

In: Oxford Handbook of Philosophy & Neuroscience. J. Bickle, ed (Oxford, 2009).

CHAPTER 14

NEUROANATOMY AND COSMOLOGY

CHRISTOPHER CHERNIAK

A bounded-resource paradigm is by now familiar in mind-brain science. The discussion here focuses not on the levels of psychology or physiology but on brain wiring, in particular, on an apparent paradox: Available connectivity in, say, the cerebrum, is stringently limited, yet deployment of interconnections shows fine-grained optimization, sometimes without detectable limits. Virtually perfect network optimization, rather than just network satisficing, sometimes applies well to neuroconnectivity architecture. Such layout cost-minimization problems are a major hurdle of microcircuit design and are known to be NP-complete (explained shortly). How does biology effectively solve such problems? Briefly, some odd wrinkles of the computational realm seem to be exploited. This line of thought suggests extensions to the anthropic principle of cosmology—that is, further possible constraints on models of a universe in which our intelligence can arise. This in turn may narrow the conventional hardware-independence thesis of computational psychology.

1. BRAIN WIRING MINIMIZATION AND OPTIMIZATION LANDSCAPES

The tension discussed here is a picture sometimes emerging in our work: seemingly almost limitless refinement in use of clearly limited connection resources (e.g., Cherniak, 1994a, 1994b; Cherniak, Mokhtarzada, and Nodelman, 2002). Prima facie,

this embodies an antagonism. The bounded-resource perspective (e.g., Cherniak, 1986) is that the mind-brain has moderate-scale resources—not zero, not perfect. The natural expectation regarding their use that harmonizes with this resource model would be that available resources are satisfied, that is, employed *moderately* well. On the one hand, they cannot be used with profligate inefficiency because they are limited. Of course, optimal use of connections in effect would maximize available connection resources. However, on the other hand, the wiring would not be anticipated to be (nearly) perfectly deployed, because that refinement itself in turn typically has its own cosmically high costs.

Network optimization theory characterizes the minimized use of limited connection resources (e.g., wirelength) in a system. A main feature of the optimization landscape is that many of the problems are nondeterministic polynomial-time complete (NP-complete) or NP-hard; that is, such problems are conjectured to require computation time typically on the order of a brute-force search of all possible solutions and are often therefore intractable.

Over the past decade or so, our laboratory has uncovered a pattern of very efficient connection use in the nervous systems of various animals (e.g., roundworm, rat, cat, primate), at different levels (from positioning of whole brain, through layout of ganglia and cortex, down to subneuronal structures like somata and arbors), and for different types of network optimization (e.g., Steiner tree, component placement optimization) (Cherniak et al., 2002). We are also beginning to find connection-minimization patterns in genomic systems. Generally, the verdict of history is that Nature just solves tractable special cases of hard optimization problems. However, we have found some odd twists. The interrelation between the neuroanatomy and cosmology in the title of this chapter is that brain structure optimization seems to mesh finely with some convenient coincidences of the computational order of the universe. These instances of biology exploiting anomalies of the computational universe operate against the background of a “nongenomic nativism” (Cherniak, 2005): We have found some complex biological structure originating directly from simple physical processes, without need of DNA involvement.

For example, we described neuron arbor approximations of Steiner tree optimization (Cherniak, Changizi, and Kang, 1999; see Chklovskii and Stepanyants, 2003). A Steiner tree is a minimum-cost arbor connecting a set of terminal loci, for example, root to synapse sites as “leaves.” The exact solution of the problem has two parts: (1) Brute force search of an exponentially exploding set of alternative tree topologies or connection patterns (the NP-hard part of the problem); (2) then embedding each topology (adjusting branch lengths for the various weights of the arbor branches, because trunks cost more than their branches), achievable via simple vector-mechanical processes. (This type of post hoc analysis applies despite the “intrinsically” driven character of typical dendrites, where leaf node loci are in fact not targets fixed in advance.) Perhaps we acquire a romantic picture from graph theory: The heart of the matter is searching the topologies (each is like a different universe, inaccessible from any other), whereas the embedding is trivial, just classical continuous mathematics.




