. Afferent connections are highly orientation-selective imprints of the elongated Gaussians presented to the photoreceptor sheet, and map development is highly unstable, indicated by an early rise in selectivity without a corresponding increase in stability. Note that the final value of each stability plot will always be 1.0; the final map compared with itself has a stability index of unity. Overall, although L develops maps, it fails to be robust to contrast, develops relatively poor quality maps, and is not very stable compared to the ferret data.

nearly all receptive fields become orientation selective, and a top-ographic organization for orientation is formed (Fig. 5). This topographic organization arises because throughout the development process, nearby neurons are similarly activated, and therefore develop similar orientation preferences. This process has been well documented in previous models (Burger and Lang, 1999; Miikkulainen et al., 2005).

L is a greatly simplified model, explicitly designed to demonstrate this process of self-organization using a minimal set of mechanisms. The afferent input to V1 is a linear rectified function of the image contrast (since the output of ON/OFF units is a linear function of the overall input from the retinal photoreceptors), as shown in Figure 6A. The activity of each LGN unit is the dot product of the photoreceptor input and a normalized difference-of-Gaussian ON or OFF receptive field, and the rela-

tionship between LGN activity and V1 afferent response is also linear. Unlike the LGN units, a V1 unit will not respond until the sum of its input activities exceeds a fixed minimum threshold, θ (Fig. 6B).

The robustness of this model to changes in input stimulus is evaluated by the behavior of the model as a function of input contrast, without changing any other parameters. Although there are many other dimensions over which these inputs could be varied (e.g., number of Gaussian patterns, size of pattern, amount of noise), robustness to contrast is an intuitive feature that can be easily evaluated and related to the visual system. The analysis of this simple model is shown in Figure 6C. For each contrast, the average stability over development (measured at intervals of 1000 iterations), the average orientation selectivity of afferent connections in the final map (iteration 20,000), and the average map quality of the final map

are computed by taking the mean results of 10 different random seeds controlling the pattern of training inputs. These three measures give a good indication of the stability, robustness, and quality of map development for each input image contrast. A fixed set of three contrasts has been chosen for illustration and analysis in all four models: the low-contrast point (Fig. 6D), the medium-contrast point (E), and the high-contrast point (F). The medium contrast level of 25% is also shown in Figure 5 and was selected as the peak of the map metric for the L model.

Figure 6C shows that low-contrast inputs to the L model (15–20% contrast) can result in higher stability but lower map quality than shown in Figure 5. At these low contrasts, neurons become moderately selective, and a small degree of stability can be achieved because, as the receptive fields of the model neurons form, connections that have been strengthened by Hebbian learning continued to strengthen throughout development, rather than being overwritten by the strengthening of a different set of connections. Figure 3D illustrates that the pinwheel densities at these contrasts are reliably distinguishable from biological maps; average results from 33 simulations of the L model at 15% contrast are shown.

In this model, contrasts of 10% and below (Fig. 6D) lead to a complete lack of map development. If the presynaptic ON/OFF activity is not strong enough to consistently exceed the necessary fixed threshold θ to activate V1 neurons, connections will not change, and receptive fields will not form. Such a developmental process is thus completely stable at the cost of achieving any degree of selectivity. Although this is a trivial property of the L model, the ability of the early visual system to ensure that V1 is sufficiently activated is an important consideration for understanding visual development.

As input image contrast increases, both the stability of map development and the overall level of map organization degrades rapidly, resulting in maps such as the one shown in Figure 6F at high contrast. The map organization of the orientation map in Figure 6F can be seen to be very poor, as reflected by the non ring-shaped structure in the Fourier power spectrum. These observations hold across all 10 randomized simulation runs, as demonstrated by the low value of the π pinwheel density metric for all high-contrast measurements (Fig. 6C).

In general, the L model is not robust to contrast. Unlike the experimental data shown in Figure 1A, selectivity is achieved rapidly, but map stability over development is not maintained. There are several key properties of this simple model that make it a poor model of robust map development in animals. First, the linearity of activity in the LGN and V1 with respect to photoreceptor activity ensures increases in contrast directly lead to corresponding increases in activity in ON/OFF units as well as in V1. Second, the lateral interactions in the model depend on the level of activation of V1 units and further amplify activity in V1 as input image contrast increases. Third, as the afferent weights in the model develop, receptive fields of V1 neurons become more orientation selective and therefore match the patterns in the input more closely, further increasing the postsynaptic responses of V1 neurons. These three properties ensure that during development, when inputs are high contrast, both ON/OFF and V1 activity will also be high.

With a fixed Hebbian learning rate, the consequence of simultaneously elevated presynaptic and postsynaptic activity (from the high ON/OFF and V1 responses respectively) ensures that the amount of Hebbian learning is increased (Eq. 10). Therefore, with high-contrast input, connections between ON/OFF and V1 cells are more likely to be overwritten as differently oriented pat-

terns are presented, leading to continual reorganization of receptive fields and a corresponding loss of stability and overall map quality.

Conversely, selectivity of receptive fields increases as contrast is increased. This is because, as contrast increases, connections begin to strengthen and weaken on a faster time scale, eventually becoming indicative of the most recent pattern on the retinal photoreceptors (which in this case will be a highly selective, elongated Gaussian pattern) rather than learning to represent longer-term correlations in the afferent input.

This simple model, although capable of forming orientation-selective receptive fields and smooth topographic organization for orientation (given a particular set of parameters), is not robust even to small changes in the input image contrast. Since experimental results indicate that map development is robust and stable over a time scale relatively consistent between animals, and that the input activity properties are likely to change during development (for example, before and after eye opening), it is necessary for the visual system to have some underlying mechanism that compensates for changes in the type of driving input.

The following sections show how feedforward gain control and automatic adaptation of the neuronal activation thresholds can improve the stability of orientation map development and robustness in response to changes in the input image properties. We then show how inclusion of both these simple and biologically motivated mechanisms results in the development of other emergent properties of V1 neurons, such as contrast-invariant orientation tuning.

Model AL: homeostatic adaptation regulates postsynaptic activity

In the L model, the fixed activity threshold across all V1 units was found to constrain robustness to contrast. For any positive threshold value, there will be a range of low contrasts at which the input to V1 units is below threshold and no development occurs. At high contrasts, the postsynaptic activity of the V1 units will be unbounded, resulting in map instabilities due to rapid Hebbian learning constantly reorganizing receptive fields across the map.

In biological V1, the firing rates of different cortical neurons cannot be characterized by a single, linear (and therefore unbounded) function of synaptic input. Unlike the linear model units with fixed threshold, real neurons are thought to have intrinsic homeostatic plasticity, whereby each neuron regulates the level of synaptic input required to generate an action potential by changing the number and distribution of its ion channels (Davis and Bezprozvanny, 2001; Daoudal and Debanne, 2003; Zhang and Linden, 2003; Schulz, 2006; Grubb and Burrone, 2010; Kuba et al., 2010). Such mechanisms allow each neuron to adapt its firing threshold to maintain a suitable average activity as incoming synaptic drive varies over time.

In the AL model, a simple model of homeostatic adaptation is introduced to the cortical sheet, allowing individual V1 units to adjust their firing thresholds to establish a suitable average target activity (Eq. 8). The response of the ON/OFF units remains linear with respect to the photoreceptor activity (Fig. 7A); the new relationship between ON/OFF response and V1 afferent activity is illustrated by Figure 7B.

