This is a repository copy of *What, if anything, are topological maps for?*.

Article:

What, if anything, are topological maps for?

Stuart P. Wilson[1]*

1. Adaptive Behaviour Research Group, Department of Psychology,

The University of Sheffield, Sheffield, S10 2TP, UK.

phone: +44 (0) 114 222 6595

email: s.p.wilson@sheffield.ac.uk

James A. Bednar[2]

2. Institute for Adaptive & Neural Computation, School of Informatics,

The University of Edinburgh, Edinburgh, EH8 9AB, UK.

phone: +44 (0) 131 651 3092

email: jbednar@inf.ed.ac.uk

What, if anything, are topological maps for?

Stuart P. Wilson & James A. Bednar

Abstract

What, if anything, is the functional significance of spatial patterning in cortical feature maps? We ask this question of four major theories of cortical map formation: self-organizing maps, wiring optimization, place coding, and reaction-diffusion. We argue that i) self-organizing maps yield spatial patterning only as a byproduct of efficient mechanisms for developing environmentally appropriate distributions of feature preferences, ii) wiring optimization assumes rather than explains a map-like organization, iii) place-coding mechanisms can at best explain only a subset of maps in functional terms, and iv) reaction-diffusion models suggest two factors in the evolution of maps, the first based on efficient development of feature distributions, and the second based on generating feature-specific long-range recurrent cortical circuitry. None of these explanations for the existence of topological maps requires spatial patterning in maps to be useful. Thus despite these useful frameworks for understanding how maps form and how they are wired, the possibility that patterns are merely epiphenomena in the evolution of mammalian neocortex cannot be rejected. The paper is intended as a non-technical introduction to the assumptions and predictions of these four important classes of models, along with other possible functional explanations for maps.

Keywords

topological map, wiring optimization, self-organization, place coding, reaction- diffusion, epiphenomenon, orientation preference.

Overview

Topological feature maps are ubiquitous in the mammalian brain. Distinct from the topographic maps in which they are embedded, which tend to correspond to the layout of sensory surfaces on the body like the skin or retina, topological maps tend to be organized in terms of similarities between more abstract features of sensory input. For example, spatial ordering of the neuronal responses to different spatial frequencies in an image, or to different directions in which a facial whisker is deflected, are topological rather than topographic maps. Much effort has been focused on explaining topological map organization, and our fascination with the continuous patterns that have been measured throughout the mammalian neocortex has driven neuroscience research forward.

However, posing the fundamental question 'what, if anything, is the spatial patterning of topological feature maps for?' reveals the limitations of current theories. Several excellent comprehensive reviews of theories and models of map formation exist (e.g., Swindale, 1996; Simpson et al., 2009; Nauhaus and Nielsen, 2014). Here, our goal is to focus specifically on what four main theories of map organization, each represented by a strong computational modeling framework, have to say about the function of the continuous cortical map patterns that we see in the majority of mammalian species.

First we outline input-driven self-organization, as a theory of cortical map development that derives from Kohonen's self-organizing map algorithm and previous work by von der Malsburg (1973). These models go a long way to explaining why we have maps, in terms of the developmental processes from which they emerge. However, we argue that only the self-organizing developmental processes, not the particular map patterns that emerge from these processes, are demonstrably useful. From an evolutionary perspective, we then outline the minimal wiring lengths hypothesis presented by Koulakov and Chklovskii (2001), before explaining with reference to place-coding theory, how considerations based on wiring length optimization can answer only one half of the question of what map continuity might be for. Finally, we consider an important class of models that explains map organization in terms of Turing pattern formation. Through this model we explain how maps could be epiphenomena, arising from a selection pressure on general-purpose developmental dynamics to generate

specific cortical circuitry, but not necessarily to generate specific spatial patterning. Thus despite the evidence that self-organization drives the evolution of topological maps, we have not found strong evidence that the spatial pattern of organization has driven this evolution.

Topological feature maps

Topological feature maps are at least locally continuous. By continuous we refer to a smooth spatial organization of feature preferences, such that similar stimulus features in some sensory and/or motor space elicit responses in nearby neurons in the brain. Continuous maps that have received extensive experimental investigation include the retinotopic maps in primary visual cortices, and the somatotopic 'homunculus' maps in primary somatosensory and motor cortices. These are continuous (or topological) maps, because stimulation of adjacent regions of the feature space, i.e., presentation of images at adjacent regions of photoreceptors on the retina, or presentation of tactile stimulation to adjacent regions of mechanoreceptors on the skin, elicit selective responses in neurons located at adjacent regions in the corresponding cortices. These particular maps are also topographic, in that distances in the input space have a well-defined (albeit often nonlinear) relationship with distances across the cortical surface.

In the case of retinotopic and somatotopic maps, the feature spaces are essentially two-dimensional, because the retina and skin can both be considered two-dimensional sensor surfaces. But feature spaces can have arbitrarily many more dimensions. For example, it has been suggested that the underlying visual feature space (or 'plenoptic function') can be captured by seven fundamental dimensions, along which visual stimuli can vary and elicit differential responses (Adelson and Bergen, 1991). The 'plenhaptic function' conveys a similar idea for the multiple dimensions along which tactile stimuli can be discriminated by neuronal activity (Hayward, 2011). We concentrate here on the representation of a particular slice of the plenoptic function, in the selective responses of neurons in primary visual cortex (V1) to the orientation of edges: orientation preference.

In a series of seminal experiments, Hubel and Wiesel (1974) found that individual neurons respond preferentially to the orientation of a bar of light presented at a fixed position on the

retina, and that the preferred orientation of neurons in cat V1 varied smoothly with respect to bars oriented in the full range of 0 to π radians, as the recording electrode was moved incrementally along the cortical surface. Like the selectivity of adjacent neurons for adjacent retinotopic locations at the coarse level, at a finer resolution the V1 orientation preference map reflects the periodic topology of the space of possible orientations (see figure 1). Subsequent recording techniques such as optical imaging revealed the two dimensional organization of orientation preferences to be locally continuous in most species tested.

Analysis of optical imaging data also showed some regions where orientation preferences vary discontinuously, across map features called pinwheels (Bonhoeffer and Grinvald, 1991; Blasdel, 1992). Pinwheels (or singularities, or point discontinuities) are sites on the cortical sheet about which orientation preferences vary continuously in a circular pattern (see, e.g., the square three down and six to the right in figure 1). When neurons are colored according to the line orientation that elicits the maximum response, the resulting image of the orientation map reveals a colorful tiling of pinwheels. Contour lines drawn so as to delineate 'iso-orientation domains', i.e., regions preferring similar orientations, radiate outwards from the pinwheel centers and connect adjacent pinwheels together. Two-photon calcium imaging has since revealed that in cats, for example, the pinwheel organization is clear even at the level of individual neurons (Ohki et al., 2005).

Insert Figure 1 around here.

However, these studies have also shown that the orientation map organization in rodent V1 is not smooth. Instead, the rodent map appears randomly organized, with the orientation preferences of adjacent neurons at most only weakly correlated, despite a well-organized retinotopic map, a full range of different orientation preferences, and high orientation selectivity (Ohki et al., 2005, 2006). Thus rodents presumably have a similar capacity to discriminate edges in the retinal image by their orientation. We therefore describe the organization in rodent V1 as an orientation map, because different stimuli elicit differential responses. We also describe this map as having good coverage of the feature space, because all orientations are represented within a local region of the retinotopic map in which it is embedded (see Swindale, 1991). However, as the topology of the feature space of orientation is not conserved we describe the rodent orientation map as

lacking continuity. In contrast, the pinwheel-dense spatial patterning of maps observed in primate V1 has been described in terms of an optimal trade-off between good continuity and good coverage (see Swindale et al., 2000; Durbin and Mitchison, 1990). Our goal here is not to explain the differences between rodent and primate maps (see Kaschube, 2014, for a recent review). Instead, inspired by the existence of both random and continuous maps in different mammalian species, our question is more general; 'what, if anything, are topological maps for?'. Henceforth we use the term 'topological feature map', as a short-hand for 'a spatial pattern that appears to trade off between good coverage and continuity'.

