



Visualizing Cortical Waves and Timing from Data

Kay A. Robbins¹
University of Texas at San Antonio
Cajal Neuroscience Research Center

Mark Robinson²
University of Texas at San Antonio

David M. Senseman³
University of Texas at San Antonio
Cajal Neuroscience Research Center

ABSTRACT

Waves are a fundamental mechanism for conveying information in many physical problems. Direct visualization techniques are often used to display wave fronts. However, the information derived from such visualizations may not be as central to an investigation as an understanding of how the location, structure and time course of the wave change as key experimental parameters are varied. In experimental data, these questions are confounded by noise and incomplete data. Recognition of waves in networks of neurons is additionally complicated by the presence of long-range physical connections and recurrent excitation. This paper applies visual techniques to analyze the structural details of waves in response data from the turtle visual cortex. We emphasize low-cost visualizations that allow comparisons across neural data sets and variables to reconstruct the choreography for a complex response.

CR Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Picture/Image Generation - Viewing Algorithms; I.3.6 [Computer Graphics]: Methodology and Techniques - Interaction Techniques.

Additional Keywords: waves, neural networks, PCA, KL decomposition, wave subspaces, flow visualization.

1 INTRODUCTION

Whitham's classic text on waves [25] defines a wave as "any recognizable signal that is transferred from one part of the medium to another with a recognizable velocity of propagation." The discussion continues "the signal may be any feature of the disturbance, such as a maximum or an abrupt change in some quantity, provided that it can be clearly recognized and its location at any time can be determined."

Much of the work on wave visualization has actually focused on the visualization of the supporting flow [6]. Visualization techniques such as arrow plots, streamlines [22], level sets [14] and advected particles [10] give a picture of the motion produced by the associated scalar or vector field. Another approach inserts spot noise textures in random places or along pixel streamlines [2]. More recently, van Wijk [24] introduced image-based flow visualization to display the distortion of an image embedded in a flow.

Other studies address the realistic rendering of waves using refraction [21] or other physical properties of the media [4] to reproduce what we would see if witnessing a wave. The study of structures within complicated dynamic fields has focused on visualization of shock waves and vortices [12, 13] because of their obvious practical importance.

While direct visualization of wave surfaces and distortion is useful, many scientific studies need to explore more detailed questions such as how structure evolves as a function of time, how different variables interact and contribute during the evolution of the response, and how parameters affect the results.

Studies of networks of neurons in the cortex have additional complexities. Assumptions of continuity and neighborhood are violated. Cortical networks are characterized by asymmetry, multiple time scales, long-range as well as neighborhood connectivity, intrinsic variability and large scale. Modelers can not hope to replicate all of the details of such systems. The natural variation between individual brains even in the same species means that comparison with experimental data for validation must be done qualitatively.

Neurovisualization work appears in several venues. Modeling environments such as Neuron [7] and Genesis [1] provide facilities for visualizing responses from individual neurons. Catacomb2 [3] provides a graphical design environment for building neural models. The environment provides basic graphical widgets for plotting curves and panning and zooming, but its emphasis is on model development rather than visualization. XPP [5] is a package for analyzing and animating behavior of small-scale dynamical systems. This package is widely used by modelers working with models on the order of tens rather than hundreds or thousands of neurons. Visu_nbc, another tool for visualizing neural activity, provides matrix visualizations for one-dimensional systems [26].

Neural systems have also been visualized using general-purpose visualization environments such as SciRun [11] or Volvis [20]. The emphasis in this work has been to reproduce the complex neuronal morphology rather than to display dynamical relationships between variables.

This paper applies visual techniques to a model of the turtle cortex. We show how a combination of simple, fast visualization techniques can be used to understand complex neuronal interactions, to qualitatively validate the model, and to answer general questions about structure. Section 2 covers some basic neuroscience background, introduces the application, and poses some questions to be answered. Section 3 analyzes wave structure by applying visual techniques such as space-time representations with different spatial sorting and contour overlays, parallel coordinate representations of response order, and projections into low-dimensional subspaces to extract structural changes. These visualizations should be compared with direct 3D representations using transparent overlays provided in the accompanying video. Section 4 discusses our experiences in using visual techniques for understanding network behavior.

¹e-mail: krobbins@cs.utsa.edu

²e-mail: mrobinso@cs.utsa.edu

³e-mail: senseman@utsa.edu

2 NETWORKS OF NEURONS IN THE TURTLE

The neuron is the fundamental building block of the brain. Neurons generally have three distinct parts --- the soma or cell body, the axon, and the dendrites. In simplified terms, the dendrites gather information from other neurons, the soma processes the information, and the axon delivers the output to other cells,

More specifically, the soma has a number of on/off switches for conductance in the form of voltage-gated ion channels (potassium, sodium or calcium). When the cell membrane potential reaches a critical threshold, the channels open and ions flow, generating an electrical event called an *action potential*. Once triggered, the action potential rapidly travels down the axon until it reaches the end or synapse, causing the release of a chemical substance called a neurotransmitter onto the dendrites or soma of the next neuron in the network.

In the turtle cortex, for example, each neuron is typically connected to hundreds of other neurons. As a neurotransmitter contacts the next cell's dendrite, it binds to various types of protein receptor molecules. The two main excitatory receptors are AMPA and NMDA; the two main inhibitory receptors are GABA_A and GABA_B. Cortical neurons are classified as either *excitatory* or *inhibitory* depending on whether they release a neurotransmitter that binds to excitatory or to inhibitory receptors, respectively. Excitatory neurons act to increase the activity of neurons to which they are connected, while inhibitory neurons decrease the activity.

