

CHAPTER 8

Brain Wiring Optimization
and Non-genomic
Nativism

Christopher Cherniak*

I will talk about combinatorial network optimization – that is, minimization of connection costs among interconnected components in a system. The picture will be that such wiring minimization can be observed at various levels of nervous systems, invertebrate and vertebrate, from placement of the entire brain in the body down to the sub-cellular level of neuron arbor geometry. In some cases, the minimization appears either perfect, or as good as can be detected with current methods – a predictive success story. In addition, these instances of optimized neuroanatomy include candidates for some of the most complex biological structures known to be derivable “for free, directly from physics” – that is, purely from simple physical energy minimization processes. Such a “physics suffices” picture for some biological self-organization directs attention to innate structure via non-genomic mechanisms, an underlying leitmotif of this Conference.

The innateness hypothesis is typically expressed in the DNA era as a thesis that some cognitive structure is encoded in the genome. In contrast, an idea of “non-genomic nativism” (Cherniak 2005) can be explored, that some biological structure is inborn, yet not genome-dependent; instead, it arises directly from simple physical processes. Not only, then, is the organism’s *tabula rasa* in fact not blank, it is “pre-formatted” by the natural order: a significant proportion of structural information is pre-inscribed via physical and mathematical law.

* Acknowledgments: I am indebted to Zekeria Mokhtarzada for his collaboration on this work. NIMH Grant MH49867 supported some of the experimental research.

In his opening remarks, Noam Chomsky described a strong minimalist thesis, that “a principled account” of language is possible: “If that thesis were true, language would be something like a snowflake, taking the form it does by virtue of natural law” (Chomsky “General Introductory Remarks,” this volume; see also 1965: 59). Of course, the snowflake reference calls to mind D’Arcy Wentworth Thompson’s *On Growth and Form* (1917), where the paradigmatic example of mathematical form in nature was the hexagonal packing array, of which snow crystals are an instance. However, even the thousand pages of the unabridged 1917 edition of Thompson’s opus contained few neural examples. Similarly, Alan Turing’s study (1952) of biological morphogenesis via chemical diffusion processes opens a conversation that needs to be continued. In effect, we examine here how far this type of idea presently can be seen to extend for biological structure at the concrete hardware level of neuroanatomy. The key concept linking the physics and the anatomy is optimization of brain wiring.

Long-range connections in the brain are a critically constrained resource, hence there seems strong selective pressure to optimize finely their deployment. The “formalism of scarcity” of interconnections is network optimization theory, which characterizes efficient use of limited connection resources. The field matured in the 1970s for microcircuit design, typically to minimize the total length of wire needed to make a given set of connections among components. When this simple “save wire” idea is treated as a generative principle for nervous system organization, it turns out to have applicability: to an extent, “instant brain structure – just add wire-minimization.” The main caveat is that in general network optimization problems are easy to state, but enormously computationally costly to solve exactly. The ones reviewed here are “NP-hard,” each conjectured to require computation time on the order of brute-force search of all possible solutions, hence often intractable. The discussion here focuses upon the Steiner tree concept and upon component placement optimization. (For a full set of illustrations, see Cherniak and Mokhtarzada 2006.) The *locus classicus* today for neuroanatomy remains Ramón y Cajal (1909).

8.1 Neuron arbor optimization

The basic concept of an optimal tree is: given a set of loci in 3-space, find the minimum-cost tree that interconnects them, for example the set of interconnections of least total volume. If branches are permitted to join at intermodal junctions (sites other than the given terminal loci, the “leaves” and “root”), the minimum tree is of the cheapest type, a Steiner tree. If synapse sites and origin of a dendrite or axon are viewed in this way, optimization of the dendrite




Fig. 8.1. Neuron arbor junction (cat retina ganglion cell dendrite). (a) Branch and trunk diameters conform to $b^3 = b_1^3 + b_2^3$, a fluid-dynamic model for minimum internal wall drag of pumped flow (laminar regime). (b) In turn, angle θ conforms to the "triangle of forces" law, a cosine function of the diameters: $\cos \theta = (b_1^2 - b_2^2)/2b_1 b_2$. This yields the minimum volume for a Y-tree junction (Cherniak et al. 1999). So, "Neuron arbor junctions act like flowing water."