In the following sections, we consider in turn what i) self-organizing maps, ii) wiring optimization, iii) place-coding, and iv) reaction diffusion models reveal about topological feature maps. It is important to note that these models do not represent mutually exclusive theories of what topological feature maps might be for. As we will see, models of class ii and iii represent largely compatible theories about the usefulness of map continuity at the local scale (i.e., between neurons), and models of class i and iv represent complementary theories about the mechanisms that additionally constrain locally continuous maps to form the millimeter-scale organizations observed in primates and carnivores (i.e., at the level of pinwheels and hypercolumns). In the final section we consider how mechanisms affecting the organization at both the local scale and the millimeter scale are responsible for the emergence of topological feature maps. The examples we consider throughout are focussed on pinwheel-dense orientation preference maps in primate V1, but the arguments that we develop apply to theories of spatial organization in the brain more generally.

Self-organization as a theory of cortical map development

Topological feature maps are spatially ordered collections of receptive fields. To say that adjacent neurons prefer similar features, e.g., similar orientations, is to say that nearby cells have receptive fields that are similar in some way. According to feed-forward neural network models of V1 simple cells, the receptive field of a neuron is closely related to the pattern of weights on the synaptic connections that it makes with peripheral neurons. In such models, maximal

responses will be elicited when the vector describing the intensity of light across the retina is most closely aligned to the corresponding vector of synaptic weight strengths, i.e., when the input pattern is a good template match to the weight pattern. In early sensory topological feature maps, nearby neurons can therefore be thought of as comprising similar templates.

The first class of models that we consider, self-organizing networks, can explain both the overall patterning of topological feature maps, and the structure of the receptive field (or weight template) represented by each point in the map, as consequences of the same relatively simple developmental algorithm.

Activity-dependent self-organization is exemplified by Kohonen's self-organizing map algorithm (Kohonen, 1982, 2007; see also Ritter et al., 1992), which has been shown to be capable of recreating the main features of cortical map organization. In this model, a population of neurons arranged on a two-dimensional sheet learns by modifying synaptic weights from a set of input units whose activity might represent, e.g., the intensity of light measured by an array of photoreceptors. The following steps are repeated for many example input patterns: (1) identify the neuron whose vector of synaptic weights is most closely aligned to the vector of input unit activations, and (2) move the vector of weights for this neuron and those nearby on the sheet towards the vector of input unit activations. Step 2 is achieved by a type of Hebbian learning. Over time, this procedure encourages the weight vectors to spread out to cover the underlying space from which the input patterns have been drawn, while ensuring that the weight vectors of neighboring neurons, and thus the input patterns that will maximally excite them, become closely aligned. As a result maps emerge that are locally continuous; maps contain discontinuities similar to those measured in primary sensory cortices, e.g. orientation pinwheels emerge in networks trained on images of oriented edges; and neuronal weight vectors end up resembling the receptive field structures of real neurons (see Obermayer et al., 1990, 1992).

An important prediction of self-organizing map models is that maps represent the underlying statistical structure in the feature space from which patterns of input are drawn (Durbin and Mitchison, 1990). Accordingly, if some region of the input space is disproportionately represented in the patterns of input presented to the network as it develops, then the resulting

distribution of cortical territory will be likewise distorted to reflect this. A map therefore reflects the personal history of developmental experiences of the animal, such that each map pattern is unique. In support of this idea, classic experiments have shown that early rewiring of projections from the optic nerve so that they drive putative auditory cortex rather than visual cortex leads to the emergence of visual orientation preference maps where normally we would expect maps to arise for auditory feature spaces (Sharma et al., 2000). This experiment suggests that some aspects of cortical regions may be equipotential, adapting to reflect their particular developmental history.

By implementing self-organizing algorithms, a given region of adult cortex is thus assured of comprising a distribution of receptive fields that match the distribution experienced during development. If we can only perceive stimuli for which we have a (reasonable) receptive field match, and self-organization provides receptive fields for the range of stimuli encountered by the organism as it develops, then it is clear that self-organization represents a highly adaptive means of ensuring that neural representations are suited to the potentially wide variety of environments faced by mammals over their evolutionary history.

Lateral interactions

The details of how self-organizing maps are implemented vary subtly between the battery of related theories of cortical self-organization, e.g., in terms of the formulation of the Hebbian learning rule used to incrementally align weight vectors to input vectors in step (2). However, one important aspect of more biologically focused models is to replace the global supervisory mechanism required to identify maximally responsive neurons in step (1), with an entirely local process based on recurrent interactions between neurons. We consider these mechanisms for lateral interactions in detail here, because the lateral interactions are responsible for the continuity and smoothness of the map.

The influential model of von der Malsburg (1973) applies Hebbian learning in localized neighborhoods of activity that have been established, without global supervision, as an emergent property of local recurrent intra-cortical dynamics (also see Dayan, 1993; Carreira-Perpin and

Goodhill, 2004). Similarly in the subsequent LISSOM model (Laterally Interconnected Synergetically Self-Organizing Map; Miikkulainen et al., 2005; Sirosh and Miikkulainen, 1994) and in the current GCAL (Gain Control, Adaptation, Laterally Connected) variant (Stevens et al., 2013), each neuron excites its local neighbors and inhibits its more distal neighbors in a series of extra steps following presentation of each new input pattern, which results in multiple neighborhoods of localized activity emerging spontaneously across the neural sheet. In these models, Hebbian modification of the synaptic weights that connect neurons laterally results in the emergence of long-range 'patchy' recurrent connections between neurons that represent similar features. In models of V1 map development trained using naturalistic image patterns as input (e.g., photographs of natural scenes), these long-range patchy connections form between neurons representing similar orientations, matching data from Bosking et al. (1997) in tree shrew V1 (see Bednar, 2012).

We can think of the localized neighborhood function used in Kohonen's algorithm, and the profile of short-range excitation and long-range inhibition in the more mechanistic approaches like LISSOM and GCAL, as representative of an assumption made by the developing cortex that statistical structure in the world is inherently spatial. The physics of our universe determine that correlated information sampled from the environment tends to be attributable to matter that is co-localized in space (indeed, before quantum physics it was difficult to imagine an alternative!). Incorporating into the developmental plan the general assumption that correlational structure in the world is locally continuous, allows a compact genetic encoding of an algorithm for extracting that correlational structure. Specifying only the gross network architecture, Hebbian learning, and local interactions requires far less genetic information than specifying the receptive field structure of each neuron individually. Exploiting spatial continuity in the environment enables a highly compressed algorithm to generate a distribution of receptive fields that is suited to the particular environment in which an organism develops.

However, although self-organizing map models have been very successful at developing the observed receptive fields and map patterns, they do not demonstrate that the map patterning itself is functionally important. One way to see this is to look at a related class of models with similar ingredients, but lacking the additional assumption of spatial continuity. For instance, clustering

algorithms closely related to Kohonen self-organizing maps but with no neighborhood function and thus no spatial ordering still generate populations of biologically realistic receptive fields (Barrow, 1987; Coates et al., 2011). There is also a large body of models based on sparse coding (assuming that the pattern of activation of cortical neurons should contain only a few active at any time) or independent component analysis (assuming that neurons should have activities statistically independent from each other) that can also explain the development of a set of receptive fields matching patterns seen during development (Olshausen and Field, 1996; Bell and Sejnowski, 1997). Biological interpretations of sparse coding models require some form of competition between neurons, as in the self-organizing map models, but for development of feature preferences do not require any spatially localized connectivity that would lead to continuous maps. Thus neither this spatially specific connectivity nor map continuity appear to be essential for how these models function.