The quality and robustness of AL across input contrast is greatly improved by the introduction of adaptive thresholds, suggesting that intrinsic homeostatic adaptation plays an important role in robust map development (Fig. 7C). With adaptation, maps have more realistic pinwheel densities, and stability drops off more slowly at high contrasts. At the low-contrast point
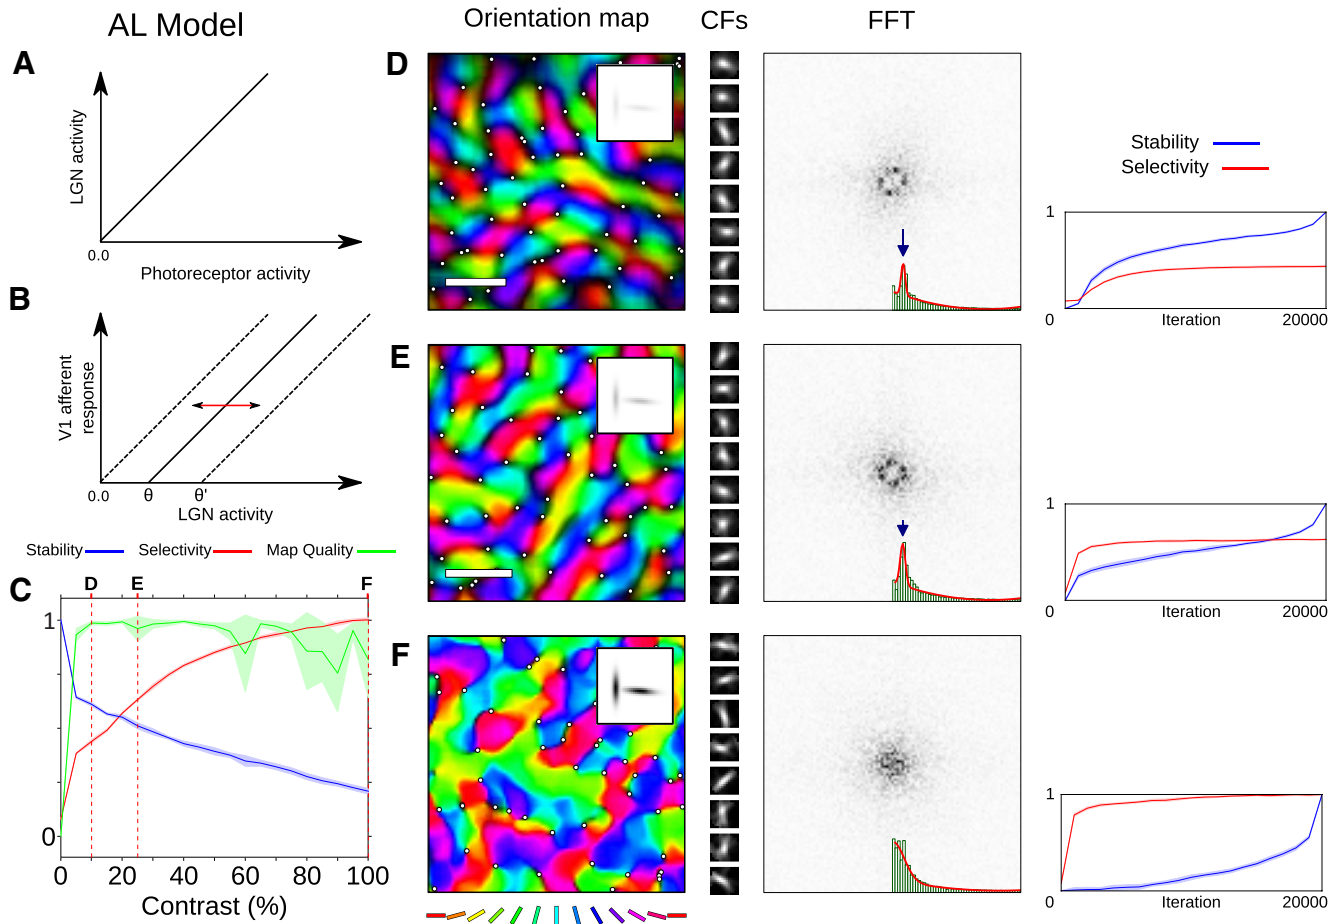
Figure 7. Model AL: adding homeostatic threshold adaptation improves selectivity, stability, and map quality across all contrasts. All plotting conventions (colors, symbols, and scale bars) are as in Figure 6. ***A***, The transfer function from photoreceptor activity to LGN activity remains unchanged. ***B***, Each V1 unit now possesses an independent adaptive threshold θ that is automatically adjusted to maintain a fixed target activity. ***C***, Maps are more stable across contrast, with higher selectivities at high contrast than the L model. Map quality is relatively high throughout, with a drop at high contrasts. ***D***, The AL model can respond and self-organize at lower contrasts by lowering the adaptive threshold of V1 units. ***E***, AL self-organizes into higher-quality maps than L at identical contrasts, though stability has suffered compared to ***D***. ***F***, Adaptation greatly improves the map quality relative to the L model at high contrasts. However, the map still suffers from sharp boundaries due to highly selective connection fields (CFs) that are imprints of the elongated Gaussians presented to the photoreceptor layer, and the FFT becomes non-ring-shaped at high contrast. Stability is also very poor (with selectivity achieved long before stability), because the afferent weights continually reorganize at high contrasts. Thus, homeostatic adaptation offers significant benefits over the L model, but is not sufficiently robust or stable to account for the animal data.

shown in Figure 7D, self-organization is able to proceed as the adaptive thresholds fall, ensuring activity in the V1 sheet and compensating for the lower input levels.

At the mid-contrast point shown in Figure 7E, selectivity and stability are higher and the maps are better organized than in the L model. This is demonstrated by the higher map metric, which reflects the greater spacing between pinwheels in the orientation map and the increasingly ring-shaped Fourier power spectrum.

At high contrasts, selectivity continues to increase but map quality and stability start to fall. The orientation map shown in Figure 7F is highly selective everywhere, including around the identified pinwheel locations. The result is a map with sharp boundaries between regions of different orientation preference instead of smooth transitions. The variance in pinwheel density increases as the ring structure in the FFT degrades, due to the hypercolumn distance becoming poorly defined in Fourier space. The high selectivity is also evident in the highly elongated weight profiles of the afferent connectivity from the ON/OFF sheets to the V1 units.

The failure of the AL model at high contrast is due to the same fundamental problem identified in the L model: rapid

Hebbian learning in the afferent connectivity from the LGN to V1. Although the postsynaptic activity is now homeostatically regulated, which keeps the intracortical circuitry working in a well-defined range, the presynaptic activity in the ON and OFF layers remains unbounded. The result is an effectively very high afferent learning rate, causing recent inputs to be memorized rather than being incorporated smoothly into the map, and destroying stability. Moderating the postsynaptic activity can only partially address this imbalance, so long as the activity in the presynaptic ON and OFF sheets remains unbounded.

Model GCL: gain control compensates for differences in input strength

Rather than compensating for high presynaptic activity by regulating postsynaptic excitability, an alternative and perhaps more direct approach would be to limit the range of possible activity levels reaching V1. This change would ensure that learning in the afferent connections between the ON/OFF sheets and V1 cannot be disrupted by high presynaptic activity.

Contrast-gain control, also known as normalization, is a well-documented phenomenon whereby neurons have a nonlinear
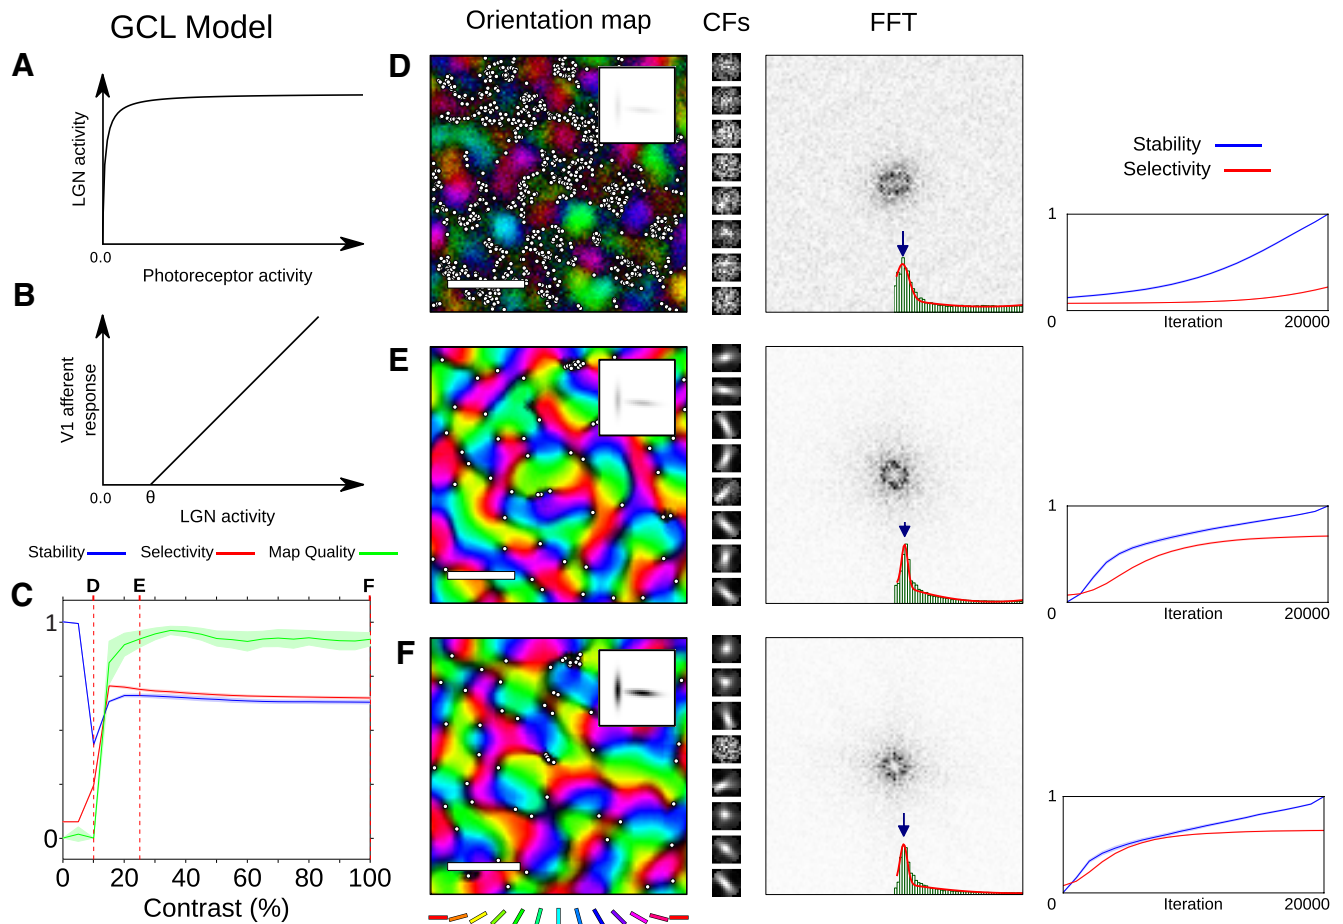
Figure 8. Model GCL: adding contrast-gain control independently improves stability and map quality. All plotting conventions (colors, symbols, and scale bars) are as in Figure 6. **A**, Contrast-gain control in the ON/OFF sheets results in a nonlinear transfer function, compressing unbounded photoreceptor inputs into a bounded range of LGN activities. **B**, V1 units share the same fixed threshold as those in the L model. **C**, Map quality is improved across all contrasts and map stability no longer degrades as contrast increases. Selectivity remains both high and stable with increasing contrast. **D**, At the low-contrast point, GCAL can just begin to self-organize as contrast-gain control boosts enough of the afferent signal over the fixed V1 threshold for some Hebbian learning to occur. **E**, GCL self-organizes into higher quality maps than L at the same 25% contrast level. Unlike the AL model, some small areas of the map fail to self-organize properly, resulting in regions of low selectivity and clusters of pinwheels. **F**, For high (100%) contrast inputs, map quality remains consistently high, connection fields (CFs) remain well formed, the map remains smooth and the FFT is appropriately ring shaped. However, small areas of the map fail to self-organize or develop selectivity, as shown by the noisy connection field. Overall, gain control supports robust and stable map development across contrasts, but some neurons are left behind as the others develop, and cannot reach threshold, leading to unevenly organized maps.

