

G. Gerstein

-1-

N00014-87-K-0766

PROGRESS REPORT FOR ONR 1 Sept 87 to 31 Aug 88.

1. The Laboratory.

During this year we have replaced our Data General MV10000 computer with a SUN4, using funds from other sources. ONR has contributed to the purchase of several programming languages and calculation packages for the new machine, and in addition has helped with the maintenance contract. The new computer is up to 8 times faster than the old one for several of our most compute intensive problems, and has been an indispensable tool in most of the research described below, both experimental and theoretical.

2. Multi neuron experiments: tools and applications.A. Rapid time variation of neuronal interactions.

We have largely completed work on development and test of the Joint PST histogram calculation, and several laboratories (ours included) are applying the calculation to real neuronal recordings. (The development built on old work of Gerstein and Perkel dating from 1972, and involved several years of collaborative work with A. Aertsen and G. Palm of Tuebingen, West Germany and M. Habib of University of North Carolina.) The new calculations allow for the first time the explicit evaluation of stimulus-time locked variations in correlation of the firing of two recorded neurons, taking into account whatever modulations there are of the individual firing rates of the neurons. A new type of significance test has been developed for evaluating the results. Details of the methods and interpretations are given in papers 7 and 8.

A few applications to experimental material are shown in paper 1. Additional papers showing applications to experiments both from our and other laboratories are in preparation. The results are that rapid, stimulus-time locked variations in correlation are frequently observed in a number of different experimental preparations and circumstances. This new observation suggests that the relations between two neurons are dynamic on a time scale of tens of milliseconds. At this stage we know neither the mechanisms nor the functional significance of such rapid changes in neuronal interaction parameters. However it would seem appropriate to investigate whether this kind of behavior occurs in theoretical models of neuronal networks as an emergent property. In other words, if such rapid changes are NOT explicitly placed into a neuronal network model, say through some special synaptic properties, will they still appear under some circumstances as a property of the network as a whole?

We have continued to evaluate experimental multi-neuron

September 8, 1988

recordings (up to 10 simultaneously and separably recorded neuronal spike trains) from both our and other laboratories. In particular we have made use of the Gravity calculation (1985 papers from our laboratory) which allow examination of the entire group of recorded neurons rather than forcing evaluation by pairs. Evaluation in pairs is of course perfectly possible in terms of computing time, but puts an impossible burden on the experimenter who must then examine literally thousands of correlograms for a typical experiment.

In all the diverse material so far examined we frequently observe changes in correlation structure associated with presentation of different stimuli or different behavioral conditions. We have traditionally expressed correlation structure among spike trains in terms of an equivalent or effective connectivity, i.e. the simplest "wiring diagram" among the observed neurons that could account for the observed correlation structure. Note that such equivalent connectivity diagrams are not necessarily related to the true anatomical connectivity: synapses that are not active could not be detected, for example.

Thus the experimental data suggest that the effective connectivity among observed neurons depends on, and varies with, the stimulus or behavioral conditions. The time scale of such changes depends on the stimulus structure (duration, repetition rate, etc.) but is generally between hundreds of milliseconds and seconds. These observations are of great interest to understanding brain function because they suggest that the neuronal systems reconfigure to deal with each different problem. The systematic rules of such reconfigurations need considerably more study. Examples of such stimulus dependent reconfiguration are shown in paper 9.

Analysis of our own auditory system multi-neuron experiments has continued along more traditional lines. Some of the excitatory neuronal interactions we have found in cochlear nucleus of the rat are described in paper 11, and further material on rapid modulation of such correlations is in the final stages of writing. (P. Gochin, J. Kaltenbach).

B. Development of new multi-neuron experimental paradigms.

One of the primary problems facing the (few) laboratories that are equipped to do multi-neuron experiments is how to design experiments that investigate neuronal assembly properties rather than just standard stimulus-response properties in parallel. At the time of this writing, G. Gerstein and P. Bedenbaugh are making a two week visit to the laboratory of Michael Merzenich at University of California, San Francisco in order to learn the details of an experimental preparation that should serve as a splendid test bed for studies at the assembly level.

