

In: Evolution, Rationality & Cognition.
A. Zilhao, ed (Routledge, 2005) 103-112.

5 Innateness and brain-wiring optimization

Non-genomic nativism

Christopher Cherniak

Our experimental work in computational neuroanatomy has uncovered distinctively efficient layout of wiring in nervous systems. When mechanisms are explored by which such "best of all possible brains" design is attained, significant instances turn out to emerge "for free, directly from physics." In such cases, generation of optimal brain structure appears to arise simply by exploiting basic physical processes, without the need for intervention of genes. An idea that physics suffices here – of some complex biological structure as self-organizing, generated without genomic activity – turns attention to limiting the role of the genome in morphogenesis. The familiar "nature/nurture" alternatives for origins of basic internal mental structure are that it arises either from the genome or from invariants of the external environment. A third alternative is explored for the neural cases here, a non-genomic nativism.

The study of minimization of neural connections reveals interrelations between the Innateness Hypothesis and theses associated with the Central Dogma of genetics. The discussion shifts from the usual focus upon abstract cognitive structure instead to underlying brain hardware structure, to hard-wired neuroanatomy.

1 Nativism

The prototypical *Innateness Hypothesis* is: abstract mental structure – e.g., relating to knowledge of language (Stich 1975) – is intrinsic to an organism, for instance genetically determined. But similarly also of course for underlying biological hardware, that is, brain structure, neuroanatomy. A brain is no more plausibly a blank slate with unlimited plasticity in response to its environment than is a mind.

As originally (narrowly) construed, the Central Dogma of genetics was: Genetic information flows one way, from the genome outward to its cellular milieu. That is, *DNA* → *RNA* → *proteins* (Crick 1958; Watson 1965). A recent example that may help to define by contrast is the challenge to the Central Dogma perceived in "prion" diseases (e.g., Kuru, Creutzfeldt-Jacobson Syndrome, Bovine Spongiform Encephalopathy (BSE)); prions are non-DNA

containing particles that appear capable of replication – and infection – in mammalian central nervous systems (Keyes 1999). To the Central Dogma we can also add a tacit “Pre-Central Dogma”: the genome information encodes an instruction-set, blueprint, representation, for construction of the organism (e.g., Watson *et al.* 1987; DePomerai 1990).

How complete can such a blueprint be? Can it approximate for biological structure-generation a kind of Maxwell’s demon micromanaging busybody? The genome instructions cannot be total. A classical regress argument from philosophy of mind can be adapted here, for example, along lines of Ryle (1949). To stop the regression, some basic level of hardware capacities must be assumed – a Lego/Meccano/Erector-set repertoire of fundamental operations, e.g., the basic physical laws of our universe.

2 Optimized neuroanatomy

However, a more extensive a-genomic domain emerges for some large-scale neuroanatomical structure. The framework is a “Best of all possible brains” hypothesis, in particular, with respect to generative principles to “Save wire.” If the brain had an unbounded supply of connections, there would be no pressure to optimize employment of wiring. But in fact connection resources are limited. And since the essential function of a brain is to connect, to make connections, connections seem to have an extremely high value, and so their use has been perfected to a very fine-grained degree. Optimized brain anatomy is detectible in minimized wiring at both micro and macro scale. In some cases, the refinement is discernible down to a best-in-a-billion level (Cherniak 2000; Cherniak *et al.* 2004). Such results begin to approach some of the most precise confirmed predictions in neuroanatomy; why such peculiarly fine-grained optimization should arise, itself, in turn, needs explaining. These wiring optimization problems, typically encountered in computer microchip design, are believed to be intrinsically computationally intractable (“NP-complete”) (Lewis and Papadimitriou 1978, Garey and Johnson 1979, Stockmayer and Chandra 1979), unavoidably requiring a combinatorially exploding amount of computation to solve.

For instance, neuron arbor morphogenesis behaves like flowing water (Cherniak 1992; Cherniak *et al.* 1999).

- i “Neural fluid mechanics”: a fluid-dynamical model for minimized wall-drag of pumped flow through a system of pipes will predict the geometry of a variety of dendrite and axon structures almost as well as it predicts configuration of river drainage networks.
- ii Water flow in branching networks in turn acts like a tree with segments composed of weights-cords-and-pulleys, that is, vector-mechanically; so also do the neuron arbors. The result is that such axons and dendrites globally minimize their total volume to about 5 percent of optimum for interconnecting their terminal loci (see Figure 5.1).




