Optimal-Wiring Models of Neuroanatomy

Christopher Cherniak, Zekeria Mokhtarzada,
and Uri Nodelman

ABSTRACT

Combinatorial network optimization appears to fit well as a model of brain structure: connections in the brain are a critically constrained resource, hence their deployment in a wide range of cases is finely optimized to "save wire". This review focuses on minimization of large-scale costs, such as total volume for mammal dendrite and axon arbors and total wirelength for positioning of connected neural components such as roundworm ganglia (and also mammal cortex areas). Phenomena of good optimization raise questions about mechanisms for their achievement: the examples of optimized neuroanatomy here turn out to include candidates for some of the most complex biological structures known to be derivable purely from simple physical energy minimization processes. Part of the functional role of such fine-tuned wiring optimization may be as a compact strategy for generating self-organizing complex neuroanatomical systems.

4.1. INTRODUCTION

How well can combinatorial network optimization theory predict structure of invertebrate and vertebrate nervous systems? The working hypothesis explored here is that brain connections are singularly limited, 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 of scarcity for expressing and solving problems of "saving wire." The question then is, how well do such concepts in fact apply to the brain? The main technique of these studies is computational experiments, the main hurdle the exponentially exploding computational requirements of optimization searches to evaluate connection-minimization of the neuroanatomy.

Good optimization findings focus attention upon possible biological mechanisms. Network optimization problems are among the most computationally intractable known; in general, only an exhaustive search of all possibilities can guarantee exact solutions. However, some "quick but dirty" probabilistic/approximation procedures

developed for microcircuit design suggest candidate models for biological mechanisms of neuroanatomy optimization. In particular, we report positive results for neural optimization via genetic algorithms and via vector mechanical energy-minimization simulations. In fact, the latter models constitute an instance of self-organizing morphogenesis of highly complex biological structure directly from simple physical processes.

4.2. CONCEPTUAL BACKGROUND

The theoretical framework of this work grew out of methodological studies of prevailing models of the agent in microeconomic, game, and decision theory *(1,2)*. The basic finding was that these models typically presupposed agents 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 that were intended to be more neurally realistic than conventional von Neumann computational architecture (cf *[3]*); again, the models tended to overestimate available resources drastically—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 then pervaded even some neuroanatomy *(4)*.

Thus, a bounded-resource philosophical critique of mind-brain science ("We do not have God's brain") focused attention on neural connections as a critically constrained neurocomputational resource. Through combinatorial network optimization theory, 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., *[5]*). In fact, minimum wiring interpretations of neuroanatomy can be traced back at least as far as Cajal's qualitative "laws of protoplasmic economy" *(6,7)* and have continued to receive attention (e.g., *[8]*).

The human brain is commonly regarded as the most complex physical structure known in the universe. In the face of such overwhelming intricacy, neuroanatomy 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 may provide a source for a "generative grammar" of the nervous system, some general principles that compactly characterize aspects of neuroanatomy. Of course, connection minimization is unlikely to be ubiquitous in the nervous system; indeed, given the many other competing desiderata driving design of a brain, the striking observation is that it should hold in even some conditions. The question, then, is characterizing where "save wire" does and does not apply.

For example, in the *Caenorhabditis elegans* ganglia case sketched below in Subheading 4.4.2., we reduced approximately one thousand pages of published anatomy diagrams *(9a-c)* to a 100-page database, which, in turn, was represented as a 10-page

connectivity matrix (see Fig. 2 in *[10]*), which we then computationally verified to conform to connection-minimizing component placement optimization better than any of the nearly 40 million alternative possible layouts (see Fig. 1). If these types of result are confirmed, they constitute a predictive success story of recent quantitative neuroanatomy.

4.3. NETWORK OPTIMIZATION THEORY

The theory of NP-completeness emerged around 1972 *(11,11a)*; the key formal concept of a computational problem being NP-complete (nondeterministic polynomial-time complete) 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. For example, a 50 component system would have 50! possible alternative layouts, which is far more than the number of picoseconds since the Big Bang 20 billion years ago. Many of the most important real-world network optimization problems (e.g., the best known, Traveling Salesman) have been proven to be NP-complete or worse in computational complexity. Steiner tree and component placement optimization, problems examined here, are of this type, having been proven to be "NP-hard".

