AA

Form Approved Through 06/30/91
OMB No. 0925-0001

DEPARTMENT OF HEALTH AND HUMAN SERVICES
PUBLIC HEALTH SERVICE

GRANT APPLICATION

Follow instructions carefully. Type in the unshaded areas only.

LEAVE BLANK FOR PHS USE ONLY.		
Type	Activity	Number
Review Group		Formerly
Council/Board *(Month, Year)*		Date Received

1. TITLE OF PROJECT *(Do not exceed 56 typewriter spaces.)*
Network Optimization in the Brain

2. RESPONSE TO SPECIFIC PROGRAM ANNOUNCEMENT ☐ NO ☒ YES *(If "YES," state RFA number and/or announcement title)*
Number: "FIRST" Title: Mathematical/Computational/Theoretical Neuroscience

3. PRINCIPAL INVESTIGATOR/PROGRAM DIRECTOR ☒ NEW INVESTIGATOR

3a. NAME *(Last, first, middle)*	3b. DEGREE(S)	3c. SOCIAL SECURITY NO.
Cherniak, Christopher G.	Ph.D M.A. B.A.	460-74-9290

3d. POSITION TITLE	3e. MAILING ADDRESS *(Street, city, state, zip code)*
Associate Professor	Philosophy Department University of Maryland College Park, MD 20742
3f. DEPARTMENT, SERVICE, LABORATORY, OR EQUIVALENT Department of Philosophy	
3g. MAJOR SUBDIVISION College of Arts & Humanities	
3h. TELEPHONE *(Area code, number, and extension)* (301) 405-5703	

4. HUMAN SUBJECTS	If "YES," exemption no. or	IRB approval date	4b. Assurance of compliance no.	5. VERTEBRATE ANIMALS	If "YES," IACUC approval date	5b. Animal welfare assurance no.
4a. ☒ NO ☐ YES				5a. ☒ NO ☐ YES		

6. DATES OF ENTIRE PROPOSED PROJECT PERIOD		7. COSTS REQUESTED FOR FIRST 12-MONTH BUDGET PERIOD		8. COSTS REQUESTED FOR ENTIRE PROPOSED PROJECT PERIOD	
From (YYMMDD)	Through (YYMMDD)	7a. Direct Costs	7b. Total Costs	8a. Direct Costs	8b. Total Costs
92/09/01	97/08/31				

9. PERFORMANCE SITES *(Organizations and addresses)*

Department of Philosophy
Computer & Space Sciences Bldg
University of Maryland
College Park, MD 20742

10. INVENTIONS *(Competing continuation application only)*
☒ NO ☐ YES If "YES," ☐ Previously reported ☐ Not previously reported

11. NAME OF APPLICANT ORGANIZATION
University of Maryland, College Pk
ADDRESS (Street, city, state, and zip code)
College Park, MD 20742

13. ENTITY IDENTIFICATION NUMBER	Congressional District
1520710851A1	5

12. TYPE OF ORGANIZATON
☒ Public: *Specify* ☐ *Federal* ☒ *State* ☐ *Local*
☐ Private Nonprofit
☐ For Profit *(General)* ☐ For Profit *(Small Business)*

14. ORGANIZATIONAL COMPONENT TO RECEIVE CREDIT TOWARDS A BIOMEDICAL RESEARCH SUPPORT GRANT
Code: 20 Identification: Other Academic

15. NAME OF OFFICIAL IN BUSINESS OFFICE	16. NAME OF OFFICIAL SIGNING FOR APPLICANT ORGANIZATION
Victor Medina	Victor Medina
TELEPHONE *(Area code, number, and extension)* (301) 405-6269	TELEPHONE *(Area code, number, and extension)* (301) 405-6269
TITLE Director, Office of Research	TITLE Administration & Advancement
ADDRESS 2100 Lee Bldg University of Maryland College Park, MD 20742	ADDRESS 2100 Lee Bldg University of Maryland College Park, MD 20742

17. PRINCIPAL INVESTIGATOR/PROGRAM DIRECTOR ASSURANCE: I agree to accept responsibility for the scientific conduct of the project and to provide the required progress reports if a grant is awarded as a result of this application. Willful provision of false information is a criminal offense *(U.S. Code, Title 18, Section 1001).*	SIGNATURE OF PERSON NAMED IN 3a *(In ink. "Per" signature not acceptable.)* Christopher Cherniak	DATE 91/09/16
18. CERTIFICATION AND ACCEPTANCE: I certify that the statements herein are true and complete to the best of my knowledge, and accept the obligation to comply with Public Health Service terms and conditions if a grant is awarded as the result of this application. A willfully false certification is a criminal offense *(U.S. Code, Title 18, Section 1001).*	SIGNATURE OF PERSON NAMED IN 16 *(In ink. "Per" signature not acceptable)* Victor Medina	DATE 9/24/91

SAI #: MD 910920-8239-3602201

PHS 398 (Rev. 10/89) (Reprinted 8/89)

AA

DESCRIPTION: State the application's broad, long-term objectives and specific aims, making reference to the health relatedness of the project. Describe concisely the experimental design and methods for achieving these goals. Avoid summaries of past accomplishments and the use of the first person. This abstract is meant to serve as a succinct and accurate description of the proposed work when separated from the application. **DO NOT EXCEED THE SPACE PROVIDED.**

The proposed computational anatomy project explores how well a set of formalisms derived from combinatorial network optimization theory fit as models for brain anatomy. The basic hypothesis is that long-range connections in the brain are a critically constrained resource, hence there is great pressure to finely optimize their deployment. The project would evaluate how accurately quantitative neuroanatomy data conforms to network optimization concepts. One concept may govern component placement in the nervous system: Given the specification of all interconnections among, for instance, the dozen ganglia of the roundworm C. elegans, are these components so positioned as to minimize total length of those connections? Our preliminary studies indicate that the actual layout of the ganglia in fact requires less connection length than any of the nearly 40,000,000 other possible placements. The planned experiments first involve compiling and/or refining appropriate neuroanatomical databases, principally from published anatomy, with initial focus including C. elegans and macaque and cat cerebral cortex. The proposed work thus meshes with current efforts toward a Neural Circuitry Database project--indeed, network optimization formalisms may prove a valuable top-down guide for the latter initiative. The experiments next involve evaluation of connection-optimality of the neuroanatomy. Network optimization problems are among the most computationally costly known; in general, only exhaustive search of all possible layouts can guarantee exact solutions. However, some probabilistic/approximation procedures developed for microcircuit chip design are promising candidates for "quick but dirty" biological mechanisms in optimization of neuroanatomy. The proposed experiments would include study of such mechanisms via parallel computer simulations. A grasp of such very general principles governing generation of the human nervous system should be a useful step in understanding how the embryology can be disrupted, and in turn can be protected from birth defects. Such basic structural principles would also be valuable in the growing field of research on therapeutic regeneration of the nervous system, particularly use of tissue transplants.

