

Applications of Computer Graphics and Image Processing to 2D and 3D Modeling of the Functional Architecture of Visual Cortex

Eric L. Schwartz, Bjorn Merker, Estarose Wolfson,
and Alan Shaw
New York University



The visual cortex of monkeys is one of the most extensively studied areas of the primate brain. Nearly half of monkey cortex is devoted to visual processing, and the twenty-odd functional areas that make up visual cortex represent, arguably, one of the most complex mechanisms in nature.

Attempts at understanding visual cortex pose a wide range of problems in computer graphics, image processing, computational geometry, and numerical methods. The term computational neuroscience has recently come into use to describe this area of study.

In this article we describe a series of studies in this area of work which illustrate a number of applications of computer graphics and image processing to the reconstruction and representation of the complex architectures that make up primate visual cortex. We demonstrate techniques for reconstructing brains in three dimensions, "peeling" them apart, and flattening the brain with minimal metric error. Finally, we show simulations of natural images as they are mapped in the brain by these architectures, including a simulation of a stereo image at the level of primary visual cortex.

Figure 1a shows the back half of one hemisphere of the brain of a macaque monkey, including the primary visual cortical area called V1.

The term cortex refers to the thin outer layer of the brain, consisting of the cell bodies of neurons. In Figure 1c we see the thickness of the cortex (about 1 mm) represented as an outer sheet, while the white matter (axons

or fibers alone) of the brain occupies most of the central volume. One of the most salient features of the brain is that cortical neurons tend to lie in thin layers, which share a common functional context. Primary visual cortex has been subdivided into at least 10 distinct layers across its thickness. Layer IV, roughly in the center of the cortical band (see the dark band in Figure 1c) is where a principal input from the retina enters the cortex.

Basic facts: Architecture as maps and columns

For the present discussion, we emphasize two basic facts about the neurons of primary visual cortex:

1. Nearby points in the visual field correspond to nearby points on the cortex.¹ The pattern of neural firing across the surface of the cortex is an image-like representation of the stimulus attributes of the image being viewed. Topographic mapping of this sort is a typical feature of sensory cortex in the auditory, somatosensory, and visual modalities. One of the goals of the present work is to be able to measure and simulate these maps.
2. Multiple maps are combined in the form of stripes or columns so that multiple stimulus parameters can be represented in a single map layer. In humans and monkeys, the left and right eyes project to layer IV of visual cortex, where they are interlaced in "ocular dominance columns." This feature of visual cortex will be extensively illustrated later. Similar

interlacing is a common architectural feature of neocortex in cats, monkeys, and humans. Examples are orientation columns (cat and monkey V1) and direction columns (monkey medial temporal cortex). Thus, we are presented with the problem of images composed of interlaced subimages, stacked up and twisting through 3-space.

Statement of the problem vis-a-vis computer graphics and image processing

There are three distinct levels of computational neuroscience applications in visual cortex:

1. At the lowest level there is the problem of reconstructing and representing the layered structure of the cortex, using computer graphics and image processing.
2. At an intermediate level there is the problem of modeling the map functions and columnar architectures that are embedded in the cortex.
3. At the highest level there is the problem of inferring the computational significance of these architectures for brain function.

In the present article we focus on the first two levels and briefly discuss the third.

Level 1: Computer-aided neuroanatomy of visual cortex

We now discuss the bottom-most problem outlined above: Given a brain, as represented by some hundreds of serial sections, we want to reconstruct a computer model of it in 3D; to "peel" apart this model, so we can view the individual layers of the cortex; and to "unroll" or "flatten" these layers to view them better and to simplify the numerical calculations to be performed on them.

A simple analogy can help clarify this problem: Imagine a stack of 10 photographs, each with a different image in it. Imagine this stack folded into a complex 3D shape, with deep convolutions and bends, and then potted in epoxy. If one's goal were to retrieve the photographs, one could slice this 3D object into many thin sections, say 400, and then begin a computer reconstruction of the sections back into the intact 3D folded stack. Then, one would have to peel apart the layers of this 3D object, to access the original photographs which were embedded within it. Finally, one would need to flatten out the crumpled photographs. We have constructed a system to effect precisely these operations. (The operations were implemented in Sun Unix in about 30,000 lines of C code.)

Experimental design

Two experimental techniques we have used in this work are based on the ability to visualize metabolic activity in the brain by two stains: cytochrome oxidase (CO) and 2-deoxyglucose (DG). (A detailed discussion of these techniques can be found in standard references.^{2,3}) The CO and DG methods can be thought of as providing a photographic record of the average activity of neurons as a result of a long-term or a short-term stimulus condition respectively. Thus, if one eye of a monkey is inactive for several weeks, a CO image of the connections of the remaining active eye will be accessible in a number of different cortical layers (see Figures 2 and 5). Similarly, if an image is shown, unmoving, to a monkey for perhaps 20 minutes, then the representation of that image can be visualized in many layers of the brain via DG imaging.

