The brain as a Darwin Machine

William H. Calvin

For parallel computers to simulate our brains, we must face the fact that human beings have a better claim on the title Homo seriatim *than* Homo sapiens — *we're more consistently serial than wise.*

AMIDST all the hyperbole about thinking machines that has accompanied the emergence of large-scale parallel computers from their serial predecessors, we have begun to contemplate the prospect of simulating some of our brain's massive parallelism. But one immediately runs into a role reversal worthy of a Mozart opera: the most distinctively human higher brain functions are surprisingly serial.

Human beings are perpetually stringing things together: phonemes into words, words into sentences, concepts into scenarios — and then fussing about getting them in the right order. Our brain uses word-order rules to create a very productive language, with an infinite number of novel messages, rather than the several dozen standard interpretations associated with the several dozen cries and grunts of any other primate species. It is not our mellifluous voices that constitute a significant advance but rather our arrangement rules, the meaningful order in which we chain our utterances.

Further, talking-to-ourselves consciousness is, among other things, particularly concerned with trying to chain together memory schemata to explain the past and forecast the future. As literary critic Peter Brooks has said[1]:

> Our lives are ceaselessly intertwined with narrative, with the stories we tell and hear told, those we dream or imagine or would like to tell, all of which are reworked in that story of our own lives that we narrate to ourselves in an episodic, sometimes semiconscious, but virtually uninterrupted monologue. We live immersed in narrative, recounting and reassessing the meaning of our past actions, anticipating the outcome of our future projects, situating ourselves at the intersection of several stories not yet completed.

It is our ability to choose between such scenarios that constitutes our free will — though, of course, our choices are only as good as our imagination in constructing a wide range of candidate scenarios. Logical reasoning also seems dependent upon the rules of reliable sequencing. Our sophisticated projection abilities are very sequential: a chess master, for example, tends to see each board configuration not just after the next move but a half-dozen moves ahead, as several different scenarios.

Not that you could get much human-like language or scenario-spinning consciousness out of the ordinary serial computer — what we are probably talking about is parallel architecture being used to create a lot of serial paths from which to choose. And, perverse though it may seem, we are also likely to make intentional use of noise, good old randomness ('stochastic process' is the polite euphemism). I can hear an incredulous voice already: "Not only does he want us to waste our precious parallel computing power simulating an old-fashioned serial device, but he wants us to make our machine intentionally *noisy*?".

> **"Technology treats noise as an unwanted impediment, darwinism as a means of exploring new avenues. But here we see it as a stimulus to evolve redundant machinery — whose secondary uses may be revolutionary."**

Yet abandoning low-level reliability (and achieving overall reliability via stable superstructures) is very useful. Noise can be creative. Every time that you think of sex, you should remember that it is all about guaranteeing some randomness — shuffling the DNA deck during crossing-over when making sperm and ova. The invention of eukaryotic sex a thousand million years ago probably prompted the great Precambrian diversification of complex life forms into the familiar tree of species.

The brain's construction of chained memories and actions is probably another tree, though a more functional metaphor might be the candelabra-shaped railroad marshalling yard, with words for cars: imagine that many trains are randomly constructed on the parallel tracks, but only the best is selected to be let loose on the 'main track' of consciousness and speech. Best is determined by memories of the fate of somewhat similar sequences in the past, and one presumes a series of selection steps[2] that shape up candidates into increasingly more realistic sequences. This selection among stochastic sequences is more analogous to the ways of darwinian evolutionary biology than to the 'von Neumann machine' serial computer. One might call it a Darwin Machine[3] instead: it shapes up thoughts in milliseconds rather than millennia, and uses innocuous remembered environments rather than the noxious real-life ones.

Before pursuing such intracerebral Darwin Machines, consider some non-biological examples. Daniel Hillis has been using massive parallelism to create some competing computer programs. They mutate, surviving on the basis of how fast they can put a list of names into alphabetical order. Just using random variations on a basic program loop, his parallel computer has re-discovered many of the known sorting algorithms[4]. Similarly, the artist Harold Cohen's computerized drawing machine AARON makes aesthetically pleasing paintings using random variations and some general selection rules[5].