KEY PERSONNEL ENGAGED ON PROJECT

NAME, DEGREE(S), SSN	POSITION TITLE AND ROLE IN PROJECT	DEPARTMENT AND ORGANIZATION
Christopher Cherniak, B.A., M.A., B.Litt., Ph.D.	Associate Professor Principal Investigator	Department of Philosophy University of Maryland

Section 2. Research Plan

Specific Aims

The broad objective of the proposed research is to examine how well a set of general principles drawn from combinatorial network optimization theory predict actual neuroanatomical structure of some vertebrate and invertebrate nervous systems. The working hypothesis is that connections in the brain, particularly long-range ones, are a stringently limited resource, both in volume and in signal-propagation times; therefore, minimizing costs of required connections strongly drives nervous system anatomy. Network optimization theory is the field in computer science that has developed formalisms for expressing and solving various problems of "saving wire." The primary question here is, do such concepts in fact apply to the brain? The proposed computational anatomy is thus cross-disciplinary, evaluating goodness of fit of abstract models to quantitative neuroanatomical data.

As an example, a key problem in design of integrated circuits is "component placement optimization:" If one knows the interconnections among a set of components, what is the physical arrangement of the components that minimizes length of the connections? In fact, our computer search of all 39,916,800 possible layouts of the ganglia of the nematode C. elegans indicates that the actual layout is the unique first-place winner requiring the minimum total length of connections [Cherniak, 1991 (Appendix 1)].

Because of the relatively recent emergence of network optimization theory, there has been little if any previous work on its applicability to neuroanatomy. Specific goals of the proposed research:

1. Extension of our initial studies of component placement optimization in the nervous system.
   a. Refinement of our neuroanatomical database for C. elegans, currently based on, e.g., Wood (1988), down to the level of individual cell body position; and evaluations of placement optimization at that level by exhaustive computer search.
   b. Extension of our databases for cortex of macaque and cat, currently based on, e.g., Van Essen (1985) and Rosenquist (1985); and statistical testing of evidence of placement optimization of cortical areas.
2. Examination of applicability of another basic optimization concept, Steiner tree, to cost-minimization of neuron dendrite and axon arbors. Again, preliminary studies [Cherniak, 1990a (Appendix 2)] suggest this formalism is a powerful predictor of local arbor structures; the planned research extends the model to global or large-scale arbor optimization.
3. Evaluation of candidates for actual biological mechanisms of network optimization in the brain. The above combinatorial optimization problems typically are highly computationally complex. However, some "quick but dirty" practical heuristics developed for microcircuit layout suggest plausible candidates for neuroanatomy optimization; some include "simulated annealing" [Kirkpatrick, et al., 1983] procedures. Study of optimization mechanisms would proceed by computer simulations, some employing the massively parallel Connection Machine [Hillis, 1985].

Background

The theoretical framework of this work grew out of methodological studies of prevailing models of the agent in microeconomic, game, and decision theory [Cherniak, 1986]. The basic finding was that these models typically presupposed an agent with unlimited computational capacities, and more realistic bounded-resource models were then developed. Subsequently, the same approach was applied in computer science, to connectionist models of massively parallel and interconnected computation [Rumelhart and McClelland, 1986; Cognitive Science, 1985] that were intended to be more neurally realistic than conventional von Neumann computational architecture; again, the models tended to drastically overestimate available resources--here, actual connectivity in the brain. At least initial connectionist models often tacitly assumed neural connections were virtually infinitely thin wires. In assembling the quantitative neuroanatomy necessary for evaluating neural feasibility of connectionist models, it became evident that a weaker but still discernible trend toward overestimation of resources pervaded even neuroanatomy [Cherniak, 1990b (Appendix 3)].

Thus, a bounded-resource philosophical critique of mind-brain science ("We do not have God's brain") focussed attention on neural connections as a critically constrained neurocomputational resource. The formalism of scarcity of connections in computer science is combinatorial network optimization theory, and so a positive research program emerged: If actual brain connections are in severely short supply, is their anatomy correspondingly optimized? The investigation thus falls in a Pythagorean tradition of seeking simple mathematical patterns in observed natural forms [e.g., Thompson, 1961].

The human brain is commonly regarded as the most complex physical structure known in the universe. In the face of such overwhelming intricacy, neuroanatomy has traditionally tended toward "descriptive geography" of the nervous system, i.e., relatively low-level ad hoc characterization of individual neural structures. The abstractive power of concepts from computation theory would aid in coping with the unparalleled complexity of the brain. In particular, network optimization theory might provide a source for a "generative grammar" of the nervous system, some general principles that compactly characterize aspects of neuroanatomy.

For example, in the case of the C. elegans ganglia sketched above, approximately a thousand pages of published anatomy diagrams were represented as a hundred-page database, which in turn was reduced to a ten-page connectivity matrix, which was computationally verified to conform to the component placement optimization rule. If this type of result is confirmed, it constitutes one of the predictive success stories of recent quantitative anatomy.

The proposed research is relevant to the Neural Circuitry Database initiative [Pechura and Martin, 1991]; it explores a class of formalisms that may provide one useful framework for anatomical aspects of the latter work. Biomedical implications of this "brain as ultimate VLSI chip" approach may include improvements in understanding

of how brain development can be disrupted and safeguarded, and also may be useful in research on therapeutic techniques for regeneration of the brain.

Network Optimization Theory

Combinatorial network optimization theory came of age about twenty years ago with the emergence of the theory of NP-completeness [Garey and Johnson, 1979]. (The best nontechnical introduction to the field remains Lewis and Papadimitriou (1978), with Stockmeyer and Chandra (1979).) The key formal concept of a computational problem being NP-complete ("non-deterministic polynomial-time complete") need not be defined here; it is strongly conjectured to be linked with a problem being intrinsically computationally intractable, that is, not generally solvable without exhaustive search of all possible solutions. Because the number of possibilities combinatorially explodes as the size of a problem-instance grows, such brute-force searches are extremely computationally costly. Many of the most important real-world network optimization problems (e.g., most notably, Travelling Salesman) have been proven to be NP-complete or worse in computational complexity. Component placement optimization and Steiner tree, problems that are the focus of the present research, are of this type, having been proven to be "NP-hard."

Component placement optimization has received the most attention in computer science recently in connection with design of very large scale integrated (VLSI) circuits [Kuh and Ohtsuki, 1990]. As mentioned earlier, the problem can be defined as: Given the connections among a set of components, find the spatial layout of components that minimizes total connection costs. The simplest cost-measure is length of connections (often represented as the sum of squares of the lengths); usually the possible positions for components are restricted to a matrix of "legal slots." As a simple example, Figures 1a and 1b diagram two of six possible configurations of components 1, 2, and 3 in slots A, B, and C; for the connections among the components, the Figure 1b placement requires the least total connection length, and Figure 1a the most.

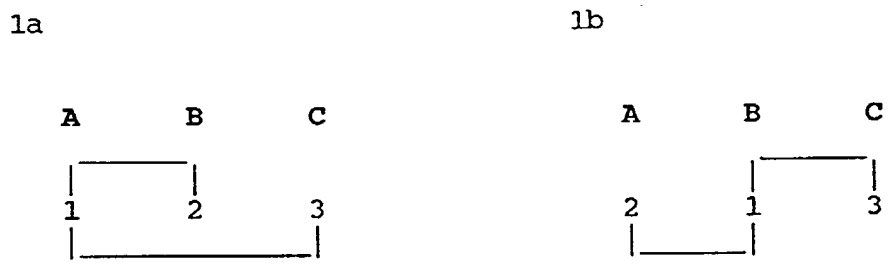

Figure 1. Component placement optimization: Two alternative placements of elements 1, 2, and 3 in positions A, B, and C. Placement 1a has greater total connection length than placement 1b.