Steiner tree has been studied in its simplest form at least since the Renaissance *(12,13)*. The most relevant version of the problem is: given a set of fixed node loci, find the set of arcs (or branch segments) between those loci that interconnects all loci and has shortest total length. The resulting 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 the length of the minimal spanning tree for the nodes. For example, Figures 2A and B show, respectively, a minimal spanning tree and a Steiner tree for five nodes on a plane. The Steiner tree is about 4% shorter than the minimal spanning tree.

Since Steiner tree is a member of the class of NP-hard problems, it is not surprising that the largest unconstrained Steiner tree problems that can currently be solved have only approx 100 nodes (cf *[14]*). However, while minimal spanning trees are equal to or longer than corresponding Steiner trees, they are not at all computationally intractable; exact algorithms for them 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? The key idea needed for such applicability is that for real-world trees, living and nonliving, not all segments are equal: the concept of an optimal tree had to be extended to include variably weighted branches and trunks *(15)*.

Component placement optimization ("quadratic assignment problem"), the other wiring problem focused upon here, has received the most attention in computer science in connection with design of large scale integrated circuits *(16, 16a)*. The problem can be defined as: given the connections among a set of components, find the spatial layout
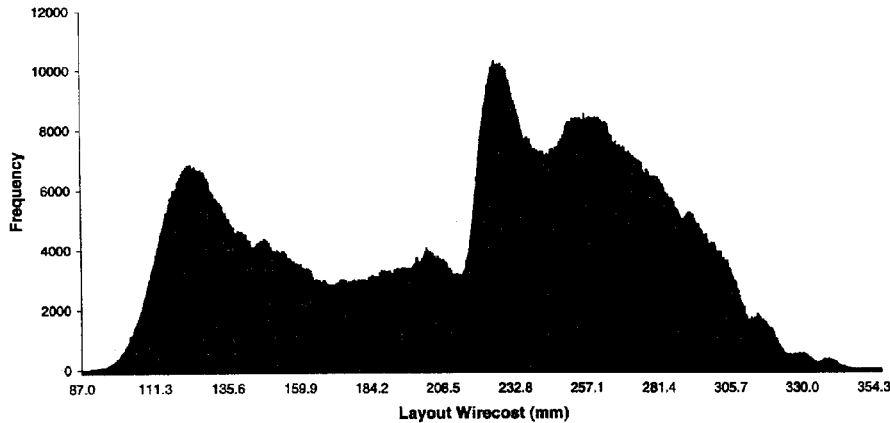
Fig. 1. Distribution of wirecosts (total wirelength) of all possible layouts of ganglia of *C. elegans*. A 10,000-bin histogram compiled from exhaustive search of all 39,916,800 alternative orderings of the 11 ganglia. Least costly and most costly layouts are rarest. In effect, the search approximates a simulation of the maximal possible history of the evolution of this aspect of the nervous system. The worm's actual layout (Fig. 2 in *[10]*) is in fact the optimal one, requiring the least total length of connecting fiber of any of the millions of possible layouts. For comparison, the last-place, "pessimal" layout would require about 4 times as much total connection fiber as the optimal one. (See *[7]*).

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 3A and B diagram two of six possible configurations of components 1, 2, and 3 in slots A, B, and C; for the connections among the components, placement 3a requires the most total connection length and 3b the least.

Again, computation costs for the exact solution of component placement optimization problems are of a magnitude not encountered in most bioscience computing, outside of gene-sequencing, and constitute one of the principal technical impediments of this research. For n components, the number of alternative possible placements is $n!$ (Size of this search space is generally unaffected by whether permissible component positions are located in 3, 2, or 1 dimensions.) Heuristic procedures that yield approx 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.