response to input strength, compressing a wide range of inputs into a smaller range of responses (Bonin et al., 2005). There is a wealth of experimental evidence supporting the idea that contrast-gain control first arises early in the visual pathway (Shapley and Victor, 1978; Derrington and Lennie, 1984; Sclar, 1987; Truchard et al., 2000; Baccus and Meister, 2002; Alitto and Usrey, 2004). A wide variety of candidate mechanisms have been proposed (Carandini et al., 2002; Geisler and Albrecht, 1997; Finn et al., 2007; Anderson et al., 2000; for review, see Carandini and Heeger, 2012).

In the GCL model, contrast-gain control is implemented using divisive inhibitory lateral connections between model ON/OFF neurons, providing an "extraclassical" supressive surround (Felisberti and Derrington, 1999; Bonin et al., 2005; Alitto and Usrey, 2008). We would expect to obtain similar results using the other mechanisms that have been proposed, but have not tested them. The addition of gain control to the L model ensures that the afferent input to V1 is no longer a linear function of the input image contrast, as shown in Figure 8A. Instead, the afferent input to V1 saturates with increasing contrast.

Figure 8C shows that gain control in the ON/OFF sheets is sufficient to greatly increase the robustness and stability of the

orientation map development process. By compensating for differences in the presynaptic input to V1, stability is greatly improved, with high map quality across most contrasts. Stability no longer degrades at high contrast, selectivity remains stable even at high contrast, and map quality remains consistently high.

At the low-contrast point shown in Figure 8D, contrast-gain control has enabled some self-organization, boosting the weak input signal over the firing threshold in the cortical layer. The orientation map is not as well organized as the corresponding AL map (Fig. 7D), but this result does demonstrate that this contrast-gain control mechanism amplifies weak afferent signals for low-contrast inputs.

At the medium contrast level shown in Figure 8E, the orientation map is comparable to the equivalent AL map. One common type of artifact observed in GCL are the pinwheel clusters found in low-selectivity regions, which do not appear in maps that have adaptive thresholds in the cortical layer (Fig. 7). These pinwheel clusters persist even in the high-contrast regime shown in Figure 8F.

The origin of these artifacts is illustrated by the afferent weight pattern shown in Figure 8F (fourth pattern down). This set of weights is noisy and has failed to self-organize, indicating areas of the map lagging in their development.
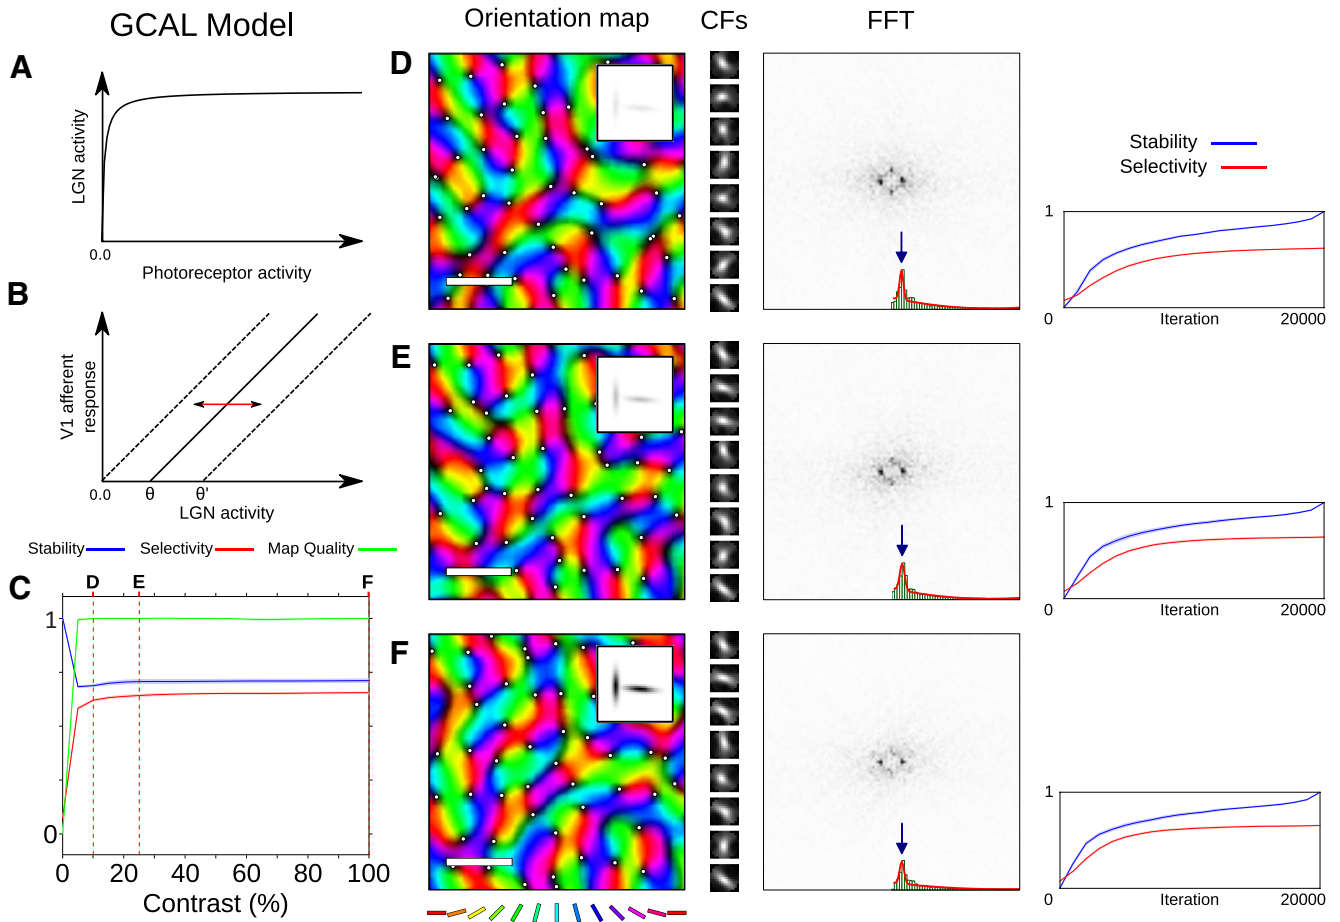
Figure 9. Model GCAL: combining homeostatic adaptation with gain control yields high-quality, selective, stable maps at all contrasts. All plotting conventions (colors, symbols, and scale bars) are as in Figure 6. **A**, The contrast-gain control mechanism introduced in GCL is retained. **B**, V1 units now have the same adaptive threshold (θ) as in the AL model. **C**, Map quality metric indicates that all GCAL maps have very close to π pinwheel density for all simulated orientation maps (also see Fig. 3D), with high stability and selectivity across all contrasts. **D–F**, Maps are properly self-organized at all contrasts, with ring-shaped FFTs. At high contrast, GCAL remains smooth and does not suffer from low-selectivity pinwheel clusters or noisy connection fields (CFs) due to the introduction of the homeostatic threshold. Stability increases alongside selectivity, indicating highly stable map development. The GCAL model robustly generates high-selectivity, highly stable, high-quality maps.

These areas of poor organization arise when neighboring cortical regions develop orientation selectivity slightly more quickly, increasing the postsynaptic activity of these surrounding regions. This increase in the activity of neighboring units increases the strength of the suppressive lateral interactions, further suppressing the development of the region already lagging in self-organization. This positive-feedback loop ensures that some areas of the map fail to develop properly. Thus, although GCL is a clear improvement over both L and AL, it does not robustly organize into biological maps that include all of the simulated neurons.