Computation costs for exact solution of component placement optimization problems are of a magnitude not encountered in most bioscience computing, and constitute one of the main technical hurdles of the proposed research. For n components, the number of alternative possible placements is n! (Size of this search space is unaffected by whether permissible component positions are located in 3, 2 or 1 dimensions.) Heuristic procedures that yield approximately optimal solutions can be much more feasible, but their performance (e.g., how close to optimality are they likely to come) is not well understood.

Steiner tree, the other network optimization concept the proposed work focuses upon, has been studied in its simplest form at least as early as the Rennaissance [Winter, 1987]. The most relevant version of the problem is: Given a set of fixed node loci, find the set of arcs (or branches) between those loci that interconnects all loci and has shortest total length. The network will always constitute a tree. When it is permitted to have branch junctions only at node sites, it is a minimal spanning tree; when branch junctions may also occur at sites that are not nodes, it constitutes a Steiner tree. The total length of the Steiner tree for a set of nodes is equal to or less than length of the minimal spanning tree for the nodes. For example, Figures 2a and 2b show, respectively, a minimal spanning tree and a Steiner tree for seven nodes on a plane. The Steiner tree is about 6% shorter than the minimal spanning tree.

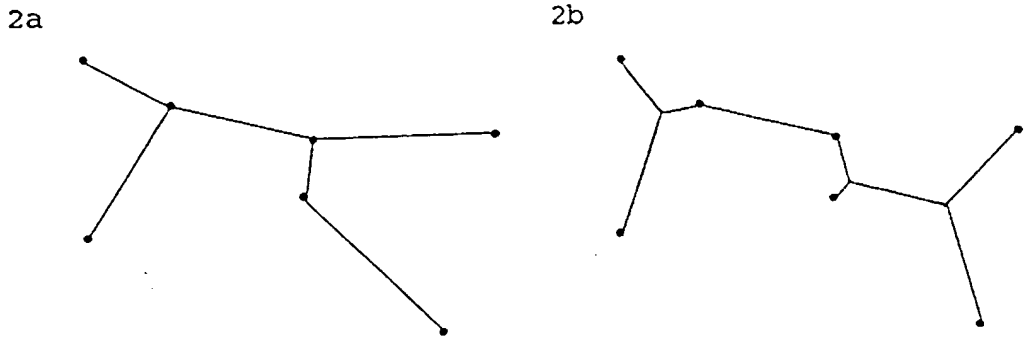

Figure 2. Minimal spanning tree, 2a, and Steiner tree, 2b, for seven nodes on a plane surface (after Lewis and Papadimitriou, 1978). The Steiner tree is shorter, but much more computationally costly to construct.

Since Steiner tree, like component placement optimization, is a member of the class of NP-hard problems, it is not surprising that the largest unconstrained Steiner tree problems that could be solved as of a few years ago had only 30 nodes [Bern and Graham, 1989]. However, while minimal spanning trees are equal to or longer than corresponding Steiner trees, they are not at all computationally intractable; exact algorithms for the former problem today perform well for quarter-million node sets. The basic question of goodness of fit of the Steiner tree concept to actual neuroanatomy is, do the dendritic and/or axonic arbors of a neuron form optimized Steiner trees interconnecting the cell body with a set of synaptic loci?

Preliminary Studies

An hypothesis that network optimization virtually unidimensionally drives some aspects of neuroanatomical connectivity--that it dominates the many other plausible orthogonal dimensions of optimization of the nervous system--is relatively simple, if not itself prima facie very plausible. This section summarizes the principal investigator's initial empirical studies of applicability of the network optimization hypothesis, funded in part by a four year, 50% appointment in the University of Maryland Institute for Advanced Computer Studies. This work yielded some interesting evidence that, in particular, component placement optimization and local Steiner tree are in fact organizing principles of nervous system anatomy. The most distinctive experimental techniques of these studies involved very large-scale searches by microcomputer; the relevant software engineering experience of the principal investigator and the programmers for these studies (and also for the planned research) originated in part from a five-year cognitive science instructional software development project [see Appendix 4], with microcomputer equipment grants from IBM's Advanced Educational Programs project.

Neural Component Placement

The hypothesis of component placement optimization in the nervous system was studied at multiple hierarchical levels, from gross to microscopic anatomy [Cherniak, 1991 (Appendix 1)]. (a) Explaining "why the brain is in the head:" At the highest level, positioning of the entire brain in the body constitutes a one-component placement problem. The simplest connection cost-measure is just total length of fiber in all sensory and motor tracts to and from the brain. Physical siting of all sensors and effectors was treated as a given, fixed "edge-constraint." The complete quantitative anatomy required for solution of this problem--in particular, for obtaining numbers of fibers in all nerve tracts--apparently has been published only for the human [Blinkov and Glezer, 1968] and nematode [Wood, 1988] nervous systems. In both cases, number of nerve fibers to and from locations anterior to the brain (or predominant concentration of the nervous system) exceeds number of fibers from locations posterior to the brain. It immediately follows that the "wire-minimizing" placement of the brain will be as far forward as possible. Actual positioning of human and nematode brains is in fact consistent with this wire-minimization prediction. In addition, inspection of gross anatomy drawings suggests that whenever anterior connections exceed posterior ones, as in the case apparently of all vertebrates and most invertebrates, the brain is correspondingly placed as far forward on the body axis as possible.

(b) Layout of functional areas of cerebral cortex: Extensive, although still incomplete, compendia of interconnections among the many cortical areas known to be involved in vision have been published recently for macaque [Van Essen, 1985] and cat [Rosenquist, 1985]. The placement optimization hypothesis for these components is that they are positioned on the 2-dimensional cortical sheet to minimize total length of their interconnections. Exhaustive search of all possible alternative layouts of even 20 components to verify

optimization would require resources of cosmic scale ($n! = 2.4 \times 10^{18}$ layouts, more than the number of seconds in the 20 billion year history of the Universe). However, if cortical components are placed to minimize interconnection lengths, one would expect to find quite computationally tractable statistical confirmation of the "adjacency rule:" If components a and b are interconnected, then they are positioned contiguous to each other, other things equal. And in fact, our connectivity and contiguity databases compiled for both macaque and cat show that each cortical layout strongly departs from random placement in favor of the adjacency rule ($p < .0001$).

(c) Layout of ganglia of C. elegans: A database giving for each of the 302 neurons of the nematode nervous system all known connections and their locations was compiled from the anatomy drawings of the Cambridge C. elegans group [e.g., Albertson and Thomson, 1976; White, et al., 1976; White, et al., 1986; Wood, 1988] and the worm computational neuroanatomy atlas of Achacoso and Yamamoto (1990). Figure 3 gives a synoptic view of all connections of the 275 ganglionic neurons ("PH," "AN,"..."LU" designate the ganglion-level components). It is probably the first complete depiction of a nervous system, at neuron-level detail, in a single, synoptic image. The adjacency rule is again strongly confirmed; however, a more striking finding can also be obtained. The ganglion-level optimization problem has 11 movable components, with 11! = 39,916,800 possible alternative orderings. A dozen microcomputers running in parallel for a few days were able to exhaustively search all of these placements, yielding the result that the worm's actual ganglion layout in fact requires less total length of connections than any of the other millions of possible layouts. In terms of methods, searches of this scale are unprecedented in computational anatomy.

