LETTER Communicated by Helge Ritter

Plasticity-Induced Symmetry Relationships Between Adjacent Self-Organizing Topographic Maps

Jared Sylvester
jsylvest@umd.edu
James Reggia
reggia@cs.umd.edu
University of Maryland, Department of Computer Science, College Park, MD 20742, U.S.A.

In many species, adjacent topographic maps in sensory neocortex are found to be oriented as roughly mirror-image copies of one another. Here we use a computational model to show for the first time that, in principle, adjacent cortical topographic maps that are mirror-image symmetric along two dimensions can arise from activity-dependent changes if the distribution radius of afferents sufficiently exceeds that of horizontal intracortical interactions. We also find that infrequently, other types of intermap symmetry and previously unexpected map relationships (such as interlocking rotation, in which two adjacent maps become intertwined) can occur. These results support the hypothesis that activity-dependent synaptic changes play a more important role in forming the orientations of adjacent cortical maps than is currently recognized.

1 Introduction

Mirror-symmetric topographic maps are common in the neocortices of many species (Schulz & Reggia, 2005), but it remains unknown to what extent such symmetries are due to genetic encoding versus thalamocortical activity and synaptic plasticity. Multiple theories have been put forward, the most widely accepted being that the initial partitioning of the cortex into different areas during development is primarily due to genetically determined chemical markers (Levitt, 2000; Sur & Leamey, 2001; Zhou & Black, 2000). However, the degree to which the subsequent reflection-symmetric orientation of maps depends on afferent activity remains unclear (Cohen-Cory, 2002; Grove & Tomomi, 2003; Karbowski & Ermentrout, 2004). Many experimental studies have demonstrated remarkable plasticity in cortical maps due to altered probabilities of sensory stimuli, peripheral nerve lesions, or cortical damage (Buonomano & Merzenich, 1998; Pantev et al., 1998; Raineteau & Schwab, 2001). Recent modeling work established that such mirror-symmetric maps can arise along a single dimension due to activity-dependent synaptic changes, leading to linear sequences of pairwise

mirror-symmetric maps, each having one or two neighbor maps (Schulz & Reggia, 2005).

Here we extend this past computational study of the selective orientation of multiple adjacent self-organizing maps along one dimension to the more general situation where adjacencies occur in two dimensions. It is important to establish this generalization because biological cortical surfaces are two-dimensional, not one-dimensional. Further, the successful generalization to two dimensions is not obvious a priori because it requires that numerous constraints between adjacent pairs of individual maps be satisfied simultaneously over long distances in four directions. Global consistency and well-formedness like this over even a single self-organizing map have been found to be difficult to achieve due to the occurrence of topological defects that arise during map formation (Aoki & Aoyagi, 2007). We show here for the first time that, in principle, multiple adjacent and mirror-symmetric topographic maps can arise in two dimensions from activity-dependent synaptic changes alone, assuming that the distribution radius of already established cortical afferents sufficiently exceeds that of horizontal intracortical interactions during map development (Brown, Keynes, & Lumsden, 2001). We also find that other types of intermap symmetries can occur, although much less frequently, as well as other unexpected types of map relationships, such as interlocking rotations in which the maps become intertwined. These results support the hypothesis that activity-dependent synaptic changes play a significant role in forming adjacent mirror-symmetric cortical maps and raise the possibility that the other types of atypically oriented intermap relationships that occasionally occur might contribute to disordered cortical information processing in childhood neurodevelopmental disorders.

2 Methods

We adopted the same multiwinner self-organizing map architecture and methodology used in the earlier study for maps occurring along one dimension (Schulz & Reggia, 2005), but now modified to permit maps to occur along two dimensions. The output or cortical layer nodes over which map formation occurs are arranged in a regular rectangular grid. The distance between two cortical nodes i and j at positions (r_i, c_i) and (r_j, c_j) is measured with the box-distance metric, $d_{lattice}(i, j) = \max(|r_i - r_j|, |c_i - c_j|)$. Conceptually each cortical node i has an afferent connection from each node in the input layer. These connections have real-valued, nonnegative weights, which are denoted w_{ij} for the connection from input node j to cortical node i. Activation of both input and output nodes is between 0 and 1.

The activation level of the input nodes is given by an input pattern $\vec{x}$ of degree P, with $\vec{x} \in [0, 1]^P$ and $\|\vec{x}\| = 1$. Each input pattern represents the stimulation of a point on a two-dimensional sensory surface. As in Schulz and Reggia (2005), coordinate encoding of input patterns was used both to maintain computational efficiency and permit unbiased comparisons with

the earlier study. Such encoding is a simplification of biological reality but has been shown in the past to be effective for topographic map formation. Training was done using 196 points, sampled uniformly from the unit square in a 14×14 grid. These patterns were then normalized by projecting them onto the unit sphere to avoid introducing biases due to unequal input vector lengths, resulting in three-dimensional input and weight vectors. Specifically, an input point from the unit square $p = (p_x, p_y)$ is projected on to $q = (\frac{p_x}{a}, \frac{p_y}{a}, \frac{b}{a})$, where $a = \sqrt{p_x^2 + p_y^2 + b^2}$ and $b = \sqrt{2} - \sqrt{p_x^2 + p_y^2}$.