There is the assumption that metabolic activity is proportional to neuronal activity, which is what we really want to visualize. In effect, we would like a kind of "magic television" on which the spatial representation, or patterns of neural firing, could be visualized. The CO and DG techniques are among the best approximations to this "magic television," but they still leave us with a massive problem of computer graphics. That is, the patterns of interest are stacked and folded in the 3D brain. Now we will illustrate the algorithms and procedures required to section, process, reconstruct, peel, and unfold them.

Digitizing, alignment, and spackling

Data is digitized with a CCD camera using a photographic color enlarger as light source. The slide to be digitized is placed in the optical path of the enlarger and imaged with an appropriate lens to fit onto the surface of a Fairchild CCD chip of resolution 488 × 380 pixels. This image is then digitized and stored on disk (see Figure 1).

To accurately align large numbers of computer images of serial sections of primate brain, a movie of the brain is shot as the brain is being cut. The movie is then digitized and used as a guide. The resulting alignment of the computer images has been satisfactory, allowing reconstruction of such detailed patterns as ocular dominance columns across the entire surface of primate visual cortex.

It might be thought that simple cross-correlation of adjacent sections would be an adequate method of alignment. However, for 3D structures that are not axially symmetric, cross-correlation is grossly unreliable. It introduces large systematic errors into the 3D brain model because one tends to "overalign" adjacent sections, that is, to ignore the fact that some small positional or rotational difference between the sections is actually correct. In our experience, shear of up to 100 percent of the dimension of the brain and torsion of up to 90

degrees are introduced by purely cross-correlational alignment of monkey visual cortex. Careful superposition of the brain sections with actual fiducial landmarks (photographs of the original frozen brain block) is necessary to achieve a veridical alignment of the data.

Three-dimensional surface tracking depends on a precise definition of the “inside” and “outside” surfaces of the cortex. They must not be connected by so much as a single voxel or the surface tracking will fail. But serial sections are often flawed. Blood-vessel holes are usually visible, and small rips or tears and other local damage are common. Thus, before performing surface tracking, it is necessary to repair serial sections. We use a combination of image processing methods including thresholding, histogram equalization, median filtering, etc., to perform this repair, which we term “spackling.”⁴

It is important to emphasize that this operation merely enhances existing image detail, and repairs small defects (for example, holes due to blood vessels) by interpolating neighboring image details. Figure 1b shows the result of image processing. The white matter of the brain has been thresholded, contrast has been enhanced, and small areas of image defect have been interpolated. We estimate that only a few percent of the image area is affected by this latter operation.

Brain peeling: A digital tangential microtome

A microtome is a fixed knife that resembles the familiar sandwich-shop meat slicer. Its purpose is to cut thin sections of brain and other organs so that the sections can be mounted on glass slides, stained, and studied. Naturally, the microtome sections are parallel to the plane of the knife. This causes major difficulties in understanding cortical anatomy, because the interesting structure of cortical systems lies in the “tangential plane” of the cortex. Because the brain surface contains many folds, it is very difficult indeed to visualize the architecture of the brain by viewing serial sections.

Tangential sections are usually produced by cutting conventional blocks of brain approximately parallel to the cortical surface. One way to ensure this is to physically flatten brain tissue (for example, striate cortex operculum) and then cut the flattened tissue tangentially. We use computer image-processing techniques to get true tangential sections from a digitized 3D reconstruction of the brain. The actual physical sectioning is in an arbitrary plane.

Given suitably aligned and prepared serial sections of brain, we apply a 3D surface-tracker⁵ to find the inner and outer boundaries of the cortex by creating an adjacency graph of all the surface elements. We close the brain by capping planes, and we use the outer detected surface as a shield to avoid “leakage.” We then output the inner surface as a set of voxel files, remove it from the 3D model, and repeat the procedure until we obtain about 25 “peels.” Each peel is both an abstract surface and a gray-scale coded (stained) map that can be displayed by




Figure 1. In (a) we see a block of monkey brain. The smooth roof (operculum) of visual cortex is the bullet-shaped region at the right. In (b) the red background shows the final result of aligning, thresholding, spackling, median filtering, histogram equalization, and further median filtering. In (c) the white background shows the original data. The dramatic increase in contrast is due to the histogram equalization. The ocular dominance columns form the periodic dark and light pattern that is most evident in the middle layers of the cortex.

conventional voxel rendering techniques.⁶

Figure 2 shows a single “brain peel” representing much of the posterior pole of a monkey brain hemisphere. This monkey had only one functional eye, so that the pattern of ocular dominance columns could be made visible by a cytochrome oxidase stain. This graphic represents a voxel reconstruction (1 voxel = 40μ 40μ × 40μ) of the pattern of cytochrome oxidase stain in the 12th peel from the top of the brain. We have thus




Figure 2. This image represents a single brain peel. It is a 40μ -thick shell, lying largely within layer IV of the posterior hemisphere of a macaque brain. The cytochrome oxidase data have been mapped onto the 3D surface. Clearly visible are the ocular dominance columns of striate cortex (heavy “zebra stripes”). Also visible are some wider V2 columns (seen best in the “flap” at the top of the brain). This image was prepared with a “laminar” brain peeler algorithm, which simply “shaved” successive thin shells from the full 3D cortex, much in the fashion of a physical microtome knife.

made visible in this display the laminar pattern from a depth 580μ below the cortical surface.