(d) Individual neuron placement: There is also evidence that placement optimization is so sensitive that it applies even to positioning of neuron cell bodies in C. elegans. (i) The following modification of the above "wire-saving" adjacency rule can be tested: If neurons a and b are interconnected, then they are placed near each other, in particular, in the same ganglion, other things equal. Again, statistical confirmation is strong. (ii) Even positioning of somata within ganglia conforms to a component placement optimization prediction. Namely, there is a strongly significant trend for somata making exclusively anterior extra-ganglionic connections to be located in the anterior half of a ganglion, while neurons with external connections only to sites posterior to the ganglion tend to be placed in the posterior half of the ganglion.

Local Neuronal Steiner Tree

The hypothesis that neuron arbors form optimized Steiner trees was studied at the local level, for internodal branch junction geometry only [Cherniak, 1990 (Appendix 2)]. (See Figure 2b for 3 such Steiner junctions.) For conventional Steiner trees, where each branch costs the same per unit length, the angles formed by each pair of branches at these isolated Steiner junctions must be 120° [Courant and Robbins, 1969]. As an initial working hypothesis to be investigated, it was supposed that a neuron cell body plus its axonic or dendritic synapses constituted the node set to be most economically interconnected. One can immediately note that both axonic and dendritic arbors include internodal branch junctions, and so are