Model GCAL: adaptation compensates for differences in V1 activation

Homeostatic adaptation and contrast-gain control independently improved map quality, selectivity, and robustness, but both mechanisms have been found to have specific shortcomings. The final GCAL model combines the ON/OFF contrast-gain control mechanism from GCL (Fig. 9A) with the adaptive threshold from AL in the cortical layer (B).

Analysis of this model in Figure 9C demonstrates that the combination of these two mechanisms results in a model that is extremely robust to contrast, while organizing much more smoothly and consistently at any contrast than the other models. All maps at all visible contrasts achieve π pinwheel density, as shown by the map metric value. As an illustration, the data

for 40 GCAL simulations at 100% contrast have been replotted in Figure 3D, where the cluster of red circles for the GCAL model clearly fits into the range of pinwheel density values seen in ferret maps. All FFT plots are ring shaped and comparable to the biological maps, and none of the maps suffer from the artifacts identified in the AL and GCL models. Selectivity and map stability are achieved rapidly once sufficient contrast is available, and remain constant across all contrasts.

The maps at the low-, medium-, and high-contrast points are qualitatively indistinguishable; all afferent weights develop properly without becoming overly selective at the expense of map quality. Combining gain-control and homeostatic adaptation achieves a trade-off between selectivity and map quality that results in highly realistic orientation preference and selectivity maps that are extremely robust to variations in contrast.

Evidently, homeostatic adaptation complements gain control by adjusting the postsynaptic target activity of each V1 unit, ensuring the long-term smooth development across the map over many visual inputs and thus preventing any particular region of the map from being poorly activated and thus lagging behind in the self-organization process. Homeostasis helps ensure consistent map development despite fluctuations in the input statistics across multiple presentations, by retaining a weighted history of previous activation.

Contrast-gain control in turn complements homeostatic adaptation by compressing the dynamic range of each individual

input, regulating the presynaptic input between the ON/OFF sheets and the cortical layer, and ensuring the rate of Hebbian learning in the afferent connections matches the rate of development of lateral cortical connectivity.

Even if the long-term average presynaptic activity in the ON/OFF sheets is appropriate for normal Hebbian learning in the afferent connections, the high dynamic range within a single visual input would continue to disrupt the afferent connectivity on time scales too rapid for the homeostatic threshold to compensate. Contrast-gain control assists the action of homeostatic adaptation over long time scales, normalizing the dynamic range of activity of individual presentations and allowing the activity threshold to respond to long-term trends in the input and not to the transient variance within individual presentations.

These two mechanisms perform two different operations on different time scales that independently increase the robustness of development to input contrast. Together these two operations complement each others' operation to adjust the average activity and dynamic range of the presynaptic and postsynaptic neural activity to ensure robust map development progresses smoothly and regularly on both long and short time scales.

Stable and robust development using realistic inputs in the GCAL model

The previous sections looked only at robustness to image contrast, but biological maps appear to be robust against not only the strength, but also the frequency and type of patterns, with similar maps developing with and without eyes open and with and without dark rearing. We now examine how the GCAL model previously tested with the abstract Gaussian input patterns can be used for other types of input patterns, without changing any of the model parameters.

When an animal opens its eyes, the driving input to the LGN and V1 changes from spontaneously generated patterns to natural visual input (Cragg, 1975; Beckmann and Albus, 1982; Tavazoie and Reid, 2000; Hooks and Chen, 2006; Huberman et al., 2006). As described above, despite this change, orientation maps measured at eye opening retain their orientation preferences, and therefore mature into a final adult map that is similar to the map measured at eye opening. We can model pre-eye-opening development in the GCAL model using input image patterns that resemble intrinsically generated activity [loosely modeled after retinal waves (Wong, 1999) and represented as noisy disk-shaped patterns randomly placed on the retinal photoreceptors at each iteration (Bednar and Miikkulainen, 2004)]. After eye opening, development can be modeled using retinal inputs extracted from a data set of natural images, from Shouval et al. (1996).

Figure 10 shows results from the GCAL model, which includes both gain control in the ON/OFF channels and homeostatic adaptation, compared with experimental results from chronic optical imaging in ferret. In the model, when the input type changes from noisy disks to natural images (at iteration 6000), there is no decrease in the map orientation stability index (Fig. 10E), even though selectivity continues to increase. The initial map structure generated before model "eye opening" is maintained after eye opening, despite the change in the nature of the driving activity. Note that no model parameters have been changed from the previous simulations with Gaussian input stimuli; the model is robust even to extreme changes in the input images without the need for any compensatory changes in parameters.

It would not be possible to use these more realistic inputs in the L or AL models without a significant retuning of the param-

eters. Without contrast-gain control, individual image-patch presentations with high contrast will disrupt the afferent connectivity from the LGN to the V1 sheets, as in Figures 6F and 7F. The GCL model is more robust to contrast changes, yet, as shown in Figure 8F, development of individual neurons is not robust, and thus only the GCAL model allows all neurons to develop robustly across a wide range of inputs.

The stability of ferret map development is measured in Figures 2, A and B, and 10D. We cannot directly compare iterations in the model with days in the experimental data, because the precise date of any initial starting point in the ferret is unclear. However, we can qualitatively compare the overall pattern of stable development for the experiment over the first 10,000 iterations, which covers the period when selectivity and stability rise before saturating in the model. In all the measured experimental cases, selectivity values increase in conjunction with stability to their maximum values over the course of development, which is also true of the simulated data. In some ferrets in which the eyes are surgically opened before the time of natural eye opening, very faint maps were already seen and were already more similar to the final map than to the control condition. Similarly, model maps that have begun to develop before eye opening are similar to the final map (Fig. 10E). These results are similar to the average development for all eight ferrets shown in Figure 2B and are consistent with binocular deprivation experiments in cat (Crair et al., 1998).

Of course, map stability could be achieved trivially by decreasing the ability of neurons to adapt once an orientation map is initially formed before eye opening (for example, by forcibly decreasing the Hebbian learning rate, as in many models). However, as shown in various experimental studies, maps continue to adapt and mature after eye opening, and ultimately come to reflect some of the underlying statistics of the natural visual input. For example, Figure 11A (reprinted from Tanaka et al., 2009) shows experimentally measured biological maps with and without continuous exposure to one particular orientation, using a goggle-rearing oriented-blurring paradigm. These experimental results suggest that when exposed to one particular orientation after eye opening, orientation maps in V1 reorganize by expanding the domains in the map that maximally respond to the goggle-reared orientation, while reducing the areas responding to the unexposed orientations. We can test this property in the GCAL model by using post-eye-opening image data sets with different orientation statistics. Such results are illustrated in Figure 11D, where the map is developed using the same set of natural image patches used in Figure 10 after they have each been convolved with a vertically elongated Gaussian kernel, approximating the anisotropy introduced by goggle rearing. The resulting orientation map also contains a much higher proportion of vertical-preferring neurons. The differences between the model map before and after eye opening, and as the map develops, show that, as in the experimental data, domains of the model V1 preferring the predominant orientation have expanded as the map has matured. The GCAL model is therefore robust and stable, but also able to adapt to changes in the input statistics.