Perhaps the most salient and daunting feature of nontrivial global optimization problems is the presence of local minima traps on the optimization landscape—that is, parameter values that yield least costs within a subregion of the search space, but not across the total space. For example, with regard to vector mechanical force-minimization treatments of the above two problems: (*i*) the dendritic tree of Figure 4C is suboptimal because of its topology, while Figure 4D shows the minimum cost topology; no vector mechanical tug of war re-embedding the suboptimal topology can ever transform it into the best topology; and (*ii*) similarly, Figure 5 shows a vector mechanical
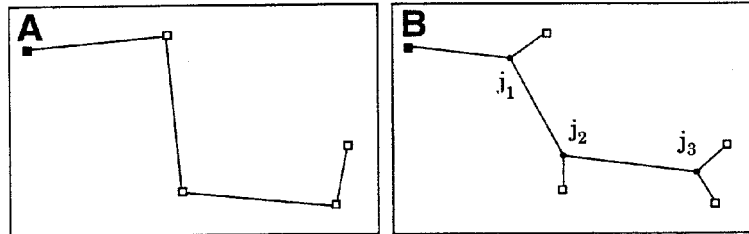
Fig. 2. Minimal spanning tree (**A**) and Steiner tree (**B**) for five nodes on a plane surface. The Steiner tree is shorter, but much more computationally costly to construct. Adapted from *(20)*.

local minimum trap for the roundworm ganglion component placement problem. The extensive modeling of cellular structures and processes in terms of compression–tension "tensegrity" by Ingber (e.g., *[17,17a]*) does not deal with local minima and, therefore, cannot account for such global optimization via evading such traps. Correspondingly, Van Essen's tension-based model of cortical folding, in terms of white matter tensegrity, also does not deal with local minimum traps and, so, will not suffice for global optimization problems of wiring minimization *(18–18b)*.

4.4. OPTIMIZATION MECHANISMS

From evaluating how well neural structures conform to minimum wiring principles of economical use of connections, we have gone on to seek biological mechanisms of the observed extremely fine network optimization. The emerging picture is that, corresponding to "Save wire" neuroanatomy optimization results, we have found neuroanatomical candidates for some of the most complex biological structures shown to be derivable purely from simple physical processes (cf *[19]*). This constitutes a further stage in developing an understanding of the generative rules that yield the highly complex anatomy of the nervous system.

4.4.1. Large-Scale Optimization of Dendrites and Axons

Some complex neuron arbor structure seems to be self-organizing, with no need of evolutionary mechanisms for its creation. While the key underlying pattern is network optimization, "Save wire" (in particular here, minimize total volume), the specific hypothesis in this case is, neuron arbor morphogenesis behaves like flowing water (see *[20]*). The volume minimizing fluid dynamic model yields two confirmed results: (*i*) it predicts diameters of branches at junctions; and (*ii*) from those diameters, branching angles and junction loci can then be predicted. The major methodological enterprise of the project centered on developing STRETCH—a package of algorithms for the computationally intractable (NP-hard) task of generating optimal trees *(13)* with variable branch-weights—as the gold standard against which to compare observed neuroanatomical trees.

Neural fluid mechanics is a simple fluid-dynamical model, for minimized walldrag of pumped flow through a system of pipes, that will predict branch diameters of some types of dendrites (e.g., of mammalian retinal ganglion cells) and axons (e.g., in rodent
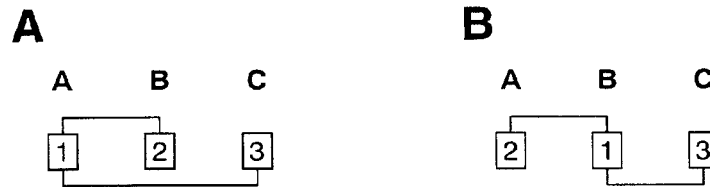
Fig. 3. Component placement optimization: two alternative placements of elements 1, 2, and 3 in positions A, B, and C. For the given interconnections, placement (**A**) has greater total connection length than placement (**B**). Adapted from *(7)*.

thalamus) almost as well as it predicts configuration of nonliving tree structures such as river drainage networks. For neurons, the fluid dynamics falls in the laminar-flow vs turbulent regime.