From a modeling perspective, there is a tradeoff between accuracy and computational efficiency. One of the simplest neuronal models, integrate-and-fire, consists only of a cell body that integrates its inputs and generates an action potential if the sum of the inputs is greater than a threshold. Dendrites and axons are sometimes modeled analytically by transmission lines or by a series of interconnected compartments (often cylinders), each of which is assumed to consist of tissue at the same potential. Compartmental models for an individual neuron can range from a few compartments to thousands [1].

The type of neurons, their modeled complexity, the time scales of the different channels and the distribution and connectivity all influence behavior. How much complexity is needed to reproduce essential function? To answer this question, one must be able to recognize essential function and to attribute factors to different aspects of behavior.

This paper uses visual analysis techniques to assess the structure of response in a model of turtle visual cortex developed by Nenadic et al. [15]. The NGU model, as it will be referred to in this paper, is a physiologically based multi-compartment network model. The overall goal of the research is to validate the model by comparison with experimental observations and then to explore the details of model behavior and information processing mechanisms.

The turtle visual cortex is considered to be a three layer cortex. Layer 2 contains the cell bodies of the excitatory neurons, which have two main types: lateral and medial pyramidal cells. Pyramidal cell dendrites extend upward to layer 1 and down to layer 3. Layer 1 contains inhibitory cells called stellate cells, and layer 3 has inhibitory cells called horizontal cells. The number of inhibitory neurons is much less than the number of excitatory neurons.

Fig. 1 shows a drawing of a cortical slice with prototypical lateral and medial cells superimposed. The black dot is the neuron's cell body (soma). The numerous fine processes extending from the soma are the cell's dendrites and axon. Lateral cells are roughly symmetric, while medial cells have a smaller tree connecting to layer 3.

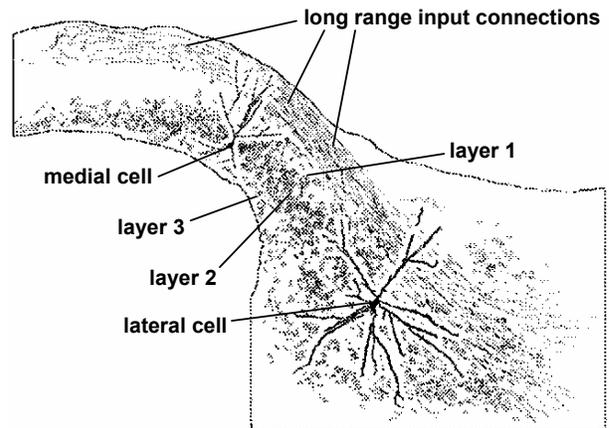


Figure 1: Section of turtle visual cortex (Ulinski, personal communication).

Although the physical volume of each cell is relatively small, its axonal arbor extends for large distances. In addition to the extensive dendrite structure, the cortex has geniculate afferents (the input lines from other cortical areas) that extend the entire length of the cortex. These lines can be seen in Fig. 1 as the band of strands at the top of the cortex.

The NGU model incorporates four types of neurons: lateral, medial, stellate and horizontal. Each neuron is represented by a compartmental model neuron whose parameters are based on extensive intracellular microelectrode recordings [23]. The cell bodies of these neurons have voltage-gated channels for potassium, sodium and calcium.

The dendrites for the neurons contain excitatory (AMPA and NMDA) and inhibitory (GABA_A and GABA_B) receptors. The neurotransmitters from the excitatory cells activate AMPA and NMDA, while those from the inhibitory neurons activate GABA_A and GABA_B. The model has heterogeneous distributions of the different types of neurons based on experimental measurements.

Fig. 2 shows the distribution of neurons in the NGU model. The four types of neurons are represented by different shapes. The excitatory neurons are displayed using gray triangles, with lateral cells being represented by upward triangles and medial cells being represented by downward triangles. The inhibitory horizontal cells are depicted as blue circles, and the inhibitory stellate cells are depicted as red squares.

While the neurons appear to be arranged in a two-dimensional square, the location of the connections to the dendrites from different cells varies by cell type, giving the model structural three-dimensionality. The dark bar at the bottom of Fig. 2 represents 200 lateral geniculate cells (LGN) that provide input to the cortex. Neurons that appear outlined in black are connected directly to the LGN input by long-range connections. The physical dimensions of the model are chosen to represent an area similar to that of Fig. 1, but the neuron density of the model is significantly less than that of the real brain.

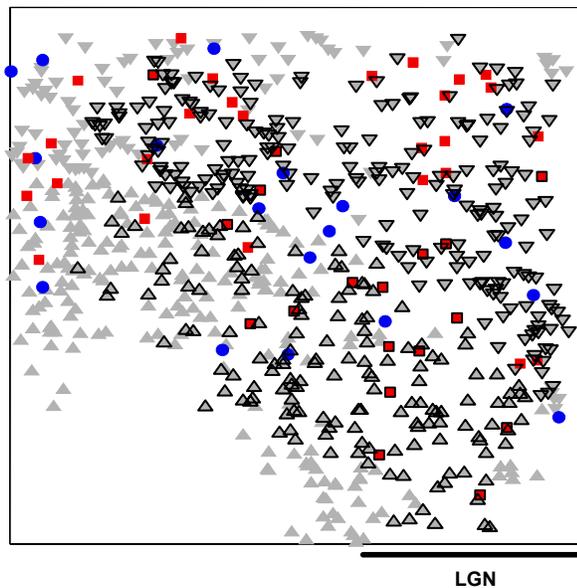


Figure 2: The NGU model of the turtle visual cortex.

3 VISUALIZATIONS

This paper applies visual techniques for model validation, response structure discovery, and organization of model variation with parameters.