or axon then can be evaluated. (Such an analysis applies despite the "intrinsically" driven character of typical dendrites, where leaf node loci are in fact not targets fixed in advance.) Approximately planar arbors in 2-space are easier to study. The most salient feature of naturally occurring arbors – neuronal, vascular, plant, water drainage networks, etc. – is that, unlike much manufactured circuitry, for each internodal junction, trunk costs (e.g., diameter) are higher than the two branch costs. The relation of branch diameters to trunk diameter fits a simple fluid-dynamical model for minimization of wall drag of internal laminar flow. Furthermore, when such micron-scale "Y-junctions" are examined in isolation, positioning of the junction sites shows minimization of total volume cost to within about 5 percent of optimal, via simple vector-mechanical processes (Cherniak 1992) (see Fig. 8.1).

This Y-tree cost-minimization constitutes local optimization. Only one interconnection pattern or topology is involved. Such small-scale optimization does not by itself entail larger-scale optimization, where local tradeoffs are often required. When more complex sub-trees of a total arbor are analyzed, the optimization problem becomes a global one, with an exponentially exploding number of alternative possible interconnection topologies. For example, a 9-terminal tree already has 135,135 alternative topologies, each of which must be generated and costed to verify the best solution. Neuron arbor samples, each with three internodal Y-junctions, minimize their volume to within around 5 percent of optimal (Cherniak et al. 1999). This optimality performance is consistent for dendrites (rabbit and cat retina cells) and also for some types of axons (mouse thalamus) (see Fig. 8.2).




Fig. 8.2. Complex biological structure arising for free, directly from physics. "Instant arbors, just add water." In each case, from micron to meter scale, neural and non-neural, living and non-living, the actual structure is within a few percent of the minimum-volume configuration shown.

8.2 Component placement optimization

Another key problem in microcircuit design is component placement optimization (also characterized as a quadratic assignment problem): Given a system of interconnected components, find the positioning of the components on a two-dimensional surface that minimizes total connection cost (e.g., wirelength). Again, this concept seems to account for aspects of neuroanatomy at multiple hierarchical levels.

"Why the brain is in the head" is a 1-component placement problem. That is, given the positions of receptors and muscles, positioning the brain as far forward in the body axis as possible minimizes total nerve connection costs to and from the brain, because more sensory and motor connections go to the anterior than to the posterior of the body. This seems to hold for the vertebrate series (e.g., humans), and also for invertebrates with sufficient cephalization to possess a main nervous system concentration (e.g., nematodes) (Cherniak 1994a, 1995).

Multiple-component problems again generally require exponentially exploding costs for exact solutions: for an n -component system, $n!$ (n factorial) alternative layouts must be searched. One neural wiring optimization result is for placement of the eleven ganglionic components of the nervous system of the roundworm *Caenorhabditis elegans*, with about 1,000 interconnections (see Fig. 8.3). This nervous system is the first to be completely mapped (Wood 1988), which enables fair approximation of wirelengths of connections (see Fig. 8.4). When all 39,916,800 alternative possible ganglion layouts are generated, the




Fig. 8.3. *C. elegans* ganglion components: their body locations and schematized shapes.

actual layout turns out in fact to be the minimum wirelength one (Cherniak 1994a). Some optimization mechanisms provide convergent support for this finding: a simple genetic algorithm, with wirecost as fitness-measure, will rapidly and robustly converge upon the actual optimal layout (Cherniak et al. 2002). Also, a force-directed placement (“mesh of springs”) algorithm, with each connection approximated as a microspring acting between components, attains the actual layout as a minimum-energy state, without much trapping in local minima (Cherniak et al. 2002) (see Fig. 8.5). This little nervous system can thereby weave itself into existence.

There is statistical evidence that this “brain as microchip” wire-minimization framework also applies in the worm down to the level of clustering of individual neurons into ganglionic groups, and even to cell body positioning within ganglia to reduce connection costs (Cherniak 1994a).

Finally, the wiring-minimization approach can be applied to placement of functional areas of the mammalian cerebral cortex. Since wirelengths of intrinsic cortical connections are difficult to derive, another strategy is to explore instead a simpler measure of connection cost, conformance of a layout to a wire-saving heuristic Adjacency Rule: If components *a* and *b* are connected, then *a* and *b* are adjacent. Exhaustive search of all possible layouts is still required to identify the cheapest one(s). One promising calibration of this approach is that the minimum wirecost actual layout of the nematode ganglia is among the top layouts with fewest violations of this adjacency rule. For seventeen core visual areas of macaque cortex, the actual layout of this subsystem ranks in the top 10^{-7} layouts best fitting this adjacency-costing; for fifteen visual areas of cat cortex, the actual layout ranks in the top 10^{-6} of all layouts (Cherniak et al. 2004).