Merzenich and his colleagues have reported that by a

September 8, 1988



A-1

number of simple manipulations they are able to alter the somatosensory cortical map. The observation is simply that the territory mapping some particular peripheral area can be enlarged on the cortex at the expense of adjacent cortical representations of other areas. The manipulations studied by the Merzenich group include various peripheral deafferentations, and also local cortical electrical micro-stimulation. We plan to use the latter, since it produces the desired changes in a matter of some four hours. This is fast enough for multiple-neuron studies of assembly organization to be feasible.

An analogy of these observations would be the structure of a magnetic material, with macroscopic domains of aligned spins. The boundaries of such domains can be shifted (at low energy costs). For our purposes, once we can reliably replicate the basic experiment of territorial enlargement of, say, the cortical representation of a particular finger tip (anesthetized rat), we propose to place an array of microelectrodes across the enlarging domain. Thus we should be able to contrast the neuronal interactions inside and outside the boundary, and to follow whatever change in organization there is as the boundary moves. This kind of reorganization is not on the fast time scale of the stimulus related reorganization we have described above. Nevertheless, similar mechanisms may be involved. A number of other interesting questions will also be addressed by these experiments. For example, in paper 10 we have examined the various proposed definitions of neuronal assembly. The new experiments will test the relations between assembly and computational function, i.e. is the domain boundary (computational) the same as the assembly boundary. Are there separate assemblies within a single domain? What is the spatial extent and density of such sub-assemblies? Are they spatially overlapping (within the larger map domain) or are they discrete sub-domains? (I find it somewhat amazing that we do not presently have answers to such fundamental aspects of nervous system organization; this type of information is vital for models of nervous systems that seek to explain how the brain does computations.)

Similar reorganization of cortical maps have been reported for auditory and visual systems. We have chosen to start with the somatosensory map because the reported experiments are cleaner and involve faster and easier manipulations. The overall property of reorganization does appear to be a general dynamic aspect of many parts of the nervous system.

3. Simulation of real neuronal networks.

A. Tools

A primary shared need of many of our projects is the use of a sophisticated and flexible neuronal network simulator. Such a tool allows purely theoretical studies of neuronal organization

September 8, 1988

or computation, but also allows the development and testing of the various analytic calculations that are necessary for experimental data. The underlying requirement is that both elements and connections allow an appropriate degree of mimicry of the real physiological situations. This type of simulator is in considerable contrast to the layered schemes with back propagation that are now commonly used to solve various computational problems, but which have little or no relation to real nervous systems.

Our initial effort in this direction has been based on an adaptation of one of the MacGregor (University of Colorado) simulator programs. Necessary modifications and additional modules were written both in Tuebingen and Philadelphia (K. Boven and P. Bedenbaugh, respectively). The current version is considerably easier to use than the original, allows various types of random and specific connections among many populations of neurons, and uses a relatively simple neuron model with potassium and calcium conductances. (Generation of action potentials is "stylized", so that sodium conductance is not explicitly modelled). This simulator does not handle each neuron as a multi-compartment model to mimic dendritic geometries, but does allow a fairly complete specification of synaptic effects, which is almost as good. In other words, this model is at a reasonable compromise between infinite detail (and hence uncomputability) and preposterous oversimplification. The resulting product can handle nets of many thousand "neurons", and on our SUN4 computes a single neuron simulation at 6 to 10 times faster than real time. As the number of "neurons" increases, so does the computation time, but in a directly proportional way with very little (less than 10%) additional overhead. Thus reasonably long runs of fairly large nets are practical; this is far from a real time tool, but works well in overnight runs.

We have evaluated the transfer of the same programs to the local supercomputer which is available to us at Princeton. The conclusion is that unless a major rewrite into vector arithmetic is undertaken, that we could expect only about a factor of 3 in computation speed. This is not worth it, given that even direct transfer involves a large amount of translation and tailoring of the program. At present, no massively parallel computer architectures seem to be available at or near the University of Pennsylvania. Rethinking this problem will be appropriate whenever additional hardware capacity appears. Another direction for development in the near future will be to acquire and test the Cal Tech neuronal simulator. This should be available sometime late this fall, but is of unknown capacity and speed. It is known to be an exceedingly detailed simulation, with an extraordinary number of available parameters.