Figure 5.1 Complex biological structure arising directly from basic physics: nerve cell anatomy behaves like flowing water, and waterflow in turn acts like a tree composed of springs. “Instant arbors, just add water.” The dendrite (input) and axon (output) branchings are portions of mammalian neurons, the river is an experimentally generated drainage network. In each of the three cases, the actual structure is within a few percent of the optimal, minimum-volume configuration shown; this is evidence of goodness of fit of the physical optimization model (Cherniak *et al.* 1999). Diagram: Mark Changizi.

As another instance, vector mechanics suffices for optimization of placement of the ganglia (neural sub-networks) of the roundworm *Caenorhabditis elegans* (Cherniak 1994a; Cherniak *et al.* 2002). “The web that weaves itself”: Some of the above self-organizing model of arbor optimization also can serve as a mechanism of ganglion placement in the roundworm nervous system. Our prior work (Cherniak 1995) had found that the actual observed positioning of the ganglia in the worm was optimal, in that it required, out of 40 million alternative layouts, the least total wire length for the animal’s interconnections. We have now constructed a force-directed placement simulator, Tensarama, where each of the worm’s connections behaves like a micro weights/cords/pulley system. This vector-mechanical net outputs the actual minimized ganglion layout by converging on energy-minimization equilibrium at the actual positioning of the ganglia – without much susceptibility to local-minima traps (see Figure 5.2).

So, the hypothesis is that optimization accounts for a significant extent of observed anatomical structure. An additional hypothesis is that simple physical processes are responsible for some of this optimization. Combined, then, the picture is:

Physics → *Optimization* → *Neural structure*

That is, self-assembling brain structure “for free, directly from physics,” i.e., generated via simply exploiting basic physical processes, without need for

and mesh of springs; a sort of plate tectonics of the brain. Some of the methodological significance here is that discrete-state processes (e.g., as in mutation of genomic sequences (Mitchell 1996)) are not required; continuous-process models can suffice for neural wiring optimization. This work falls in the tradition of seeking simple underlying mathematical form in complex aspects of Nature, ranging from Pythagoras through D'Arcy Wentworth Thompson (1917/1961).

3 Non-genomic nativism

Hence, contrary to a "nature/nurture" dichotomy, a third possibility-zone – "nature" here includes not only biology, but also mathematics and physics. *A via media* between two extremes: (a) brains of course cannot grow like crystals, entirely non-genomically; (b) yet life must still play by the rules of the game, subject to mathematical and physical law. So, there is a division of labor between the genome and simple physical processes. The organism's genome is written not upon a *tabula rasa*, but upon a specifically pre-formatted, pre-printed form or slate, already inscribed with a significant proportion of structural information.

Perhaps such a picture of interpenetrating domains clashes with some metaphysics hardwired in human beings, with a category structure that draws a bright line between the animate and the inanimate. Also, the expression "non-genomic nativism" may sound like a bit of a solecism to sharper ears. (Piaget: "That which is inevitable does not have to be innate" (1970).) Of course, label choice is not of interest. The underlying point is that some complex biological structure – comparable in extent to genome-specified structure – is intrinsic, inborn, yet not genome-dependent.

One rationale for organisms to exploit such free anatomy is the vast mismatch of scale between brain structure and genome structure. The human brain is commonly characterized as the most complex physical structure known, yet the total information representation capacity of the human genome is comparatively small. After allowing for noncoding introns (about 95 percent of the total), the amount of brain-specific DNA available might amount to as little information as is contained in a desk dictionary (about 50,000 entries, \cong 100 Mb total). Hence, information on brain structure must pass through a "genomic bottleneck" constraint on DNA information-representation capacity (Cherniak 1988, 1992). Brain structure for free lowers this genome information-carrying load.