We will not further investigate the details of the clustering, sparse-coding, and independent-component-analysis (ICA) models here, using them only as examples of self-organization without spatial patterning. It is interesting to consider how these models may be extended to generate topological map patterns, e.g., as the topographic-ICA model of Hyvärinen et al. (2001) extends ICA to create continuous maps by grouping according to the remaining higher-order dependencies between linearly independent neurons. However, in the context of the self-organizing map models, what is important is that these examples suggest that self-organizing map algorithms may have been selected because of the component they have in common; competitive interactions between neurons. This competition ensures that neurons develop different receptive fields, and hence adaptively give rise to appropriate distributions of receptive fields across the population. In self-organizing map algorithms the competition is enforced by the neighborhood function (either explicitly as in Kohonen's algorithm, or implicitly as in LISSOM and GCAL), and so it is inherently spatially localized. Spatial localization of these competitive interactions may represent a saving of genetic information required to ensure that emerging receptive fields compete to represent features typical in visual scenes. And thus self-organization can provide a sound explanation for why we have maps. But the point to emphasize is that at no stage in formulating this explanation are we required to attribute function to the spatial patterning, and so these models do not provide evidence for why maps should need to be ordered

in the way that they are.

Continuity is optimal if continuity was optimized

Apart from developmental algorithms, another explanation for why topological feature maps exist is that they optimize wiring lengths: Representing nearby points in feature space by the activity of nearby neurons in the brain minimizes structural and metabolic costs associated with connecting neurons over larger distances. This wiring optimization principle is perhaps the most popular explanation for the existence of spatial organization in the brain provided by authors of textbooks, when motivating the importance of studying neural maps. But it is an incomplete explanation. This was apparent even to early researchers who articulated the 'minimal wiring lengths hypothesis'; "[W]hy is there a map at all? [. . . The] answer rests on an assumption, which is that interactions between cortical neurones are much more important for cells representing points close together in visual space" (Cowey, 1979; see also Mitchison, 1991). We explain here how making this assumption is problematic for the minimal wiring lengths hypothesis. We do so by deconstructing its most explicit statement, in the computational model of Koulakov and Chklovskii (2001), whose original discussion of the limitations is often neglected.

Koulakov & Chklovskii used simulations to determine the pattern of orientation preferences across a cortical sheet that minimizes connection lengths, when each neuron is required to make a proportion of its synaptic connections to others that depends on the similarity of their orientation preferences. The algorithm they used is meant as an abstraction of the kind of evolutionary process by which cortical maps may have been selected, and proceeds as follows.

All simulations begin by randomly assigning orientation preferences to neurons arranged as a regular lattice on a two-dimensional sheet. For a given simulation, the first step is to define a target histogram a priori, which relates the proportion of connections that each neuron should make to others based on the similarity of their orientation preference. For example, we might specify that each neuron should be connected to five others with a difference in orientation less than 10°, to three others with a difference of 10-20°, and so on. The next step is to connect every

neuron to the nearest set of neurons on the sheet that result in it having the same distribution of connections as that defined by the target histogram. Note that the shape of the target histogram for all neurons is identical. Once the connections for every neuron have been (independently) made, the total length of connections in the network is determined. The process of wiring length optimization then begins by randomly perturbing the orientation preference of the first neuron, and accepting the change in orientation preference if the overall wiring length is reduced, else by rejecting the change (with some probability) and restoring the original orientation preference. The procedure of randomly rejecting perturbations that increase a cost function is a type of Monte Carlo method, where the cost function here is simply the overall wiring length. The method is applied to each neuron on the sheet in turn, such that by the end of one iteration through all neurons the arrangement of orientation preferences will have changed to reduce the overall wiring length. Over many such iterations, a map structure will emerge from this process, which makes the wiring length as short as possible given the target distribution.

Koulakov and Chklovskii found that using different distributions of target connectivity yields different patterns of orientation preference across the cortical sheet. If each neuron is required to make equal numbers of connections to others that differ in preferred orientation over the full range of orientation differences, i.e., to make equal numbers of connections with others of similar and dissimilar orientation tuning, then the pattern that makes wiring lengths as short as possible is a random 'salt-and-pepper' arrangement like that in rodent V1 (figure 2; left panels). For a narrower target distribution, where neurons are required to make more connections to others with a similar orientation preference, the pattern that minimizes wiring lengths displays continuity, such that nearby orientations end up represented by neurons that are next to each other. As the width of the target distribution is made smaller the colored images of the resulting map patterns start to resemble bands like rainbows, and as the width is further reduced these rainbows twist into a configuration that starts to resemble a staircase. For very narrow target distributions, where almost all connections are to be made between neurons with similar orientation preferences, the patterns twist further until they resemble maps measured in cat and primate primary visual cortex, which are punctuated by pinwheel point discontinuities, about iso-orientation domains that radiate outwards from the pinwheel centers (figure 2; right panels). Hence, optimizing for a narrow target distribution can lead to the observed local trade-off

between map continuity and coverage in topological feature maps.

Insert Figure 2 around here.

The implication of this result is that an evolutionary process, that likewise minimizes the overall wiring lengths across a sheet of neurons tuned to a periodically repeating feature space like visual edge orientation, should result in the locally continuous pinwheel-rich maps that we measure in the primary visual cortices of most mammalian species. In similar terms we might explain why orientation preferences in rodents are organized as random maps across the cortical sheet, despite being well tuned for orientation; perhaps these species did not undergo the same selection pressures to preferentially connect neurons with similar feature tuning, or perhaps organization with respect to some other dimension of the visual feature space carried a greater selection pressure.

Thus an intuitive answer to the question, why do we have topological maps?, is that our maps have been selected to minimize the clear structural and metabolic costs associated with having to connect neurons over long distances. Yet this is only one half of the full picture. It is the other half of the question that is most important for us to answer; why preferentially connect similarly tuned neurons, i.e., why should evolution have optimized for a narrow target connectivity distribution? Wiring minimization explains map organization in terms of the cost to having long wires, but it does not explain map organization in terms of the benefit to preferentially connecting similarly tuned neurons. Only once we assume that preferentially connecting similarly tuned neurons is beneficial, does wiring minimization open the door to an explanation for map organization via evolution by natural selection.

The minimal wiring hypothesis therefore leaves us asking whether there is a good computational reason to put adjacent regions in the feature space next to each other in the cortical tissue. Let us frame this question the other way around. What are the implications of taking a topological map with a set of connections that result from optimizing a narrow target distribution, and then moving every neuron to a random location on the cortical sheet while keeping the connectivity the same? This would be equivalent to taking the initial pattern of connectivity in Koulakov &

Chklovskii's networks before having run any optimization; neurons with similar orientation preference are preferentially connected, and we are likely to have very long overall wiring lengths. Metabolic costs aside, what functional capacity have we lost by having long wires between similarly tuned neurons? Clearly, if no quality of a signal traveling from one neuron to another with which it makes connections depends on the distance traveled, then the network is functionally equivalent to the network obtained after wiring optimization. The same is true if we permit the signal propagation time to vary randomly with distance. But the two networks are not functionally equivalent if some quality of the signal varies non-randomly with the distance traveled. In the following section, we consider a computational mechanism that can only work in the latter case, where physical distances between points on the map correspond to distances in the feature space. We ask whether such a mechanism can complete the missing half of the explanation for map spatial patterning that is left open by the tautology of the minimal wiring lengths hypothesis.

Are topological maps computational maps?