B. Simulation studies.

Initial studies with the modified MacGregor simulator have been aimed at replicating several types of experimental

September 8, 1988

multi-neuron results. In particular we have been interested in possible mechanisms for (a) the production of broad (20-50 ms) peaks in cross-correlograms, and (b) the stimulus dependent effective connectivity (or correlation structure) described above for real neurons.

These studies (P. Bedenbaugh) have demonstrated the existence of a whole new class of mechanism for both rapid and slow changes of correlation structure among observed neurons. The relevant simulation studies involve several interconnected "neurons" which mimic the set of neurons actually observed in a typical experiment, and in addition, a large pool of "neurons" that is unobserved in the experiment. If these pool neurons project upon the observed neurons, it is possible to produce certain types of stimulus dependence in the correlations, i.e. to mimic the experimentally observed reorganizations. Parameters that are relevant to the details of the pool influence are strength of projection (both numbers and individual synaptic values), the connectivity within the pool, and modulations of pool activity (whether spontaneous or externally imposed). The connectivity influences synchrony in the pool firings, and hence in the projected effects. Early results of this project are being presented at two workshops ("Neural Network Models and Their Relevance to Neurobiology", Jerusalem, May 24 - June 8, 1988 and "Principles of Cortical Function", Irvine, Sept 18-21, 1988) and at the 1988 Neurosciences meeting in Toronto. A paper which reviews various concepts of neuronal assemblies as well as giving a demonstration of some of these new pool effects is in press (paper 10).

5. Gabor filter computational models for texture segmentation and for perception of slanted textured surfaces. (M. Turner)

Application of Gabor filter models to segmentation of natural and artificial pictures according to regions of different textures is described in paper 2. Included also is a review of neuronal assembly properties, and an examination of assemblies to realize the segmentation computations.

Most of the textured surfaces perceived by humans and animals have a non-perpendicular orientation to the viewer's line of sight. The most common example of this is the textured surface of the ground as perceived during normal activity. It is important, therefore, that the animal have a perception of the surface orientation in order to plan and facilitate his movements within the environment. A number of models for recovery of textured surface orientation (estimation of slant and tilt of a textured plane) have been proposed within a machine perception context. However, these models have not addressed the mechanisms by which such perception would occur in the visual cortex. The current research has addressed this question in three sub-projects. Three separate papers are currently in preparation describing

September 8, 1988

respectively parts 5A, 5B, and 6B of this report.

First, we have created an algorithmic model which uses as its low level filter the 2D Gabor function found to be an accurate and economical description of simple cell receptive fields. This model, although it operates by evaluation of equations, does not embody principles incompatible with current physiological understanding. Development of such a model affords us a mathematical and programmatic base upon which to extend our investigations.

Second, we have developed a model to explain psychophysical data about human perception of texture surface orientation. It has been known for almost 40 years that humans tend to underestimate the slant of a textured surface, particularly when the texture is "irregular". Before this research, however, there were no models to explain this phenomenon. Using the algorithmic model previously mentioned, we have shown that such underestimations are the inevitable result of local/spectral receptive field types such as many of the current descriptions for simple cells.

Third (part 6B below), we have developed a higher order receptive field model which is capable of performing the same function as the algorithmic model in a more physiologically realistic manner. These receptive fields are orientation detectors in a 4 dimensional space with 2 spatial and 2 spatial frequency dimensions. To the best of our knowledge, such receptive fields have never been described. However, it is also unlikely that anyone would have performed experiments capable of detecting them. This phase of the research, then, makes a prediction for a new receptive field type and describes the experimental techniques which could be used to discover them.

A. An algorithmic model for recovery of texture surface slant using 2D Gabor functions.

A textured surface having a non-perpendicular orientation to a viewer exhibits systematic frequency shifts across its image. When measured by 2D Gabor functions these frequency shifts induce changing energy distributions. An algorithm has been developed which recovers the slant and tilt parameters of the original surface from these energies. The heart of this recovery process is two routines which are iteratively executed. The first adjusts local sets of parameters to reduce the error between predicted and measured energies. The second propagates the local parameters to neighboring regions to consolidate the estimates of slant and tilt. The algorithm is capable of operating in parallel on any number of regions in the image and with a diverse set of filter inputs. While this algorithm operates by evaluations of equations on local patches of the image, its form is not inconsistent with physiological principles. It uses as its low level operator a receptive field type experimentally found to be a good description of simple cells in V1 and utilizes

September 8, 1988

propagation of values between neighboring computing elements not inconsistent with the inter and intra columnar connections found in cortex.