The next step is to “flatten” the brain to view the map structures within it more easily, and, more importantly, to be able to measure and characterize them in planar representation.

To perform this brain flattening, we need an additional data structure on the brain, in the form of a 3D triangulation. We will now briefly describe methods we have developed to achieve automatic 3D triangulation of such complex surfaces as the brain.

Voxel-to-polyhedral and polyhedral-to-voxel transformations

A voxel data structure is required for brain peeling and detailed imaging, but it is far too detailed to allow convenient numerical manipulation, such as flattening. We therefore need to construct a conventional polyhedral model of the same surface.

There is a large body of literature on the problem of triangulating 3D surfaces from serial sections, or contours.⁷⁻¹² These methods work well on segments of

“generalized cylinder”—that is, runs of sections containing single loops—but they often fail when attempting to process highly convoluted surfaces, because of changes in 2D topology from section to section. Laborious and error-prone human interaction is then needed.

Our method¹³ uses the full description of a surface provided by a voxel model to classify all the topological transformations and thereby guide the triangulation process, automatically ensuring a correct representation of the surface. The voxel surface is viewed as a discrete sampling of an orientable differentiable manifold, whose topological properties admit of only four kinds of transformations: birth, death, merging, and splitting of contours. We detect the “critical levels” where such changes occur, and localize the transitions at “critical points.” Where contours merge or split, the critical point is used to link together the several contours in the critical level that participate in the transition, and hence the generalized cylinders that terminate at these contours. Figure 3a shows such an automatically produced polyhedral model, which corresponds to the voxel model shown in Figure 2.

The reverse operation, transformation from a polyhedral to a voxel model, is achieved easily by scan conversion. Figure 3b shows a polyhedral model of the surface of primary visual cortex which was scan-converted into voxels and rendered by standard voxel display techniques (“back-to-front”). We are able to go freely back and forth between the polyhedral and voxel data structures.

Brain flattening

The mapmaker’s problem is to find a flat representation of a curved surface. Classical mapmaking has been restricted to the relatively simple spherical surface of the earth. In the case that the surface of interest is complex, and possibly nonconvex, there are no known methods of finding quasi-isometric planar representations.

In previous work we described an early version of our brain flattening algorithm.¹⁴ Briefly, we compute the “geodesic” distances along the surface between all nodes of the polyhedral model of the brain. Finding minimal distances on a polyhedral surface is difficult. Recently, an algorithm with polynomial complexity has been described to perform this function.¹⁵ To our knowledge, this rather difficult algorithm has not yet been implemented. We have developed a simpler algorithm, with exponential complexity, but which we have implemented and used successfully for computing the metric structure of brain surfaces.¹⁶

Once we have the matrix (or some subset of it) representing interpoint distances in the surface, we perform a gradient descent on the goodness of fit of the 3D distances to a set of (initially random) 2D distances, as described in the earlier *CG&A* article.¹⁴ The result is a planar model of the brain configuration such that the metric properties of the original 3D model are optimally




Figure 3. Figure 3a represents the back half of a monkey brain hemisphere, triangulated by an automatic algorithm, which required no human intervention. This algorithm uses a full voxel surface tracking of the cortex to effect a veridical triangulation. In 3b we see a computer-generated reconstruction of the entire striate cortex, viewed from the bottom. The opercular surface is away from the viewer, while the convoluted “calcarine” cortex appears as a flower-petal shape closer to the viewer. This rendering was constructed by scan-converting a polyhedral model of the cortical surface consisting of about 2,500 nodes, followed by a “back-to-front” voxel rendering algorithm.

preserved. Figure 4 shows an example of primary visual cortex (the polyhedral model of Figure 3), which has been flattened by this procedure. The local “error” of the flattening averages 5 percent, and it is color coded in Figure 4.

Brain painting from 3D to 2D and from 2D to 3D

Having constructed voxel and polyhedral models of the same brain and flattened the polyhedral model, we have one final task: to texture-map the gray-scale values from 3D into the 2D flattened model. We perform this texture-mapping by deriving a bilinear warp function from the coordinates of the vertices of each triangle of the 3D polyhedral model and those of its corresponding flattened 2D shape. Then we scan-convert the destination triangle and use the derived bilinear warp function to find the source point in 3D from which to transfer the value to each pixel. The result of this operation is shown in Figure 5, which indicates the ocular dominance column pattern of macaque V1.