The qualities of signal propagation that we might imagine to vary predictably with communication distances in a topological feature map include the signal to noise ratio, the amplitude of the signal, and the signaling delay. In the third class of model that we consider, place coding, we are concerned with the latter, the signaling delay. Let us assume that for a finite signal propagation speed the signaling delay varies monotonically with the physical distance separating communicating cells. For propagation velocities slower than a meter per second, an action potential may take tens of milliseconds to register at a cell located several millimeters away. The idea of place coding is that in such a scheme it might be possible for relative signaling delays to subserve computation of the relative timing of external events.

The place coding model proposed by Jeffress (1948) provides a sketch of how this could be done. Jeffress considered that if two events (A and B) register at two sufficiently distant neurons (a and b), and signals travel at a finite speed between them, then those signals will coincide at a specific location between a and b that depends on the relative timing of events A and B. If events A and B are registered simultaneously then signals propagating from a and b will coincide at a

neuron located exactly half way between them. If event A precedes B then the two signals will coincide at a point that lies closer to b, and vice versa. For an interval longer than the time it takes for the signal to travel from a to b, the two will never coincide, but within this range longer intervals yield coincidences that occur closer to the second neuron. Hence, in an array of neurons between a and b, with thresholds high enough that they respond only when multiple signals arrive coincidently, the identity of the active neuron uniquely reports the inter-stimulus time delay. Wiring length might therefore be useful for deriving from a lower level topological map representing a spatial feature like edge orientation, a higher-order topological map representing the stimulus velocity.

Jeffress' place-coding model has most successfully been used to explain the selectivity of midbrain auditory neurons to the relative time at which sound arrives at two ears, as determined by the azimuth angle of the sound source relative to the head orientation of an auditory specialist like the barn owl (reviewed in Yoris and Yin, 2006). Claims that such computation could occur in mammalian cortices have been somewhat controversial, but we have recently provided evidence that such a scheme could operate between barrel columns in rodent primary somatosensory cortex, given the relatively large distances separating the barrels and axonal conduction velocities as slow as ten centimeters per second (Wilson et al., 2011). This mechanism could render supragranular neurons sensitive to a psychophysically relevant range of inter-whisker deflection intervals (a range of up to ten milliseconds).

At first, this idea seems very powerful. It is not hard to imagine that the ability to compute spatial-temporal derivatives for an appropriately mapped feature space could place a species at a selective advantage. The barn owl can use such a scheme to localize prey, the rat could use it to direct biting, the bat could use it to echolocate, and so forth. Indeed, consideration of the model later formalized by Jeffress led Boring (1933) to argue that we should "search for a place theory for all dimensions of consciousness". It is also noteworthy that in the review that popularized the term 'computational map' (Knudsen et al., 1987), the only concrete example provided for a computation for which map continuity is required, is Jeffress' place coding theory.

However, there is a fundamental limitation to the generality of this model. It can only work for

maps of one and two dimensional feature spaces, such as the topographic maps of sensor surfaces, not many-dimensional topological maps. Consider that the analogous computation with respect to visual orientation preference would be the abstraction of the rotation velocity of image contours. For this to be possible, signals between sequentially activated iso-orientation domains would have to travel not isometrically across the two cortical sheet dimensions, but in circles around the pinwheel centers.

To accept the minimal wiring hypothesis, we must escape its inherent tautology by asking what advantage the specific map patterning yields, and so we must ask what trading off between continuity and coverage is good for. Through place coding, a reasonable computational argument for continuity can be made in terms of isotropic mechanisms allowing the extraction of spatial-temporal derivatives. The justification for having good coverage is also clear, because ensuring that each orientation is represented at a given point in V1 is necessary for detecting oriented stimuli at all retinal locations (Durbin and Mitchison, 1990). The problem comes when we attempt to trade continuity for coverage, as in smooth topological feature maps. Promoting coverage at any cost to continuity causes isotropic map decoding mechanisms to break down. Anisotropic mechanisms that can compensate for degraded continuity are not hard to imagine, particularly in the context of the self-organizing map models where circuitry is shaped by experience. But if cortical circuitry is free to build anisotropic mechanisms, then why promote continuity at all? Indeed, patchy connectivity between like-tuned neurons suggests that V1 circuits operate despite continuity, not because of it, potentially allowing activation to spread between similar iso-orientation domains by leap-frogging from pinwheel to pinwheel. This discontinuous spread of activation does not appear to exploit the continuity in map patterns, so why then has pressure for coverage not taken over entirely to generate random maps, which promote coverage at all spatial scales?

A plausible answer is that discontinuous leap-frogging is useful. It may be that orientation maps actually maximize the retinotopic distance over which orientation-specific activity is able to spread, and hence that orientation map patterning has been selected to maximize interaction lengths over the retinotopic map in V1 (as far as is metabolically and structurally viable). Topological feature maps may therefore maximize feature-specific interaction distances, for

minimal wiring lengths. Or perhaps the cortex simply has to compute with respect to whatever pattern it is given, and it is given continuity because of its efficiencies in some non-computational regard, as we will consider in the next section.

Hence, despite the intuitive notion that mirroring distance in feature space by distance in a topological map in the brain should be useful, we are aware of no concrete cortical computation that demonstrates this to be the case in a mammal. We are therefore no closer to explaining a selective advantage to preferentially connecting similarly tuned neurons, and thus we are no more motivated to accept wiring length optimization as a complete explanation for topological feature maps.

Developmental dynamics optimize circuitry; patterning comes for free

The final main class of models that we consider in detail derives from the reaction-diffusion models of pattern formation (morphogenesis) introduced by Turing (1952). At a conceptual level, and paraphrasing the tutorial presented by Kauffman (1993), the basic ingredients of reaction-diffusion models are as follows. For each cell in a dense array, chemical A increases production of chemicals A and B at that location and in the immediately surrounding cells, while chemical B decreases production of A and B with a more diffuse spatial profile that spans a greater distance. With concentrations of 'activator' A and 'inhibitor' B equal at each cell, nothing interesting happens. But increasing A even slightly at (arbitrary) point X leads activation to build up and form into a sharp peak at X, which in turn causes the more diffuse inhibition to build up around point X, but with a flatter peak of lower amplitude. The ratio of A to B will become largest at X, but in surrounding regions, beyond the spatial extent of A, B dominates and the ratio of A to B becomes low. Further away from X, just beyond the spatial extent of B, any other slight increases in A will give rise to similar regions of activation flanked by inhibition. In the ratio of activation to inhibition we see a tiling of Mexican hats emerge, each with the same spatial profile, repeating at a particular spatial frequency. Given noisy initial concentrations across the array, it can be possible for a given profile of activator versus inhibitor to amplify an infinite number of spatial frequencies, but boundary conditions that stop A and B diffusing beyond the edges of the cell array yield attractors permitting only wavelengths that are integer

fractions of the array length (in a given direction) to dominate. Thus, the patterns that will form from noisy initial conditions during Turing instabilities are affected primarily by the profile of activator and inhibitor diffusion (relative to the array length if boundaries are enforced).