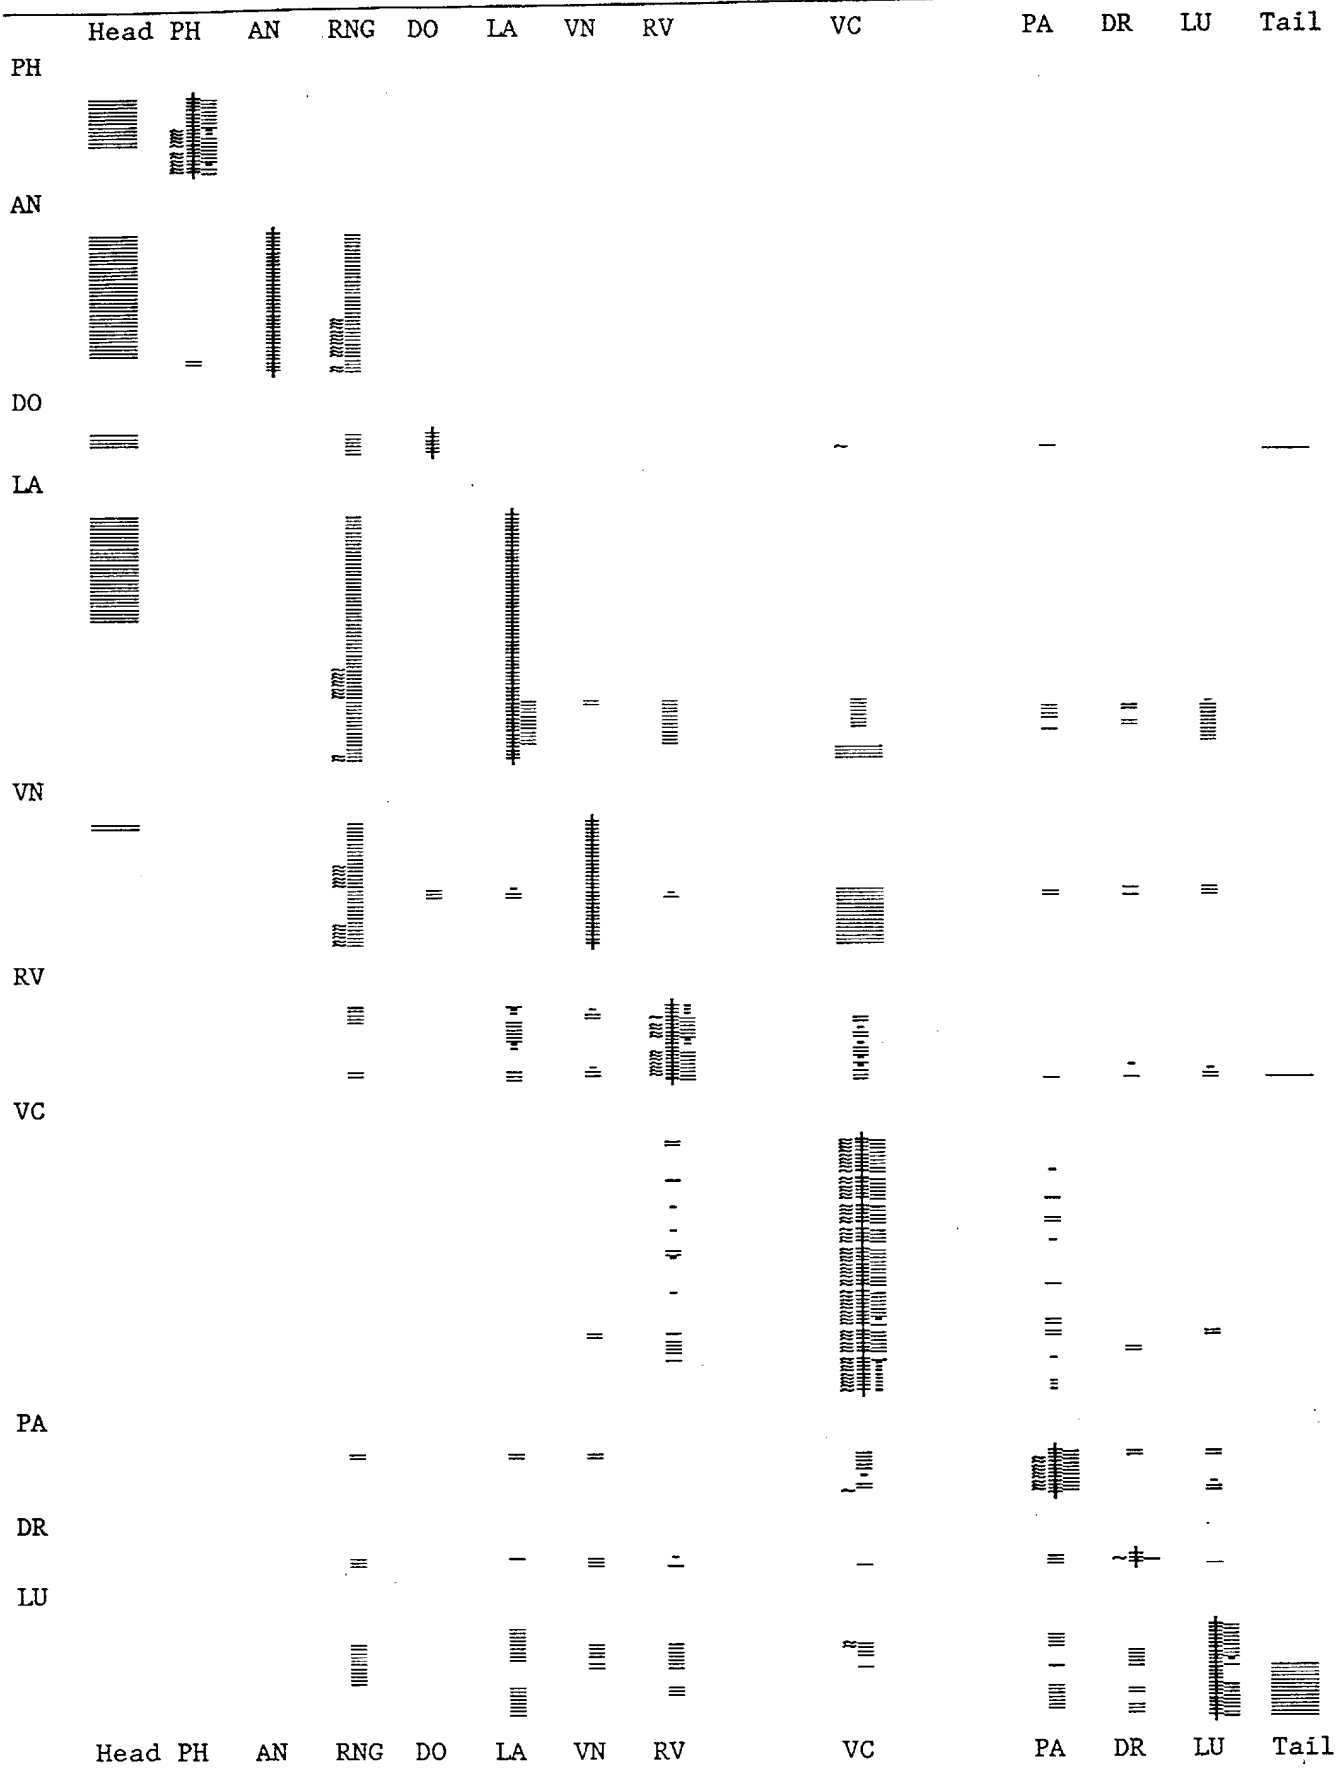

Figure 3. Ganglion-level connectivity map for C. elegans. Each row represents one neuron: +, soma; —, assymmetrical synapse; -, symmetrical gap synapse; ~, muscle connection; ——, sensor. From [Cherniak, 1991].

candidates for less costly Steiner trees, rather than minimal spanning trees.

However, these internodal junctions almost never have the 120° angles that are a signature, a necessary condition, of uniform branch-cost Steiner trees. Instead, throughout nature, for nonliving as well as living tree structures (e.g., river deltas as well as plant roots), there is a robust trend toward internodal junction angles of about 60°. Perhaps the simplest proposal to account for this is that neuronal trees do not have branches of constant cost. In some sense, a "trunk" (the branch nearer the soma) costs more per unit length than the other two branches of an internodal junction, hence the angle between the latter, cheaper branches is decreased below 120° in order to decrease the amount of the costlier trunk (see Figure 4).

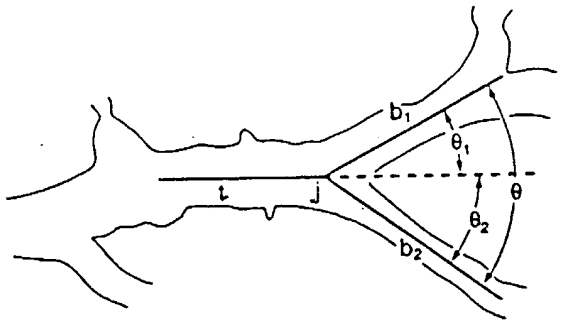

Figure 4. Bifurcating junction in a neuron arbor. $\underline{t}$, trunk; $\underline{b}_1$, $\underline{b}_2$, branches; j, internodal junction. Internal angle $\theta = \theta_1 + \theta_2$. The "Y-tree" diagram is superimposed upon a simplified outline of a junction in the dendritic arbor of a cat cortical stellate neuron.

A simple model can be derived that is a function from cost per unit length of each branch and trunk to the internal angle θ that minimizes cost of the internodal junction. Let $\underline{w}_0$ be the cost per unit length of trunk, and $\underline{w}_1$, $\underline{w}_2$ be the corresponding costs of the branches:

$$\cos\,\theta = \frac{\underline{w}_0^{\,2} - \underline{w}_1^{\,2} - \underline{w}_2^{\,2}}{2\underline{w}_1\underline{w}_2} \quad .$$

The next question is, what is the cost that is supposed to be minimized? Perhaps the simplest hypothesis to begin with is that the relevant cost is just total volume of the internodal junction structure; as with component placement optimization, this hypothesis of course ignores other possible roles of cell architecture, such as electrophysiological signal processing at the branch junctions.

Neuron junction data on branch diameters and angles was obtained principally by measurements from published camera lucida drawings, plus some calibration data sets from direct 3-d computer microscopy [cf. Nevin, 1989]. The data sets included both dendrite and axon arbors, both vertebrates and invertebrates, as well as being drawn

from a wide range of brain structures. Goodness of fit studies for local optimization models similar to the above one have been reported for arteries [Zamir, 1976] and river junction geomorphology [Roy, 1983]; in these studies, none of the alternative cost assumptions was supported significantly more strongly than the others. However, two sensitivity enhancements had been introduced into the present goodness of fit analysis: (a) The widest possible range of branch data was collected; if there in fact were a correlation between observed and predicted branch angles for each junction, a narrow range of data values would tend to miss it. (b) The mean value of the aggregated data for each related junction group, rather than the individual junction data, was used in the test for the observed-predicted correlation; this is a conventional strategy for coping with high levels of random variance, as are found in the junction data.

The overall mean difference between predicted junction angle and observed angle for over 200 neuronal junctions was quite small, below 1^{o}. In addition, correlation between mean predicted and observed angles for the set of arbor samples was strong, exceeding $\underline{r} = 0.7$. Alternative cost assumptions, such as that minimized cost = surface area of the Y-tree, consistently yielded mean "predicted vs observed" errors more than an order of magnitude greater than that for cost = volume. Thus, the local Steiner tree model appears to be a good predictor of junction anatomy.

Some clues concerning the mechanism of volume-minimization of the local neuron arbor Y-trees: Branch junctions of non-living tree structures, such as river deltas and electric discharge patterns, conform to the volume-minimization hypothesis about as well as neuron arbors; so the mechanism of neural optimization might involve simple physical processes like those presumably at work in non-living structures. In addition, the above cosine law for determining the cost-minimizing branch angle turns out to be identical with the "Triangle of Forces" law of vector mechanics, which gives the equilibrium angle of three interconnected cords pulled by three weights. Thus, the volume-minimizing angle at neuronal internodal junctions might be set by a similar "tug of war" process.

One rationale for so simple an anatomy-generating process can be seen in a dilemma nature confronts: Human brain wiring is among the most complex structures known in the universe, yet its layout information must pass through the "genomic bottleneck" of very limited DNA information-representation capacity [Cherniak, 1988]. The harmony of local neuroanatomy and vector mechanics suggested by the above account would lower this hereditary information load by giving the genome local network optimization automatically and comparatively cheaply directly from the physics.

Experimental Design

The above preliminary work suggests the network optimization hypothesis yields a research program worth further pursuit. The planned investigation continues to focus on the two network optimization concepts that were the subject of the earlier studies.