Level 2: Simulating high-resolution monocular and stereo images mapped through the ocular dominance column system

Through the series of steps outlined above, we have attained control over the fine detail of the entire cortical surface, as represented by a set of stained serial tissue sections. This software, together with appropriate experimental techniques, allows us to measure cortical map functions with precision.




Figure 4. This shows the flattening of the 3D polyhedral model of Figure 3a. A small section of the most peripheral extent of V1 in this data (about 1-2 mm) was lost, and is not shown. The error is color coded in units of percent.

But what is the mathematical structure of these maps? In this section we show some recent simulation work from our lab, using image processing methods to simulate the map structure of primary visual cortex on high-resolution images.

We know from previous work that the complex logarithm function provides a good approximation to the visual map of primate cortex.¹⁷⁻²⁰ Figure 6 shows a




Figure 5. The borders of V1 and V2 were determined by inspection, and V1 was isolated from the “peel” of Figure 2. This V1 model was then flattened. This figure shows the cytochrome oxidase data of a 3D peel (as in Figure 2) “image-mapped” into the flattened V1 model. This peel was produced by an “anatomical peeler”; it corresponds to the bottom of layer IV of striate cortex. The ocular dominance column pattern is visible.

simulation of the cortical representation of a natural scene. The original scene of our lab hallway, digitized at an effective resolution of $16,000 \times 16,000$ pixels, is in Figure 6a, while Figure 6b shows an eye chart at the center of the scene. The eye chart was located 20 feet from the camera, so that the 20/20 line (line 8) would be readable in this simulation (and it was in the original).

Figure 6d shows a complex-log simulation of this 60-degree field, with one minute of arc resolution at the center. Note that the eye chart occupies almost a third of the entire cortical surface, whereas in the visual field, the eye chart occupies less than 1 percent of the solid angle. This is a graphic demonstration of the strongly space-variant structure of the human retina and cortex.

Humans have a very small foveal region of very high resolution (about one minute of arc) and resolution falls off in approximately inverse proportion to angle away from the fovea.

But there is an additional complication in cortex, caused by the existence of ocular dominance columns. The left and right eyes are “interlaced,” as shown in Figures 5 and 7. The question naturally arises, given a simulation of the map of a single eye (via the complex log, or any other map function), what graphics and image processing problems are raised by the columnar interlacing? We have developed a solution to this problem, based on the properties of generalized Voronoi regions, which we will briefly outline here.

Generalized Voronoi regions, ocular dominance columns, and proto-columns

Historically, little attention has been paid to the details of how an interlacing of data from the entire left and right eyes might be produced in a single map. We now describe a simple algorithm that can interlace an arbitrary number of complete maps into a single layer, via columns. The columns can have an arbitrary polygonal shape, and each pixel of each input map is assigned to a specific column.

Full “proto-maps” exist before birth

Rakic has shown that, before birth, monkey visual cortex contains two full visual maps, one each for the left and right eyes, at the level of striate cortex.²¹ We call these maps “proto-maps.” We can visualize the cortical proto-columns as existing in each of the cortical proto-maps (see Figure 7). Slightly before birth the cells in the two visual maps migrate together into a single map, forming the ocular dominance column system.

How does a given cell “know” which column to migrate toward? The simplest rule is based on proximity: The cell would simply migrate to the nearest column-to-be. This assumption is the basis of our “proto-column” algorithm.

Earlier attempts to simulate ocular dominance column patterns

Two previous simulations of ocular dominance column patterns have been published.^{22,23} In both these simulations the actual cortical ocular dominance column pattern (not the proto-column pattern) was mapped back to the retina. This resulted in the loss of half the visual field from each eye. That is, wherever a left-eye column occurs in the cortex, there was no right-eye input to map back to the retina, and vice versa for the right-eye columns.

However, when we close one eye, we do not lose that half of the visual field that corresponds to the cortical pattern of the closed eye’s columns. To take account of the true complexity of the situation, we must construct the proto-column boundaries of the left and right eyes as well as the actual column boundaries.




Figure 6. Figure 6a shows a wide-angle fish-eye view of a scene in the hall of our laboratory. A ladder is to the right; an eye chart is in the very center of the frame (almost invisible). The original version of this scene was digitized to an effective resolution of $16,000 \times 16,000$ pixels by a polar-coordinate mosaic technique. In 6b we see a blowup of the central region of this original frame. This is an eye chart, at a distance of twenty feet. In the original, line 8 of this chart could be easily read, indicating an effective acuity of 20/20, or about one minute of arc. In 6c this scene is blurred by a space-variant filter, which is modeled

after human visual acuity. In 6d we see the image of 6a modeled in terms of a complex-logarithmic model¹⁷ of human visual cortex. The eye chart occupies almost half of the surface of visual cortex, although it occupies a tiny fraction of the original scene. The ladder and the windows of the original are compressed to almost the same size as the centrally fixated letters of the eye chart. This illustrates the tremendous space-variant compression of human vision. Variations in linear size of about 100:1 (10^4 in solid angle) occur from the center to the periphery of the human visual system.