An instructive demonstration of how reaction-diffusion can be used to study pattern formation in cortical maps is provided by Ermentrout et al. (2009). They analyzed a Turing system using equations for the profile of activator and inhibitor that specify similar Mexican-hat interactions to those described above, but in this case defined across a two-dimensional disk of neurons that correspond to a barrel column. Using the ratio of chemical attractor to inhibitor to specify the density of thalamocortical axons across the barrel, they were able to recreate sub-barrel patterning characteristic of that revealed by cytochrome oxidase staining in real barrels, resembling a coffee bean, a Mercedes car badge, a baseball, and then a bullseye pattern, as the size of the disk, and therefore the number of solution modes, was increased. As well as being able to predict the shape of a pattern given the size of the column boundary, numerical simulations with boundary conditions constrained by the outlines of real barrels generated strong matches to the patterning measured in those barrels, providing strong support that specific patterning found in cortical maps may arise via Turing instabilities. Models of this form allow the relative contribution of terms comprising the profile of activator versus inhibitor to be investigated analytically, such as the rates of axon production, pruning, and diffusion in the example from Ermentrout et al. (2009). Recognizing that short-range excitatory interactions and long-range inhibitory interactions in the recurrent activity of neurons might serve as the activators and inhibitors in a Turing-like model of feature map formation, we are primed to understand the elegant reaction-diffusion model of orientation map development considered by Wolf (2005).

Wolf (2005) starts by assuming that orientation map development can be modeled in terms of Turing pattern formation (thereby implicitly assuming map continuity), and by considering a form for the profile of activator and inhibitor that guarantees that all orientations will be represented (thereby explicitly assuming map coverage). Under the assumption of Mexican-hat-like lateral interactions (see Reichl et al., 2012; Hein and Kaschube, 2014), many different types of continuous map organization were found to be possible, including rainbow-like

and pinwheel-dense maps like those found by Koulakov and Chklovskii (2001). Wolf showed that what distinguishes between these various map solutions is the number of modes that become selectively amplified as the dynamics of the model unfold. When Turing instabilities selectively amplify a large (finite) number of modes, maps have pinwheel densities in the range observed across real cortices. The model of Wolf assumes no cortical boundary effects, but asks instead what extra constraints added to the Mexican-hat-like lateral interactions will lead to the amplification of many modes, and thus the emergence of topological feature maps.

Wolf's analysis shows that many modes will tend to be amplified when the recurrent inhibitory interactions are additionally constrained to be strongest between similar orientations over large cortical distances (far beyond a hypercolumn). This additional constraint stabilizes the developmental dynamics, such that Turing instabilities yield pinwheel-dense continuous maps like those in primate primary visual cortex. Importantly, the additional interactions that lead to primate-like orientation maps are required to be based primarily on orientation similarity rather than on the proximity of neurons. Hence, like the model of Koulakov and Chklovskii (2001), orientation-specific interactions appear to be key.

However, for our discussion the key insight from the Wolf (2005) model is that orientation preference maps might well reflect an optimization process; but an optimization based on the strength of interaction between like-tuned neurons over larger distances rather than an optimization based on reducing the distance between like-tuned neurons. Moreover, this process takes place through the developmental dynamics that play out within each organism, rather than through the evolutionary dynamics that play out across generations. The account of topological feature map organization is thus largely in agreement with that provided by the self-organizing map models considered earlier. The specific spatial patterning of primate orientation maps may result from pressure for interactions between neurons to be orientation specific. It is interesting to note that feature-specific long-range interactions emerge spontaneously through Hebbian learning in mechanistic models like GCAL, due to correlational structure in the images used to train them, hence these models help explain where the feature-specific constraints, shown by Wolf (2005) to be so important for primate-like map development, might come from. Mechanistic models and reaction-diffusion models of map self-organization might therefore

agree, that evolution constrained general-purpose, pattern-forming, developmental dynamics, to promote feature-specific recurrent interactions, and in doing so incidentally created the pinwheel-dense maps that we observe.

For the natural selection of maps to have begun, as Wolf did, with general-purpose pattern-forming dynamics, makes evolutionary sense because it means that genes need only specify the rules of recurrent interaction that will guarantee neuronal representation of (and hence the ability to perceive) all stimulus combinations, rather than specifying the receptive field structure of each individual neuron. Selective suppression between the correlates of similar orientations over larger distances (over the cortex and thus over the retinotopic map) may have emerged since, because it facilitates visually relevant integrative operations like contour completion and pop-out effects across the visual scene.

Insert Figure 3 around here.

Kaschube et al. (2008) later derived from the model of Wolf (2005) the additional theoretical prediction that the pinwheel density, defined as the number of pinwheels per hypercolumn, should approach the mathematical constant π. Here the hypercolumn is defined in terms of the distance separating adjacent iso-orientation domains, hence the π-density metric is unitless. Kaschube et al. (2010) then tested this prediction, by measuring pinwheel densities in the maps of many animals taken from multiple mammalian lineages (including tree shrew, bush baby, and ferret), representing branches of the phylogenetic tree that diverged up to 65 million years ago. The pinwheel densities in each lineage were indeed found to converge to π. This remarkable result, indicating the presence of a fundamental mathematical constant in the mammalian brain, not only validates the theory that topological maps emerge as Turing instabilities, but also suggests that nature has converged upon this solution multiple times. In line with the conclusion of Kaschube et al. (2010), the emergence of maps by Turing instability may thus represent a deeply canalized attractor in the 'epigenetic landscape' (see Striedter, 1998).

The Koulakov and Chklovskii (2001) model suggests that evolution may have optimized maps for the length of connections between the neural correlates of similar orientation. The Wolf

(2005) model instead suggests that development may optimize map-generating processes for the strength of interaction between the neural correlates of similar orientation. As an explanation of what maps might be for, the former is a tautology, amounting to a formal description of what the map is; i.e., like-tuned connectivity for minimal wiring cost. The latter model takes us beyond re-stating what the map is, to hint at what the connectivity behind the map might be for; i.e., suppression of responses to similar features across the visual scene.

Concluding discussion

We asked of four important classes of model, what, if anything, are topological maps for? Self-organizing maps explain why we have maps in terms of efficient competitive mechanisms between neurons that ensure that the distribution of receptive fields within the population match the distribution of features experienced as an organism develops. In these terms we were able to explain the emergence of maps, but to do so we did not require their patterning to be useful. Considering next the principle of wiring optimization we obtain only a tautology; maps with good coverage and local continuity like we see in primate visual cortex minimize wiring lengths only if we first assume that preferentially connecting like-tuned neurons yields a selective advantage: Continuous maps are demonstrably optimal if we start with the assumption that map continuity has been optimized. The place-coding proposal that continuity is optimized to mirror in cortical space distances in feature space is viable for topographic maps, with one or two cortical sheet dimensions isometric with one or two dimensions of the feature space, as in the Jeffress model for computing stimulus velocity with respect to cortical distance. However, it fails to provide an advantage for continuity in the case of topological maps with multiple feature dimensions, as in the visual cortex. Finally, the reaction-diffusion model of Wolf (2005) suggests that orientation-specific long-range connectivity may be the key constraint that discriminates the pinwheel-dense maps we see from the many other possibilities that satisfy continuity and coverage constraints.

We interpret these analyses as support for a two-part evolutionary model: i) General pattern-forming mechanisms based on reaction-diffusion principles were originally chosen for their efficiency in generating maps with good coverage of the feature space, in terms of the

genetic information required to specify the necessary competition between emerging receptive fields, ii) additional pressure to generate feature-specific cortical circuitry for long-range integration across the sensor surface is responsible for the specific pinwheel-dense patterning of topological feature maps, which are themselves epiphenomena of the underlying developmental process. This interpretation is consistent with the minimal wiring hypothesis as well as the mechanistic accounts of map development provided by self-organizing map models like GCAL.