The input to each cortical node is given by $h_i = \vec{w}_i^T \vec{x}$, and multiple winners are determined by a one-step selection process in which all nodes that have a greater activation than their N closest neighbors are selected as winners. Here N is the number of cortical nodes within a distance r_{comp} from a node at the center of the lattice. For our computational experiments, we used $r_{comp} = 6$, which gives each node $N = 168$ other nodes to compete against. If cortical node i is determined to be a winner, it is taken to be the peak of an island of activation. That is, node i is taken to be maximally active ($y_i = 1$), and the activations of its neighboring nodes are taken to decrease exponentially according to their distance from i. Specifically, the activation of an arbitrary node j will be $y_j = \gamma^{d(i,j)}$, with $\gamma \in (0, 1)$, and i being closer to j than any other winning node.

Individual weights are initialized to uniform random values on $[0, 1]$, followed by a normalization of each node's weight vector to unit length. Inputs are presented in a randomized order each epoch, and Hebbian weight updating occurs after each, according to $\vec{w}_i = \vec{w}_i + \mu y_i \vec{x}$ with a learning rate $\mu \in (0, 1]$, followed by a normalization step $\vec{w}_i = \frac{\vec{w}_i}{\|\vec{w}_i\|}$ to induce competition between weights. As is often done with self-organizing maps, learning parameters are adjusted as the learning process progresses. Training begins with larger values of γ and μ, and these values gradually decrease. We vary our parameters according to $\gamma(t) = \gamma_{fin} + \frac{\gamma_{init} - \gamma_{fin}}{1 + e^{(t - \gamma_{infl})/\gamma_\sigma}}$ where t is the proportion of epochs elapsed, the initial value $\gamma_{init} = 0.9$, the final value $\gamma_{fin} = 0.0$, $\gamma_\sigma = 0.1$ controls the rate of decrease, and the inflection point between the coarse and fine stages of learning is $\gamma_{infl} = 0.33$. The same function determines μ, using $\mu_{init} = 0.5$, $\mu_{fin} = 0.0$, $\mu_{infl} = 0.5$, and $\mu_\sigma = 0.1$.

In assessment of map formation, common metrics are ill suited for situations like ours where multiple maps are present (Schulz & Reggia, 2005). For this reason, as with the earlier study, we adopt the M metric, which is defined as the mean of the smallest 2% of all pairwise dot products of adjacent cortical node weight vectors. M is inversely proportional to the average distance between receptive field centers of adjacent cortical elements in those regions where such distances are greatest. Because weight vectors are normalized, $0 \leq M \leq 1$, with larger values of M indicating better map formation.

To visualize map formation, we imagine a spiral pattern to be superimposed on the 14×14 grid of input patterns (see Figure 1a). Each input