The proto-column algorithm: Generalized Voronoi regions and generalized Voronoi diagrams

The concept of Voronoi regions is familiar from both crystallography and computational geometry. "The Voronoi diagram of a finite set S of points in the plane

is a partition of the plane so that each region of the partition is the locus of points which are closer to one member of S than to any other member."²⁴ If we start with polygons, such as ocular dominance columns, instead of points, we can construct a generalization of the Voronoi diagram. The locus of points that are nearest to a given polygon is the generalized Voronoi region of the polygon. In our terminology, it is also the *proto-column*.




Figure 7. In (a) the image of the ocular dominance column pattern of macaque striate cortex is reprinted from LeVay et al.²⁵ The black-and-white areas correspond to cortical cells that receive input from the left eye and the right eye respectively. In (b) we see the construction of “proto-columns.” Each colored area represents the territory associated with a left-eye column in Figure 7a.

In the case of cortical systems, we perform the following construction: We take a given cortical area, which by assumption is a continuous 2D sheet. If n column systems are interlaced in that area we make n copies of the cortical sheet. In the i^{th} sheet, we retain the i^{th} column system and erase the others. The generalized Voronoi diagram corresponding to the i^{th} column system corresponds to the proto-map of that column parameter. The union of all the generalized Voronoi regions in that proto-map will smoothly cover the area of the proto-map, with no overlap, and every pixel in the proto-map will be assigned to one of the proto-columns.

An Algorithm to find Voronoi regions

Let us start with a definition of a column system in terms of labeled pixels. We assume the black regions represent left-eye afferents for cortical ocular dominance columns²⁵ (see Figure 7).

We label each column with a unique integer or color. We then apply a standard contour-follower to isolate the boundaries of the columns. Each white pixel is then assigned the color of the nearest boundary pixel. After we have examined all such intercolumn pixels, every pixel in the image has been assigned to one or another of the original columns and has been marked with the color of its column. Each set of marked pixels forms the proto-column of the correspondingly colored column.

Figure 7a, reprinted from LeVay et al.,²⁵ shows the ocular dominance column system of a part of monkey visual cortex, similar to the larger example shown in Figure 5. The proto-column system generated from this data is shown in Figure 7b. Each of the proto-columns corresponds to one of the actual (left-eye) ocular dominance columns of Figure 7a.

In Figure 8 we show the application of this proto-column map, together with the earlier demonstration of the cortical topographic map, to simulate a natural scene, at the level of striate cortex, contingent on the pattern of

ocular dominance columns. We are thus able to simulate realistically the stereo image projected onto one of the map surfaces of the brain when an animal scans a 3D scene with its two eyes.

Discussion

The proto-column algorithm performs a necessary extension of the concept of a “cortical map.” Because most cortical layers consist of two or more “copies” interlaced together in complex fashion, the simple concept of a “regular” or “continuous” map of $R^2 \rightarrow R^2$ is not sufficient.

A simple nearest-neighbor rule for multicolumn systems defines the proto-column regions associated with the observable columnar boundaries. We have developed an algorithm that can find these regions. Without knowledge of these proto-column regions, accurate simulations of multicolumn systems are impossible. Conversely, however, with good knowledge of a global, regular topographic map, such as the complex logarithm, and knowledge of proto-column boundaries, accurate simulations can be performed.

Level 3: The brain as map machine

We have reviewed, in a bottom-up fashion, a series of computer graphics applications which were developed to aid research into the nature of cortical computation. There are a number of challenging algorithmic, implementational, and experimental issues related to reconstructing cortical systems from serial sections. There are also challenging problems related to modeling cortical architectures, which are indicated by the kinds of data that emerge from the cortical reconstruction. But finally, the algorithmic implications of these map structures must be addressed. And this is the motivation for the entire process.

The very fact that intricate maps of the visual field exist




Figure 8. In (a) the entire figure from (b) is mapped through the proto-column construction into the left eye columns alone. Each proto-column was painted onto its associated column as in the brain painter described in the text, using the proto-column construction to determine the warp function. In (b) the simple cortical simulation of the eye chart from Figure 6 is shown. In (c) and (d), performing a similar analysis for the right eye and then “oring” together the two frames provides a true “binocular” simulation of the eye-chart scene as it is mapped to the surface of layer IV of cortex for each of the two eyes. We show this simulation for zero minutes of arc disparity, and for five minutes of arc disparity—(c) and (d) respectively.

in great profusion in the brain raises a computational question: What is the function, if any, of this phenomenon? Broadly speaking, there are two possible answers. First, the existence of maps and columns may be a developmental or anatomical epiphenomenon. It may simply be easier or more efficient for neurons to arrange themselves in interesting spatial patterns, but these patterns themselves may have little significance. This might be termed the “topological” or “rubber-sheet connectionist” approach. One could argue that as long as the actual connections between neurons remain the same, their location is irrelevant, as is the case for the conventional von Neumann architecture.