Figure 14.1. Mouse thalamus axon subarbor (wireframe representation). (A) Actual multijunction tree in broken lines; optimal embedding of actual topology (with respect to volume minimization) in solid lines. Volume error of actual tree is 2.2 percent. (B) "Best of all possible topologies" (with respect to volume); volume error is 2.5 percent. Only 10 of the 10,395 alternative topologies here have lower total volume costs, when optimally embedded, than the actual topology. Topology search yields little improvement compared with embedding. Embedding can be accomplished by a simple fluid mechanical process; topology cannot.




Figure 14.2. Nine-terminal arbor of mouse thalamus axon: distribution of volume costs of all 135 possible topologies, each optimally embedded. The histogram shows the usual pattern for natural arbors, living and nonliving—more costly topologies are more common, cheapest ones are rarest. The most costly optimally embedded "pessimal" layouts have only about 12 percent greater volume than cheapest one. Hence, for optimization, "topology does not matter."

However, the unintuitive story we found is that roughly, "Topology does not matter." Figures 14.1 and 14.2 illustrate that in fact relatively little tree cost-minimization seems to be at stake in selecting a topology; much more can be gained in just refining the embedding. (So also for classical uniform-cost Steiner trees, with all branch junctions at 120° .) Thus, Nature seems to "solve" the Steiner tree problem via a lucky if counterintuitive factoid of computational geometry, that embedding dominates

over topology; topology can in effect be ignored at comparatively little cost. (This strategy may also be worth exploring for rectilinear Steiner tree design in microchip engineering.) Neuron arbors solve their minimization problem to within a few percent of optimality, about as well as nonliving arbors, such as river drainage networks; the neurons seem to attain their volume minimization via the same basic fluid dynamic processes as do water networks.

Another exploitation of special cases emerges in mechanisms for optimization of layout of ganglia in the nematode nervous system. This is a component placement optimization problem, involving positioning of a system of interconnected components to minimize total connection costs. Again, the general quadratic assignment problem (QAP) here is NP-hard. Yet the worm's actual layout is the unique cheapest one in total wirecost (Cherniak, 1994a). As we have reported, a very simple genetic algorithm, with total wirelength of nervous system as fitness measure, will robustly and reliably find that optimal layout (Cherniak et al., 2002). Yet so will a forced-directed placement algorithm (FDP) or vector-mechanical "mesh of springs." See figures 14.3 and 14.4 (see also Chklovskii, 2004).

Such energy-minimization approaches have not proven very successful for microchip layout because of trapping in local minima. Why do they work for worm brain layout? One strange feature we found is that the FDP algorithm has a knife-edge sensitivity to the worm connectivity matrix: We found some changes of adding or cutting a *single* worm brain connection that can paralyze the algorithm in local minima. So the moral seems to be that the worm brain may indeed constitute a




Figure 14.3. GenAlg, a simple genetic algorithm, rapidly and reliably finds the optimal (minimum wirelength) layout of *C. elegans* ganglia among 11! alternatives. The initial population of this run was only 10 individuals, all with reverse of actual ordering of ganglia. (For another GenAlg run with the same initial population, see Cherniak et al., 2002, figure 5.)

Input: rid.add

T E N S A R A M A

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

PH (152.000000)
 AN (352.000000)
 RNG (492.000000)
 DO (26.000000)
 LA (616.000000)
 VN (796.000000)
 RV (1000.000000)
 VCa (1908.000000)
 VCP (3908.000000)
 PA (4778.000000)
 DR (4862.000000)
 LU (4936.000000)

Final layout popped out after: 150,000 iterations
 Tension Constant: 0.0100
 Total Wirecost: 90,743.7500 um

Figure 14.4. Tensarama, a force-directed placement (FDP) algorithm that optimizes layout of *C. elegans* ganglia. This “tug of war” vector-mechanical energy-minimization simulation represents each of the worm’s approximately 1,000 connections acting on the movable ganglia PH, AN, and so on. The key feature of Tensarama performance for the actual worm connectivity matrix is its low susceptibility to local minima traps—unlike Tensarama performance for small modifications of the actual connectivity matrix (a “butterfly effect”), and unlike FDP algorithms in general for circuit design. Here Tensarama is locked in a nonoptimal layout by a “killer” connectivity matrix that differs from the actual matrix by only one added connection. See also Cherniak et al. (2002, figure 6) and Cherniak (2005, figure 5.2).

special tractable case of QAP; but the picture may be that the abstract connectivity matrix itself (the “mind”) co-evolves with its wirecost-minimized physical layout. Instead of the top-down design methodology familiar in computer engineering, this is a cart-before-horse (cart-beside-horse?), tail-wags-dog picture.