One caveat in interpreting genome information limitations concerns data compression. The key idea, from algorithmic information theory (Li and Vitanyi 1997; Chaitin 1987), is: for a given symbol sequence, what is the smallest program of a given format that will generate the sequence (and then halt)? A related concept (intimately connected with unsolvability of the Halting Problem) is that of, e.g., a five-state "Busy Beaver" Turing machine – a ten-line Turing machine program that generates the largest number of

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
PH (100.000000)	
AN (300.000000)	
RNG (440.000000)	
DO (506.000000)	
LA (564.000000)	
VN (744.000000)	
RV (948.000000)	
VCa (1856.000000)	
VCp (3856.000000)	
PA (4726.000000)	
DR (4810.000000)	
LU (4884.000000)	

Final layout popped out after: 100,000 iterations [final 100]

Tension Constant: 0.010000

Total Wirecost: 87802.750000 um

Figure 5.2 Runscreen for "Tensarama," a force-directed placement algorithm for optimizing layout of ganglia of the nematode *Caenorhabditis elegans*: that is, minimizing total length of interconnections. This vector mechanical simulation represents each of the roundworm's \sim 1,000 interconnections as a sort of micro-spring acting upon the horizontally movable ganglia (nervous system clusters) "PH," "AN," etc. (Connections themselves do not appear on runscreen, nor fixed components such as sensors and muscles.) The above screendump shows the final configuration of the system after 100,000 iterations (re-update cycles for forces and locations): the system has found the global minimum-cost positioning of the ganglia (with about 8.7 cm total of wire) – which is also the actual layout. In this way, physics suffices to generate this neuroanatomical structure, out of \sim 40 million alternative possible configurations (Cherniak 1995; Cherniak *et al.* 2002).

intervention of genes. Physics suffices; complex biological structure as self-organizing. So, besides conserving connections, another possible explanation for the extreme level of optimization observed might be just that it is a side-benefit of complex biostructure hitching a free ride from physics – in particular, for minimal-wiring neuroanatomy for dendrites, axons, and roundworm ganglia (Cherniak *et al.* 2002).

Now also we have observed even finer wiring optimization of cerebral cortex layout of cat and macaque monkey (Cherniak 2000; Cherniak *et al.* 2004). Thus, a harmony of physics and neuroanatomy via fluid-mechanics

1's for its size on an initially all-blank tape and halts; the most productive such program presently known outputs 4,098 1's (Dewdney 1993) (see Figure 5.3). Six-state Turing Machines are now known with much vaster output. Thus, one can ask, what is the minimum such program that emits the text of *Anna Karenina*, or the bitstream soundtrack of a Beethoven quartet performance, or that generates a given genome sequence? One could imagine such information compression occurring on the genome to help pack in extensive brain structure specification. However, it seems a mistake to go on to conclude that this data compression is in practice virtually unlimited; for compression/decompression itself entails computation costs. For instance, the above five-state busy beaver candidate requires over 11 million steps to generate its string of 1's. So, no free lunch here: storage space in effect is traded off for encode/decode time.

4 Misgivings

An uneasy "biology from physics" coda may help in articulating nongenic nativism. Proceeding from static structure to behavior, physiology: that is, let us explore instances of complex biological functions similarly originating directly from simple physical processes. For example, as explained above, Tensarama yields the neuroanatomical layout of *C. elegans* by simulating each of the worm's thousand connections as approximately a

[CST	INP]	→	[OUT	MCV	NST]
1	0	1	L	2	
	1	1	L	1	
2	0	1	R	3	
	1	1	R	2	
3	0	1	L	1	
	1	1	R	4	
4	0	1	L	1	
	1	1	R	5	
5	0	1	R	0	
	1	0	R	3	

Figure 5.3 Turing machine program that has been the contender for title of five-state "busy-beaver" – maximally productive TM program – without challenge for over a decade (Dewdney 1993). It takes 11,798,826 steps to generate 4,098 "1"s on an initially blank tape before it halts. Thus, the program illustrates ~1:100 data-compression, which genome encoding might exploit. However, the program also illustrates the high computation cost of message "decompression" (a human being with pencil and paper would take about a year to complete the computation, a TM simulator (Cherniak 1990) takes a few hours).

microspring; as the system proceeds to vector-mechanical equilibrium, it weaves the ganglia into the actually observed layout – a one-in-ten-million achievement. Of course, such complex anatomy-generation can in itself be pictured instead as a slow-motion intelligent action.

A corresponding behavioral instance would be to derive realistic locomotion – similarly from a mesh of springs. In this connection, the "Sodaconstructor" package yields strikingly lifelike crawling motion for a simulated "worm" device, walking for a genetic quadruped, etc. (Sodaplay.com). Yet all this animation is generated from what amounts to again just a mesh of springs (with masses), driven by a single simple sine wave. Another case is a random-path searcher device: a simple smooth sphere containing a motor-driven counterweight will robustly outperform robots with conventional obstacle detection and evasion chip software. Simulated annealing is an example of a system using only basic physical processes (thermodynamic temperature-schedule models) to solve search problems with interesting structure that have local minima traps (Kirkpatrick *et al.* 1983). Also suggestive is the idea of mechanisms that exploit chaos-theoretic (nonlinear dynamical) phenomena – e.g., intricate behavior from a simple device, such as a jointed pendulum (Thompson and Stewart 1986). One question is how far such a for-free-from-physics approach can proceed; what is the most complex behavior so derivable? (For examples of intricate behavior from simple rotational dynamics, see Walker 1985.)