Waterflow in branching networks, in turn, acts like a tree composed of a weights-pulleys-cords system (non-Hooke's Law), that is, vector mechanically; so also do the neuron arbors. As a result, the trees globally minimize their total volume to about 5% of optimum for interconnecting their terminals (see Fig. 4). One unanticipated moral that emerges is that, in a sense, "Topology does not matter", that is, the worst or "pessimal" connection pattern typically costs only relatively little more than the optimal pattern, compared to the wide corresponding possible range of costs for embedding a given topology. The conclusion here is only that this minimum volume configuration is the default neuron arbor structure, probably often modified in many complex ways (cf, e.g., *[21]*).

4.4.2. C. elegans *Ganglion Placement Optimization*

We have extended the above results on large-scale optimization of individual neuron arbors to the entire *C. elegans* nervous system. The basic picture is indeed that vector mechanics suffices for optimization of placement of the ganglia of *C. elegans*. As mentioned earlier, our prior research had found that the actual placement of the ganglia in the worm was optimal, in that it required the least total length for the animal's (approx 1000) interconnections, out of roughly 40 million alternative possible ganglion orderings (see Fig. 1). We had also reported a related set of optimization results for rat, cat, and macaque cortex, in terms of placement of connected Brodmann areas that conforms to an "Adjacency Rule" *(7,10,22,23)*. (Nonetheless, as we have noted *(22)*, the majority of connections in the actual worm are not to nearby components; therefore, merely positioning components so their connections tend to go to nearby components will not in itself suffice to yield the minimum wirecost layout.) As mentioned, if this 1-in-10-million type of result is replicated, it begins to approach some of the most precise confirmed neuroanatomy predictions (see *[24, 24a]*); hence, we sought convergent support by finding feasible mechanisms for such fine-grained optimization.

We have constructed Tensarama, a force-directed placement simulator, in which each of the worm's connections behaves like a microweight-and-pulley system (Fig. 5 shows a screendump). Analog hardware devices of this type have been used to solve simple (noncombinatorial) placement optimization problems for over a century *(25, 25a)*. (Similar "mesh of springs" simulations have become a focus of current modeling
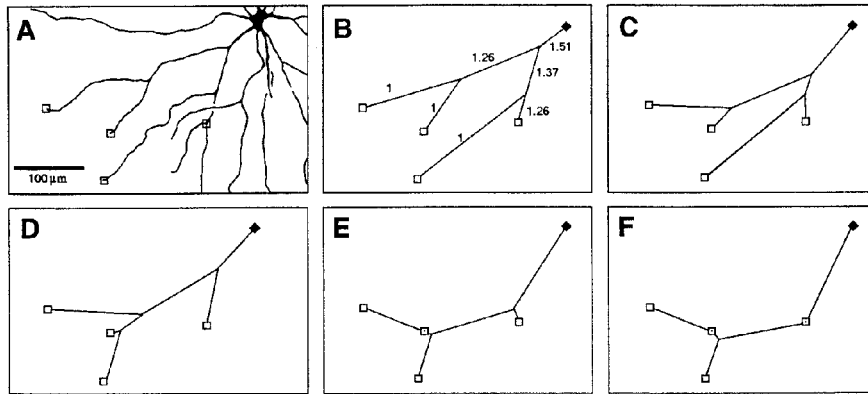
Fig. 4. Optimization analysis of a 5-terminal subtree from dendritic arbor of an α ganglion cell in rabbit retina. (**A**) A quadrant of the original camera lucida drawing containing the subtree (after *[35]*); soma is in upper right corner. "Leaf terminals" of the analysis are boxed (note that one of them is not a branch termination); "root-terminal" is at soma. (**B**) Wireframe representation of actual tree, with branch segments straightened between loci of terminals and internodal junctions. The labels give diameters assigned to the branch segments via the power law of the laminar-flow model. (**C**) Optimal re-embedding of the topology of the actual tree, with respect to total volume cost, via the STRETCH algorithm; this minimum volume embedding of the actual topology is 1.06% cheaper than the volume of the actual tree in panel **B**. (**D**) Optimal embedding of the optimal topology for the given terminal loci, with respect to volume-cost. It can be seen to differ from the actual topology of panels **A–C**; it is only 2.64% cheaper in volume than the actual topology in its actual embedding, in panel **B**. (**E**) Optimal embedding of the optimal topology, with respect instead to total tree surface area; actual vs optimal error is now 27.22%, much greater. (**F**) Optimal embedding of the optimal topology, with respect to total tree length; actual vs optimal error is now 60.58%, even greater. Thus, this dendritic arbor best fits a minimum volume model. Adapted from *(20)*.