At the other extreme is an optical computer. Consider one performing Fourier transforms via a lens. In this case, the computation and the physical structure of the machine are identical: The form is the function.

Which of these limiting cases represents the brain? A prudent choice, given the current state of almost complete ignorance about the bases of neural computation, would be to take the middle road, and assume “a little of both.” Curiously, however, the dominant contemporary paradigm of neural computation is almost purely connectionist: “Neural networks,” which are essentially independent of the space-time structure of the nervous system, represent the vast majority of attempts at understanding neural computation.

But there is very little data available about nontrivial neural networks. One can observe networks in the simplest invertebrates, but the behavioral repertoire of these creatures is limited. Large-scale network observations

(more than 2 neurons at one time) in cats or monkeys are essentially nonexistent. But, as we have indicated, there is a vast amount of data (and technology) available to observe and characterize architecture at the map and column level. In fact, if one considers that such basic phenomena as lateral inhibition, spatial filtering, and receptive field structure are intrinsically map-level (spatial) phenomena, then it would be fair to claim that the majority of neural data about the visual system is in the domain of spatial mapping.

What is a map, and how many might there be in the visual cortex of monkeys? We take as a general definition for map any spatial representation of stimulus variables. Classical receptotopic maps, column systems, and more general “cognitive” maps, if they exist, fall into this category. Often, a map is restrictively defined as the representation of a sensory surface (retina, basilar membrane, skin receptors) onto a central neural surface. But we generalize this to include the spatial representation of any neural surface onto another. The algorithmic construction of proto-columns in this article gives us a well-defined means of discussing multiple columnar maps interlaced into single maps. This provides a notion of “map” that is more general than the conventional one, but is still precisely defined.

The number of maps in the visual system is thus related to the number of physiologically distinct layers in the visual cortex. There are more than twenty distinct visual cortical areas currently known, and the number has been growing steadily over the past few years. There are five classical neural layers in neocortex, but V1 has

at least 10 distinct sublayers. Thus, there are perhaps 100 to 200 neural layers in the visual system, and there may be comparable numbers in other areas of primate neocortex. We have not counted any of the thalamic or midbrain (or cerebellar or brain stem) structures that have a maplike architecture. A number of maps in the range of 1,000 is easy to imagine.

How can we view maps as computational entities? One direct way is to consider the familiar computational metaphor of data structures and algorithms. Good data structures allow simpler algorithms to be invoked: As the task becomes more complex, we have often observed that the quality of the data structure becomes increasingly important. If one interprets the maplike architecture of the brain as a form of data structure or representation, then one could view the "network" component as algorithm. This forms the basis for a "map machine" point of view.

As an example of this approach, the ocular dominance column system of V1 provides a very simple algorithm for stereo vision: If one passes a stereo image (as shown in Figure 8) through a simple space-domain filter (a cepstral filter^{26,27}), then a very fast stereo segmentation may be performed.²⁸ In other work we have investigated the size-, rotation-, and projective-invariant aspects of the complex-logarithmic cortical map,¹⁷ and we have developed algorithms to choose a "scan path" automatically for a space-variant system (such as the human), and to blend multiple space-variant scans into a single "stable" visual percept.²⁹ We have also begun to develop a more general "calculus of maps" which allows cognitive classifications on networks of maps, so that we may combine some of the combinatorial aspects of network approaches with a space-domain or map approach to cognition.³⁰

In summary, computer graphics has at least three areas of major contribution to understanding the nature of vision. Simply analyzing the data of animal visual systems makes sophisticated demands on our ability to use computer graphics and computational geometry. But the more exciting possibility is that the brain itself is a kind of elaborate parallel frame buffer which warps, transforms, and interlaces multiple map systems in large number, to effect the real-time performance of the only working visual machine that currently exists. ■

Acknowledgments

This research was supported by AFOSR 85-0341, System Development Foundation, and the Nathan S. Kline Psychiatric Research Center.