Validation of individual neuron models based on structural measurements and voltage-clamp experiments is well-documented [1]. However, very little work has been done on either qualitative or quantitative validation of network behavior. Three types of experimental measurements are available for the turtle visual cortex: electrode measurements of single neurons in slice [23], voltage sensitive dye imaging of whole brain-eye preparations *in vitro* [19], and voltage sensitive dye imaging *in vivo* [16]. The latter two types of experiments provide spatially averaged information on a 24×24 grid at a time resolution of 1 ms. Because the imaging data has low spatial resolution and the spatial geometry varies significantly for measurements on different animals as well as between the types of experiments, it is only possible to qualitatively compare population characteristics.

Membrane potential or voltage is the only observable response in the imaging experiments. Because state information is completely available in the model, one can use the model to infer detailed answers to structural questions, once the model has been validated. For example, one might ask how closely soma voltage tracks current. When does inhibition due to $GABA_A$ become significant? Which excitatory receptors are important in the generation of the initial response wave? What is the role of excitatory receptors that don't contribute to the initial response? What part of the response is due to the input (LGN excitation) as opposed to recurrent excitation by excitatory neurons?

More integrative questions can be also asked. How does the response change as the stimulus is varied? Can the change be parameterized? Is this change related to some internal cortical mapping? Other classes of questions address the influence on cortical response of structural change in the model or the experiment. What happens to the response, for example, if the NMDA receptors are weakened or removed entirely? What if a chemical that blocks a particular receptor type, e.g., $GABA_A$, is applied to the cortex?

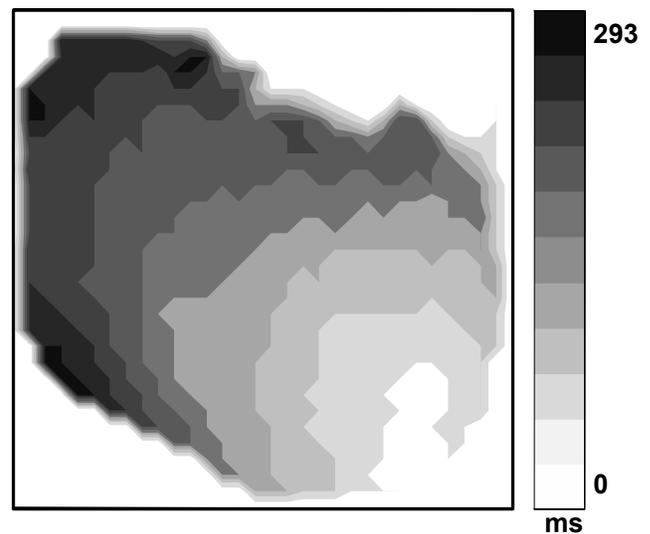


Figure 3: A half-height latency map of response of excitatory (lateral and medial) neurons to a diffuse light flash in the NGU model.

3.1 Latency maps

We begin by applying latency maps to understand the structure of the initial response to a stimulus. Latency maps are computed by finding the time at which each neuron first reaches a specified fraction of its maximum height. The times are then interpolated on a rectangular grid to produce a contour map [19].

Fig. 3 shows the half-height latency map for a response to a diffuse light flash in the NGU model. The response takes approximately 293 ms to sweep across the cortex. The direction of maximal speed is aligned at approximately a 135° line. The roughly equally spaced contours imply that the response sweeps across the cortex with a relatively uniform speed. The modeled brain size is 1.6×1.6 mm giving a velocity of 7.7 mm/sec, which is in the range of 4 to 10 mm/sec observed experimentally. The direction, speed and uniformity of the response wave are consistent with experimental measurements obtained from voltage sensitive dye imaging.

Some aspects of the model's latency map disagree with experimental observations. The model has no visual pathway, so the wave begins almost immediately after the stimulus is applied. In the real cortex, information about the stimulus is generated in the eye, travels the visual pathway, and arrives at the cortex after considerable delay. Also, the visual cortex of real animals is embedded in a much larger cortical structure.

Latency maps are effective in this setting because the initial response is a rapidly rising front of activation that sweeps across the cortex. Generally the time to half-height is used for experimental data because it avoids problems associated with baseline noise and oscillations near the peak response.

The latency visualization reveals the structure of the initial response and is in qualitative agreement with experimental data, but the behavior is filtered and interpolated. We lose the detailed behavior of individual neurons and structure of overall response.

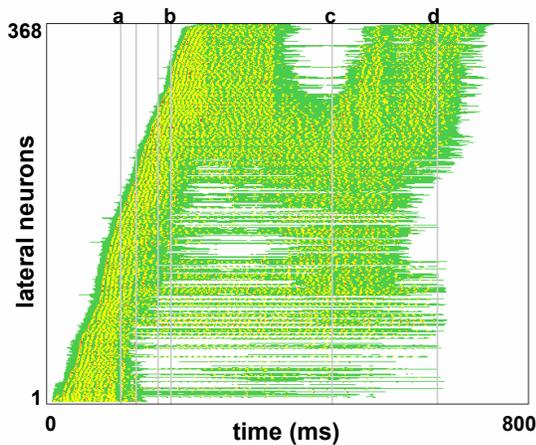


Figure 4: Space-time diagram of soma voltage of the lateral pyramidal cells in response to a diffuse flash. Neurons are ordered by time to first reach half of maximum height.

3.2 Space-time diagrams

Another classical method for displaying waves in one spatial dimension is the space-time diagram. Space-time diagrams display time along the horizontal axis and spatial position or spatial element number along the vertical axis. The value of the response at each time and spatial position is color-coded to represent magnitude.

Sampled data sets with multiple spatial dimensions can be represented by space-time diagrams if an appropriate total ordering of the spatial points can be found. Neurons can be ordered by the latency of the initial response, using actual spatial coordinates in succession to break ties.

Fig. 4 displays a space-time diagram for the lateral cells in the NGU model for a diffuse flash stimulus. The sorting is based on the time each neuron takes to first reach half of its maximum height. The response values are offset by their resting potential and are always positive. Response is linearly mapped to five colors (white, green, yellow, orange, red) in order of increasing response. Thus, response levels less than 20% of the maximum are set to the background color of white, allowing the front to be more clearly distinguished.