Fig. 8.4. Complete ganglion-level connectivity map for *C. elegans* nervous system (apparently, the first depiction of approximately complete connectivity of a nervous system down to synapse level). Each horizontal microline represents one of the 302 neurons. Horizontal scaling: $\sim 100\times$. This actual ganglion layout requires the least total connection length of all ~ 40 million alternative orderings (Cherniak 1994a).

In general, a Size Law seems to apply to cases like macaque and cat (and worm) with such local-global tradeoffs: The larger the proportion of a total system the evaluated subsystem is, the better its optimization. We have observed this Size Law trend recently also for rat olfactory cortex and for rat amygdala (Rodríguez-Esteban and Cherniak 2005). For the largest systems studied (visual, auditory, plus somatosensory areas of cat cortex), there is evidence of

Input: rad.cue

T E N S A R A M A

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

Teprons

```

PH <189.000000>
AN <300.000000>
BNC <492.000000>
DO <425.000000>
LA <615.000000>
UN <795.000000>
RV <1000.000000>
UCa <1998.000000>

UCP <3908.000000>
Pa <4779.000000>
DR <4882.000000>
LU <4936.000000>

Final layout popped out after 150000 iterations
Tension Constant = 0.010000
Local Ultecost = 80954.250000

```

Fig. 8.5. Tensarama, a force-directed placement algorithm for optimizing layout of *C. elegans* ganglia. This “mesh of springs” vector-mechanical energy-minimization simulation represents each of the worm’s ~ 1,000 connections (not visible here) acting upon the moveable ganglia PH, AN, etc. The key feature of Tensarama performance for the actual worm connectivity matrix is its low susceptibility to local minima traps (Cherniak et al. 2002) – unlike Tensarama performance for small modifications of the actual connectivity matrix (a “butterfly effect”), and unlike such force-directed placement algorithms in general for circuit design. Here Tensarama is trapped in a slightly sub-optimal layout, by a “killer” connectivity matrix that differs from the actual matrix by only one fewer connections.




Fig. 8.6. Cerebral cortex of cat. (Lateral aspect; rostral is to right.) Placement of 39 interconnected functional areas of visual, auditory, and somatosensory systems (in white). Exhaustive search of alternative layouts suggests this actual layout ranks at least in the top 100 billionth of all possible layouts with respect to adjacency-cost of its interconnections (Cherniak et al. 2004). – “Best of all possible brains”?

optimization approaching limits of current detectability by brute-force sampling techniques (see Fig. 8.6). A similar Size Law pattern also appears to hold for Steiner sub-tree optimization of neuron arbor topologies.

8.3 Optimization: mechanisms and functional roles

The neural optimization paradigm is a structuralist position, postulating innate abstract internal structure – as opposed to an empty-organism blank-slate account, without structure built into the hardware (structure is instead vacuumed up from input). The optimization account is thereby related to Continental rationalism – but for brain structure, rather than the more familiar mental structure.

The picture here is of limited connections deployed very well – a predictive success story. The significance of ultra-fine neural optimization remains an open question. That is, one issue raised by such “best of all possible brains” results is, what is the function of minimizing, rather than just reducing, neural connection costs? Wiring optimization is of course subject to many basic constraints, and so cannot be ubiquitous in the nervous system; the question is where it does in fact occur, and how good it is. Tradeoffs of local optimality for better cost minimization of a total system are one way in which global optimization can be obscured.

The high levels of connection optimization in the nervous system seem unlike levels of optimization common elsewhere in organisms. Optimization to nearly absolute physical limits also can be observed in human visual and auditory sensory amplitude sensitivities, and in silk moth olfactory sensitivity to pheromones (Cherniak et al. 2002) – that is, at the very meniscus of the neural with its environment. Why should the neural realm sometimes demand optimization, rather than the more familiar biological satisficing? (For some biological optimization phenomena elsewhere, see Weibel et al. 1998).