So, the obvious leap to a conclusion would be a sunny picture, of complex neuroanatomical structure "for free, directly from physics." And similarly also of complex behavior derived directly from basic physical processes, the latter again partaking of the Pythagorean tradition of D'Arcy Thompson (1917/1961), of simple mathematical form in Nature. However, a darker picture also crystallizes, of a retrogression back to a familiar type of neo-reductionism: rather than progress, a recantation of a generation of philosophy of mind threatens.

Let us contemplate the development of mind/brain sciences in the post-behaviorist era, since c.1970. Functionalist/computationalist theories of mind emerged, opening the conceptual possibility of mental states not having to be identified just with physical hardware states, but instead with abstract software states (Putnam 1964). The corresponding idea in genetics was of course where we began the essay, the picture of "genome as program." However, by now some of us may wonder whether AI has perhaps not fulfilled its early promise. For instance, compare the 1966 "Eliza" conversationalist program (an anti-AI project, really) (Weizenbaum 1976) with the remarkably unimpressive "Alice" – recent 3-time winner of the Loebner Turing Test prize (ALICE 2004). AI and von Neumann-architecture machines both emerged a half century ago: if hardware had developed as has AI, we would still be using abacuses and sliderules – computers would merely be exotic laboratory confusions. Hardware and software engineering differ profoundly (Cherniak 1988). Perhaps faith here in the future of progress may begin to waver.

The above five-state Busy Beaver candidate, although of course itself a program, illustrates another potential Pyrrhic victory for computationalism. This TM table was discovered by a huge constrained brute-force search of many possible TM tables by a program written by Heiner Marxen, of the Technical University of Berlin. When one contemplates the TM's operation, it indeed displays a kind of inhuman, unintelligible elegance – it seems repeatedly to re-embroider segments of its tape, to re-use parts of its own code palindromically, etc. Rather than incomprehensibly huge, it is incomprehensibly compact (a “nano-kluge”). And so an idea emerges of unintelligible intelligence, of programs that might generate intelligent behavior, but without identifiable representations of their world, rules for using such models, etc. This corresponds to the above picture of intelligence for free from physics. What turns out to work efficiently may not be – indeed, may be antagonistic to – what is humanly comprehensible (Cherniak 1988). (The satirist Emo Phillips proposed, in another context, a universal nondenominational prayer: (approximately) “Lord, please arrange the Universe for my convenience.”)

Yet I suspect anyone who actually experienced the aridity of the empty-organism era could not regard its revival in any form with equanimity. The unintelligibly compressed programs recall the Ramsey sentences and Craig interpolations of 1950s instrumentalist philosophy of science (Ramsey 1931; Craig 1956). I conclude, or fail to conclude, with this anomie, this uneasy equivocation. A variation upon Marthe Arnold's lines that Henry Adams (1918/1973) took as motto comes to mind: To find oneself caught between worldviews, one declining, one struggling to bring itself into existence.

Acknowledgments

Some of the experimental work described here was supported by NIMH grant MH49867.

Bibliography

- Adams, H. (1918/1973) *The Education of Henry Adams*, Boston: Houghton.
- ALICE (2004) [2000, 2001 and 2004 Loebner “Turing Test” Prize winner] www.alicebot.org.
- Chaitin, G. (1987) *Algorithmic Information Theory*, New York: Cambridge University Press.
- Cherniak, C. (1988) “Undebuggability and cognitive science,” *Communications of the Association for Computing Machinery*, 31: 402–412.
- Cherniak, C. (1990) MindSet CogSci Courseware. PCOMP Logic Engine Package: TM* Turing Machine Modeller. www.glue.umd.edu/~cherniak/philcomp/.
- Cherniak, C. (1992) “Local optimization of neuron arbors,” *University of Maryland Institute for Advanced Computer Studies Technical Report*, No. 90–90 (1990); *Biological Cybernetics*, 66: 503–510 (1992).
- Cherniak, C. (1994a) “Component placement optimization in the brain,” *University*