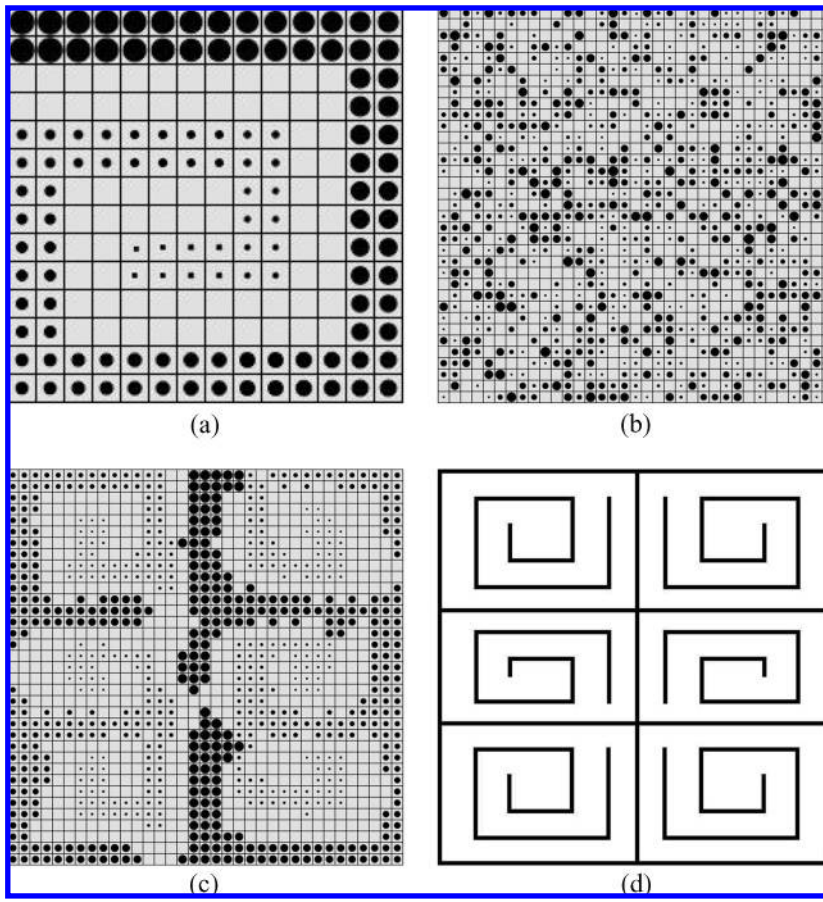

Figure 1: (a) Spiral label pattern used for map formation visualizations, depicted over the sensory surface. (b) Initial cortical map of an untrained network of size 35×35, which is highly disorganized due to random initial weights. (c) The same network as in *b*, after training. Six maps have formed (three rows, two columns) with mirror symmetry between every adjacent pair. (d) Schematic of *c*, showing the orientations of each map and intermap relationships.

location is either represented by a black circle of a specific size (with smaller circles in the middle of the region) or is left blank. The labeling of input locations is for visualization purposes only and does not affect the activation dynamics or learning process used in map formation. Because the spiral pattern is asymmetric, the relative orientation of any two adjacent representations of it on the cortical surface can be unambiguously categorized. When a map is visualized, each cortical node i is associated with the

Table 1: Number of Maps and Symmetries for Different Sized Networks.

	Number of Maps			Pairwise Symmetries			
Size	Mean	Minimum	Maximum	Mirror	Glide	Rotate	Interlock
15×15	1.00	1	1	—	—	—	—
20×20	2.00	2	2	.000	.000	.000	1.000
25×25	2.00	2	2	.000	.000	.000	1.000
30×30	3.35	2	7	.672	.033	.164	.131
35×35	6.25	2	9	.737	.125	.125	.013
40×40	8.25	6	9	.834	.009	.152	.005
45×45	10.10	7	13	.811	.081	.105	.004
50×50	13.00	10	17	.826	.049	.118	.008
55×55	16.15	12	19	.883	.049	.068	.000
60×60	18.40	15	24	.840	.063	.092	.005
65×65	24.40	16	29	.848	.074	.079	.000
Overall (size $\geq 30 \times 30$)				.806	.060	.113	.021

input pattern $\vec{x}_j$, which maximizes its activation. Whatever label (circle or blank) is associated with the input pattern j is placed on the map in the location of node i. For well-formed topographical maps, the input spiral will show up as a slightly distorted but topologically equivalent image on the cortical surface. An example of map visualization before and after training can be seen in Figures 1b and 1c. Figure 1d gives a schematic representation of the six maps shown in Figure 1c. In this example, each individual map of the input surface is a mirror image of each adjacent map.

3 Results

Computational experiments were done using 11 different output lattice sizes, ranging from 15×15 to 65×65. Twenty runs lasting 2500 epochs each, differing only in their initial random weights but otherwise equivalent, were done for each network size. A typical result can be seen in Figure 1c, where six maps, each mirror-symmetric with its neighbors, formed during learning. Informally these model maps correspond to adjacent maps observed in various biological somatosensory cortices (reviewed in Schulz & Reggia, 2005). Table 1 shows the average, minimum, and maximum number of maps formed for each lattice size. In general, the larger the network, the more maps that formed. Maps sizes tended to be roughly $(2r_{comp} + 1) \times (2r_{comp} + 1)$ or occasionally $(2 \times (2r_{comp} + 1)) \times (2r_{comp} + 1)$. This is intuitive, since the map size should correspond to the area of the lattice in which a node must compete for the highest activation in order to be declared a "winner." Averaged across all network sizes, each map occupied 200.83 output nodes, which corresponds to a square with sides 14.17 nodes across. This is equivalent to a radius of 6.58 nodes around the central

node with peak activation, which is close to the value of $r_{comp} = 6$ that was used. If we discount networks that formed only a single pair of interlocking maps, the value falls to 6.43, becoming even closer to the expected r_{comp} value.

For networks 30×30 and larger in which multiple adjacent maps formed consistently, four types of relationships between adjacent maps were observed. Three of these were symmetries (isometries): mirror reflection, glide reflection, and rotation. The fourth was a pattern that we refer to as *interlocking rotation*. These four relationships are illustrated in Figure 2. For networks larger than 30×30, the great majority of intermap relationships observed (about 81% overall) were mirror symmetries (see Figures 1c and 2a), in which one map is reflected in an axis that lies along the maps' mutual border. A further 6% exhibited a distorted type of mirror-image symmetry that we labeled *glide reflection* (see Figure 2b). Glide reflection symmetry is similar to mirror symmetry, except the reflection occurs along an oblique axis and is followed by a translation, so about 87% of adjacent maps exhibited reflection symmetries (mirror and glide reflection). Another 11.3% of neighboring maps exhibited rotational symmetry with their neighbors (see Figure 2c). Approximately 90% of these rotations were 180 degrees, with the remainder being 90 degrees. Table 2 compares these observed percentages of intermap relationships to the theoretically expected percentages of the eight possible relationships between two maps when map orientations are determined randomly and independently (assuming a square tiling and absence of interlocking rotations). As seen in Table 2, the observed percentages of reflections (mirror plus glide) are remarkably high compared to what would be expected. (The 6% observed reflect + rotate relationships are all glide reflections, which can be viewed as an orthogonal reflection followed by a rotation.) The proportion of mirror symmetries increases progressively up to maps of size 40×40 and then levels off at about 85% (or about 91% mirror and glide reflection), rising only slightly as map size increases further. Also, for 7.5% of networks, there were some very small areas of the cortical surface that did not organize into a map of the input space.