The speckled appearance of Fig. 4 is due to the spiking behavior of the individual neurons. A front is visible, propagating from neuron 1 at time 0 to neuron 368 at about 220 ms. The vertical gray lines labeled **a**, **b**, **c** and **d** mark times 124 ms, 208 ms, 475 ms and 650 ms. The two unlabeled gray lines correspond to 150 ms (end of the diffuse flash stimulus) and 186 ms. These times are referenced in later visualizations.

The U-shaped region at the top of the figure shows that the wave reflects and then another front propagates after the line marked **c**. The response dies out somewhat abruptly throughout the cortex. Notice the streaks of color ahead of the front. Because the coordinate sorting is based on the half-height latency, the early activity of these neurons is ignored in the sorting.

Although the neuron sort-order given on the vertical axis does not directly represent spatial position, the linear structure of the front gives an indirect representation of velocity. A steeper slope represents faster propagation.

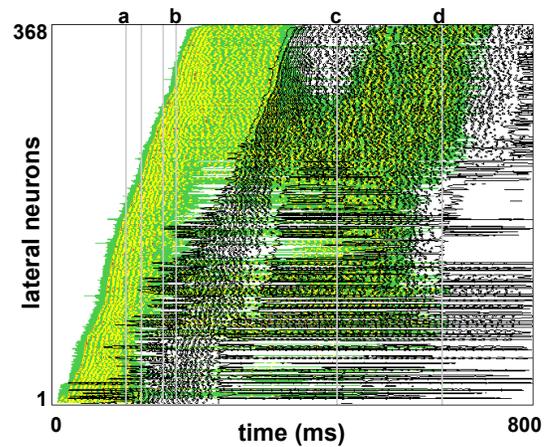


Figure 5: Space-time diagram of Fig. 4 overlaid with a contour map (in black) of the response to a spot stimulus.

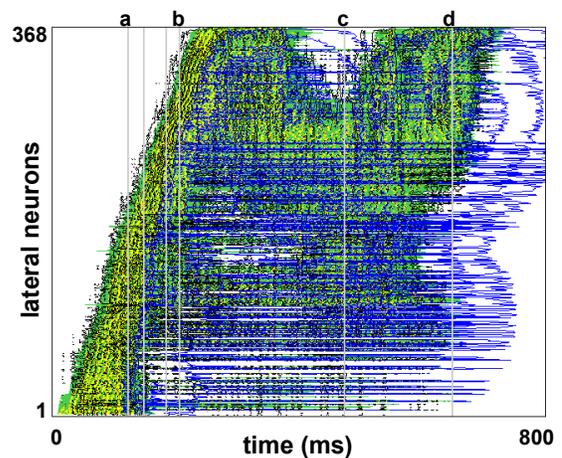


Figure 6: Space-time diagram of Fig. 4 overlaid with contour maps of AMPA (in black) and NMDA (in blue) currents.

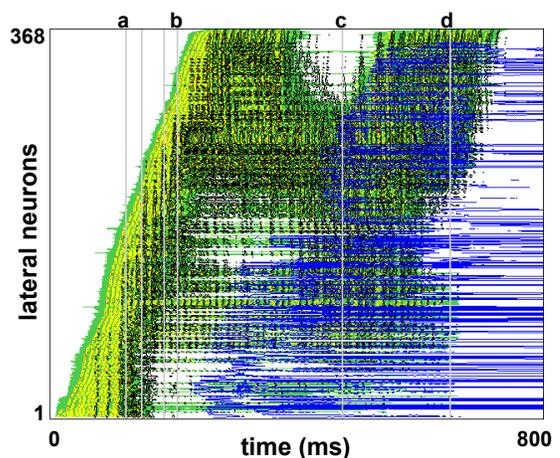


Figure 7: Space-time diagram of Fig. 4 overlaid with contour maps of GABA_A (in black) and GABA_B (in blue) currents.

One can overlay space time diagrams to compare the evolution of two data sets. The diffuse flash stimulus activates all 201 neurons of the LGN input for 150 ms. Fig. 5 compares this response to that when only LGN neurons 1 through 20 are stimulated. The neurons for the spot stimulus response are displayed in the same order as the diffuse stimulus. The overlaid space-diagram is represented by unfilled black contours.

Several features are immediately apparent. The low numbered neurons (e.g., 1 through 20) are activated at the same time as those for the light flash response. However, the overall response develops much more slowly and there are considerable differences in the order of activation among neurons in the mid-range. An organized front sweeps across the high numbered neurons (which in fact correspond to neurons that are spatially far away from the input lines). This front is temporally delayed in the spot response by roughly 150 ms and dies more slowly than the front in the diffuse response. The spot response does not show reflection or secondary front propagation.

Overlaid space-time diagrams can also be used to evaluate the different factors contributing to the overall response. Fig. 6 shows the space-time diagram of Fig. 4 overlaid with the space-time diagrams of the AMPA and NMDA currents. These currents are caused by input from other excitatory neurons (lateral and medial). The space-time diagram for the voltage is treated as a surface and displayed with filled contours, while the space-time diagrams for the AMPA and NMDA currents are displayed with unfilled black and blue contours respectively.

Many observations can be made about Fig. 6. Except for a very small region near the origin, the contours of the AMPA current closely match those of the voltage. Since the contributions from the LGN are not included in the plot, the match suggests that recurrent excitation from other excitatory neurons, rather than outside input is responsible for most of the response. The AMPA contours follow the voltage through the entire time course, reproducing the reflection and secondary front.