This review has focused on single–cortical-level models, such as those for simple cells in the visual cortex, to see if topological map patterns can be justified at that level. One interesting direction for future work is to consider how these maps could interact across a sensory hierarchy, which might provide more constraints on possible functions. For instance, several models of maps of complex (phase-insensitive) cells in V1 rely on a multi-stage architecture, with an orientation map of simple cells as a first layer, followed by complex cells pooling indiscriminately over multiple nearby simple cells with the same orientation but different spatial phases (Antolik and Bednar, 2011; Hyvärinen et al., 2009; Weber, 2001). This approach simultaneously requires a continuous organization for orientation (so that local pooling will not destroy orientation selectivity) and a disordered organization for spatial phase (so local pooling will sample from multiple phases). The models differ on how they explain the local phase disorder (either mechanistically due to variability in long-range projections, Antolik and Bednar, 2011, or as a mathematical convenience based on squaring of putatively negative activations, Hyvärinen et al., 2009; Weber, 2001), and thus these models do not yet offer a clear functional explanation for why such an organization may occur. But because complex cells are thought to be useful for visual computations, these models suggest that further study of how they develop may offer stronger evidence for a functional role for continuity (and discontinuity, in the case of spatial phase) in topological maps.

Common to the self-organizing models discussed, which includes the reaction- diffusion model of Wolf (2005), is the assumption of lateral interactions that are excitatory at short ranges and inhibitory over longer ranges. Interestingly, the anatomical data reveals a much more complex architecture—-the longest lateral connections are from excitatory to excitatory cells, but the net effect of surround modulation is typically suppressive at high input contrasts (Gilbert et al.,

1990; Hirsch and Gilbert, 1991; Weliky et al., 1995; Ren et al., 2007; Somers et al., 1998), probably due to polysynaptic inhibition. Thus it is likely that the aspects of the circuit that dominate map development are only part of the story of how the circuit operates.

In any case, the assumption of Mexican-hat-like lateral interactions is a clear link between the mechanistic and reaction-diffusion approaches (as shown formally by Wolf and Geisel, 1998, and Keil and Wolf, 2011; see also Hein and Kaschube, 2014). In the absence of a formal analysis, it would seem that map organization in LISSOM and GCAL-like models is due to similar principles of pattern formation. As strong (albeit circumstantial) support for this claim, maps from the GCAL model of Stevens et al. (2013) have a pinwheel density that reliably converges to π. Therefore we might think of such models as being mechanistic implementations of reaction-diffusion models, or at least consistent with their predictions. The mechanistic models are difficult to analyze compared to the Wolf (2005) model, because every neuron is a non-linear system of equations, and it would have been difficult to first recognize and second understand the π pinwheel density prediction had Wolf's analytically tractable model not so elegantly paved the way.

An important distinction between these two modeling approaches is that the reaction-diffusion formalism explicitly assumes feature-specific lateral interactions between neurons during development, whereas feature-specific interactions in the mechanistic models emerge from the same Hebbian processes that underlie map self-organization, because unsupervised learning at intra-cortical synapses captures the long-range statistical structure of naturalistic image patterns. The former explains why feature-specific interactions are a key ingredient in generating observed topological feature maps, and the latter reveals how such interactions can be driven by statistical structure in the environment.

Whether by accepting only short connections, assuming Mexican-hat-like connectivity, or utilizing distance-dependent delays, our explanations for the spatial patterning of primate V1 maps each require neurons to interact preferentially over short distances. In addition, models based on the principles of wiring optimization and pattern formation both rely on the additional assumption that representations of similar stimulus features should preferentially interact. In both

cases this additional assumption is problematic; explaining precisely why similar features should preferentially interact is deceptively difficult.

The mechanistic self-organizing map models go a step further to reveal how preferential interactions based on the similarity of stimulus features can emerge through the same Hebbian learning mechanisms as those responsible for establishing map patterning itself. Mechanistic algorithms like GCAL, which rely on Hebbian learning of environmental statistics, generate rather than assume preferential interactions between similar features, insofar as they recreate observed patchy connectivity between like-tuned neurons without explicit instruction to do so. Perhaps then the optimization constraint represented by the narrow target histogram in the Koulakov and Chklovskii (2001) model, and by the orientation-specific nonlinearities in the Wolf (2005) model, is Hebbian learning itself. I.e., perhaps Hebbian learning of environmental statistics (rather than a Jeffress-like decoding mechanism) is the criterion against which wiring lengths are optimized, and no further computational constraint is required to escape the minimal wiring lengths tautology and explain spatial patterning in topological feature maps.

Again, this explanation does not depend on spatial patterning being useful. Thus so far the default answer to 'what, if anything, are topological maps for?' remains that topological map patterns may not serve any purpose, beyond providing a common reference point for validating different theories of cortical evolution and development.

Acknowledgements

References

Adelson, E. H. and Bergen, J. R. (1991). The plenoptic function and the elements of early vision. Computational models of visual processing, 1(2):1–20.

Antolik, J. and Bednar, J. A. (2011). Development of maps of simple and complex cells in the primary visual cortex. Frontiers in Computational Neuroscience, 5:17.

Barrow, H. G. (1987). Learning receptive fields. In Proceedings of the IEEE First International Conference on Neural Networks (San Diego, CA), volume IV, pages 115–121. Piscataway, NJ: IEEE.

Bednar, J. A. (2012). Building a mechanistic model of the development and function of the primary visual cortex. J Physiol Paris, 106(5-6):194–211.

Bell, A. J. and Sejnowski, T. J. (1997). The "independent components" of natural scenes are edge filters. Vision Research, 37:3327.

Blasdel, G. G. (1992). Orientation selectivity, preference, and continuity in monkey striate cortex. The Journal of Neuroscience, 12:3139–3161.

Bonhoeffer, T. and Grinvald, A. (1991). Iso-orientation domains in cat visual- cortex are arranged in pinwheel-like patterns. Nature, 353:429–431.

Boring, E. G. (1933). The physical dimensions of consciousness. The Century Co.

Bosking, W. H., Crowley, J. C., and Fitzpatrick, D. (2002). Spatial coding of position and orientation in primary visual cortex. Nature Neuroscience, 5(9):874–882.

Bosking, W. H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. J Neurosci, 17(6):2112–2127.

Carreira-Perpin, M. A. and Goodhill, G. J. (2004). Influence of lateral connections on the structure of cortical maps. J Neurophysiol, 92(5):2947–2959.

Coates, A., Ng, A., and Lee, H. (2011). An analysis of single-layer networks in unsupervised feature learning. Journal of Machine Learning Research, Workshop and Conference Proceedings: AISTATS 2011, 15:215–223.

Cowey, A. (1979). Cortical maps and visual perception the grindley memorial lecture. Quarterly Journal of Experimental Psychology, 31(1):1–17. PMID: 424501.

Dayan, P. (1993). Arbitrary elastic topologies and ocular dominance. Neural Computation,

5:392–401.

Durbin, R. and Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. Nature, 343(6259):644–647.

Ermentrout, B., Simons, D. J., and Land, P. W. (2009). Subbarrel patterns in somatosensory cortical barrels can emerge from local dynamic instabilities. PLoS Comput Biol, 5(10):e1000537.

Gilbert, C. D., Hirsch, J. A., and Wiesel, T. N. (1990). Lateral interactions in visual cortex. In The Brain, volume LV of Cold Spring Harbor Symposia on Quantitative Biology, pages 663–677, Cold Spring Harbor, NY. Cold Spring Harbor Laboratory Press.

Hayward, V. (2011). Is there a 'plenhaptic' function? Philos Trans R Soc Lond B Biol Sci, 366(1581):3115–3122.

Hein, B. and Kaschube, M. (2014). Weakly coherent neuronal interactions are sufficient to explain highly coherent orientation preference maps. In Society for Neuroscience abstracts, 155.02.

Hirsch, J. A. and Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat's visual cortex. The Journal of Neuroscience, 11:1800– 1809.

Hubel, D. H. and Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. J Comp Neurol, 158(3):267– 293.

Hyvärinen, A., Hoyer, P. O., and Inki, M. (2001). Topographic independent component analysis. Neural Comput, 13(7):1527–58.