Finally, because the mechanisms added to make the GCAL model were not arbitrary, but were chosen based on a wide range of well-established experimental work, it is possible to make specific links between the behavior of the GCAL model's neurons and experimental data. For example, contrast-gain control mechanisms have been proposed to be involved in contrast-invariant tuning, whereby the tuning curves for V1 neurons retain the same width across contrasts (Sclar and Freeman, 1982; Skottun et al.,
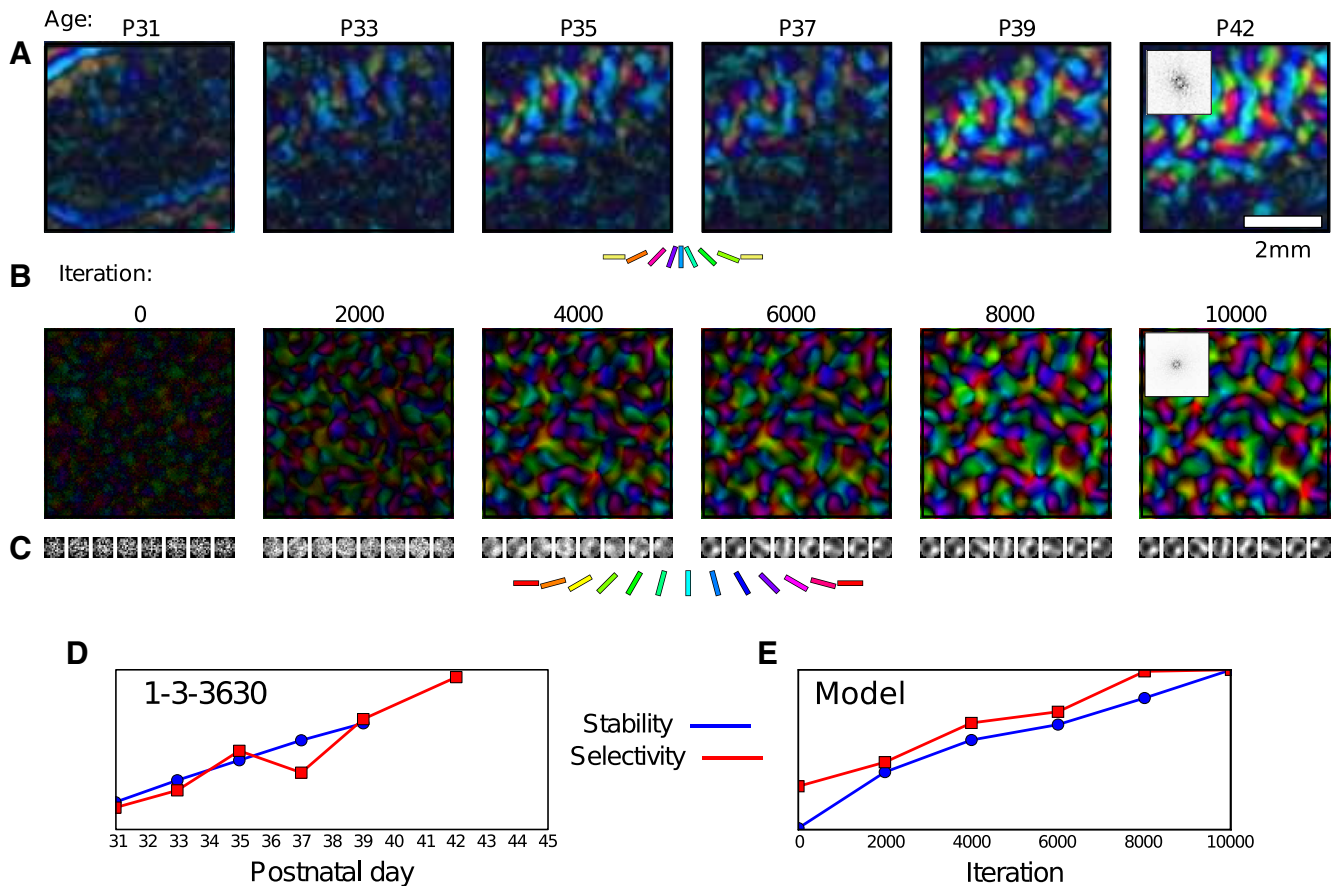
Figure 10. Stable map development before and after eye opening. **A**, Polar orientation maps recorded using chronic optical imaging at different ages in one ferret [Fig. 1A, ferret 1-3-3630; reprinted from the study by Chapman et al. (1996)]. Scale bar, 2 mm. **B**, Simulated GCAL polar orientation maps. Noisy disk patterns drive the map development until 6000 iterations, after which natural images are presented to the model retina. **A**, **B**, Both the ferret and model map have a ring-shaped FFT (inset in the final map plot of each). **C**, Afferent connections from the ON sheet to V1 are shown for an arbitrary set of model V1 units throughout development, to illustrate how neurons become more selective over time. **D**, Orientation stability indices (Eq. 1) across development for the ferret from **A**, replotted from Figure 2B for comparison. **E**, SI for the simulation shown in **B**. As selectivity develops in both the ferret and GCAL, the map smoothly increases in stability, indicating a highly stable process of development. Stability can also be seen by tracking individual features of the orientation map across time, again for the ferret and for GCAL. These results show that the mechanisms in GCAL are sufficient to account for the observed levels of stability in ferret map development.

1987). Figure 12 shows that the GCAL model's self-organized orientation-selective neurons have achieved robust contrast-invariant tuning. This property arises in the model because the contrast-gain control in the ON/OFF channels increases the responses to low-contrast relative to high-contrast stimuli. The recurrent lateral interactions in V1 then ensure that differences in the responses to preferred and nonpreferred orientations are amplified (Antolík, 2010). Responses to nonpreferred orientations will therefore always be smaller than those to preferred orientations. Similar methods of achieving contrast-invariant orientation tuning have been described in previous nondevelopmental models (Geisler and Albrecht, 1997; Anderson et al., 2000; Carandini et al., 2002; Finn et al., 2007), but to our knowledge GCAL is the first model to show how contrast-invariant tuning using this method can emerge robustly over development for a full map of V1 neurons.

Discussion

We have shown that the integration of a small number of simple, well-known, biologically realistic mechanisms is sufficient to reproduce stable and robust map development. Contrast-gain control in the early visual pathway and homeostatic maintenance of activity levels have the effect of regulating both the presynaptic activity and postsynaptic activity of

V1 neurons. This regulation can ensure that throughout development, activity levels, and therefore the rate of Hebbian learning, remain stable. In this way, receptive fields of individual neurons can approach a final organized structure without being overwritten many times, thus achieving the stability and robustness observed in experimental studies. These mechanisms also potentially underlie contrast-invariant orientation tuning of single neurons.

Other activity-dependent models such as the elastic net have previously demonstrated stable development for orientation and for ocular dominance under certain conditions (Keil and Wolf, 2011; Keil et al., 2010). However, these models are formulated at a more abstract level not suitable for identifying important mechanisms like homeostasis and gain control. In particular, they do not simulate responses to individual input patterns such as natural images or retinal waves, and thus cannot address robustness with respect to changes in the input pattern types.

In the real biological system, the differences between input types may be more (or less) extreme than those presented here, leading to two possibilities. First, if the differences in the types of inputs that drive the real visual system are more extreme, further adaptation mechanisms may be required. For instance, if the den-
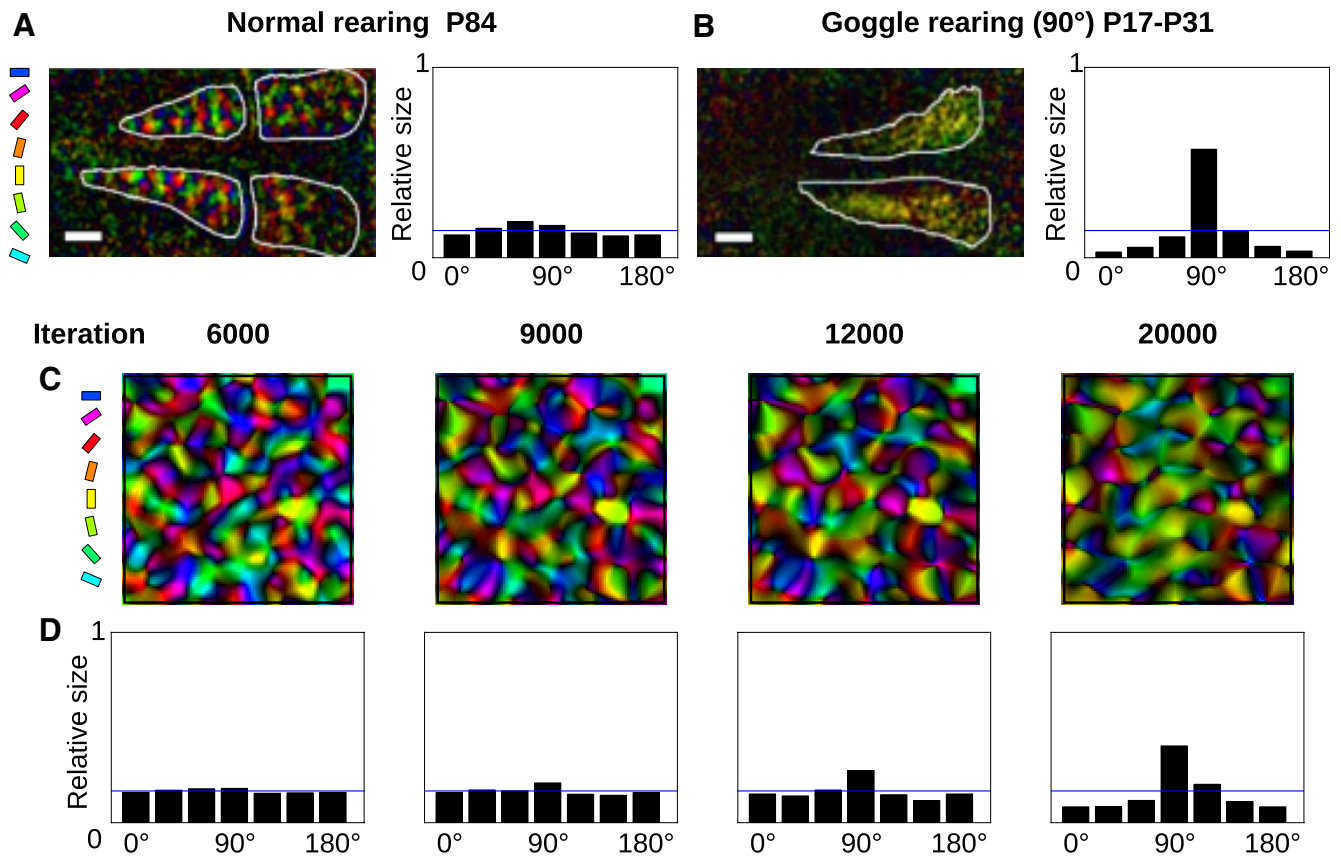
Figure 11. Experimental and GCAL model orientation maps developed in orientation-biased environments. ***A***, Orientation map measured in a normally reared kitten at postnatal day 84 and orientation histogram. The color and brightness indicate the preferred orientation and the orientation magnitude, respectively. The color code for preferred orientations is shown to the side. Scale bars: 2 mm. ***B***, Orientation map measured in a kitten at postnatal day 42 after 13 d of goggle rearing with vertical lines. Orientation histogram now shows a strong bias toward vertical orientations. Reprinted from Tanaka et al. (2009). ***C***, Orientation preference maps during development from a GCAL simulation driven by noisy disk patterns until 6000 iterations, after which natural images that have been anisotropically blurred vertically are presented to the model retinal photoreceptors. ***D***, Histograms show an expansion of yellow (vertically preferring) regions in the orientation after eye opening, reflecting the statistics of the natural image input and reproducing the results observed in biological map development.
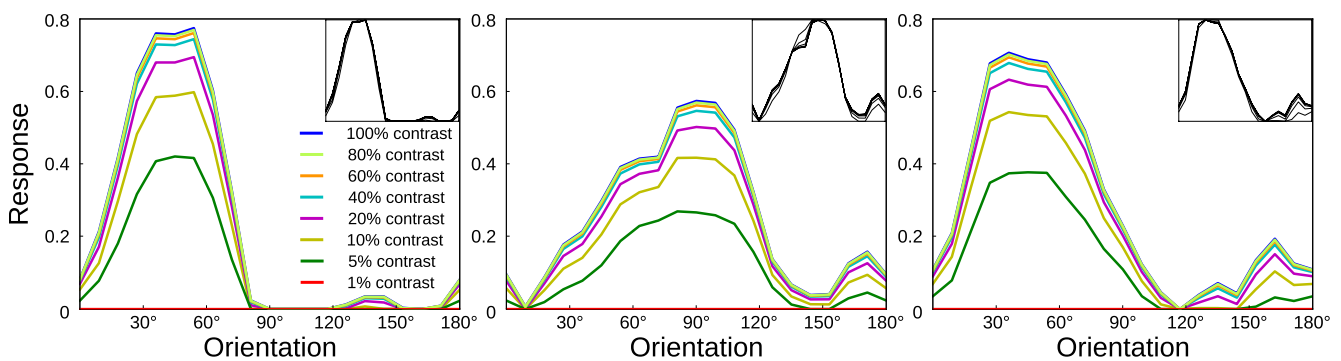