We also observed a previously unknown and unexpected type of relationship between pairs of maps, which we refer to as interlocking rotation. An example can be seen in Figure 2d. Interlocking maps arose only rarely, appearing in less than 1% of map pairs in networks 35×35 and larger (but in 100% of networks of size 20×20 and 25×25). For any of the three previous types of symmetries, the border between two maps comprises nodes that are representative of stimuli near one edge of the input space. However, two maps joined by interlocking symmetry have a common border, which represents stimuli from the interior of the input space. Such arrangements seem to arise when there is a surplus of space on the cortex that is needed for one map to form, but not enough space for two whole maps side-by-side (networks of size 20×20 and 25×25). In all of these runs, two maps

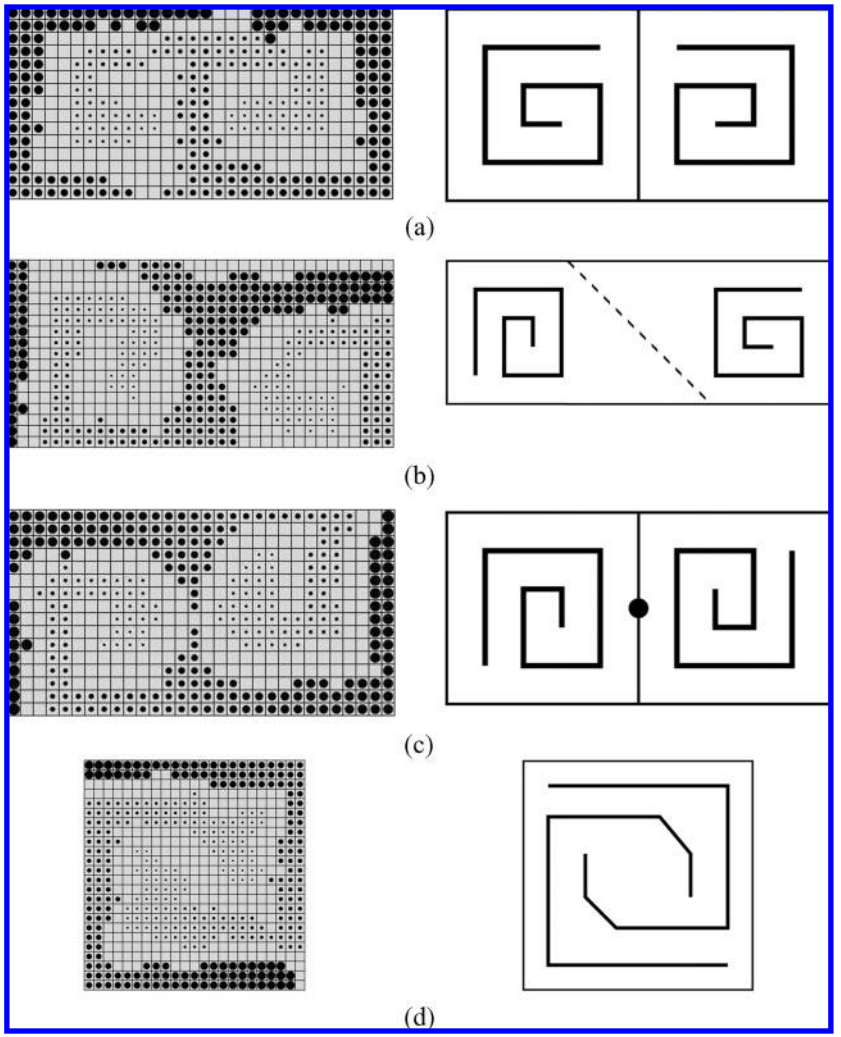

Figure 2: The four types of relationships observed between adjacent maps (left), along with a schematic diagram of each (right). Each display shows a single pair of adjacent maps indicated by how a single spiral "painted" on the sensory surface (Figure 1a) is represented. (a) Mirror symmetry, with a vertical reflection axis in the middle of the cortical surface. (b) Glide reflection symmetry—a distorted mirror symmetry in which the left spiral can be viewed as translated along an oblique axis (dotted line) and then reflected to match the right spiral. (c) Rotational symmetry, where the right spiral matches the left after a 180 degree rotation about a rotation point (small, solid black circle in center of schematic). (d) Interlocking rotations in which the interiors of the two spirals are intertwined.