of protein folding of amino acid chains *[26]*.) Over a wide range of initial input configurations of the ganglia, our vector mechanical net outputs the actual layout via tug-of-war, converging upon equilibrium at the actual minimum wirecost positioning of the ganglia—without major susceptibility to local minima traps. We have also constructed Genalg, a genetic algorithm (cf *[27]*) package that stably outputs the actual minimum wirecost placement (see Fig. 5); it is, in effect, a demonstration that evolutionary processes suffice for worm wiring optimality. (A caveat on interpretation of the vector mechanical models: while actual physical forces appear to drive neuron arbor optimization, it is likely in the case of nematode ganglion layout that the forces involved should instead be viewed more abstractly as governing natural selection processes; neuron somata need not, in fact, move during development of the individual organism. We have similarly argued *(7)* concerning the simplest neural component placement problem, of brain positioning, that the brain's sensory-motor connections, of course, do not behave literally vector mechanically over evolutionary history.)

But the bottom line here once more seems to be that, in a sense, "Physics suffices": since no genome is required for this self-organization, some interesting limits may thereby emerge on the Central Dogma of genetics. (Cf also the related picture regard-
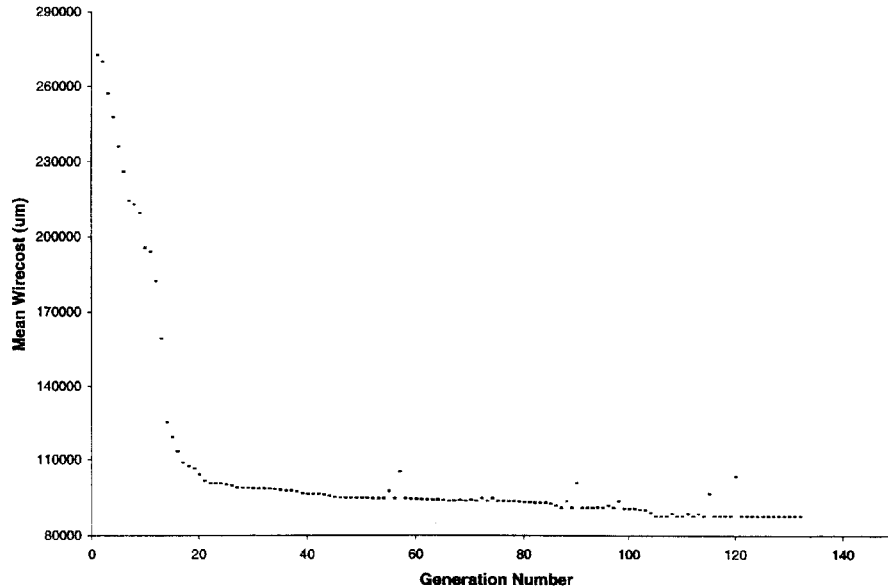
Fig. 5. GenAlg, a simple genetic algorithm, rapidly and reliably finds the optimal (minimum wirelength) layout of *C. elegans* ganglia. The initial population in this run is small, 10 individuals, each here with a reverse ganglion ordering of that found in the actual worm; the algorithm converges upon the minimum total wirecost layout (87,803 µm) in only 130 generations. The evolution of wirecost shows the usual pattern: a very rapid initial improvement of fitness (about 90% during the first 20 generations), followed by a much longer slower fine-tuning phase to optimality. Some of the random mutations cause the half-dozen brief "blips" of increased mean wirecost of the population during the later phase. The robust performance of this genetic algorithm, and also of our force-directed placement algorithm, is further converging support for the hypothesis that the actual layout of *C. elegans* is in fact perfectly optimized.

ing prions *[28]*.) A discrete-state process like a genetic algorithm is not needed to generate highly complex types of biological structure. One rationale for nongenomic anatomy-generating processes, as well as for such simple generative rules as "Save Wire," is apparent in a dilemma that 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 *(3,15)*. The harmony of neuroanatomy and physics suggested here would lower this hereditary information load by accomplishing network optimization without required participation of the genome.