Finally, we have found a peculiar fine-tuning of a quick-and-dirty heuristic for connection cost minimization of layouts of worm ganglia (also perhaps of mammalian cerebral cortex). The adjacency rule wiring heuristic is: If two components are connected, then place them next to each other. It can be used for laying out a system and also as an easily applicable wirecost measure for such layouts (instead of measuring actual wirelengths, which is often unfeasible). We have described a calibration of performance of this heuristic for worm ganglia (Cherniak, Mokhtarzada, Rodriguez-Esteban, and Changizi, 2004; see figures 14.5 and 14.6). In general, the adjacency heuristic does not yield especially low-wirecost ganglion layouts. However, again, a very narrow range of special cases behaves quite well, namely, layouts that conform best to the adjacency rule do have the cheapest wirecosts. The




Figure 14.5. Adjacency rule conformance, versus total wirecost, of a sample of 100,000 *C. elegans* ganglion layouts. (Adjacency rule: “If components *a* and *b* are connected, then *a* and *b* are adjacent.”) Generally, the adjacency rule is not an effective heuristic for good wirecost. However, the small set of layouts best fitting the adjacency rule (points at far left) behave markedly differently: They correlate well with the best wirecost layouts.




Figure 14.6. Parcellation of functional areas of macaque cerebral cortex. Component placement optimization analysis of layout of 17 core areas (white) of visual cortex, along with immediately contiguous edge areas (dark gray). Lighter straight lines indicate reported interconnections among core areas. Rostral (forward) is to right. In a connection cost analysis (with respect to the adjacency rule), this actual layout of the core visual system ranks in the top one-millionth of all alternative layouts (Cherniak et al., 2004). For 39 areas of cat cortex, actual layout may rank in top 100 billionth of all possible layouts. “Best of all possible brains?”

computational order is such that for arbor optimization, topology does not matter; yet for component placement optimization, it suffices. So once more, the leitmotif seems to be exploitation of happy accidents of the fine scale microtopography of the optimization landscape.

What these cases convey is the need to develop means of systematically mapping such "chaotic" problem terrains—concepts beyond simple sampling.

2. BRAIN WIRING OPTIMIZATION AND THE ANTHROPIC PRINCIPLE

The simplest relation of the brain to the universe, stemming from the Laplacean picture of the cosmos as a vast clock, comes from a tradition of computing the maximum total computational capacity of the cosmos. A recent estimate, based on an entropic theory of information, and taking account of quantum phenomena, is 10^{50} operations/second (Lloyd, 2002). With the familiar speculation (e.g., Donofrio, 2001) that one human brain is capable of around 10^{16} flops (10 petaflops, or 10 quadrillion floating-point arithmetic operations/second), the universe then has the computing capacity of 10^{34} human brains. Or, with $\sim 10^9$ human brains presently alive on Earth, the universe-as-computer is equivalent to 10^{25} Earth human brain-populations.

Another connection of neuroanatomy to cosmology would proceed via the anthropic principle (Barrow and Tipler, 1986). A weak form of the anthropic principle asserts that any model of the universe must meet the adequacy condition that it permits the development of life and intelligence—as has in fact occurred. (Although the idea is much older, the phrase "anthropic principle" is credited to Carter, 1974.) In particular, the physical constants—the number of dimensions, energy states of the electron, strength of the weak nuclear force, and so on—must be fine-tuned to permit carbon-based life, and human intelligence, to arise.

For instance, in the search for extrasolar planets capable of supporting life, the focus is on ones in the habitability range, not too hot or cold. Similarly, on the largest scale for the universe: Its physics must not yield just Big Bang conditions with too high temperatures for stable structure, nor only heat death conditions with too little energy exchange for information-processing—neither bang nor whimper.

In the previous section, we reported brain-wiring optimization in a number of nervous systems, at a number of levels. If such extreme neural network optimization is not just gratuitous, but in fact somehow is a prerequisite for brain functioning (Cherniak et al., 2002), we can now contemplate a possible further set of brain-enabling conditions of the universe.