The NMDA contours shown in blue in Fig. 6 reveal that the effects of the slower NMDA receptors first appear close to the time marked **a** at 124 ms. The low-numbered neurons are nearly simultaneously activated. A second simultaneous activation of NMDA occurs near the next gray line at 150 ms, when the stimulus is shut off. Here an additional middle group of neurons is simultaneously activated by NMDA. NMDA activation persists long after AMPA has died. The NMDA contours are much smoother than the AMPA, indicating less of a spiked response. The NMDA contours appear to be more circular.

Fig. 7 displays a space-time diagram that reveals the interplay between excitation and inhibition for the response to a diffuse light flash. The space-time diagram for the voltage for the lateral neurons is overlaid by contours (black) representing currents generated from GABA_A receptors and by contours (blue) representing currents generated from GABA_B receptors. Both GABA_A and GABA_B receptors are activated by neurotransmitters from inhibitory neurons, neurons whose role is to diminish response.

Fig. 7 shows that GABA_A currents do not become significant until shortly before the line marked **a**, which appears at 124 ms. Notice that there is considerable synchronization in the GABA_A response as evidenced by the vertical lines of black dots corresponding to synchronized spikes. Once GABA_A has been activated, its time course closely follows that of the voltage, with both disappearing at the same time. In contrast, GABA_B is much slower to initiate. Once activated, the front of GABA_B follows roughly the same order of neuron activation as that of the voltage. However, the GABA_B front propagates more slowly as indicated by the smaller slope of the blue front in Fig. 7. GABA_B activity continues long after the 800 ms shown in the diagram.

The timing of the two types of inhibition clearly shows that only GABA_A is involved in the dynamics of the initial response and that inhibition does not become important until quite far into the development of the response. The latter observation is in conflict with detailed experimental measurements that show the presence of inhibition from the beginning of the response. This observation caused the modelers to add another type of inhibitory neuron to the model, the subpial cell, which contributes to very early inhibition. A related observation is that the initiation of NMDA currents in the pyramidal cells is closely timed with GABA_A initiation. It appears that NMDA provides sufficient driving for the pyramidal cells to activate the inhibitory cells.

3.3 Latency ordering diagrams

We now return to the issue of early low-level activity which is observed ahead of the front in Fig. 4. Space-time diagrams with neuron ordering determined by half-height latency capture the primary activity of the front, and one could easily ignore the early active neurons as “noise”. However, the NGU model has no external noise.

After careful viewing many data sets as movies using both a two-dimensional representation available in Davis [Robbins, et al., 2004] and as three-dimensional evolving surfaces (see accompanying video), we realized that the early-active neurons were distributed spatially throughout the cortex and reflected the delivery of the input from the LGN over the long-range connections in the cortex. Thus, these early-active neurons were, in fact, biologically significant.

One approach to understanding this early activity is to study how the selection of the latency threshold affects the response ordering of the neurons. The space-time diagram of Fig. 4 shows that when the latency threshold is 0.5 (half-height), the ordering closely follows the front. We expect that lower thresholds will give orderings that are more strongly influenced by the delivery of LGN input.

The order of the response can be visualized by treating the ordering of the neurons in the front as a permutation of the integers 1 to *n* and using parallel coordinate visualization [9] to compare the orderings. Fig. 8 shows a parallel coordinate visualization of the ordering of each neuron at different latency thresholds. The vertical axis threshold of 0.5 corresponds to the latency map of Fig. 3.

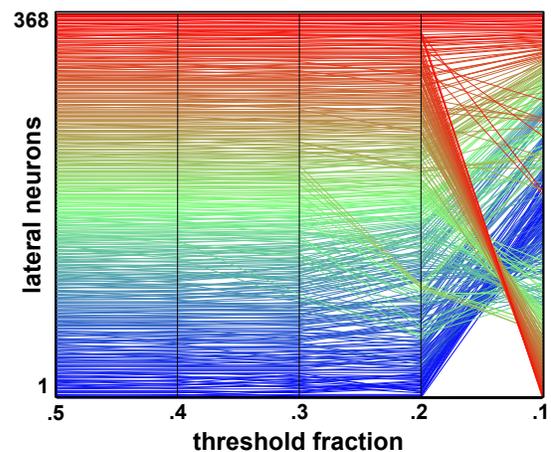


Figure 8: Ordering of response latency time to reach the specified fraction of max height for a diffuse stimulus.

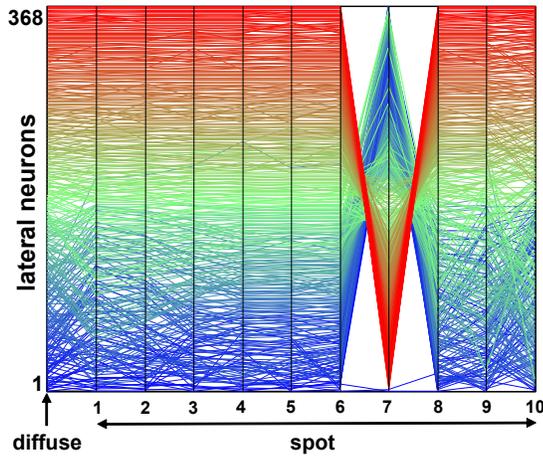


Figure 9: Comparison of response order for diffuse stimulus with 10 non-overlapping spot stimuli using a latency threshold of 0.5.

In the parallel coordinate visualization of Fig. 8, neurons in the base case (threshold of 0.5) are assigned values 1 through 368. Each neuron is connected to the position on the next line corresponding to its order as determined by the latency at the next threshold. The lines representing the neurons are color-coded using the rainbow map, with the neurons labeled by low numbers at a threshold value of 0.5 displayed in blue.