Table 2: Approximate Percentages of Adjacent Map Relationships.

	Translate	Rotate 180°	Rotate 90°	Reflection	Reflect + Rotate
Expected	12.5	12.5	25	12.5	37.5
Observed (size $\geq 30 \times 30$)	0	10	1	81	6

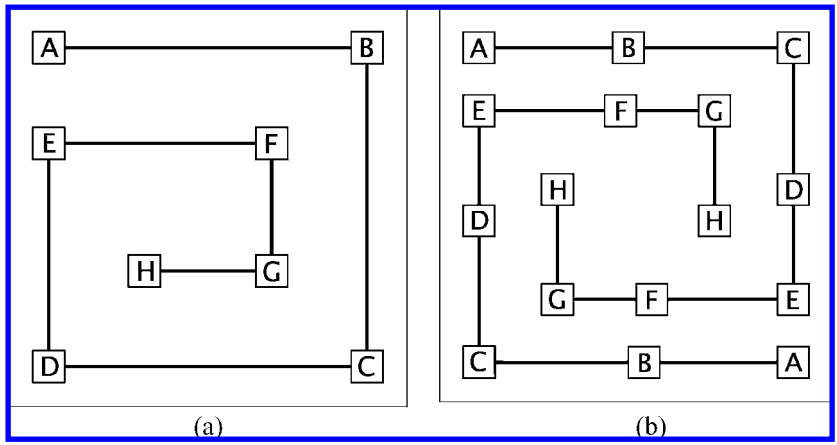

Figure 3: Schematic of interlocking rotation. (a) The input space, with eight sample points labeled A through H. (b) Where the peaks of activation might fall for each of these labeled points in the interlocking maps that form on the cortical surface. The representations of the input space have been "unwound" and compressed so that each copy of the perimeter of the sensory surface, comprising points from all four sides, covers just two sides of the cortical surface. Despite the apparent distortion brought about by interleaving the maps, each point on the cortex maintains the appropriate topographic relationship with its neighbors. For example, B on the input surface, situated between A, F, and C, remains surrounded by A, F, and C in the cortical map. Although the two maps have become unraveled and intertwined, the topographic relationships are still preserved locally.

formed in each network and aligned themselves in an interlocking pattern. Figure 3 provides a more detailed characterization of this phenomenon.

To clarify the conditions most likely to lead to interlocking map formation, we systematically carried out additional experiments with networks ranging in size from 15×15 to 30×30 and with r_{comp} values varying between 4 and 8. Ten random runs were done for each combination of network size and r_{comp}. Table 3 shows how many of these trials resulted in a pair of interlocking maps. These results confirm that for all values of r_{comp}, very small networks do not have enough space on the output layer

Table 3: Number of Times Two Interlocking Maps Formed in 10 Trials.

	Side Length of Square Cortical Layer																				
r	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35
4	4	7	8	10	9	10	8	8	4	3	0	0	0	0	0	0	0	0	0	0	0
5	0	1	2	10	10	10	10	10	9	8	8	9	9	5	3	2	0	0	0	0	0
6	0	0	0	4	7	10	10	10	10	10	10	10	10	6	6	5	4	7	2	1	1
7	0	0	0	0	1	6	7	10	10	10	10	10	10	10	9	7	8	1	4	2	5
8	0	0	0	0	0	0	5	6	9	8	10	10	10	10	10	10	10	10	6	6	4

to support more than a single map, while the larger-sized networks have enough space to support multiple maps with the usual mirror, glide, and rotation symmetries. In between this range of network sizes, a single pair of interlocking maps was the dominant pattern. The networks most suitable for interlocking map formation were those with a side length between approximately $1.5 \times (2r_{comp} + 1)$ and $2 \times (2r_{comp} + 1)$.

We also ran five simulations having much larger 120×120 output layers to verify that the results in Table 1 scale up to bigger networks in which any edge effects will be less influential. The number of maps that formed ranged from 73 to 87 (mean: 77.2), with an average of 86.4% mirror reflections, 5.1% glide reflections, and 8.5% rotations, consistent with extrapolation from Table 1. Figure 4 shows the maps that formed during one of these runs. All map relationships in Figure 4 are mirror-symmetric except the eight rotational symmetries labeled R and three glide reflections labeled G. We interpret this result to indicate that mirror-symmetric relationships between adjacent maps are to be expected in that they globally minimize the distance between weight vectors of neighboring nodes. The rotational and glide relationships generally occurred along "defects" or "fault lines" separating different uniform domains that have all mirror-symmetric maps. This is illustrated in Figure 4, where a uniform domain in the uppermost right quadrant (e.g., the inner "u" of each spiral opens either up or down) can be distinguished from a different uniform domain throughout the other three quadrants (inner "u" of all spirals opens left or right), these two domains being separated by a diagonal boundary zone of maps having reflection and glide relationships. We hypothesize that such intermap defect lines are the analog of intramap topological defects (so-called folds or bumps) often observed in single-map studies (Aoki & Aoyagi, 2007), and that they occur for the same reason.