- of Maryland Institute for Advanced Computer Studies Technical Report*, No. 91–98 (1991); *Journal of Neuroscience*, 14: 2418–2427 (1994).
- Cherniak, C. (1994b) “Philosophy and computational neuroanatomy,” *Philosophical Studies*, 73: 89–107.
- Cherniak, C. (1995) “Neural component placement,” *Trends in Neurosciences*, 18: 522–527.
- Cherniak, C. (2000) “Network optimization in the brain: NIMH Grant Application,” *University of Maryland Institute for Advanced Computer Studies Technical Report*, UMIACS-TR-2001-28.
- Cherniak, C., Changizi, M., and Kang, D. (1999) “Large-scale optimization of neuron arbors,” *University of Maryland Institute for Advanced Computer Studies Technical Report*, No. 96-78 (1996); *Physical Review E* 59: 6001–6009 (1999). <http://pre.aps.org/>.
- Cherniak, C., Mokhrzadza, Z., and Nodelman, U. (2002) “Optimal-wiring models of neuroanatomy,” in G. Ascoli (ed.), *Computational Neuroanatomy: Principles and Methods*, Totowa, NJ: Humana.
- Cherniak, C., Mokhrzadza, Z., Rodriguez-Esteban, R., and Changizi, B. (2004) “Global optimization of cerebral cortex layout,” *Proceedings National Academy of Sciences*, 101: 1081–1086.
- Craig, W. (1956) “Replacement of auxiliary expressions,” *Philosophical Review*, 65: 38–55.
- Crick, F. (1958) “On protein synthesis,” *Symposia of the Society for Experimental Biology*, 12: 138–167.
- DePomerai, D. (1990) *From Gene to Animal*, New York: Cambridge University Press.
- Dewdney, A. (1993) *The New Turing Omnibus*, New York: W.H. Freeman, ch. 39.
- Garey, M. and Johnson, D. (1979) *Computers and Intractability: A Guide to the Theory of NP-Completeness*, San Francisco, CA: W.H. Freeman.
- Keyes, M. (1999) “The prion challenge to the ‘Central Dogma’ of molecular biology, 1965–1991,” *Stud. Hist. Biol. and Biomed. Sci.*, 30: 1–19.
- Kirkpatrick, S., Gelatt, C., and Vecchi, M. (1983) “Optimization by simulated annealing,” *Science*, 220: 671–680.
- Lewis, H. and Papadimitriou, C. (1978) “The efficiency of algorithms,” *Scientific American*, 238: 96–109.
- Li, M. and Vitanyi, P. (1997) *Introduction to Kolmogorov Complexity and its Applications*, New York: Springer.
- Mitchell, M. (1996) *An Introduction to Genetic Algorithms*, Cambridge, MA: MIT Press.
- Piaget, J. (1970) *Structuralism*, New York: Basic.
- Putnam, H. (1964) “Robots: Machines or artificially created life?,” *Journal of Philosophy*, 61: 668–691.
- Ramsey, F. (1931) “Theories” in his *The Foundations of Mathematics and other Logical Essays*, ed. by R.B. Braithwaite, London: Routledge & Kegan Paul, 212–236.
- Ryle, G. (1949) *The Concept of Mind*, New York: Barnes & Noble.
- Sodaconstructor locomotor simulator package: www.sodaplay.com.
- Stich, S. (ed.) (1975) *Inmate Ideas*, Berkeley, CA: University of California Press.
- Stockmeyer, L. and Chandra, A. (1979) “Intrinsically difficult problems,” *Scientific American*, 240: 140–159.
- Thompson, D. (1917/1961) *On Growth and Form*, New York: Cambridge University Press.

- Thompson, J. and Stewart, H. (1986) *Nonlinear Dynamics and Chaos*, New York: Wiley.
- Walker, J. (1985) *Roundabout: The Physics of Rotation in the Everyday World*, New York: W.H. Freeman. (Cf. "tippetop" inverting top, "rattleback" spin-reverser, boomerang.)
- Watson, J. (1965) *Molecular Biology of the Gene*, 1st edn, New York: W.A. Benjamin.
- Watson, J. et al. (1987) *Molecular Biology of the Gene*, 4th edn, Menlo Park, CA: Benjamin/Cummings.
- Weizenbaum, J. (1976) *Computer Power and Human Reason*, San Francisco, CA: W.H. Freeman.