Mechanisms of neural optimization are best understood against the background mentioned earlier, that the key problems in general are computationally theory are NP-complete, hence exact solutions in general are computationally intractable. For example, blind trial-and-error exhaustive search for the minimum-wiring layout of a 50-component system (such as all areas of a mammalian cerebral cortex), even at a physically unrealistic rate of one layout per picosecond, would still require more than the age of the Universe (Cherniak 1994b). Thus, to avoid universe-crushing costs, even evolution instead must exploit “quick and dirty” approximation or probabilistic heuristics.

One such possible strategy discernible above is optimization “for free, directly from physics.” That is, as some structures develop, physical principles

cause them automatically to be optimized. We reviewed above some evidence for arbor optimization via fluid dynamics, and for nematode ganglion layout optimization via “mesh of springs” force-directed placement simulation. As could be seen for each of the neural optimization examples above, some of this structure from physics depends in turn on exploiting anomalies of the computational order (Cherniak, 2008). While neuron arbors seem to optimize on an embryological timescale, component placement optimization appears to proceed much slower, on an evolutionary timescale. For component placement optimization, there is the chicken-egg question of whether components begin in particular loci and make connections, or instead start with their interconnections and then adjust their positions, or some mix of both causal directions. It is worth noting that both a force-directed placement algorithm for ganglion layout, and also genetic algorithms for layout of ganglia and of cortex areas, suggest that simple “connections → placement” optimization processes can suffice.

If the brain had unbounded connection resources, there would be no need or pressure to refine employment of wiring. So, to begin with, the very fact of neural finitude appears to drive “save wire” fine-grained minimization of connections. Another part of the functional role of such optimization may be the picture here of “physics → optimization → neural structure.” Optimization may be the means to anatomy. At least our own brain is often characterized as the most complex structure known in the universe. Perhaps the harmony of neuroanatomy and physics provides an economical means of self-organizing complex structure generation, to ease brain structure transmissibility through the “genomic bottleneck” (Cherniak 1988, 1992) – the limited information carrying-capacity of the genome. This constitutes a thesis of non-genomic nativism, that some innate complex biological structure is not encoded in DNA, but instead derives from basic physical principles (Cherniak 1992, 2005).

The moral concerns not only “pre-formatting” for evolutionary theory, but also for modeling mind. Seeing neuroanatomy so intimately meshed with the computational order of the universe turns attention to constraints on the computationalist thesis of hardware-independence of mind; practical latitude for alternative realizations narrows.

Discussion

PARTICIPANT: I am a biologist and I'm interested in this concept of minimality or perfect design in terms of language. Coming from immunology, we have a mixture of very nice design and also huge waste. That is to say, every day you make a billion cells which you just throw in the bin because they make

antibodies you don't need that day. And I am wondering whether in the brain there is a combination of huge waste in terms of enormous numbers of cells, and beautiful design of the cell itself and the way it copes with incoming information. Some neurons take something like 40,000 inputs, and there doesn't seem to be any great sense in having 40,000 inputs unless the cell knows how to make perfect use of them. And that seems to be something that very little is written about. The assumption is that the cell just takes inputs and adds them up and does nothing much with them. But I would suggest that there may be something much more interesting going on inside the cell, and that focusing on the perfect design of the cell might be more attractive and more productive than looking at perfect design in terms of the network as a whole, which is hugely wasteful in having far too many cells for what is needed. I wonder if you would like to comment on that.

CHERNIAK: Just to start by reviewing a couple of points my presentation garbled: anyone around biology, or methodology of biology, knows the wisdom is that evolution *satisfices* (the term “satisfice” is from Herbert Simon 1956). The design problems are so crushingly difficult that even with the Universe as Engineer, you can't optimize perfectly; rather, you just satisfice. And so, I remember literally the evening when we first pressed the button on our reasonably debugged code for brute-force search of ganglion layouts of that worm I showed you, to check on how well minimized the wiring was; I certainly asked myself what I expected. We had already done some of the work on neuron arbor optimization, and so I figured that the nematode (*C. elegans*) wiring would be doing better than a kick in the head, but that it would be like designing an automobile: you want the car to go fast, yet also to get good mileage – there are all these competing desiderata. So when our searches instead found perfect optimization, my reaction was to break out in a cold sweat. I mean, quite happily; obviously the result was interesting.