For purposes of further analysis, we divided the composite maps formed with the 160 networks of Table 1 that were of size 30×30 and larger into four mutually exclusive categories, based on the types of symmetries present. Because it was not possible to measure M for particular isolated pairs of maps, we grouped the networks together into categories based on the dominant type of symmetry present over the entire cortical area. The interlocking

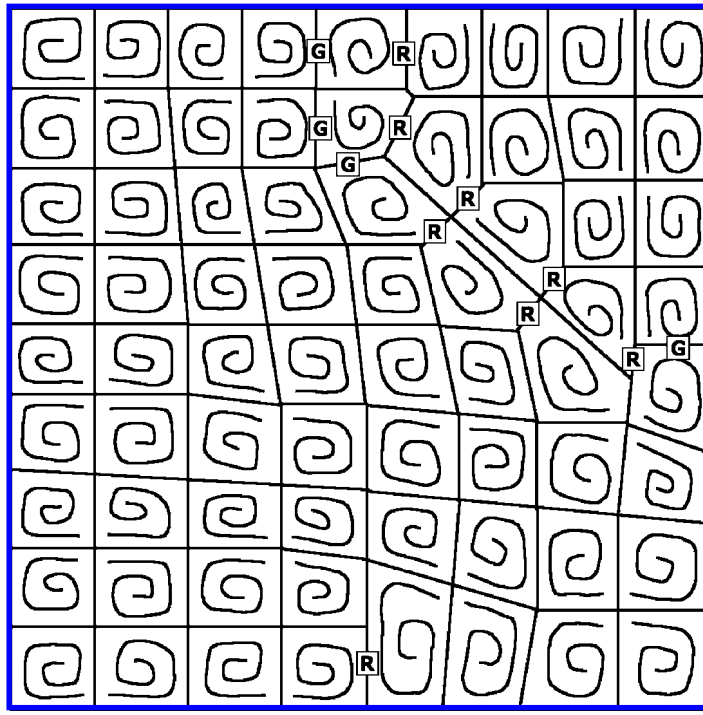

Figure 4: Maps forming on a large 120 × 120 output layer. All are mirror-symmetric except those designated G (glide) or R (rotation).

category are those networks for which the entirety of the cortical surface was covered by a single pair of interlocking maps, which occurred only for networks of size 40 × 40 and below. The mainly-mirror category comprises networks in which all or all but one of the map pairs were mirror-symmetric. The mixed category are networks that displayed more than one instance of glide, rotational, or interlocking symmetries, although mirror-symmetric maps were also present simultaneously. Finally, the unorganized category is composed of networks that had any unorganized regions on the cortical surface.

Table 4 shows the mean μ and variance σ^2 of the M measure across all the networks in each of these four categories. The interlocking maps exhibited the highest M values, followed by mainly mirror, then mixed, and finally the unorganized category. Using a t-test, we found that the mean M value for the interlocking class was significantly higher than that of the mirror class with $p < 0.0001$, that the mirror class mean was significantly higher than that of the mixed class with $p = 0.0054$, and that for the mixed class was higher than for the unorganized class but only with $p = 0.0331$. The higher

Table 4: Values of the Measure M for Symmetry Categories.

Symmetry Category	μ	σ^2	N^a
Interlocking rotation	.9881	.0043	50
Mainly mirror	.9665	.0091	43
Mixed	.9622	.0077	95
Unorganized	.9572	.0076	12

[a]Number of networks in the category.

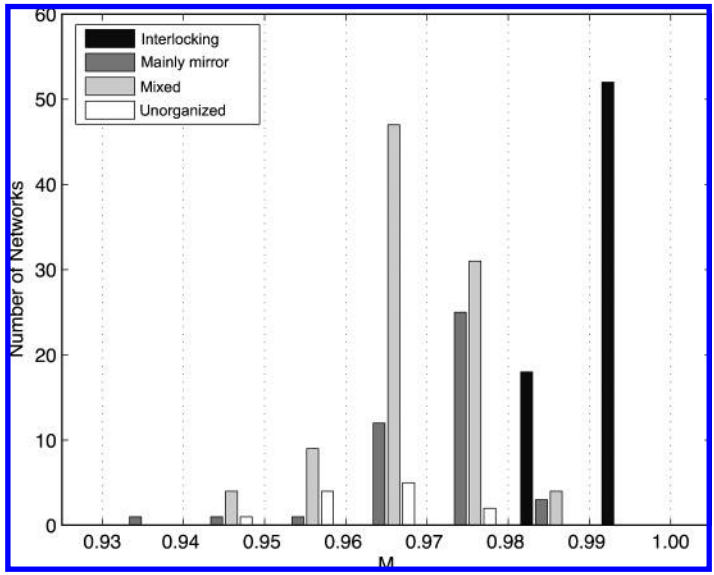

Figure 5: Histogram of M values for four different map categories.