General Methods

The main technique of the contemplated research is again a series of computational experiments, employing a hierarchy of computational resources:

(a) A DOS-based 80486-chip workstation, to be purchased for the laboratory. In addition, the University of Maryland can provide several consecutive days of use of sets of 10 to 20 such machines for simple types of parallel execution.

(b) A Sun engineering-type workstation, with UNIX programming environment. Time on this machine to be obtained from University of Maryland.

(c) Time on a VAX mini-mainframe. Available for purchase from University of Maryland.

(d) Time on the San Diego Cray supercomputer system. University of Maryland has a direct satellite hookup to this machine, and an arrangement including grants of time blocks on it.

(e) Some time on the University of Maryland's Connection Machine, a massively parallel computer with 16,000 separate processors.

Our basic "minimum computational force necessary" strategy in choice of computational environment for a given task entails, in particular, tying up as little funding as possible in acquisition of hardware (which rapidly obsolesces, requires costly maintenance regimes, etc.). A typical data stream for Part I below: From published diagrams, input is obtained by tracing tablet, simple image-handling software is then applied to extract dimensions of interest, and results are encoded in a low computational overhead database via, e.g, a commercially available spreadsheet package such as Lotus 3. The exhaustive searches of both the earlier and some of the planned studies involve perhaps unusually long computer runs, but on conventional microcomputer hardware, with quite small-sized programs. For instance, as described above, the full-scale search of worm ganglion layouts required about two days when executed in parallel on a dozen 80386 machines; our estimates suggested one such search run would have required at least one hour on the San Diego X-MP/48 Cray , not a dramatically more desirable use of computational resources. Some of the studies below, however, specifically require Cray or parallel Connection Machine time. As in the past, all source code developed by this project will be made publically available.

I. Neural Component Placement

This study is a straightforward, incremental continuation of our earlier work, and so is the most natural starting point for the planned research.

A. C. elegans

The first project is extension and refinement of our C. elegans

neuroanatomical database. To begin with, the "classical" worm anatomy summarized in Wood (1988) has recently been augmented by independent projects such as Hall and Russell (1991), and Schatz (1991) which need to be incorporated into our database. In addition, the structure of our database ought to be upgraded: The present data matrix for component placement evaluation is all 302 neurons x 11 locatable ganglion-level components; Figure 3, with some compression of the vertical axis, includes most of this 302 x 11 matrix. The finest-scale connection matrix would be 302 x 302, that is, representing every connection from a neuron to a specific neuron (rather than only connections from a neuron to a ganglion site). Also, more sensitive cost-measures than just length of fibers should be tested; the most promising weighting would be the number of distinct synapses that a fiber makes to a given neuron. This information is accessible in the published anatomy (e.g., White, et al., 1986), and could be entered in each cell of the full 302 x 302 matrix, instead of just in effect "yes" or "no" with regard to a possible connection. Finally, the specific physical location of each individual neuron should be obtained, from scanned images of the anatomy diagrams; our earlier studies have chiefly used a simpler "quasi-black box" physical model of each ganglion with siting of its cell bodies only specified as in either the anterior or posterior half of the ganglion. At least initial work with this database can be continued on a conventional high-end DOS machine.

With this more realistic model of connections and component placement in the worm nervous system, the prior finding of optimality of the actual ganglion layout can be checked. In addition, the entire search space of 40,000,000 possible layouts should be more closely characterized (e.g., in terms of cost-distribution of runner-up layouts). However, even with each layout and its total cost specified in a single alphanumeric line, storage of all layouts requires one gigabyte; analysis of just some more interesting regions of the total search space would still by itself require over 100 megabytes of hard disk storage. Finally, with the fine-scale database, it would be feasible to verify optimal placement of individual neurons within medium-sized ganglia by exhaustive search of all layouts, thereby extending the earlier statistical analyses of this level of component placement.

C. elegans neurons tend to make connections within fiber bundles [White, et al., 1983]. Another relevant network optimization problem encountered in VLSI circuit design is "wire routing"--in effect, a component placement problem involving the minimization of costs of positioning of interconnecting wires. Thus, a question of "fiber routing" optimization in neuropil bundles can be posed. In particular, a major proportion of all interconnections in the C. elegans nervous system are made in the ring, a kind of white-matter crossbar network in the anterior portion of the nervous system; some analysis of component placement in this structure should also be tractable.

B. Cerebral Cortex

Our connectivity database for cerebral cortex is restricted to cortical areas implicated in vision, and was constructed principally

from Van Essen (1985) for macaque and from Rosenquist (1985) for cat. A similar analysis for all cortical areas--auditory, somatosensory, olfactory, motor, association--appears now to be within reach. For instance, for macaque, Pandya and Yeterian (1985) with Felleman and Van Essen (1991) constitute sources for a good starting point. The developing Neural Circuitry Database initiative suggests more cortical connectivity data should be soon available. (It should be noted that exhaustive search of all possible layouts in the total cortex problem is out of the question; for example, for 50 Brodman-type areas, there are $50! = 3 \times 10^{64}$ alternative placements, far more than the number of attoseconds (10^{-18} sec) since the Big Bang.) Our database should also be upgraded to encode some measure of strength or density of an intracortical connection; a rule related to the adjacency rule discussed earlier that is worth investigation is, "If a and b are strongly interconnected, then a and b are more likely to be contiguous."

Of course, similar component-placement studies should then be pursued for other vertebrates and invertebrates where there is sufficiently complete neuroanatomy. Rat cerebral cortex (which has been proposed for inclusion in the Neural Circuitry Database [Pechura and Martin, 1991]) is a promising next candidate. Anticipated timetable for all of the above component placement studies is about two years.

II. Global Neuronal Steiner Trees

Our initial studies found good volume-minimization of internodal branch junctions in dendrite and axon arbors. These exploratory studies sampled as wide a variety of arbors--in terms of brain region, type of animal, and developmental stage--as possible. The natural follow-up project is an in-depth investigation of a large, uniform data set of state of the art quality; R. Nevin (Berkeley) has obtained such a data set for us, by 3-d computer microscopy, for dendritic arbors of cricket interneuron and of rat hippocampal granule cell. It should be emphasized, as mentioned earlier, that, because of the high levels of variance associated with the model, a narrowly focussed, homogeneous data set presents special problems for analysis. However, consultation with statisticians P. Hanges and S. Grimshaw (College Park) has suggested an approach that deals with these "noise-swamping" problems.

A deeper issue is that this local optimization is not sufficient for solution of the complete, global Steiner tree problem. The combinatorial question of where to place the internodal junctions (e.g., see Figure 2b) must also be answered. As mentioned earlier, the full-scale global Steiner tree problem is in the category of computational complexity known as "NP-hard;" no exact solution procedure is known, nor expected to exist, that does not require exhaustive search of all possible trees interconnecting a given node set. Such searches are not currently feasible for node sets much larger than around 30, while a typical neuron has thousands of synapses. So, even if some types of neuron arbors were forming perfect Steiner trees, there would be no way to directly verify this.

Certainly many kinds of neuron arbors do not seem to be even

candidates for Steiner trees. As a typical example, in cat cerebral cortex some intrinsic axons show double hairpin turns [Winer and Morest, 1983]. One possibility is that, with component placement and local Steiner tree so finely optimized, not globally optimizing neuron arbors (particularly axonic ones) provides the degree of freedom, the lack of constraint, needed for neural wiring to be able to satisfy the functional requirements of an operating nervous system.