Hyvärinen, A., Hurri, J., and Hoyer, P. O. (2009). Natural Image Statistics: A Probabilistic Approach to Early Computational Vision. Springer, Berlin.

Jeffress, L. A. (1948). A place theory of sound localization. J Comp Physiol Psychol, 41(1):35–39.

Kaschube, M. (2014). Neural maps versus salt-and-pepper organization in visual cortex. Curr Opin Neurobiol, 24(1):95–102.

Kaschube, M., Schnabel, M., Lwel, S., Coppola, D. M., White, L. E., and Wolf, F. (2010). Universality in the evolution of orientation columns in the visual cortex. Science, 330(6007):1113–1116.

Kaschube, M., Schnabel, M., and Wolf, F. (2008). Self-organization and the selection of pinwheel density in visual cortical dev. New Journal of Physics, 10:015009.

Kauffman, S. A. (1993). Origins of order: Self-organization and selection in evolution. Oxford University Press.

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

Knudsen, E. I., du Lac, S., and Esterly, S. D. (1987). Computational maps in the brain. Annu Rev Neurosci, 10:41–65.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. Biological Cybernetics, 43:58–69.

Kohonen, T. (2007). Kohonen network. Scholarpedia, 2(1):1568.

Koulakov, A. A. and Chklovskii, D. B. (2001). Orientation preference patterns in mammalian visual cortex: A wire length minimization approach. Neuron, 29:519–527.

Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (2005). Computational maps in the visual cortex. Springer, Berlin.

Mitchison, G. (1991). Neuronal branching patterns and the economy of cortical wiring. Proc Biol Sci, 245(1313):151–158.

Nauhaus, I. and Nielsen, K. J. (2014). Building maps from maps in primary visual cortex. Curr Opin Neurobiol, 24(1):1–6.

Obermayer, Blasdel, and Schulten (1992). Statistical-mechanical analysis of self- organization and pattern formation during the development of visual maps. Phys Rev A, 45(10):7568–7589.

Obermayer, K., Ritter, H., and Schulten, K. J. (1990). A principle for the formation of the spatial structure of cortical feature maps. Proceedings of the National Academy of Sciences, USA, 87:8345–8349.

Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., and Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature, 433(7026):597–603.

Ohki, K., Chung, S., Kara, P., Hubener, M., Bonhoeffer, T., and Reid, R. C. (2006). Highly ordered arrangement of single neurons in orientation pinwheels. Nature, 442(7105):925–928.

Olshausen, B. A. and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature, 381:607–609.

Reichl, L., Heide, D., Lwel, S., Crowley, J. C., Kaschube, M., and Wolf, F. (2012). Coordinated optimization of visual cortical maps (ii) numerical studies. PLoS Comput Biol, 8(11):e1002756.

Ren, M., Yoshimura, Y., Takada, N., Horibe, S., and Komatsu, Y. (2007). Specialized inhibitory synaptic actions between nearby neocortical pyramidal neurons. Science, 316(5825):758–761.

Ritter, H., Martinez, T., and Schulten, K. (1992). Neural computation and self- organizing maps- An introduction. Addison-Wesley, New York.

Sharma, J., Angelucci, A., and Sur, M. (2000). Induction of visual orientation modules in auditory cortex. Nature, 404(6780):841–847.

Simpson, H. D., Mortimer, D., and Goodhill, G. J. (2009). Theoretical models of neural circuit development. Curr Top Dev Biol, 87:1–51.

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

Somers, D. C., Todorov, E. V., Siapas, A. G., Toth, L. J., Kim, D. S., and Sur, M. (1998). A local circuit approach to understanding integration of long-range inputs in primary visual cortex. Cereb Cortex, 8(3):204–217.

Stevens, J.-L. R., Law, J. S., Antolik, J., and Bednar, J. A. (2013). Mechanisms for stable, robust, and adaptive development of orientation maps in the primary visual cortex. Journal of Neuroscience, 33:15747–15766.

Striedter, G. F. (1998). Stepping into the same river twice: homologues as recurring attractors in epigenetic landscapes. Brain Behav Evol, 52(4-5):218– 231.

Swindale, N. V. (1991). Coverage and the design of striate cortex. Biol Cybern, 65(6):415–424.

Swindale, N. V. (1996). The development of topography in the visual cortex: A review of models. Network: Comput Neural Syst, 7(2):161–247.

Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., and Hubener, M. (2000). Visual cortex maps are optimized for uniform coverage. Nature Neuroscience, 3(8):822–826.

Turing, A. M. (1952). The chemical basis of morphogenesis. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 237(641):37–72.

von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex. Kybernetik, 14(2):85–100.

Weber, C. (2001). Self-organization of orientation maps, lateral connections, and dynamic receptive fields in the primary visual cortex. In Proceedings of the International Conference on Artificial Neural Networks, Lecture Notes in Computer Science 2130, pages 1147–1152, Berlin. Springer.

Weliky, M., Kandler, K., Fitzpatrick, D., and Katz, L. C. (1995). Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. Neuron, 15(3):541–552.

Wilson, S. P., Bednar, J. A., Prescott, T. J., and Mitchinson, B. (2011). Neural computation via neural geometry: a place code for inter-whisker timing in the barrel cortex? PLoS Computational Biology, 7(10):e1002188.

Wolf, F. (2005). Symmetry, multistability, and long-range interactions in brain development. Phys Rev Lett, 95(20):208701.

Wolf, F. and Geisel, T. (1998). Spontaneous pinwheel annihilation during visual development. Nature, 395(6697):73–78.

Yoris, P. and Yin, T. C. T. (2006). A matter of time: Internal delays in binaural processing. Trends Neurosci, 30:70–78.

Figure captions

Figure 1: *Retinotopic and orientation map in V1.* Given a particular fixation point (marked with a red + symbol above), the visual field seen by an animal can be divided into a regular grid, with each square representing a 1°×1° area of visual or retinal space. In cortical area V1 of mammals, neurons are arranged into a retinotopic map, with nearby neurons responding to nearby areas of the retina. As an example, the image on the right shows the retinotopic map on the surface of V1 of a tree shrew for an 8°×7° area of visual space (adapted from Bosking et al., 2002; scale bar below is 1mm). A stimulus presented in a particular location in visual space (such as the thick black bar shown) evokes a response centered around the corresponding grid square in V1 (3° right, 5.5° down). Which specific neurons respond within that general area, however, depends on the orientation of the stimulus. The V1 map is color coded with the preferred orientation of neurons in each location; e.g. the black bar shown at left will primarily activate neurons colored in purple in the corresponding V1 grid squares.

Figure 2: *Map organization predicted by wiring-length minimization.* Panels on the top row show the target histogram of connectivity that each neuron should make to others based on the similarity between their orientation preferences. Target histograms are defined a priori to guide the process of wiring optimization. By the algorithm of Koulakov and Chklovskii (2001), perturbations of each orientation preference are iteratively accepted if reconnecting neurons to the nearest set of neighbors satisfying the target histogram reduces the overall wiring length. The corresponding map organizations that emerge are shown in corresponding panels below. As the target histograms increasingly promote connections between like-tuned neurons, from left to right, the resulting maps can be described as 'salt-and-pepper', 'rainbow', 'staircase', and 'pinwheel-tiling', respectively. Color-coding of orientation preferences as in figure 1. Adapted from figures 2 and 3 of Koulakov and Chklovskii (2001).