The horizontal portions of the lines signify that the ordering is the same, meaning that the order in which the response first reaches that fraction of the maximum response is the same as in the previous fraction. The order of response changes little for thresholds 0.5, 0.4 and 0.3. When the threshold fraction falls below 0.2, a dramatic reordering of the neurons occurs. Some neurons that are very distant from the LGN (color-coded in red) receive some of the earliest responses.

One can also use parallel coordinate visualizations to compare response order when different stimuli are applied. Fig. 9 compares the response order for a diffuse stimulus (all LGN neurons are activated for 150 ms) with spot stimuli at different positions. Spot k in Fig. 9 corresponds to activation of 20 neurons starting at position $20*(k-1) + 1$ in the LGN. The mapping uses a latency threshold of 0.5, so the diagram reflects differences in the propagation of the front.

The most apparent feature of Fig. 9 is the anomalous ordering in the response to spot 7. Further investigation shows that this part of the LGN corresponds to a problem “blind spot” in the model and results in no front propagation at all.

The other spot data sets show some variation in order in the lower numbered neurons but little reordering in the neurons far from the LGN. This observation is consistent with the space-time diagram of Fig. 5 which compares the response to the diffuse stimulus with the response to the spot 1 stimulus using an ordering imposed by the diffuse stimulus at a latency threshold of 0.5. The black ragged front edge for the low number neurons reflects the reordering of the front for the spot response.

Fig. 10 displays the same data sets as Fig. 9, using a latency threshold of 0.1. The lower threshold should reflect the influence of different inputs. There is dramatic difference in the ordering of the inputs as the stimulus varies. Some adjacent spots first activate neurons on opposite ends of the cortex.

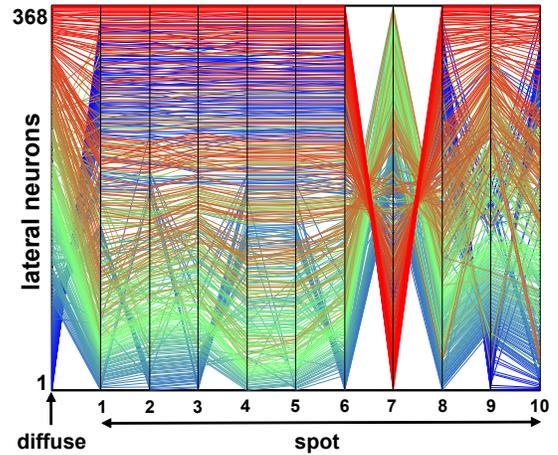


Figure 10: Comparison of response order for diffuse stimulus with 10 non-overlapping spot stimuli using a latency threshold of 0.1.

3.4 Subspace projections

This section discusses the representation of these data sets by their projection into a low-dimensional subspace. One can then relate landmarks or features of the resulting projection curve to specific details of the overall response. One can also compare responses from different stimuli by projecting their data sets into the same subspace. We use KL decomposition or principal component analysis (PCA) to find a low-dimensional subspace that captures correlated behavior [8].

Suppose the data set has N time snapshots, each consisting of L spatial samples. The data set can be arranged in an $L \times N$ matrix $\mathbf{U}=(u_{x,t})$. Each column of \mathbf{U} represents a time snapshot of the data set, and each row of \mathbf{U} represents a time series for a particular spatial sampling point. The KL spatial basis functions, ϕ^k , are the eigenvectors of the $L \times L$ matrix $\mathbf{U}\mathbf{U}^T$. The $L \times 1$ vectors ϕ^k are usually sorted by the size of their associated eigenvalues, with ϕ^1 capturing the most data set energy. For the example in this paper we plot $\mathbf{U}^T\phi^1$ (horizontal axis) versus $\mathbf{U}^T\phi^2$ (vertical axis) to obtain a low-dimensional subspace representation.

Fig. 11 shows the subspace representation for the diffuse stimulus data set. Since we are projecting into a two-dimensional subspace, each time sample is represented by a point in the plane. The black curve connects the time samples in succession.

The seven dots on the diagram in Fig. 11 mark specific time events. The *stimulus onset* and *stimulus end* events delineate the time course of the visual stimulation from 0 to 150 ms. After *stimulus onset*, the projection curve is relatively linear, indicating that the contributions of ϕ^1 and ϕ^2 maintain the same proportion as they grow. The *end of linear growth* event at 124 ms occurs when the projection $\mathbf{U}^T\phi^2$ reaches its maximum value. At this time, labeled **a** on the space-time diagrams in the paper, the curve makes a sharp change in direction and linear growth ends.

The *prototypical event* represents the time at which $\mathbf{U}^T\phi^2$ first returns to zero after stimulation and the spatial structure most closely resembles ϕ^1 . The prototypical event occurs at 208 ms and is labeled **b** on the space-time diagrams.

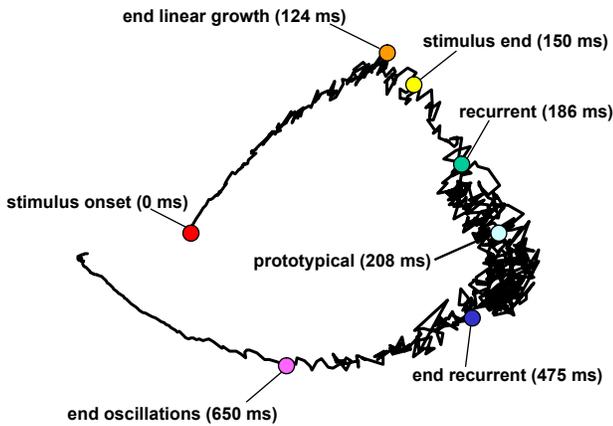


Figure 11: Projection of the dataset on ϕ^1 plotted versus the projection on ϕ^2 for diffuse stimulus. Shape landmarks correspond to vertical gray lines on earlier space-time diagrams.