Another observation of robustness worth further study is that, for both the global arbor and ganglion neural optimization problems, random noise injection (e.g., as in simulated annealing *[29]*) generally was not needed to evade local minima traps— unlike for typical network optimization problem instances.

Mapping "chaotic" optimization landscapes: we have found that both a genetic algorithm like GenAlg and a force-directed placement (FDP) algorithm like Tensarama

```
                                                        Input:  actual.mtx

                    T E N S A R A M A

Head                                              Tail
0     0     1     1     2     2     3     3     4     4     5
0     5     0     5     0     5     0     5     0     5     0   Tetrons

    PH    (100.000000)
      AN    (300.000000)
        RNG   (440.000000)
          DO    (506.000000)
          LA    (564.000000)
            VN    (744.000000)
              RV    (1096.000000)
                    VCa   (2004.000000)

                                            VCp   (4004.000000)
                                                  PA      (4874.000000)

        DR    (854.000000)
        LU    (928.000000)

Final layout popped out after:   1,000,000 iterations
Tension Constant:   0.001000
Total Wirecost:   88485.250000 um
```

Fig. 6. Runscreen for Tensarama, an FDP algorithm for optimizing layout (minimizing total wirecost) of *C. elegans* ganglia. This vector mechanical simulation represents each of the worm's approx 1,000 connections (cf Fig. 2 in *[10]*) as a weight-and-pulley (non-Hooke's Law) element acting upon the movable ganglia "PH", "AN", etc. (Connections do not appear on the runscreen, nor do fixed components such as sensors and muscles). At each iteration, the program computes net horizontal force on each ganglion, and correspondingly updates its left/ right position; the cycle is repeated a given number of times. (Ganglion locations are in "tetrons", or quarter-microns, to decrease round-off errors.) The most striking feature of Tensarama performance for the actual worm's connectivity matrix is its comparatively low susceptibility to local minima traps—unlike Tensarama performance for minor modifications of the actual connectivity matrix and unlike FDP algorithms, in general, for circuit design. However, the above screendump shows the final configuration of the system for an identified "killer" initial configuration input of the actual matrix: Tensarama has frozen in a local minimum with ganglia in positions (notably, ganglia DR and LU in head, rather than tail) that yield a final layout wirecost of 88,485.25 μm, about 0.8% more than the actual layout. The fatal initial layout here (ganglion left edges at 0 tetrons) differs only slightly from a quite innocuous initial layout (ganglion centers at 0).

perform notably well in optimizing ganglion placement for the actual connectivity matrix (Fig. 2 in *[10]*) of *C. elegans*. However, this good performance turns out to be interestingly narrow-tuned: (*i*) adding or removing as little as a single connection (of approx 1,000 total) in some cases can change the actual matrix into a "killer matrix" input that is highly prone to paralyzing an FDP algorithm in local minima traps; (*ii*) similarly, some "killer layout" initial input positionings of the ganglia of the actual matrix will paralyze the FDP algorithm (Fig. 6). Each of these instances of discontinuous, very sharply tuned performance prima facie suggests chaotic structure (e.g., *[30]*) and seems worthy of systematic exploration. Each exhibits a Butterfly Effect: some quite small changes of input conditions, but only in a limited range, yield drastic

changes in behavior. We need to compare these natural neuroanatomical connection matrices with some typical benchmark micro circuits (e.g., *[31]*). Such studies entail basic mapping of the optimization terrain. For instance, the "neighborhood" around actual layouts (i.e., the subregion of nearby layouts that differ from the actual one by only a small number of component swaps) appears to be a particularly good one, which is richer in lower cost layouts than randomly sampled zones.