References

1. D.M. Daniel and D. Whitteridge, "The Representation of the Visual Field on the Cerebral Cortex," *J. Physiology*, Vol. 159, 1961, pp. 203-221.
2. M.T.T. Wong-Riley, "Changes in the Visual System of Monocularly Sutured or Enucleated Cats Demonstrable with Cytochrome Oxidase Histochemistry," *Brain Research*, Vol. 171, 1979, pp. 11-28.
3. L. Sokoloff et al., "The 14C-Deoxyglucose Method for the Measurement of Local Cerebral Glucose Utilization. Theory, Procedure and Normal Values in the Conscious and Anesthetized Albino Rat," *J. Neurochemistry*, Vol. 28, 1977, pp. 897-916.
4. W. Kaplow and E.L. Schwartz, "Preparing and Repairing Serial Sections of Brain for 3D Reconstruction: The Brain Spackler," *Comp. Neurological Tech. Report*, Vol. CNS-TR-17-86, 1986, NYU Medical Center/Computational Neuroscience Laboratories, N.Y.
5. C. Frederick and E.L. Schwartz, "Exploiting Scan Plane Order To Construct a Fast Brain Peeler," *Comp. Neurological Science Tech. Reports*, Vol. CNS-TR-5-87, 1987.
6. G. Frieder, D. Gordon, and R.Q. Reynolds, "Back-to-Front Display of Voxel-Based Objects," *CG&A*, Jan. 1985, pp. 52-60.
7. E. Keppel, "Approximating Complex Surfaces by Triangulation of Contour Lines," *IBM J. Research Development*, Jan. 1975, IBM, Yorktown Heights, N.Y., pp. 2-11.
8. H. Fuchs, S.M. Kedem, and S.P. Uselton, "Optimal Surface Reconstruction from Planar Contours," *CACM*, Vol. 20, 1977, pp. 693-702.
9. H. Christiansen and M. Stephenson, *MOVIE.BYU*, University Press, Brigham Young Univ., Provo, Utah, 1983.
10. O.D. Faugeras, M. Hebert, D. Mussi, and J.D. Boissonnat, "Polyhedral Approximation of 3-D Objects Without Holes," *Computer Vision, Graphics, and Image Processing*, Vol. 25, 1984, pp. 169-183.
11. J.D. Boissonnat, "Geometric Structures for Three-Dimensional Shape Representation," *ACM Trans. Graphics*, Jan. 1987, pp. 266-286.
12. K. Anjyo, T. Ochi, Y. Usami, and Y. Kawashima, "A Practical Method of Constructing Surfaces in Three-Dimensional Digitized Space," *The Visual Computer*, Feb. 1987, pp. 4-12.
13. A. Shaw, W. Light, and E.L. Schwartz, "Finding, Labeling, and Parsing 2.5D Contours of a Surface in 3D: A Semantics and Syntax for Automatic Triangulation," *Comp. Neurological Tech. Report*, Vol. CNS-TR-6-86, NYU Med.Ctr./Computation Neuroscience Laboratories, N.Y., 1986.
14. E.L. Schwartz and B. Merker, "Computer-Aided Neuroanatomy: Differential Geometry of Cortical Surfaces and an Optimal Flattening Algorithm," *CG&A*, Mar. 1986, pp. 36-44.
15. J.S.B. Mitchell, D.M. Mount, and C.H. Papadimitriou, "The Discrete Geodesic Problem," *SIAM J. Computation*, Aug. 1987, pp. 647-668.
16. E. Wolfson and E.L. Schwartz, "Computing Minimal Distances in Arbitrary Two-Dimensional Polyhedral Surfaces," *CIMS Tech. Report*, Vol. 255, NYU Courant Inst., N.Y., 1987.
17. E.L. Schwartz, "Computational Anatomy and Functional Architecture of Striate Cortex: A Spatial-Mapping Approach to Perceptual Coding," *Vision Research*, Vol. 20, 1980, pp. 645-670.
18. E.L. Schwartz, "On the Mathematical Structure of the Retinotopic Mapping of Primate Striate Cortex," *Science*, Vol. 227, 1985, p. 1066.
19. B.M. Dow, A.Z. Snyder, R.G. Vautin, and R. Bauer, "Magnification Factor and Receptive Field Size in Foveal Striate Cortex of Monkey," *Experimental Brain Research*, Vol. 44, 1981, pp. 213-228.
20. D.C. Van Essen, W.T. Newsome, and J.H.R. Maunsell, "The Visual Representation in Striate Cortex of the Macaque Monkey: Asymmetries, Anisotropies, and Individual Variability," *Vision Research*, Vol. 24, 1984, pp. 429-448.
21. P. Rakic, "Principles of Neuronal Migration," *Handbook of Sensory Physiology* W.M. Cowan, ed., American Physiological Society, Bethesda, Md., 1987.