Figure 3: *Self-organizing orientation preference maps and the π-pinwheel density prediction.* The left panel shows an example of a realistic orientation preference and selectivity map with approximately π pinwheel density generated using the self-organizing GCAL model of Stevens et al. (2013). The right panel shows the pinwheel density of three species (diamonds) and simulated maps (circles) as a function of hypercolumn size. Animal data is replotted from Kaschube et al. (2010). Horizontal lines indicate median values of each cluster, with the medians of animal maps and of realistic model maps (GCAL) clustered around π, and those of relatively poorly organized maps (here from the L model, a simplified version of LISSOM) typically being much larger and more variable. Reprinted from Stevens et al. (2013).
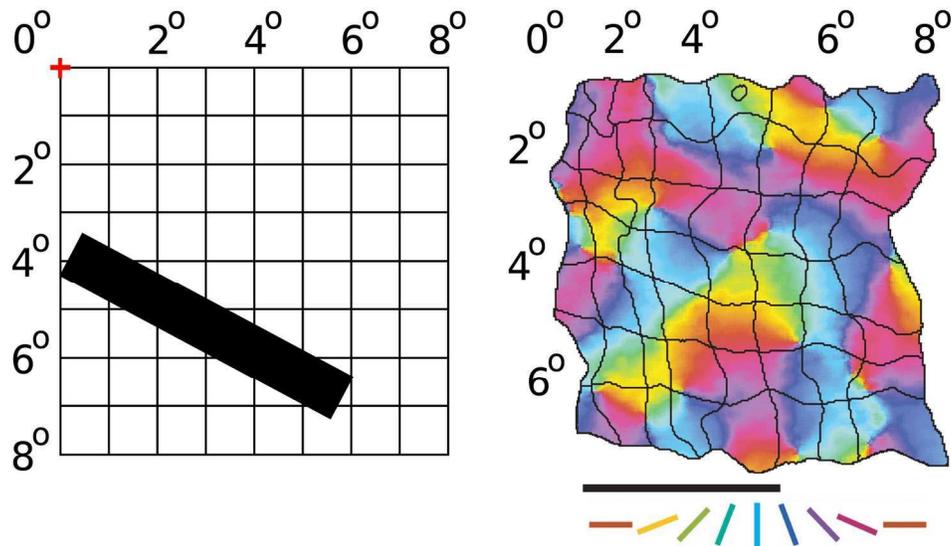
Figure 1: Retinotopic and orientation map in V1. Given a particular fixation point (marked with a red + symbol above), the visual field seen by an animal can be divided into a regular grid, with each square representing a 1°×1° area of visual or retinal space. In cortical area V1 of mammals, neurons are arranged into a retinotopic map, with nearby neurons responding to nearby areas of the retina. As an example, the image on the right shows the retinotopic map on the surface of V1 of a tree shrew for an 8°×7° area of visual space (adapted from Bosking et al., 2002; scale bar below is 1mm). A stimulus presented in a particular location in visual space (such as the thick black bar shown) evokes a response centered around the corresponding grid square in V1 (3° right, 5.5° down). Which specific neurons respond within that general area, however, depends on the orientation of the stimulus. The V1 map is color coded with the preferred orientation of neurons in each location; e.g. the black bar shown at left will primarily activate neurons colored in purple in the corresponding V1 grid squares.
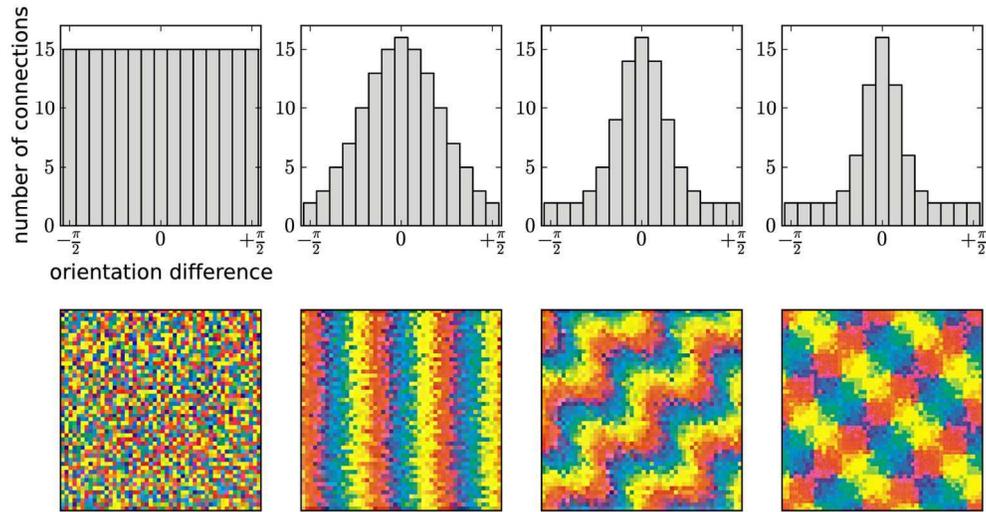
166x95mm (300 x 300 DPI)

Figure 2: Map organization predicted by wiring-length minimization. Panels on the top row show the target histogram of connectivity that each neuron should make to others based on the similarity between their orientation preferences. Target histograms are defined a priori to guide the process of wiring optimization. By the algorithm of Koulakov and Chklovskii (2001), perturbations of each orientation preference are iteratively accepted if reconnecting neurons to the nearest set of neighbors satisfying the target histogram reduces the overall wiring length. The corresponding map organizations that emerge are shown in corresponding panels below. As the target histograms increasingly promote connections between like-tuned neurons, from left to right, the resulting maps can be described as 'salt-and-pepper', 'rainbow', 'staircase', and 'pinwheel-tiling', respectively. Color-coding of orientation preferences as in Figure 1. Adapted from figures 2 and 3 of Koulakov and Chklovskii (2001).
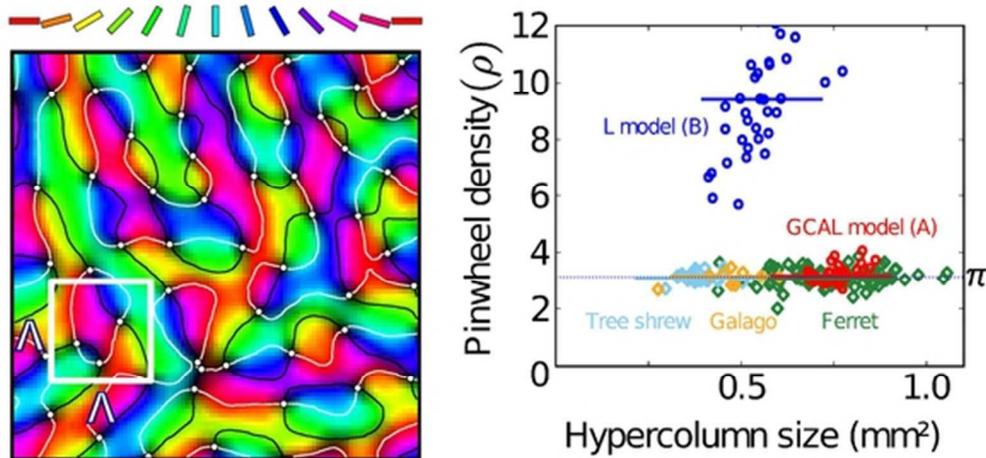133x70mm (300 x 300 DPI)

Figure 3: Self-organizing orientation preference maps and the π-pinwheel density prediction. The left panel shows an example of a high-quality, realistic, orientation preference and selectivity map with approximately π pinwheel density generated using the self-organizing GCAL model of Stevens et al. (2013). The right panel shows the pinwheel density of three species (diamonds) and simulated maps (circles) as a function of hypercolumn size. Animal data is replotted from Kaschube et al. (2010). Horizontal lines indicate median values of each cluster, with the medians of high quality model maps (GCAL) clustered around π, and those of relatively poorly organized maps (here from the L model, a simplified version of LISSOM) typically being much larger and more variable. Reprinted from Stevens et al. (2013).
49x23mm (300 x 300 DPI)