4.5. FUNCTIONAL ROLE OF NEURAL OPTIMIZATION

Finally, a larger question: if one takes seriously the above instances of distinctively fine-grained neural optimization, a larger question emerges: why is such extreme connectivity minimization occurring? Of course, "Save wire" has obvious fitness value as explained earlier—in reducing volume of a delicate metabolically costly tissue and in reducing signal propagation delays in a notably slow transmission medium. For instance, neuron volume minimization directly minimizes the significant metabolic "pumping" costs of maintaining ion concentrations across cell membranes *(32)*. Nonetheless, such optimization nearly to absolute physical limits is only rarely encountered in biology (for instance, human visual and auditory system amplitude sensitivities under good conditions and the silk moth olfactory system, which can detect single molecules of mating pheromones *[33–33b]*). The usual view (e.g., *[34]*) is that nature cannot afford to optimize, but instead, like any finite resource engineer, only satisfices with a compromise among competing desiderata that is "good enough". Natural selection almost never gets to begin with a clean slate, but instead must design organisms as a prisoner of prior evolutionary history.

Thus, the type of striking neural optimization we are observing, in itself, needs explanation regarding its functional role: it could be a clue about basic brain mechanisms that require such extraordinary connectivity minimization, and/or a sign of some unexpectedly feasible means of attaining such optimization, and/or, as mentioned above, part of an economical scheme for generation of complex structure. Neuroanatomical cases, where such optimization is not present, become as diagnostically significant as cases where it is present. "Why", thus, becomes as important as "how" here. Attention thereby turns to issues of neural function as well as structure—in any case, the two really mesh seamlessly. Just as a real brain does not consist of infinitely thin wires, its connections do not have virtually infinite signal propagation velocity. Hence, the methodological approach we started with for brain structure volume, and the issue of stringency of limits upon it, needs in turn to be recapitulated for brain function and its temporal constraints.

ACKNOWLEDGMENTS

REFERENCES

1. Cherniak C. Minimal Rationality. MIT Press, Cambridge, MA, 1986.
2. Cherniak C. Philosophy and computational neuroanatomy. Philos Stud 1994; 73:89–107.
3. Cherniak C. Undebuggability and cognitive science. Communica Assoc Comput Mach 1988; 31:402–412.
4. Cherniak C. The bounded brain: toward quantitative anatomy. J Cognit Neurosci 1990; 2:58–68.
5. Thompson D. On Growth and Form. Cambridge University Press. New York, 1917/1961.
6. Cajal S. Ramon y. Histology of the Nervous System of Man and Vertebrates, Vols. I & II. Swanson N, Swanson L., transl. Oxford University Press, New York, 1909/1995.
7. Cherniak C. Neural component placement. Trends Neurosci 1995; 18:522–527.
8. Mitchison G. Neuronal branching patterns and the economy of cortical wiring. Proc R Soc Lond B 1991; 245:151–158.
9. Albertson D, Thomson J. The pharynx of *Caenorhabditis elegans*. Philos Trans R Soc Lond B 1976; 275:299-325.
9a. White J, Southgate E, Thomson J, Brenner S. The structure of the ventral cord of *Caenorhabditis elegans*. Philos Trans R Soc Lond B 1976; 275:327–348.
9b. White J, Southgate E, Thomson J. Brenner S. The structure of the nervous system of the nematode *Caenorhabditis elegans*. Philos Trans R Soc Lond B 1986; 314:1–340.
9c. Wood W (ed.) The Nematode *Caenorhabditis Elegans*. Cold Spring Harbor Laboratory Press, Cold Spring Harbor, NY, 1988. See also Achacoso T, Yamamoto W. *AY's* Neuroanatomy of *C. elegans* for Computation. CRC Press, Boca Raton, 1992.
10. Cherniak C. Component placement optimization in the brain. University of Maryland Institute for Advanced Computer Studies Technical Report, No. 91–98, 1991; J Neurosci 1994; 14:2418–2427.
11. Garey M, Johnson D. Computers and Intractability: A Guide to the Theory of NP-Completeness. W.H. Freeman, San Francisco, 1979. Lewis H, Papadimitriou C. The efficiency of algorithms. Sci Am 1978; 238:96–109.
11a. Stockmeyer L, Chandra A. Intrinsically difficult problems. Sci Am 1979; 240:140–159.
12. Courant R, Robbins H. What is Mathematics? Oxford University Press, New York, 1941/1969.
13. Hwang F, Richards D, Winter P. The Steiner Tree Problem. North-Holland, Amsterdam, 1992.
14. Bern M, Graham R. The shortest-network problem. Sci Am 1989; 260:84–89.
15. Cherniak C. Local optimization of neuron arbors. University of Maryland Institute for Advanced Computer Studies Technical Report. No. 90-90, 1990; Biological Cybernetics 1992; 66:503–510.
16. Kuh E, Ohtsuki T. Recent advances in VLSI layout. Proc IEEE 1990; 78:237–263.
16a. Sherwani N. Algorithms for VLSI Physical Design Automation, 3rd ed. Kluwer, Boston, 1995. Mead C. Analog VLSI and Neural Systems. Addison-Wesley, Reading, MA, 1989.
17. Ingber D. Cellular tensegrity: defining new rules of biological design that govern the cytoskeleton. J Cell Sci 1993; 104:613–627.
17a. Ingber D. The architecture of life. Sci Am 1998; 278:48–57.
18. Van Essen D. A tension-based theory of morphogenesis and compact wiring in the central nervous system. Nature 1997; 385:313–318.
18a. Armstrong E, et al. Cortical gyrification in the rhesus monkey: a test of the mechanical folding hypothesis. Cereb Cortex 1991; 1:426–432.
18b. Armstrong E, et al. The ontogeny of human gyrification. Cereb Cortex 1995; 5:56–63.
19. Kauffman S. At Home in the Universe. Oxford University Press, New York, 1995.
20. Cherniak C, Changizi M, Kang D. Large-scale optimization of neuron arbors. University