1. For the generation of approximations of optimal Steiner trees, search of topologies does not have a significant downside risk, compared to embedding. This is a lucky accident of computational geometry, because exhaustive topology search in fact is intractable, whereas embedding is straightforward.
2. There are connectivity matrices for the "mind" of the nematode, and of other creatures (that implement required brain functions) for which energy-minimization placement processes are not trapped in local minima—as FDPs often are, and as very small variations from the actual worm brain matrix are.
3. Regarding worm nervous system layouts, in the vast majority of cases, adjacency heuristic conformance does not correlate with cheap wirecost. However, the few cases of *best* adjacency rule performance do correspond to the cheapest wirecost cases. Worm ganglia, and rat, cat, and macaque cortex areas in fact do optimize well with respect to the adjacency heuristic, with a Size Law included (Cherniak et al., 2004).

In concluding, we refrain from intelligent design speculations about such fortunate harmony of universe and brain. There are traditions of haruspicing messages in natural patterns, for example, from extraterrestrials in bacteriophage DNA (Yokoo and Oshima, 1979), and from Someone in the cosmic microwave background radiation (Hsu and Zee, 2006). As a caveat, we cite a universal non-nominal prayer that social commentator Emo Phillips proposed, in another context: Approximately, "Please arrange the universe for my convenience." Reality is the ultimate Rorschach.

A modest moral concerns modeling mind: Some of the landscape of philosophy of mind begins to shift. Namely, seeing neuroanatomy so intimately meshed with the computational order of the universe prompts reconsideration of constraints on the computationalist thesis of hardware-independence of mind. To the extent that wiring optimization seems both requisite, and a hostage to brain-friendly computational anomalies of the universe, practical latitude for alternative realizations of mind narrows.

ACKNOWLEDGMENT

A version of this chapter was presented as an invited contribution to a symposium on Computational Complexity Analyses of Cognitive Models at the Society for Mathematical Psychology annual meeting, July 2006, Vancouver.

REFERENCES

- Barrow, J., and Tipler, F. (1986). *The Anthropic Cosmological Principle*. Oxford: Oxford University Press.
- Carter, B. (1974). Large number coincidences and the anthropic principle in cosmology. *International Astronomical Union Symposium* 63: 291.
- Cherniak, C. (1986). *Minimal Rationality*. Cambridge, Mass.: MIT Press.
- Cherniak, C. (1994a). Component placement optimization in the brain. *Journal of Neuroscience* 14: 2418–2427.
- Cherniak, C. (1994b). Philosophy and computational neuroanatomy. *Philosophical Studies* 73: 89–107.
- Cherniak, C. (2005). Innateness and brain-wiring optimization: Non-genomic nativism. In A. Zilhao (Ed.), *Evolution, Rationality and Cognition*. London: Routledge, 103–112.
- Cherniak, C., Changizi, M., and Kang, D. (1999). Large-scale optimization of neuron arbors. *Physical Review E* 59: 6001–6009.
- Cherniak, C., Mokhtarzada, Z., and Nodelman, U. (2002). Optimal-wiring models of neuroanatomy. In G. Ascoli (Ed.), *Computational Neuroanatomy: Principles and Methods*. New York: Humana.
- Cherniak, C., Mokhtarzada, Z., Rodriguez-Esteban, R., and Changizi, B. (2004). Global optimization of cerebral cortex layout. *Proceedings of the National Academy of Sciences USA* 101: 1081–1086.
- Chklovskii, D. (2004). Exact solution for the optimal neuronal layout problem. *Neural Computation* 16: 2067–2078.
- Chklovskii, D., and Stepanyants, A. (2003). Power-law for axon diameters at branch point. *BMC Neuroscience* 4: 18.
- Donofrio, N. (2001). Turing Memorial Lecture. *Computer Journal* 44: 67–74.
- Hsu, S., and Zee, A. (2006). Message in the sky. *Modern Physics Letters A* 21: 1495–1500.
- Lloyd, S. (2002). Computational capacity of the universe. *Physical Review Letters* 88: 1–4.
- Yokoo, H., and Oshima, T. (1979). Is bacteriophage ϕ X174 DNA a message from an extraterrestrial intelligence? *Icarus* 38: 148–153.