Toolmaking can operate the same way, and perhaps did so even two million years ago when hominids had ape-sized brains. The late Glynn Isaac used to demonstrate early toolmaking techniques during his archaeology lectures by pounding together two potato-sized rocks, not delicately but furiously: chips would soon be scattered all over the floor. After a minute, he would stop and sort through the dozens of stone flakes. And he would pick up some excellent analogues of the single-edged razor blade, just the thing for incising the tough hide of a savannah animal, or amputating a leg at the joint. This stochastic toolmaking is one round of a Darwin Machine: make lots of random variants by brute bashing about, then select the good ones. Perhaps another round of bashing resulted in a flake splitting, two sharp edges intersecting in a point. Careful craftsmanship probably developed where the raw materials were scarce.

Darwinian schemes[6] (some of which have the successive rounds of randomness plus shaping-up selection that mark them as members of the class I am calling Darwin Machines) have also emerged as partial explanations for

- bacterial food-finding[7] (*E. coli* intermittently randomize their swimming path by tumbling, but suppress tumbling when the nutrient concentration is increasing, and so select random paths that lead towards high concentrations);
- long-term memory consolidation[8] (selective retention, during a culling procedure, of randomly generated but subsequently used synapses, analogous to

photographic development retaining exposed silver grains[8]); and

- perceptual categorization via shaping up cortical interconnections (Gerald Edelman[9] notes that, in consequence of the random element, "we must look at all acts of perception as acts of creativity"). Given that half the cortical synapses are disconnected during childhood[10], there is again much opportunity for darwinian editing.

And, for at least a century[11], it has been recognized that even the highest-known biological function, human thought, involves random generation of many alternatives and is only shaped up into something of quality by a series of selections. Like the elegant eyes and ears produced by biological randomness, the Darwin Machine's final product (whether sentence or scenario, algorithm or allegory) no longer appears random because of many millisecond-long generations of selection shaping up alternative sequences off-line.

The Darwin Machines of particular interest here are the ones associated with chaining together actions (sequencing). Although they are often useful, command queues for detailed preplanning are seldom essential: goal-plus-feedback usually suffices, as when raising a cup to one's lips and getting progress reports from the joints and muscles. Where planned chains become essential, and thus likely to evolve rapidly, is where feedback becomes impossible, yet a linked series of moves must be precisely executed. Reaction time becomes a problem in brief ballistic movements such as hammering, throwing, clubbing or kicking: the progress reports will usually arrive too late for corrections to be made. For organisms that need to be both large (metres of conduction distance) and fast, one often needs the neural equivalent of an old-fashioned roll for a player piano. During 'get set', we carefully plan to act without feedback.

Grammar

Sequencing may involve much of the left cerebral hemisphere in mammals[12]. Left premotor cortex tends to program linked movements for not only the right hand and arm, but the left as well[13]. Left hemisphere is best at deciphering rapid sound sequences — and so it apparently became a natural home for many language-related abilities. Indeed, the core of human language cortex is a sequencing area for both incoming sounds and outgoing movements[14], just what grammar needs.

Comparison of grammars shows that the typical subject–verb–object word order of an English sentence is not biologically determined: Japanese syntax uses subject–object–verb, while classical Arabic puts the verb first. What the biology may provide is the serial buffer to hold the phrase while it is analysed according to the learned rules (though more subtle grammatical linkages are perhaps constrained by buffer branchings, corresponding to Chomsky's deep structure). There are some suggestions that the capacity of one important serial buffer is about a half-dozen items, judging from phenomena such as chunking of memory[15].

Yet there is surely more than one serial buffer: the human brain seems to orchestrate many sequences in parallel. Most are subconscious, with only one entering our 'main line' awareness (as in the railroad yard's parallel-to-serial bottleneck); traditional lines of evidence are scene-shifting in dreams, subconscious problem-solving and how subconscious scenarios sometimes pathologically intrude into speech.

Another suggestion of parallel sequencers has arisen from a very different direction. Human beings often hunt with projectiles; faster and farther throws are always better, provided accuracy can be maintained. A biophysical model for throwing[16] has emphasized the need for sub-millisecond timing precision, far in excess of what one can expect from noisy neurons[17] in a single command buffer. This suggests that the precision of the release of a projectile must arise from the Law of Large Numbers (the same rationale as why, in order to halve a standard deviation, one averages four times as much data).

Thus there must be many sequencers which, at least temporarily, can be ganged in parallel (imagine multiple columns of horses pulling a single wagon) during the occasions demanding peak performance in one-shot timing. To hit a rabbit-sized target reliably from twice the distance requires that the jitter in rock release time must be narrowed by a factor of eight, and the only known way of accomplishing this feat is, as one gets set to throw each time, to assign 64 times as many noisy neurons to the task and then average their recommendations for the release time.