M values for the interlocking class can be seen more clearly in Figure 5, which gives a histogram of the networks' individual M values for each class. These high M values may arise because the two maps do not form a strict border between themselves, but "share" the output nodes that represent the center of the input space. For comparison, 15×15 networks with a single map had $M = 0.9930$. Figure 5 also shows that while their mean M values are close, the distributions of M values for the mirror and mixed class are substantially different (e.g., they have different mode values).

To test the effect of uneven distributions of input stimuli on the model, the input space was divided into three approximately equally sized regions. Region I was composed of the 64 input points located closest to the center, region III was composed of the 64 input points around the perimeter, and region II was composed of the remaining 68 points between regions I and III.

Table 5: Map Formation with Varying Spatial Distributions of Stimuli.

	Number of Maps			Pairwise Symmetries			
Type	Mean	Minimum	Maximum	Mirror	Glide	Rotate	Interlock
60×60 rim heavy	20.80	16	29	.914	.035	.051	.000
60×60 uniform	18.40	15	24	.840	.063	.092	.005
60×60 center heavy	23.42	15	29	.696	.125	.178	.001
Overall	20.87	15	29	.817	.074	.107	.002

Using the same procedures and parameter values, we ran two additional sets of 20 simulations, each having networks with 60×60 node output layers. In the first set of rim-heavy simulations, points in region III were stimulated three times as often as those in region I, and points in region II were stimulated twice as often as points in region I. The second set of center-heavy simulations reversed these ratios.

As seen in Table 5, uneven sampling that was either rim heavy or center heavy was associated with a small increase in the average number of maps that formed when compared to the uniform simulations presented earlier. Rim-heavy simulations had the highest proportion of map pairs exhibiting mirror symmetries. Conversely, center-heavy simulations formed fewer mirror symmetries than the uniform-stimuli baseline. The mean M values for rim-heavy, uniform, and center-heavy networks were 0.9583, 0.9644, and 0.9570, respectively. The score of the uniform category is significantly higher than either of the others at a 99% confidence level, but there is no statistically significant difference between center-heavy and rim-heavy networks. If one examines all neighbors' pairwise dot products rather than just the smallest 2% used in the M-metric, the rim-heavy network had more pairs of nodes with very similar weights, while the center-heavy network's distribution showed nodes that are farther apart on average, and the uniform network fell in between. Recalling that there were more mirror symmetries in the rim-heavy network, followed by uniform, and finally center heavy, thus offers confirmation that mirror symmetries allow more adjacent nodes to minimize the distance between their weight vectors.

We also explored the possibility of magnification effects resulting from nonuniform stimuli. More frequent stimulation or greater innervation of a sensory region often causes that region's representation on the cortical map to be disproportionately large (Dykes & Ruest, 1984; Kohonen, 2001). The proportion of the cortex representing the three sensory regions, for uniform and nonuniform stimuli distributions, is shown in Table 6. For uniform stimulations, we find that central region I is overrepresented on the cortical surface, and rim region III is underrepresented compared to the roughly equal distributions one might expect. Under the rim-heavy distribution, the most stimulated area, region III, increased to 40.4% of the cortex, consistent with magnification effects seen in past single-map studies (Grajski &

Table 6: Proportion of Cortical Nodes Representing Each Sensory Region.

Type	Region I	Region II	Region III
60×60 rim heavy	.283	.313	.404
60×60 uniform	.452	.344	.204
60×60 center heavy	.474	.259	.267
Overall	.403	.305	.292

Merzenich, 1990; Sutton, Reggia, Armentrout, & D'Autrechy, 1994), while region I decreased to 28.3% (significantly different from the uniform case with $p < 10^{-6}$). An opposite magnification effect is also evident for the center-heavy distribution ($p < 10^{-6}$).

4 Discussion

The model of map formation presented here, while greatly simplified from biological reality, produces individual maps with many features that are reminiscent of topographic map formation in nature and many past computational models of single map formation. Multiple maps emerged over the cortical surface solely as a by-product of the tendency of Hebbian learning to have adjacent output nodes take on similar weight vectors. The limited radius at which nodes compete with one another relative to the size of the network allowed the multiple maps that formed to arrange themselves with four different types of intermap relationships. Of these, the mirror, glide, and rotational symmetries were previously observed in a similar study limited to maps with one-dimensional adjacencies that form along a narrow cortical strip (Schulz & Reggia, 2005), and in similar proportions. The key point is that in both this and the earlier study, adjacent maps became oriented with reflection symmetry in the vast majority of cases: 87% and 93%, respectively, counting mirror and glide reflections, indicating that map formation in one region strongly influences is influenced by that in adjacent regions. Further, in the largest networks, the rotational and glide relationships tended to occur along fault lines between regions where maps were all uniformly mirror-symmetric, suggesting that such relationships are the intermap analog of topological defects sometimes observed inside single maps (Aoki & Aoyagi, 2007). We also found that changing the distribution of stimuli occurring over the sensory surface resulted in magnification effects in the cortical maps and that when the points near the edges of the sensory surface were stimulated more frequently than other points, mirror symmetries became more common, presumably due to the increased similarity between edges of adjacent maps that this produced.