Figure 12. GCAL model: contrast-invariant orientation tuning. Example orientation tuning curves for three representative model neurons, measured at the indicated contrasts. Orientation tuning width remains constant despite changes in contrast (except for small deviations at the lowest contrasts), as confirmed by the normalized tuning curves (inset).

sity of edges in visual scenes varies dramatically across different rearing environments (e.g., between desert and jungle environments), it may be necessary for the adaptation mechanism to take both presynaptic and postsynaptic activity into account together within each cell. We have implemented such a mechanism in our modeling, but found that it was not necessary to handle the range of input types considered here, while being less obvious how to implement locally in V1 or LGN neurons. Alternatively, if real input differences are less extreme, the mechanisms introduced here may not both be required, although this is an unlikely pos-

sibility given that there is substantial experimental evidence for both of the simple mechanisms involved. The observation that these mechanisms are also consistent with many other single-neuron properties, such as contrast-invariant orientation selectivity and automatic maintenance of average activity, suggests that the visual system has achieved many different properties using similar underlying mechanisms.

The ability of neurons to learn and adapt requires changes in synaptic connectivity, yet it is also essential that neurons are neither excessively active or silenced. Both excessive activity and

silence would lead to a loss of information at each level of the neural pathway (Turrigiano and Nelson, 2004), and so homeostatic adaptation not only contributes to stable and robust map development, but also helps to maintain the transfer of visual information. More complex homeostatic mechanisms may also stabilize map development, for example, by regulating the distribution of activity rather than just the target average value (Triesch, 2005) or by operating on the synaptic weights directly via a process such as synaptic scaling (Tetzlaff et al., 2011). In any case, the idea that neurons use some measure of their average activity over time to trigger homeostasis and regulate their responses is widely accepted, and the specific simple homeostatic mechanism used in the GCAL model is sufficient to demonstrate how such regulation improves stability and robustness.

In the visual pathway, it is possible that both homeostatic adaptation and contrast-gain control are present at every stage of processing. In the GCAL model, we have chosen to apply the explicit contrast-gain control mechanism only to the ON/OFF sheets and homeostatic adaptation only to the V1 units. It would have been possible to add both of these mechanisms to all of the stages of the model, but this would have made the effect of each mechanism harder to analyze and interpret, and we do not believe such a change would affect our results significantly.

The GCAL model depends heavily on lateral interactions. Lateral interactions in V1 arise because of lateral recurrent connections between neurons in V1 (Gilbert et al., 1990; Bosking et al., 1997; primarily in layer 2/3). In other simple self-organizing map algorithms, these specific lateral interactions are replaced by a more convenient and computationally efficient but highly abstract mechanism. For example, in the self-organizing map model (Obermayer et al., 1990; Kohonen 1982; Farley et al., 2007), the algorithm finds the maximally active neuron in the cortex based on the distance between the input pattern and each neuron's set of afferent weights. Afferent weights are then adapted only in a specified circular (Gaussian) neighborhood around the maximally active unit. There are therefore no specific lateral interaction strengths in the SOM algorithm, weight changes depend only on the position of the maximally active unit, and the cortical activity does not depend on the contrast of the input stimulus. As shown in this study, we believe that it is important to include specific lateral interaction strengths in mechanistic models of V1 to understand the constraints on the real system.

Although lateral interactions are important for the operation of the GCAL model, the architecture is deliberately simplified and does not reflect the detailed anatomy of cortical connectivity. In animals, lateral interactions involve a complex circuit based on long-range excitation and disynaptic inhibition, with different effects at different contrasts (Hirsch and Gilbert, 1991; Weliky et al., 1995). A model related to GCAL demonstrates how this more elaborate circuit could give similar results (Law, 2009), but requires many more parameters and more complicated analysis methods. Similarly, a related model shows how simulating multiple V1 laminae can explain both simple and complex cells in V1, allowing neurons with random phase preferences to develop in the simple cells and realistic orientation maps in the complex cells (Antolík and Bednar, 2011). These other models improve on the realism of the GCAL model, but because of their greater complexity and the larger associated parameter spaces, they are more difficult to analyze and understand, and so the GCAL model is more useful for studying phenomena that do not require those additional circuit elements.

Apart from the lateral connectivity, another feature of the SOM and other algorithms that is not shared by the GCAL model is that each SOM V1 neuron begins with initial connections to the whole of the retina. This starting point requires that both retinotopic and orientation-selective receptive fields develop simultaneously. Achieving smooth development of both these features in a computational model requires several global ad hoc mechanisms. For example, the processes of lateral neighborhood decrease and learning rate decrease in the SOM algorithm help ensure that self-organization does not fall into a local minimum where the mapping is not retinotopic (i.e., they provide an "annealing" of the map). These ad hoc mechanisms themselves result in model orientation map instability, such as a decrease in the size of orientation domains and a large reorganization of the map over time. However, real V1 maps are arguably formed sequentially, with a coarse retinotopic map already in place before the orientation map (for review, see Huberman et al., 2008). In the model algorithm presented here, there is an initial coarse retinotopic map in V1, which avoids the need for global ad hoc mechanisms to achieve smooth map development.

The removal of ad hoc mechanisms and addition of mechanisms explicitly regulating activity levels makes the GCAL model easier to understand with numerical simulations (as we have done here), and we expect the same fundamental theoretical principles of operation to apply to the GCAL model as for other self-organizing map models (von der Malsburg, 1973; Kohonen, 1982). These properties have been analyzed extensively in previous work with these related models. In particular, the GCAL model achieves good coverage and continuity in its representation of the input feature space (Swindale et al., 2000) by forming sparse, decorrelated activity bubbles in response to each input, just as in the LISSOM model (Miikkulainen et al., 2005). This process has been shown to fold a high-dimensional feature space onto a two-dimensional surface, as a discretized approximation of the principal surfaces of the input (Ritter et al., 1992), and is a form of dimensionality reduction (Durbin and Mitchison, 1990). We argue that the GCAL model should be seen as a mechanistic implementation of the well-established principles of self-organization, allowing them to proceed robustly for realistic inputs of variable strengths and types.

Apart from the specific models introduced, we showed how pinwheel density can be used to assess the quality of simulated orientation maps. Of course, pinwheel density is just one of many factors that could be measured and compared between maps, but we argue that it is a useful metric for several reasons. First, there is a unique, cross-species reference value, which allows objective decisions about map quality. Having a pinwheel density close to π is a necessary, if not always sufficient, criterion for a map to be considered biological. Second, reliable high performance on this measure requires that the maps be smooth (thus implicitly incorporating alternative metrics like LHI; Nauhaus et al., 2008) and periodic (thus implicitly incorporating the ringness in the Fourier power spectrum). Third, the reference value of π was measured empirically through careful analysis of the raw imaging data that ensured that π pinwheel density is a genuine, structural property of the underlying orientation map organization, not simply an artifact of the particular filtering process used. This reference value can thus be compared directly against values from simulated maps, which do not need filtering or other preprocessing that requires human judgment.