of Maryland Institute for Advanced Computer Studies Technical Report, No. 96–78, 1996; Phys Rev E 1999; **59**:6001–6009. (http://pre.aps.org/).

21. Purves D, Lichtman J. Principles of Neural Development. Sinauer, Sunderland, MA, 1985.
22. Cherniak C. Reply to Letter to Editor. Trends Neurosci 1996; **19**:414–415. (www.wam.umd.edu/~cherniak/).
23. Young M. Objective analysis of the topological organization of the primate cortical visual system. Nature 1992; **358**:152–155.
24. Brush S. Feynman's success: demystifying quantum mechanics. Am Sci 1995; **83**:476–477.
24a. Kinoshita T. New value of the alpha^3 electron anomalous magnetic moment. Phys Rev Lett 1995; **75**:4728–4732.
25. Quinn N. The placement problem as viewed from the physics of classical mechanics. In: Proceedings of the 12th IEEE Design Automation Conference, 1975, 173–178.
25a. Francis R, McGinnis L, White J. Facility Layout and Location, 2nd ed. Prentice-Hall, Englewood Cliffs, NJ, 1992.
26. Berendsen H. Protein folding: a glimpse of the holy grail? Science 1998; **282**:642–643.
27. Mitchell M. An Introduction to Genetic Algorithms. MIT Press, Cambridge, MA, 1996.
28. Keyes M. The prion challenge to the 'Central Dogma' of molecular biology, 1965–1991. Stud Hist Biol Biomed Sci 1999; **30**:1–19.
29. Kirkpatrick S, Gelatt C, Vecchi M. Optimization by simulated annealing. Science 1983; **220**:671–680.
30. Thompson J, Stewart H. Nonlinear Dynamics and Chaos. Wiley & Sons, New York, 1986.
31. (http://vlsicad.cs.ucla.edu/~cheese/ispd98.html).
32. Laughlin S, de Ruyter van Steveninck R, Anderson J. The metabolic cost of neural information. Nat Neurosci 1998; **1**:36–41.
33. Cornsweet T. Visual Perception. Academic Press, New York, 1970.
33a. Green D. Introduction to Hearing. Lawrence Erlbaum, Hillsdale, NJ, 1976.
33b. Shepherd G. Neurobiology, 2nd ed. Oxford, New York, 1988.
34. Gould S. The Panda's Thumb. Norton, New York, 1980.
35. Peichl L, Buh, E, Boycott B. Alpha cells in rabbit retina. J Comp Neurol 1987; **263**:25–41.