In any case, very small arbors of some prima facie promising types of dendritic arbors can be directly checked for Steiner optimality by just constructing the Steiner tree for the node set. In addition, various indirect strategies can be pursued: As one example, Du and Hwang (1990) recently proved that, for a node set on the plane, length of its Steiner tree must be greater than or equal to $\sqrt{3} / 2$ of the length of its minimal spanning tree. It is also conjectured that length of the Steiner tree may be equal to this value. Thus, for a neuron arbor that is approximately planar, we can construct the corresponding minimal spanning tree (a quite tractable problem, as mentioned earlier), obtain its length, and then apply the Hwang and Du estimator to determine a lower bound on the Steiner tree length. Actual length of the arbor can then be measured, to compare how close it approaches this lower bound--and, indeed, whether the neuron arbor is in fact any shorter than even its minimal spanning tree. (One manageable complication is that the Hwang and Du estimator, like most Steiner tree results, is for trees of uniform branch cost; as discussed above, our preliminary work established that neuron arbors have varying branch costs.) Our first studies will be of retinal ganglion cell dendrites, which in fact approximate planar trees relatively well. Other computationally tractable Steiner tree length-estimators will also be investigated.

Applicability to neural structure and function of a number of other related tree-optimization concepts should also be explored, such as shortest path through a network, and flow optimization through a network [Garey and Johnson, 1979]. Timetable for the Steiner tree studies will require at least one year.

III. Biological Mechanisms

Finding phenomena of network optimization in the brain immediately raises questions concerning the actual processes by which the optimization occurs.

A. Component Placement

Even our preliminary studies raised the possibility that neural component placement optimization may be very precise; for example, the second-best layout of C. elegans ganglia--after the actual layout--required well under 1% more total extra-ganglionic connecting fiber. Yet, a "perfect" blind trial and error evolutionary process of exhaustive search for even the 50-element cortical area problem is not remotely feasible, even with major parallelization; the calculations earlier indicate that such a phylogenetic mechanism would exceed resources of cosmic scale. Instead, presumably evolution "blindly" selects quick but dirty heuristic procedures for design of good brain

layouts. That is, these layout optimizing procedures involve a tradeoff of guaranteed exact but computationally intractable solutions for much faster probabilistic/approximate results--a familiar strategy in software engineering [Cherniak, 1986, ch. 4]. Thus, one question is, what are plausible, biologically feasible candidates for such quick but dirty layout processes? The range of candidates is constrained by the fact that the mechanisms must be quick but, as noted, apparently cannot be very dirty.

One of the categories of process that is most worthy of investigation involves "tug of war" vector-mechanical phenomena. As sketched earlier, our preliminary studies identified such a process as a good candidate for the actual mechanism of local Steiner optimization of neuron arbors. An identical process suffices for efficient solution of the 1-component placement problem (analog devices of this type were in fact employed at the turn of the century [Francis and White, 1974]). Thus, phylogenetically, for the simple brain-placement problem, each sensory and motor nerve fiber could be represented as behaving over generations in effect like a spring anchored to a fixed sensor or muscle site and attached to a movable brain.

However, this type of process cannot suffice for multiple-component problems; it will tend to become trapped in highly non-optimal "local minimum" solutions. An extension of the tug of war idea that is used in VLSI layout is the "mesh of springs" concept: Each interconnection in a 2-dimensional array of circuit elements is again represented as a spring. However, the system is not taken to equilibrium in a single step; rather, energy minimization proceeds by a relaxation method, in small stages. (See Kuh and Ohtsuki (1990) for a review of this procedure implemented as a resistive network.)

The mesh of springs method, and any other feasible approach, is only a quick and dirty procedure--it is guaranteed to fail, the only questions are how often, when, and how badly. Unfortunately, performance evaluation of quick and dirty heuristics is particularly difficult conceptually, because probability distribution of the bad problem-instances is hard to formally define [Cherniak, 1988].

We plan a series of simulations of a mesh of springs mechanism for the worm ganglion and cell placement problems, beginning with pilot studies on conventional DOS-environment machines. The available parallelism of the 16,000 processor Connection Machine should then be especially appropriate for realistic modelling of such a layout heuristic, e.g., with an ensemble of processors representing each neuron and a separate processor for each connection of a given neuron. One question the simulations will focus upon concerns interpretation of the adjacency rules discussed above. The chicken-egg question of direction of causation arises: Does connection between two neural components tend to lead to contiguous placement of them, or vice versa, or do both coevolve together, or are there different causal directions in different cases?

B. Global Arbor Generation

What are some neurally feasible candidates for mechanisms for growth of full neuron arbors? For a sense of the constraint of neural feasibility, suppose contrary to the earlier discussion that neurons were in fact just constructing arbors that were minimal spanning trees (see Figure 2a). One of the most elegant, and quite computationally cheap, procedures for generating minimal spanning trees is the "nearest neighbor" algorithm: Choose an arbitrary node from the set to be interconnected (e.g., the cell body site), and connect it to the node that is its nearest neighbor; now, connect the unattached node that is nearest an attached node to the latter node; repeat this step until no unattached nodes remain [Hillier and Lieberman, 1986]. The question is how this type of process would be ontogenetically realizable. For, it requires in effect much message-passing between the connected nodes at any stage of development to determine which free node is currently the nearest neighbor; one needs to envisage how sufficient, and sufficiently rapid, information-transmission of this type could actually occur on an embryological time scale.

If in fact some kinds of neuron arbors turn out to be Steiner trees, the argument of the previous section again applies: the Steiner tree problem is in the same category of computational complexity as component placement optimization, hence the most plausible mechanisms of arbor generation will be quick and dirty probabilistic/approximation procedures. Perhaps the first category of tree-construction heuristics to investigate includes a "simulated annealing" process [Kirkpatrick, et al., 1983]. The idea involved in this type of search procedure was inspired by simulations of solid-state thermodynamic phenomena; it involves random, stochastic adjustment of solutions according to a "temperature schedule"--an initial "hot" stage where large-scale random modifications are introduced, followed by "warm" stages where only smaller-scale refinements are generated, terminating with "quenching" and the solution frozen in place. Procedures based upon the simulated annealing metaphor have been applied to network optimization (e.g., the Travelling Salesman problem), although issues of performance evaluation along the lines sketched above have made their effectiveness a matter of debate.

The most natural implementation of simulated annealing as a tree-generation heuristic would involve interpreting temperature as probability of making longer, costlier connections. That is "hot" would here mean, make many such connections, while "warm" would mean, make and/or change fewer, cheaper connections. The somewhat paradoxical phenomenon of multilevel die-back in the normal development of the nervous system has been frequently noted for, e.g.: cells [Williams and Herrup, 1988], dendrite branches [Wassle, 1988], axon branches [LaMantia and Rakic, 1990], and synapses [Zevcevic, et al., 1989]. One familiar interpretation of it has been that it represents a kind of fine-tuning of the anatomy--here, in particular, a pruning back of neuron arbors. In terms of a simulated annealing concept, die-back would represent "cooling down" of the solution process; in this way, a simulated annealing arbor-refining procedure has some prima facie appearance of neural realizability. A network-

optimization paradigm is of course an alternative to conventional approaches to neural development [e.g., Purves and Lichtman, 1985], but may nonetheless be compatible and complementary with the latter approaches. For example, the above simulated annealing process would mesh with a trophic account where fiber growth tips "home" on appropriate targets when they come within a critical recognition range.