The proposed metric could be extended to include additional criteria for map quality, but those of which we are aware have

significant issues that make them unsuitable for characterizing model maps. For instance, systematic bias in the orientation histograms has been quantified for biological maps (Coppola et al., 1998; Müller et al., 2000). However, our results and those of Tanaka et al. (2009) in Figure 11 show that these biases depend crucially on the visual input statistics in both models and animals, rather than revealing inherent properties of the maps. Because the visual statistics have not yet been measured for laboratory environments, and are likely to vary significantly between labs, the levels of map bias found in a particular study are not a suitable reference value for a general model. Similarly, Müller et al. (2000) and Kaschube et al. (2010) show that biological maps have slightly different average distances between pinwheels of the same polarity (with orientation increasing clockwise or counterclockwise) or opposite polarities, which is also true of the high-quality maps from the GCAL model (data not shown). However, there is no clear reference value for this difference, which varies across maps and species. Calculating a meaningful value for polarity biases also requires precise localization of independent pinwheels, which is difficult to automate for any maps other than high-quality maps like those from GCAL. Thus, the pinwheel polarity distribution gives relatively little information for distinguishing between maps of varying quality. However, each of these analyses can be useful alongside the proposed metric, and whenever reference values that hold across laboratories and across species become available, the metric can be extended to include these.

With recent advances in calcium imaging (e.g., two-photon imaging), we anticipate novel, high-resolution experimental results concerning the time course of map development that we can relate to the GCAL model. Although rodent species are most commonly used in existing calcium-imaging studies, these techniques are beginning to yield new results for species that develop orientation maps, such as the macaque monkey (Nauhaus et al., 2012). Using genetically encoded calcium indicators, it has also become possible to use calcium imaging for chronic recordings (Lütcke et al., 2013). Chronic calcium recordings of orientation map development would complement the existing optical imaging results, offering insights into the process of map formation at cellular resolution.

The GCAL model may also be applicable for understanding cortical development in species that do not form continuous orientation maps, such as rodents Ohki et al. (2005). The homeostatic adaptation and contrast-gain control mechanisms remain applicable even with salt-and-pepper organization of preferences, and it has been shown that this type of organization can emerge in models related to GCAL (Law, 2009). It may be possible to adapt the GCAL model to account for other developmental phenomena involving orientation-selective cells, such as the convergence of ipsilateral eye and contralateral eye orientation preference in binocular cells of mice (Wang et al., 2010). The GCAL model's robust mechanisms should allow the transition between intrinsic activity and visually driven activity to be simulated with little or no retuning of model parameters.

Finally, the realistic features of the GCAL model make this model a useful starting point for investigating properties of map development using a wide range of inputs. Previous state-of-the-art high-dimensional models of topographic map development, such as the LISSOM model (Miikkulainen et al., 2005), required a complete retuning of a large number of parameters depending on the input patterns used. However, there is a relatively small number of free parameters of the GCAL model (Hebbian learning rate, target V1 activity, strength of RGC/LGN inhibition, and

V1 lateral interaction strengths). Moreover, these values hold for a wide range of inputs without any tuning.

Fundamentally, we believe that contrast-gain control and homeostatic adaptation are important basic principles underlying topographic map development, even though they have primarily been considered only at the single-neuron and small-network levels in previous work.

Notes

Supplemental material for this article is available at www.topographica.org. The models are implemented in the Topographica simulator, freely available at www.topographica.org. This material has not been peer reviewed.

References

Alitto HJ, Usrey WM (2004) Influence of contrast on orientation and temporal frequency tuning in ferret primary visual cortex. J Neurophysiol 91:2797–2808. CrossRef Medline

Alitto HJ, Usrey WM (2008) Origin and dynamics of extraclassical suppression in the lateral geniculate nucleus of the macaque monkey. Neuron 57:135–146. CrossRef Medline

Anderson JS, Carandini M, Ferster D (2000) Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. J Neurophysiol 84:909–926. Medline

Antolík J (2010) Unified developmental model of maps, complex cells and surround modulation in the primary visual cortex. PhD thesis, School of Informatics, The University of Edinburgh, Edinburgh, UK.

Antolík J, Bednar JA (2011) Development of maps of simple and complex cells in the primary visual cortex. Front Comput Neurosci 5:17. Medline

Baccus SA, Meister M (2002) Fast and slow contrast adaptation in retinal circuitry. Neuron 36:909–919. CrossRef Medline

Barrow HG, Bray AJ, Budd JM (1996) A self-organizing model of 'color blob' formation. Neural Comput 8:1427–1448. CrossRef Medline

Beckmann R, Albus K (1982) The geniculocortical system in the early postnatal kitten: an electrophysiological investigation. Exp Brain Res 47:49–56. Medline

Bednar JA, Miikkulainen R (2004) Prenatal and postnatal development of laterally connected orientation maps. Neurocomputing 58–60:985–992.

Bienenstock EL, Cooper LN, Munro PW (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. J Neurosci 2:32–48. Medline

Blakemore C, Cooper GF (1970) Development of the brain depends on the visual environment. Nature 228:477–478. CrossRef Medline

Blakemore C, Van Sluyters RC (1975) Innate and environmental factors in the development of the kitten's visual cortex. J Physiol 248:663–716. Medline

Blasdel GG (1992a) Differential imaging of ocular dominance and orientation selectivity in monkey striate cortex. J Neurosci 12:3115–3138. Medline

Blasdel GG (1992b) Orientation selectivity, preference, and continuity in monkey striate cortex. J Neurosci 12:3139–3161. Medline

Blasdel GG, Salama G (1986) Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. Nature 321:579–585. CrossRef Medline

Bonin V, Mante V, Carandini M (2005) The suppressive field of neurons in lateral geniculate nucleus. J Neurosci 25:10844–10856. CrossRef Medline

Bosking WH, Zhang Y, Schofield B, Fitzpatrick D (1997) Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. J Neurosci 17:2112–2127. Medline

Burger T, Lang EW (1999) An incremental Hebbian learning model of the primary visual cortex with lateral plasticity and real input patterns. Z Naturforschung C 54:128–140. Medline

Burger T, Lang EW (2001) Self-organization of local cortical circuits and cortical orientation maps: a nonlinear Hebbian model of the visual cortex with adaptive lateral couplings. Z Naturforsch C 56:464–478. Medline

Carandini M, Heeger DJ (2012) Normalization as a canonical neural computation. Nat Rev Neurosci 13:51–62. CrossRef Medline

Carandini M, Heeger DJ, Senn W (2002) A synaptic explanation of suppression in visual cortex. J Neurosci 22:10053–10065. Medline

Chapman B, Bonhoeffer T (1998) Overrepresentation of horizontal and vertical orientation preferences in developing ferret area 17. Proc Natl Acad Sci U S A 95:2609–2614. CrossRef Medline

Chapman B, Gödecke I (2000) Cortical cell orientation selectivity fails to develop in the absence of ON-center retinal ganglion cell activity. J Neurosci 20:1922–1930. Medline

Chapman B, Stryker MP (1993) Development of orientation selectivity in ferret visual cortex and effects of deprivation. J Neurosci 13:5251–5262. Medline

Chapman B, Stryker MP, Bonhoeffer T (1996) Development of orientation preference maps in ferret primary visual cortex. J Neurosci 16:6443–6453. Medline

Coppola DM, White LE, Fitzpatrick D, Purves D (1998) Unequal representation of cardinal and oblique contours in ferret visual cortex. Proc Natl Acad Sci U S A 95:2621–2623. CrossRef Medline

Cragg BG (1975) The development of synapses in the visual system of the cat. J Comp Neurol 160:147–166. CrossRef Medline

Crair MC, Gillespie DC, Stryker MP (1998) The role of visual experience in the development of columns in cat visual cortex. Science 279:566–570. CrossRef Medline

Daoudal G, Debanne D (2003) Long-term plasticity of intrinsic excitability: learning rules and mechanisms. Learn Mem 10:456–465. CrossRef Medline

Davis GW, Bezprozvanny I (2001) Maintaining the stability of neural function: a homeostatic hypothesis. Annu Rev Physiol 63:847–869. CrossRef Medline

Derrington AM, Lennie P (1984) Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. J Physiol 357:219–240. Medline

Durbin R, Mitchison G (1990) A dimension reduction framework for understanding cortical maps. Nature 343:644–647. CrossRef Medline

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. CrossRef Medline

Farley BJ, Yu H, Jin DZ, Sur M (2007) Alteration of visual input results in a coordinated reorganization of multiple visual cortex maps. J Neurosci 27:10299–10310. CrossRef Medline

Felisberti F, Derrington AM (1999) Long-range interactions modulate the contrast gain in the lateral geniculate nucleus of cats. Vis Neurosci 16:943–956. Medline

Finn IM, Priebe NJ, Ferster D (2007) The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. Neuron 54:137–152. CrossRef Medline