The dots in Fig. 11 labeled *recurrent*, *end recurrent*, and *end oscillations* occur at 186 ms, 475 ms and 650 ms, respectively. These hand-placed dots mark visual changes in the character of the project curve. Around 186 ms, oscillations of the projection curve become more pronounced and disorganized. At 475 ms, the response appears to reorganize and smoothly decays at 650 ms.

3.5 Direct views and animation

Direct visualizations display snapshots or animations of the response associated with their spatial position. For example, Fig. 12 represents neurons by the same shapes as in Fig. 2 and uses color for response.

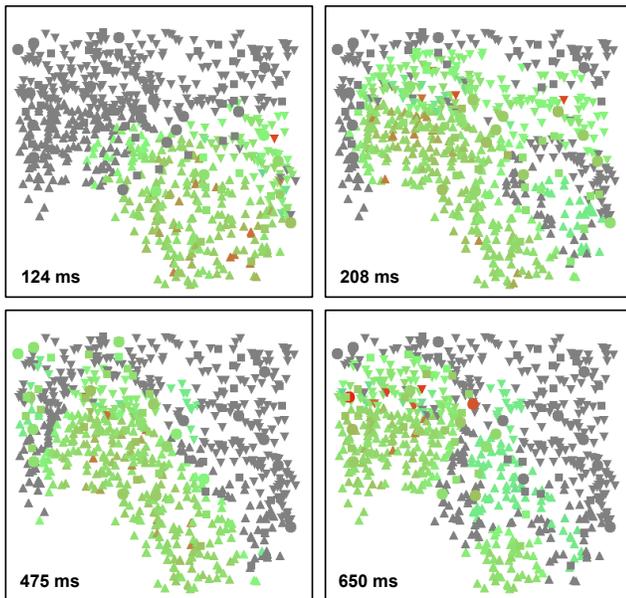


Figure 12: Snapshots of the soma voltage in response to a diffuse flash stimulus for selected times.

We observe from Fig. 12 that only about half of the neurons have been activated at 124 ms when the linear growth period ends. The prototypical state at 208 ms appears to have more activation in the central regions of the cortex. The hill shape of this state is apparent in three-dimensional visualizations of the response surface shown in the accompanying video.

3.6 Timing summary diagram

Biologists often use block diagrams to summarize the interacting cortical components in a neural network. The different cortical components are represented by text blocks with labeled arrows representing the types of interactions. A block diagram, such as the one shown in Fig. 13 can also be used to organize timing information as observed from the visualizations. The red arrows mark inhibitory interaction and the green arrows mark excitatory connections. Notice that the horizontal cells have no feedback inhibition. The medial cells are not shown in the diagram because the visualizations revealed they had little effect on the initial response.

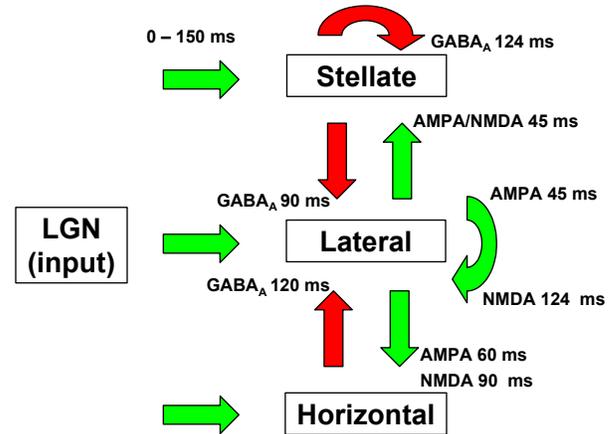


Figure 13: Block diagram for the interacting factors contributing to the initial response to a diffuse flash.

The space-time diagram of Fig. 7 shows that GABA_A inhibition of the lateral neurons appears around 90 ms into the response and grows as more inhibitory neurons become active. Direct viewing of visualizations, such as those of Fig. 12, reveals that the two columns of synchronous spikes of GABA_A that appear prior to 90 ms are produced by a single stellate neuron located near the LGN in the lower left corner of the domain. The end of linear growth event at 124 ms coincides with GABA_A reaching a significant size, suggesting that inhibitory neurons are acting to control the growth of the response.

The space-time diagram of excitation for the stellate neurons (not displayed due to space limitations) shows that both AMPA and NMDA receptors are activated about 45 ms after the stimulus begins, but that this does not result in significant voltage (and hence downstream inhibition of the lateral neurons) until after 100 ms. The space-time diagram for the horizontal neuron membrane potential (not shown) has first activation as the result of a single horizontal neuron receiving excitatory input from neighboring lateral cells. This neuron is located farther from the initial activation and occurs later than the response of the first stellate neuron. The sparse distribution of stellate and horizontal cells definitely affects the details of the response.

4 DISCUSSION

The visualizations presented in this paper focus on understanding of the initial response to diffuse flash in the NGU model of the turtle visual cortex, but many more complex questions can be addressed. The visualizations provide model validation as well as detailed comparisons among datasets.

Investigations typically begin with direct viewing of the data. The displays of Fig. 12 were snapshots taken from Davis, a data viewing tool [17]. Davis allows users to single step through frames, play them automatically, or to control the progress through the frames with a slider. The latter capability proved to be essential for discovering and verifying timing details.

The color displays representing response, such as in Fig. 12, were helpful for understanding patterns, but we only discovered the importance of the low-level input “priming” after viewing three-dimensional surface representations of the response. These surfaces represent response by height, making it easier for the eye to distinguish small changes over a large range. Unfortunately the data must be interpolated and filtered for this visualization, so detailed information about particular neurons is lost.

It is almost impossible to recognize structurally coherent changes from directly viewing. One would not mark 186 ms or 475 ms as special times without the low-dimensional projections of Fig. 11. Having identified times associated with structural changes in the low-dimensional projections, we were able to correlate these times with events of biological significance.