Technology treats noise as an unwanted impediment, darwinism as a means of exploring new avenues. But here we see it as a stimulus to evolve redundant machinery — whose secondary uses may be revolutionary. There may even have been a 'noise window' in hominid evolution: lacking sufficient neuron noise to overcome, Ice Age hominids might have become proficient projectile predators without the massively serial scheme[18]. While timing precision is the argument for why so many parallel planning tracks were evolved in the first place, the really interesting things are the possible spare-time uses — if those extra buffers are capable of randomly sequencing other things when not needed for throwing–hammering–clubbing muscle commands.

If the separate tracks can also be unhitched to operate independently, then one might expect a Darwin Machine to emerge. By providing many candidate queues, it might foster stringing words together into more sophisticated sentences, or schemata into more credible scenarios. Rather than our productive language and planning-for-the-future consciousness arising gradually through their own selective advantages, they could have emerged as novel spare-time uses of neural machinery originally under selection for more mundane forelimb movements — much as a novelty called bird flight probably emerged willy-nilly as a consequence of natural selection for keeping warm via forelimb feathers (because it takes a lot of feathers to begin flying).

Serendipity

Neural-like networks[19–21], once they become capable of stochastic generation of sequences, then successive selections by remembered environments, do offer an obvious route to machine intelligence — though, should we succeed, we shall surely have to cope with machine imagination and machine 'free will'. We do not yet know how much of our own mental life might be explained by serendipitous secondary benefits of stochastic sequencers. But just as darwinian gradualism has been supplemented with notions of sexual and group selection, isolation and speciation, stasis and 'fast tracks', so we might expect a fuller understanding of our mental life to identify additional processes that regulate and elaborate the stochastic shaping-up of novel constructs in our intracerebral Darwin Machines. □

1. Brooks, P. *Reading for the Plot* 3 (Knopf, New York, 1984).
2. Dawkins, R. *The Blind Watchmaker* (Longman, London, 1986).
3. Calvin, W.H. *Whole Earth Review* **55**, 22–28 (1987).
4. Hillis, W.D. personal communication (1987).
5. McCorduck, P. *Whole Earth Review* **55**, 45–51 (1987).
6. Young, J.Z. *J. R. Soc. Med.* **72**, 801–814 (1979).
7. Segall, J.E., Block, S.M. & Berg, H.C. *Proc. natn. Acad. Sci. U.S.A.* **83**, 8,987–8,991 (1986).
8. Calvin, W.H. & Ojemann, G.A. *Inside the Brain* 65–67 (NAL, New York, 1980).
9. Edelman, G.M. in *How We Know* 24 (Nobel Conference, 1985).
10. Rakic, P., Bourgeois, J.-P., Eckenhoff, M.F., Zecevic, N. & Goldman-Rakic, P.S. *Science* **232**, 232–235 (1986).
11. Campbell, D.T. in *The Philosophy of Karl Popper* (ed. Schilpp, P.A.) 413–463 (Open Court, La Salle, Illinois, 1974).
12. Bradshaw, J.L. & Nettleton, N.C. *Behav. Brain Sci.* **4**, 51–91 (1981).
13. Kimura, D. *Phil. Trans. R. Soc.* **B292**, 135–149 (1982).
14. Ojemann, G.A. *Behav. Brain Sci.* **6**, 189–230 (1983).
15. Simon, H.A. *Models of Thought* **41** (Yale University Press, New Haven, 1979).
16. Calvin, W.H. *J. Theoret. Biol.* **104**, 121–135 (1983).
17. Calvin, W.H. & Stevens, C.F. *J. Neurophysiol.* **31**, 574–587 (1968).
18. Calvin, W.H. *The River That Flows Uphill: A Journey from the Big Bang to the Big Brain* 407 (Macmillan, New York, 1986).
19. Rumelhart, D.E., McClelland, J.L. & the PDP Research Group *Parallel Distributed Processing* (MIT Press, 1986).
20. Dehaene, S., Changeux, J.-P. & Nadal, J.-P. *Proc. natn. Acad. Sci. U.S.A.* **84**, 2,727–2,731 (1987).
21. Kleinfeld, D. *Proc. natn. Acad. Sci. U.S.A.* **83**, 9,469–9,473 (1986).

William H. Calvin is in the Biology Program, University of Washington, NJ-15, Seattle, Washington 98195, USA.