Pilot simulations of the temperature-schedule arbor growing procedure would begin on conventional DOS-based machines. In fact, we have already developed THERMO, an instructional software environment for simple simulated annealing experiments that is presently in pilot distribution [Appendix 4]. Again, the Connection Machine architecture is a natural choice for more realistically parallel simulations of the process, for example, with each internodal junction and its sprouting daughter branches represented by a set of processors.

Another source of clues regarding the actual mechanisms of neuronal arbor generation will be benchmark comparisons of how well natural non-neural trees approximate global optimality--in particular, river deltas and drainage basins (with exclusively non-living processes), and plant arbors.

A timetable for obtaining adequate concrete results concerning biological mechanisms of neural network optimization is, of course, the most difficult to specify, but about two years is anticipated for these last studies.

Literature Cited

Achacoso, T., and Yamamoto, W. (1990). The AY Neuroanatomic Atlas of C. Elegans for Computation. (Washington, D.C.: George Washington University Medical School).

Albertson, D., and Thomson, J. (1976). The pharynx of Caenorhabditis elegans. Phil. Trans. R. Soc. Lond. B 275: 299-325.

Bern, M. and Graham, R. (1989). The shortest-network problem. Scientific American 260: 84-89.

Blinkov, S., and Glezer, I. (1968). The Human Brain in Figures and Tables: A Quantitative Handbook. (New York: Plenum).

Cajal, S. Ramon y (1972). Histologie du Systeme Nerveux de l'homme et des Vertebres, vols. I & II. (Madrid: Consejo Superior de Investigaciones Cientificas).

Cherniak, C. (1986). Minimal Rationality. (Cambridge, Mass.: MIT Press).

Cherniak, C. (1988). Undebuggability and cognitive science. Communications of the Association for Computing Machinery 31: 402-412.

Cherniak, C. (1990a). Local optimization of neuron arbors. University of Maryland Institute for Advanced Computer Studies Technical Report, No. 90-90. [= Appendix 2]

Cherniak, C. (1990b). The bounded brain: Toward quantitative anatomy. Journal of Cognitive Neuroscience 2: 58-68. [= Appendix 3]

Cherniak, C. (1991). Component placement optimization in the brain. University of Maryland Institute for Advanced Computer Studies Technical Report, No. 91-98. [= Appendix 1]

Cognitive Science (1985). 9.

Courant, R., and Robbins, H. (1969). What is Mathematics? (New York: Oxford University Press).

Du, D., and Hwang, F. (1990). A proof of Gilbert-Pollak conjecture on Steiner ratio. Foundations of Computer Science (FOCS) Proceedings.

Felleman, D., and Van Essen, D. (1991). Distributed hierarchical processing in the primate cerebral cortex. Cerebral Cortex 1: 1-47.

Francis, R., and White, J. (1974). Facility Layout and Location. (Englewood Cliffs, N.J.: Prentice-Hall).

Garey, M., and Johnson, D. (1979). Computers and Intractability: A Guide to the Theory of NP-Completeness. (San Francisco: W.H. Freeman).

Hall, D., and Russell, R. (1991). The posterior nervous system of the nematode Caenorhabditis elegans: Serial reconstruction of identified neurons and complete pattern of synaptic interactions. Journal of Neuroscience 11: 1-22.

Hillier, F., and Lieberman, G. (1986). Introduction to Operations Research, 4th ed. (Oakland, Calif.: Holden-Day).

Hillis, D. (1985). The Connection Machine. (Cambridge, Mass.: MIT Press).

Kirkpatrick, S., Gelatt, C., and Vecchi, M. (1983). Optimization by simulated annealing. Science 220: 671-680.

Kuh, E., and Ohtsuki, T. (1990). Recent advances in VLSI layout. Proceedings of the IEEE 78: 237-263.

LaMantia, A., and Rakic, P. (1990). Axon overproduction and elimination in the corpus callosum of the developing rhesus monkey. Journal of Neuroscience 10: 2156-2175.

Lewis, H., and Papadimitriou, C. (1978). The efficiency of algorithms. Scientific American 238: 96-109.

Nevin, R. (1989). Morphological analysis of neurons in the cricket cercal system. Ph.D. dissertation. University of California, Berkeley.

Pandya, D., and Yeterian, E. (1985). Architecture and connections of cortical association areas. In A. Peters and E. Jones, eds. Cerebral Cortex, vol. 4. (New York: Plenum).

Pechura, C., and Martin, J., eds. (1991). Mapping the Brain and its Functions: Integrating Enabling Technologies into Neuroscience Research. (Washington, D.C.: National Academy Press).

Purves, D., and Lichtman, J. (1985). Principles of Neural Development. (Sunderland, Mass.: Sinauer).

Rosenquist, A. (1985). Connections of visual cortical areas in the cat. In A. Peters and E. Jones, eds. Cerebral Cortex, vol. 3. (New York: Plenum).

Roy, A. (1983). Optimal angular geometry models of river branching. Geographical Analysis 15: 87-96.

Rumelhart, D., and McClelland, J., eds. (1986). Parallel Distributed Processing, vols. I & II. (Cambridge, Mass.: MIT Press).

Schatz, B. (1991). Building an electronic scientific community. IEEE Proceedings, 24th International Conference on System Sciences 3: 739-748.

Stockmeyer, L., and Chandra, A. (1979). Intrinsically difficult

problems. Scientific American 240: 140-159.

Thompson, D. (1961). On Growth and Form. (New York: Cambridge University Press).

Van Essen, D. (1985). Functional organization of primate visual cortex. In A. Peters and E. Jones, eds. Cerebral Cortex, vol. 3. (New York: Plenum).

Wassle, H. (1988). Dendritic maturation of retinal ganglion cells. Trends in Neuroscience 11: 87-89.

White, J., Southgate, E., Thomson, J., and Brenner, S. (1976). The structure of the ventral cord of Caenorhabditis elegans. Phil. Trans. R. Soc. Lond. B. 275: 327-348.

White, J., Southgate, E., Thomson, J., and Brenner, S. (1983). Factors that determine connectivity in the nervous system of Caenorhabditis elegans. Cold Spring Harbor Symposia on Quantitative Biology 48: 633-640.

Williams, R., and Herrup, K. (1988). The control of neuron number. Annual Review of Neuroscience 11: 423-453.

Winter, P. (1987). Steiner problem in networks: A survey. Networks 17: 129-167.

White, J., Southgate, E., Thomson, J., and Brenner, S. (1986). The structure of the nervous system of the nematode Caenorhabditis elegans. Phil. Trans. R. Soc. Lond. B 314: 1-340.

Wood, W. (1988). The Nematode Caenorhabditis Elegans. (Cold Spring Harbor, N.Y.: Cold Spring Harbor Laboratory).

Zamir, M. (1976). Optimality principles in arterial branching. Journal of Theoretical Biology 62: 227-251.

Zevcevic, N., Bourgeois, J., and Rakic, P. (1989). Changes in synaptic density in motor cortex of rhesus monkey during fetal and postnatal life. Developmental Brain Research 50: 11-32.