We also noted a new type of relationship between maps that was not present in previous work, which we refer to as *interlocking rotation*. These were likely not formed in the previous study (Schulz & Reggia, 2005)

because the cortical regions of the networks being considered were all 11-node-wide strips, while interlocking maps seem to occur only when there is more than enough room for one map but less than enough for two. More specifically, our simulations showed a minimum size for two interlocking maps to form and a maximum size at which they are common, even though they have higher M values than mirror, glide, or rotational symmetric maps. Judging from the M metric, interlocking symmetries are a more efficient arrangement of the cortex, but because it does not have clear distinctions between where one map begins and the other ends, it may not lend itself to the improved downstream processing and parallelization proposed as a possible benefit of multiple maps (Kaas, 1988; Levitt, 2000; Zhou & Black, 2000).

Finally, we speculate on the significance of non-mirror-image map relationships, such as rotation symmetric and interlocking rotation maps. Given the remarkably common finding of mirror-image adjacent maps in the cortical regions of many species and that they form the vast majority of intermap relationships in our computational model, it seems reasonable to view mirror-image maps as the norm. Presumably the genetically determined topographic organization of interregional connections and pathways has evolved to be consistent with such mirror-image orientations. Assuming this is so raises the question of what would happen if, as occurred with a small percentage of our computational cases, two adjacent maps organized during development to have rotational or interlocking rotational relationships. We hypothesize that such atypically oriented adjacent maps, in the context of normal connectivity between cortical regions, would be expected to cause abnormal cortical information processing and that this might account for some of the cognitive deficits and functional imaging changes observed in childhood neurodevelopmental disorders such as dyslexia, dyscalculia, autism spectrum disorder, and specific language impairment.

Acknowledgments

This work was supported by NINDS grant NS35460 and NSF grant IIS-0325098.

References

Aoki, T., & Aoyagi, T. (2007). Self-organizing maps with asymmetric neighborhood function. *Neural Computation*, *19*, 2515–2535.

Brown, M., Keynes, R., & Lumsden, A. (2001). *The developing brain*. New York: Oxford University Press.

Buonomano, D., & Merzenich, M. (1998). Cortical plasticity. *Ann. Rev. Neurosci.*, *21*, 149–186.

Cohen-Cory, S. (2002). The developing synapse: Construction and modulation of synaptic structures and circuits. *Science*, *298*, 770–776.

Dykes, R., & Ruest, A. (1984). What makes a map in somatosensory cortex? In E. Jones & A. Peters (Eds.), *Cerebral cortex* (vol. 5, pp. 1–29). New York: Plenum Press.

Grajski, K., & Merzenich, M. (1990). Neural network simulation of somatosensory representational plasticity. In D. Touretzky (Ed.), *Advances in neural information processing systems, 2* (pp. 52–59). San Francisco: Morgan Kaufmann.

Grove, E., & Tomomi, F. (2003). Generating the cerebral cortical area map. *Annual Review of Neuroscience*, *26*, 355–380.

Kaas, J. (1988). Why does the brain have so many cortical areas? *J. Cognitive Neuroscience*, *1*, 121–134.

Karbowski, J., & Ermentrout, G. (2004). Models of the early development of thalamocortical connections and area patterning via signaling molecules. *J. Computational Neurosci.*, *17*, 347–363.

Kohonen, T. (2001). *Self-organizing maps* (3rd ed.). New York: Springer.

Levitt, P. (2000). Molecular determinants of regionalization of the forebrain and cerebral cortex. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 23–43). Cambridge, MA: MIT Press.

Pantev, C., Oostenveld, R., Engelian, A., Ross, B., Roberts, L., & Hoke, M. (1998). Increased auditory cortical representation in musicians. *Nature*, *392*, 811–814.

Raineteau, O., & Schwab, M. (2001). Plasticity of motor systems after incomplete spinal cord injury. *Nature Reviews Neuroscience*, *2*, 263–273.

Schulz, R., & Reggia, J. (2005). Mirror symmetric topographic maps can arise from activity-dependent synaptic changes. *Neural Computation*, *17*, 1059–1083.

Sur, M., & Leamey, C. (2001). Development and plasticity of cortical areas and networks. *Nature Reviews Neuroscience*, *2*, 251–262.

Sutton, G., Reggia, J., Armentrout, S., & D'Autrechy, C. (1994). Cortical map reorganization as a competitive process. *Neural Computation*, *6*, 1–13.

Zhou, R., & Black, I. (2000). Development of neural maps. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 213–236). Cambridge, MA: MIT Press.

Received April 21, 2008; accepted April 25, 2009.