B. A model for the underestimation of slanted, textured surfaces by human observers.

It has been known since the experiments of Gibson in 1950 that, in the absence of other cues, a human observer will tend to underestimate the slant of a textured planar surface. When shown with a perspective projection, observers tend to perceive the textured surface as having a steeper or more perpendicular orientation than it actually has. Moreover, it has also been found that the amount of underestimation tends to be greater for textures which are perceived as being more "irregular" both in texture element shape and placement. Although this perceptual tendency has been replicated in a variety of psychophysical studies, as yet no models have been offered in explanation. In this work we have used gradients in the output of 2D Gabor filters to recover the orientation of a textured surface. However, for "irregular" textures these gradients of filter outputs are inconsistent with curves for the correct amount of slant. The presence of closely spaced multiple spectral components which comprise the irregular textures tends to flatten the gradient curves, making the gradients more consistent with simpler textures at a lower angle of slant. Therefore, this tendency toward underestimation will be a characteristic of any visual system which uses a local spectral operator as its low level filter.

6. Layered networks and receptive field models for perception of slanted, textured surfaces.

Layered networks and learning algorithms like back propagation are not intended to mimic real neuronal networks or their computations. Nevertheless, such model studies can be very interesting in exploring or even finding algorithms for well defined transformations between inputs and outputs.

A. Tools

We have purchased and installed the Rochester connectionist simulation package. This tool has been particularly organized for SUN type computers, has excellent input and output facilities, and has so far required little customizing. It includes the back propagation algorithm as an option for layered structure studies. Other programs are in principle available (i.e. from the Rumelhard group), but we have so far made no effort to compare advantages or performance with the Rochester program.

B. Application to orientation of a textured plane.

The algorithm described in part 5A. operates by finding

solutions to a complex system of non-linear equations using Gabor function energy values. A system such as this is more likely to find realization in cortex as sets of cells representing, in their receptive fields, specific instances or templates of the equations at different slant values. We have generated and examined such receptive fields in a layered model with back propagation; the general approach is similar to the Anderson-Zipser oculomotor study or the Sejnowski shape from shading study.

The three layer network was set up (M. Salganicoff) to handle a selection of Gabor filter energies as input, and six possible orientations of slant (one dimension only) as output. Input was precalculated as filter energies (at several frequencies and orientations) spatially averaged over six regions along a vertical centered stripe of the learning set of pictures. In each of the six regions we used filters of 5 frequencies and 2 orientations, i.e. 10 spatially averaged energy values. Thus there were 60 input "neurons", each corresponding to one of the filters in one of the regions. The network used three "neurons" in the hidden layer, and six in the output layer.

Initial training was on a set of 5 textured pictures (some natural and some artificially constructed), each presented 1000 times in random order for 3 different values of slant. Learning was rapid, and the final configuration of the layered network attained a high performance level. Additional work is in progress on performance with larger learning sets and with regard to interpolation and generalization.

The most interesting aspect of these results so far is the description of the receptive fields of the hidden layer "neurons" after the network has stabilized. These fields can be described as a central, elongated excitatory region flanked by two elongated inhibitory regions (very familiar) BUT in a space with coordinates y (spatial picture region) and f (spatial frequency). Note that such receptive fields are extremely common in an x - y space, i.e. with regard to picture regions. Such an x - y description matches the Gabor fields that are found physiologically for simple neurons in visual cortex, and which were the starting point for all this modeling activity.

In the current situation of the y - f space, the "orientation" of the elongated axis determines the spatial frequency gradient, and hence the slant angle of the picture material to which the neuron is tuned. The width of the excitatory and inhibitory regions will be related to the sharpness of tuning.