22. E.L. Schwartz, "The Development of Specific Visual Projections in the Monkey and the Goldfish: Outline of a Geometric Theory of Receptotopic Structure," *J. Theoretical Biology*, Vol. 69, 1977, pp. 655-685.
23. D.H. Hubel and D.C. Freeman, "Projection into the Visual Field of Ocular-Dominance Columns in Macaque Monkey," *Brain Research*, Vol. 122, 1977, pp. 336-343.
24. F.P. Preparata and M.I. Shamos, *Comp. Geometry: An Introduction*, Springer-Verlag, N.Y., 1985, 234.
25. S. LeVay, D.H. Hubel, and T.N. Wiesel, "The Pattern of Ocular-Dominance Columns in Macaque Visual Cortex Revealed by a Reduced Silver Stain," *J. Comp. Neurology*, Vol. 159, 1975, pp. 558-576.
26. B.P. Bogert, W.J.R. Healy, and J.W. Tukey, "The Frequency Analysis of Time Series for Echoes: Cepstrum Pseudo-Autocovariance, Cross-Cepstrum and Saphe Cracking," *Proc. Symp. Time Series Analysis*, John Wiley and Sons, N.Y., 1963, pp. 209-243.
27. A.V. Oppenheim and R.W. Shafer, *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, N.J., 1975.
28. Y. Yeshurun and E.L. Schwartz, "Cepstral Filtering on a Columnar Image Architecture: A Fast Algorithm for Binocular Stereo Segmentation," To be published in *IEEE Trans. Pattern Analysis and Machine Intelligence*.
29. Y. Yeshurun and E.L. Schwartz, "Shape Description with a Space-Variant Sensor: Algorithms for Scan-Path, Fusion and Convergence over Multiple Scans," *IEEE Trans. Pattern Analysis and Machine Intelligence*, to appear 1988.
30. A. Rojer and E.L. Schwartz, "A New Pattern Classifier Motivated by Brain Architecture: Towards a Calculus of Maps," *Courant Inst. of Math. Sciences, Tech. Rpt.*, Vol. 203, 1988, NYU Courant Inst., N.Y.



Eric L. Schwartz is associate professor at the Computational Neuroscience Laboratories, Department of Psychiatry, of New York University Medical Center. He is a research scientist at Nathan S. Kline Research Institute and an adjunct associate professor of computer science at the Courant Institute of Mathematical Sciences of New York University. His current research interests are the anatomy and physiology of the nervous system. A major theme of this research is the application of mathematical and computer models to describe the various architectures that have been experimentally observed in primate visual systems, and to infer the possible computational significance of these architectures for both machine and biological visual computation.

Schwartz received the AB degree in physics and chemistry in 1967 and MA, MPhil, and PhD degrees in high energy physics from Columbia University in 1973.



Bjorn Merker is research assistant professor at the Brain Research Laboratory of the NYU Medical Center. Prior to his current work on the primate visual system, he worked on behavioral consequences of brain lesions in hamsters, neuroanatomical staining and lesioning methods, and the control of eye movements in the cat. His interests include the functional significance of spatial and geometric aspects of brain structures in general and the computational significance of topographic mappings in particular.

Merker received an AB in psychology from Queens College of the City University of New York in 1975 and a PhD in psychology and brain science from MIT in 1980. He is a Swedish citizen.



Estarose Wolfson is on the research staff of the Computational Neuroscience Laboratories located at New York University Medical Center and the Courant Institute of Mathematical Sciences. Formerly, she worked at CIMA, Université des Beaux Arts de Architecture VI, Paris, doing research in computer graphics and architecture, and at the College de France, Paris, and Brookhaven National Laboratories, New York, in high-energy and solid-state physics respectively. Her research interests involve the combining of art, mathematics, and computer science by using computational geometry, computer graphics, and image processing.

Wolfson received a BS in mathematics from Queens College of the City University of New York and a masters degree in mathematical sciences from the Université de Paris VII.



Alan Shaw is an assistant research scientist at the Computational Neuroscience Labs of New York University Medical Center. Prior to his current work in graphics and image processing, he worked in the motion picture industry and as a translator.

Shaw received the BA degree in mathematics from Washington Square College of New York University, MA degree in Chinese studies from St. John's University, and the MS in computer science from the Courant Institute, New York University.

The authors can be reached at the Computational Neuroscience Laboratories, Dept. of Psychiatry, New York University School of Medicine, 550 First Ave., New York, NY 10016.

University of Berne, Switzerland

Institut für Informatik und angewandte Mathematik
(Dept. of Computer Science and Applied Mathematics)

The department invites applications for the
faculty position of a

PROFESSOR OF COMPUTER SCIENCE.

The candidate's basic duties and responsibilities include—teaching courses for graduate and undergraduate students in the fields of computer architecture and operating systems—research in one or more areas from the following list; computer architecture, operating systems, software engineering, man machine interface, computer networks.

Qualified applications should be sent to:

Erziehungsdirektion des Kantons Bern
Abt. Hochschulwesen (1600.27/86)
Sulgeneckstrasse 70
CH-3005 Bern, Switzerland

and must be received not later than 15.9.1988.

Further information is available from:

Institut für Informatik
und angewandte Mathematik
der Universität Bern
Länggassstr. 51
CH-3012 Bern, Switzerland
Tel.: 031/65 86 81