One open question, of course: it is easy to see why you would want to save wire; but why you would want to save it to the *n*th degree is a puzzle. One pacifier or comfort blanket I took refuge in was the work Randy Gallistel referred to on sensory optimality (see “Foundational Abstractions,” this volume). Just in the course of my own education, I knew of the beautiful Hecht, Schlaegel, and Pirenne (1942) experiments showing the human retina operating at absolute quantum limits. And the similar story, that if our hearing were any more sensitive, we would just be hearing Brownian motion: you can detect a movement of your eardrum that is less than the diameter of a hydrogen atom. A third sensory case (obviously, I'm scrambling to remember these) is for olfactory sensitivity – the Bombyx silk moth, for example. Romance is a complicated project; the

moths' "antennas" are actually noses that are able to detect single pheromone molecules. If you look at the titration, males are literally making Go/No-go decisions on single molecules in terms of steering when they are homing in like that. However, these are all peripheral cases of optimality, and they don't go interior; so that is one reason why I wanted to see if we could come up with *mechanisms* to achieve internal wiring minimization. Another reassurance we sought was to look at other cases of possible neural optimization. The claim cannot be that everywhere there is optimization, we cannot say that on the basis of what we are seeing. Rather, the issue is whether or not there are other reasonably clear examples of this network optimization. Now, some of the work that got lost in my talk improvisation is on cortex layout; so you are moving from the nematode's approximately one-dimensional nervous system, to the essentially two-dimensional one of the cerebral cortex (which is much more like a microchip in terms of layout). And cortex results are similar to the worm. For cortex, you need more tricks to evaluate wiring optimality. But still, when we search alternative layouts, we can argue that the actual layout of cat cortex is attaining wiring-minimization at least somewhere in the top one-billionth of all possible layouts. As an outside admirer, I find the single cell a prettier, less messy world than these multi-cellular systems. I would point out that the work I showed you on arbor optimization is at the single-cell level – actually at the sub-cellular level, in the sense that it is for the layout of single arbors. (The one caveat is that those arbors are approximately two-dimensional. The mathematics is somewhat simpler than for 3D.)

HAUSER: I may not have the story completely right, but I was reading some of the work of Adrian Bejan (Bejan and Marden 2006), an engineer at Duke, who has made somewhat similar kinds of arguments as you have about tree structure, and especially about the notion of optimal flow of energy or resources. In a section of one of his books, he makes the argument that there is a necessary binary bifurcation in many tree structures at a certain level of granularity. This is probably a leap, but in thinking about some of the arguments concerning tree structure in language, is it possible that there is more than mere metaphor here? In other words, could the fact that trees, lightning, neurons, and capillaries all show binary branching indicate that this is an optimal solution across the board, including the way in which the mind computes tree structures in language? Could this be the way language *had* to work?

CHERNIAK: Yes, that is a classic sort of inter-level connection, and I don't think it is just metaphorical. When we went into this field, all the network optimization theory, all the graph theory for arbors, had been done for what are called Steiner trees. (The usual history of mathematics story, misnamed after

Jacob Steiner of the nineteenth century; but in fact you can find work on the idea going back to the Italian Renaissance, within the research program of Euclidean geometry.) The classical models assume trunks cost the same as branches, and so we had to retrofit four centuries of graph theory to cover cases where trunks cost more than branches – as they usually do in nature. So that is the one caveat on this. But if you go back to the classic uniform wire-gauge models, then the usual theorems are in fact that optimal trees will have such bifurcating nodes; this is a completely abstract result. A caution I hasten to add is: there is another type of tree, the minimal spanning tree. With Steiner trees, you are allowed to put in internodal junctions, and you get a combinatorial explosion of alternative topologies. The largest Steiner trees that have been solved by supercomputer have perhaps around a hundred nodes. There are more towns than that in Tennessee, so the computational limits on Steiner trees are very much like the traveling salesman problem. But if you instead look at this other type of tree ("minimal spanning tree" probably approximates a standard name), in this case junctions are only permitted at nodes or terminals, which is not of course what you see for neuron arbors. However, minimal spanning trees are incredibly fast to generate, and indeed the most beautiful algorithms in the universe that I know of are for generating minimal spanning trees. You see quarter-million-node sets being solved. Anyway, if you look at the neuron cell body, you can treat that one case as a local minimal spanning tree, and the theorem there is: Not two, but six branches maximum. And indeed micrographs of retinal ganglion cells show six branches from the soma. Anyway, again, regarding your query, it's a theorem of graph theory that optimal Steiner trees have binary bifurcations. And, yes, I agree, this is germane to theorizing about tree structures in linguistics.