Unfortunately, the KL decomposition used for the low-dimensional projections in this paper is relatively costly for large data sets. Furthermore, experimental data sets contain noise and may often have non-stimulus related artifacts that contaminate the KL subspace. Robbins and Senseman [18] have shown that a two-dimension subspace called a wave subspace can be computed efficiently from sampled data and results in a more consistent two-dimensional representation of wave structure in experimental data than straight KL decomposition. For model data, noise and confounding do not occur, so we used straight KL decomposition.

The combination of direct viewing, structural visualization provided by low-dimensional projections, and summary information provided by overlaid space-time diagrams allow detailed investigation of this complex system. These visualizations can be computed quickly and scale well to larger data sets.

Acknowledgements

Our sincere thanks to Philip Ulinski, Bijoy Ghosh and Zoran Nenadic. This work was partially supported by NIH (G12 RR13646), NSF (ACI-9721348; EIA-0217884) and ONR (N00014-97-0029).

REFERENCES

- [1] J. Bower and D. Beeman. *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural Simulation System*. 2nd ed. Springer-Verlag, 1998.
- [2] B. Cabral and L. Leedom. Imaging vector fields using line integral convolution. *Proc. of ACM SIGGRAPH 1993, Computer Graphics*, 27, 263–272, 1993.
- [3] M. Cannon, M. Hasselmo, and R. Koene. From biophysics to behavior: Catacomb2 and the design of biological plausible models for spatial navigation. *Neuroinformatics*, 1, 1, 3–42, 2003.
- [4] D. Enright, S. Marschner, and R. Fedkiw. Animation and rendering of complex water surfaces. *ACM Trans. on Graphics*, 21, 3, 736–744, 2002.

- [5] B. Ermentrout. *Simulating, Analyzing, and Animating Dynamical Systems: a Guide to XPPAUT for Researchers and Students*. SIAM Software Environments and Tools 14, 2002.
- [6] J. Helman and L. Hesselink. Representation and display of vector field topology in fluid flow data sets. *Computer*, 22, 8, 27–36, 1989.
- [7] M. Hines and N. Carnevale. NEURON: A tool for neuroscientists. *The Neuroscientist*, 7, 123–135, 2001.
- [8] M. Kirby. *Geometric Data Analysis*. John Wiley & Sons, 2001.
- [9] A. Inselberg and B. Dimsdale. Parallel coordinates: a tool for visualizing multi-dimensional geometry, *IEEE Visualization '90*, 361–378, 1990.
- [10] B. Jobard, G. Erlebacher, and M. Hussaini. Lagrangian-Eulerian advection for unsteady flow visualization, *IEEE Visualization '01*, 53–60, 2001.
- [11] C. Johnson, S. Parker, D. Weinstein and S. Heffernan. Component-based problem solving environments for large-scale scientific computing, *J. on Concurrency and Computation: Practice and Experience*, 14, 1337–1349, 2002.
- [12] C. Lawrence, L. Lee, L. Pao and R. Novoselov. Shock and vortex visualization using a combined visual/haptic interface. *IEEE Visualization '00*, 131–137, 2000.
- [13] K.-L. Ma, J. Van Rosendale, and W. Vermeer. 3D shock wave visualization on unstructured grids. *Proc. 1996 Symp. on Volume Visualization*, 87–94, 104, 1996.
- [14] R. Malladi, J. A. Sethian, and B. C. Vermuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 17, 2, 158–175, 1995.
- [15] Z. Nenadic, B. Ghosh, and P. Ulinski. Propagating waves in visual cortex: A large-scale model of turtle visual cortex. *J. of Computational Neuroscience*, 14, 2, 161–184, 2003.
- [16] J. Prechtl, L. Cohen, B. Pesaran, P. Mitra, and D. Kleinfeld. Visual stimuli induce waves of electrical activity in turtle cortex. *PNAS*, 94, 7621–7626, 1997.
- [17] K. Robbins, I. Grinshpan, K. Allen, and D. Senseman. Synchronized views for exploring populations of neurons. *Proc. Visualization and Data Analysis*, San Jose, Jan. 21–22, 2004.
- [18] K. Robbins and D. Senseman. Extracting wave structure from biological data with application to responses in the turtle visual cortex. *J. Computational Neuroscience*, 16, 267–298, 2004.
- [19] D. Senseman and K. Robbins. High-speed VSD imaging of visually evoked cortical waves: Decomposition into intra- and intercortical wave motions. *J. Neurophysiology*, 87, 3, 1499–1514, 2002.
- [20] L. Sobierjski, R. Avila, D. O'Malley, S. Wang, and A. Kaufman. Visualization of calcium activity in nerve cells. *IEEE Computer Graphics and Applications*, July, 55–61, 1995
- [21] P. Ts'o and B. Barsky. Modeling and rendering waves: Wave-tracing using beta-splines and reflective and refractive texture mapping. *ACM Trans. on Graphics*, 6, 3, 191–314, 1987.
- [22] S.-K. Ueng, C. Siorski and K.-L. Ma. Efficient streamline, streamribbon, and streamtube constructions on unstructured grids. *IEEE Trans. Vis. and Computer Graphics*, 2, 2, 100–110, 1996.
- [23] P. Ulinski. Neural mechanism underlying the analysis of moving visual stimuli. In Ulinski, P., Jones, E., and Peters, A., Eds. *Cerebral Cortex Vol. 13: Models of Cortical Circuitry*, Plenum Press, New York, 382–399, 1999.
- [24] J. Van Wijk. Image based flow visualization. *ACM Tran. on Graphics*, 21, 3, 745–754, 2002.
- [25] G. Whitham. *Linear and Nonlinear waves*, Wiley Interscience, 1974.
- [26] XNBC Homepage: www.b3e.jussieu.fr/xnbc/.