Geisler WS, Albrecht DG (1997) Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. Vis Neurosci 14:897–919. CrossRef Medline

Gilbert CD, Hirsch JA, Wiesel TN (1990) Lateral interactions in visual cortex. Cold Spring Harb Symp Quant Biol 55:663–677. CrossRef Medline

Gödecke I, Kim DS, Bonhoeffer T, Singer W (1997) Development of orientation preference maps in area 18 of kitten visual cortex. Eur J Neurosci 9:1754–1762. CrossRef Medline

Goodhill GJ (2007) Contributions of theoretical modeling to the understanding of neural map development. Neuron 56:301–311. CrossRef Medline

Grubb MS, Burrone J (2010) Activity-dependent relocation of the axon initial segment fine-tunes neuronal excitability. Nature 465:1070–1074. CrossRef Medline

Hirsch JA, Gilbert CD (1991) Synaptic physiology of horizontal connections in the cat's visual cortex. J Neurosci 11:1800–1809. Medline

Hooks BM, Chen C (2006) Distinct roles for spontaneous and visual activity in remodeling of the retinogeniculate synapse. Neuron 52:281–291. CrossRef Medline

Huberman AD, Speer CM, Chapman B (2006) Spontaneous retinal activity mediates development of ocular dominance columns and binocular receptive fields in V1. Neuron 52:247–254. CrossRef Medline

Huberman AD, Feller MB, Chapman B (2008) Mechanisms underlying development of visual maps and receptive fields. Annu Rev Neurosci 31:479–509. CrossRef Medline

Kaschube M, Schnabel M, Löwel S, Coppola DM, White LE, Wolf F (2010) Universality in the evolution of orientation columns in the visual cortex. Science 330:1113–1116. CrossRef Medline

Keil W, Wolf F (2011) Coverage, continuity, and visual cortical architecture. Neural Syst Circuits 1:17. CrossRef Medline

Keil W, Schmidt KF, Löwel S, Kaschube M (2010) Reorganization of columnar architecture in the growing visual cortex. Proc Natl Acad Sci U S A 107:12293–12298. CrossRef Medline

Keil W, Kaschube M, Schnabel M, Kisvarday ZF, Lwel S, Coppola DM, White LE, Wolf F (2012) Response to comment on universality in the evolution of orientation columns in the visual cortex. Science 336:413. CrossRef

Kohonen T (1982) Self-organized formation of topologically correct feature maps. Biol Cybern 43:59–69. CrossRef

Kuba H, Oichi Y, Ohmori H (2010) Presynaptic activity regulates Na(+) channel distribution at the axon initial segment. Nature 465:1075–1078. CrossRef Medline

Law JS (2009) Modeling the development of organization for orientation preference in primary visual cortex. PhD thesis, School of Informatics, The University of Edinburgh, Edinburgh, UK.

Linsker R (1986) From basic network principles to neural architecture: Emergence of orientation columns. Proc Natl Acad Sci U S A 83:8779–8783. CrossRef Medline

Löwel S, Schmidt KE, Kim DS, Wolf F, Hoffsümmer 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. CrossRef Medline

Lütcke H, Margolis DJ, Helmchen F (2013) Steady or changing? long-term monitoring of neuronal population activity. Trends Neurosci 36:375–384. CrossRef Medline

Miikkulainen R, Bednar J, Choe Y, Sirosh J (2005) Computational maps in the visual cortex. New York: Springer.

Miller KD (1994) A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. J Neurosci 14:409–441. Medline

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 Comput 12:2573–2595. CrossRef Medline

Nauhaus I, Benucci A, Carandini M, Ringach DL (2008) Neuronal selectivity and local map structure in visual cortex. Neuron 57:673–679. CrossRef Medline

Nauhaus I, Nielsen KJ, Disney AA, Callaway EM (2012) Orthogonal microorganization of orientation and spatial frequency in primate primary visual cortex. Nat Neurosci 15:1683–1690. CrossRef Medline

Obermayer K, Ritter H, Schulten K (1990) A principle for the formation of the spatial structure of cortical feature maps. Proc Natl Acad Sci U S A 87:8345–8349. CrossRef Medline

Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature 433:597–603. CrossRef Medline

Ritter H, Martinetz T, and Schulten KJ (1992) Neural computation and self-organizing maps: an introduction. Reading, MA: Addison-Wesley.

Rochester N, Holland JH, Haibt LH, Duda WL (1956) Tests on a cell assembly theory of the action of the brain, using a large digital computer. IRE Trans Inform Theory 2:80–93. CrossRef

Schulz DJ (2006) Plasticity and stability in neuronal output via changes in intrinsic excitability: it's what's inside that counts. J Exp Biol 209:4821–4827. CrossRef Medline

Schummers J, Mariño J, Sur M (2004) Local networks in visual cortex and their influence on neuronal responses and dynamics. J Physiol Paris 98:429–441. CrossRef Medline

Sclar G (1987) Expression of "retinal" contrast gain control by neurons of the cat's lateral geniculate nucleus. Exp Brain Res 66:589–596. CrossRef Medline

Sclar G, Freeman RD (1982) Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. Exp Brain Res 46:457–461. Medline

Sengpiel F, Stawinski P, Bonhoeffer T (1999) Influence of experience on orientation maps in cat visual cortex. Nat Neurosci 2:727–732. CrossRef Medline

Shapley RM, Victor JD (1978) The effect of contrast on the transfer properties of cat retinal ganglion cells. J Physiol 285:275–298. Medline

Shouval H, Intrator N, Law CC, Cooper LN (1996) Effect of binocular cor-

tical misalignment on ocular dominance and orientation selectivity. Neural Comput 8:1021–1040. CrossRef Medline

Sirosh J, Miikkulainen R (1994) Cooperative self-organization of afferent and lateral connections in cortical maps. Biol Cybern 71:66–78.

Sirosh J, Miikkulainen R (1997) Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. Neural Comput 9:577–594. CrossRef Medline

Skottun BC, Bradley A, Sclar G, Ohzawa I, Freeman RD (1987) The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. J Neurophysiol 57:773–786. Medline

Sullivan TJ, de Sa VR (2006) Homeostatic synaptic scaling in self-organizing maps. Neural Netw 19:734–743. CrossRef Medline

Swindale NV (1996) The development of topography in the visual cortex: a review of models. Network 7:161–247. CrossRef Medline

Swindale NV, Shoham D, Grinvald A, Bonhoeffer T, Hübener M (2000) Visual cortex maps are optimized for uniform coverage. Nat Neurosci 3:822–826. CrossRef Medline

Tanaka S, Ribot J, Imamura K, Tani T (2006) Orientation-restricted continuous visual exposure induces marked reorganization of orientation maps in early life. Neuroimage 30:462–477. CrossRef Medline

Tanaka S, Tani T, Ribot J, O'Hashi K, Imamura K (2009) A postnatal critical period for orientation plasticity in the cat visual cortex. PLoS ONE 4:e5380. CrossRef Medline

Tavazoie SF, Reid RC (2000) Diverse receptive fields in the lateral geniculate nucleus during thalamocortical development. Nat Neurosci 3:608–616. CrossRef Medline

Tetzlaff C, Kolodziejski C, Timme M, Wörgötter F (2011) Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. Front Comput Neurosci 5:47. Medline

Triesch J (2005) A gradient rule for the plasticity of a neuron's intrinsic excitability. In: International Conference on Artificial Neural Networks (ICANN) 2005, Lecture Notes in Computer Science, Vol 3696 (Duch W, et al., eds.), pp 65–70 Berlin: Springer-Verlag.

Truchard AM, Ohzawa I, Freeman RD (2000) Contrast gain control in the visual cortex: monocular versus binocular mechanisms. J Neurosci 20:3017–3032. Medline

Turrigiano GG (1999) Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. Trends Neurosci 22:221–227. CrossRef Medline

Turrigiano GG, Nelson SB (2004) Homeostatic plasticity in the developing nervous system. Nat Rev Neurosci 5:97–107. CrossRef Medline

von der Malsburg C (1973) Self-organization of orientation-sensitive cells in the striate cortex. Kybernetik 14:85–100. CrossRef Medline

Wang BS, Sarnaik R, Cang J (2010) Critical period plasticity matches binocular orientation preference in the visual cortex. Neuron 65:246–256. CrossRef Medline

Weliky M, Kandler K, Fitzpatrick D, Katz LC (1995) Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. Neuron 15:541–552. CrossRef Medline

White LE, Fitzpatrick D (2007) Vision and cortical map development. Neuron 56:327–338. CrossRef Medline

White LE, Coppola DM, Fitzpatrick D (2001) The contribution of sensory experience to the maturation of orientation selectivity in ferret visual cortex. Nature 411:1049–1052. CrossRef Medline

Wiesel TN, Hubel DH (1963) Single-cell responses in striate cortex of kittens deprived of vision in one eye. J Neurophysiol 26:1003–1017. Medline

Wong RO (1999) Retinal waves and visual system development. Annu Rev Neurosci 22:29–47. CrossRef Medline

Zhang W, Linden DJ (2003) The other side of the engram: experience-driven changes in neuronal intrinsic excitability. Nat Rev Neurosci 4:885–900. CrossRef Medline