The significance of the new observation of elongated, oriented, excitatory-inhibitory receptive fields in a y - f space is its prediction for experimentation. We should look, in the real system, for neurons that have a systematically varying preferred spatial frequency across their receptive field. In other words, we are raising the suggestion that somewhere in the visual

system there should be large receptive fields that exhibit clear gradients of some tuning properties across their extent. The idea should not only be related to spatial frequency, as in the current example, but might include time as the second coordinate. Such receptive fields would be velocity (vector) sensitive. As far as we know, neurons with tuning gradients across receptive fields have neither been sought nor reported. Experimental determination of such receptive field properties is time-consuming, but perfectly feasible -- it just has never been tried.

In summary, the back propagation model has suggested a new solution to the slanted texture problem which can, in principle, be easily attained by physiologically realistic (weighted) summation of Gabor filter (simple area 17 neurons) elements. In addition the model predicts some receptive field properties that should be sought among real neurons in real brains.

September 8, 1988

PUBLICATION LIST FOR ONR PROGRESS REPORT SEPT. 1987 to AUG. 1988.Papers published or in press.

1. Information flow and state in cortical neural networks: Interpreting multi-neuron experiments. G.L. Gerstein. In: Organization of Neural Networks. , Eds.: W. v. Seelen, G. Shaw, R. Leinhos. VCH Verlagsgesellschaft, Weinheim, 1988.
2. Neuronal Assemblies as building blocks of cortical computation. G.L. Gerstein and M. Turner. In: Computational Neuroscience. Ed: E. Schwartz, MIT Press, 1988. (in press)
3. Cortical auditory neuron interactions during presentation of 3-tone sequences: effective connectivity. I. Espinosa and G.L. Gerstein, Brain Research 450 : 39-50, 1988.
4. Detecting spatio-temporal firing patterns among simultaneously recorded single neurons. M. Abeles and G.L. Gerstein, J. Neurophysiol. (in press) 1988.
5. Unsupervised waveform classification for multi-neuron recordings: A real-time software based system. I. Algorithms and Implementation. M. Salganicoff, M. Sarna, L. Sax, and G.L. Gerstein, J. Neurosci. Methods (in press) 1988.
6. Unsupervised waveform classification for multi-neuron recordings: A real time software based system. II. Performance comparison to other sorters. M. Sarna, P. Gochin, J. Kaltenbach, M. Salganicoff and G.L. Gerstein, J. Neurosci. Methods (in press) 1988.
7. Dynamics of neuronal firing correlation: modulation of "effective connectivity". A.M.H.J. Aertsen, G.L. Gerstein, M. Habib and G. Palm with the collaboration of P. Gochin and J. Kruger, J. Neurophysiol. (in press) 1988.
8. On the significance of correlations among neuronal spike trains. G. Palm, A.M.H.J. Aertsen and G.L. Gerstein, Biological Cybernetics 59 : 1-11, 1988.
9. Interactions within neuronal assemblies: theory and experiment. G. L. Gerstein, In: Brain Organization and Memory: Cells, Systems and Circuits, Eds.: J.L. McGaugh N.M. Weinberger and G. Lynch., Oxford University Press, (in press) 1988.
10. Neuronal Assemblies. G.L. Gerstein, P. Bedenbaugh and A.M.H.J. Aertsen, IEEE Trans. Biomed. Eng. (in press) 1989.

September 8, 1988

Papers submitted or in preparation:

11. Coordinated Activity of Neuron Pairs in Anesthetized Rat Dorsal Cochlear Nucleus. P.M. Gochin, J.A.Kaltenbach and G.L. Gerstein. (submitted)
12. Neuronal assemblies as observed in multi-neuron experiments: Dynamic organization depending on stimulus. G.L. Gerstein, A.M.H.J. Aertsen et al.
13. An Algorithmic Model for Recovery of Textured Surface Slant Using 2D Gabor Functions. M. Turner, R. Bajczy, G. Gerstein.
14. A Model for the Underestimation of Slanted, Textured Surfaces by Human Observers. M. Turner, G. Gerstein, R. Bajczy.
15. Layered Networks: Application to Orientation of a Textured Plane. M. Turner, M. Salganicoff, G. Gerstein.

September 8, 1988