DAN RYDER

REVIEW ESSAY

MEDITATIONS ON FIRST NEUROSCIENCE: CRITICAL NOTICE OF MARK CHANGIZI'S *THE BRAIN FROM 25,000 FEET*

Mark A. Changizi, *The Brain from 25,000 Feet: High Level Explorations of Brain Complexity, Perception, Induction and Vagueness,* Kluwer Academic Publishers, Dordrecht, 2003, Hardcover, $115.00.

Neuroscience today is burdened by a massive accumulation of data, and a relative paucity of high-level theory relating brain to mind. One of the reasons for this may be the average neuroscientist's conviction that the brain must be studied from the bottom-up. One must first uncover the structural and causal details of the mechanism at the molecular, cellular, and network level. Gradually this will reveal what the mechanism does – but we are only getting the first glimpses of this.

In *The Brain from 25,000 Feet*, Mark Changizi argues that the average neuroscientist has it all bass-ackwards, or at least that this methodology is bound to miss something essential for understanding the brain. The mechanistic details, he suggests, won't make any sense in the absence of a "high-level" understanding of the brain, by which he means an understanding of the principles and constraints that must apply to *any* brain-like object, i.e., any thinking thing.

It would be lovely if Changizi could, in neo-Cartesian form, derive the function and structure of the fundamental cortical circuit from *a priori* mathematical truths. Of course, he does nothing of the sort, but I think he is successful in demonstrating how high-level understanding can enrich the sciences of the mind, often in surprising ways. I think he is less successful in illuminating neuroscience in particular, but that is perhaps the minor failing of a misleading title rather than a serious failure to support his main thesis. Chapter 1 is primarily neuroscientific, but the remainder of the book is about the brain only in the weak sense that *any* branch of cognitive science is ultimately about the brain, at least in part. Chapter 2 (on "perceiving the present" and optical illusions) *could* potentially be useful as a guide to

Synthese **141:** 277–285, 2004.

mechanistic discovery, but I fear that Chapters 3 (on induction) and 4 (on vagueness) would leave the practicing neuroscientist regretting the time spent reading them. (I don't mean to question their value in general, just their value for neuroscience.) Changizi makes a weak attempt to show that these chapters fit the book's title, but even he doesn't sound too convinced.

Changizi's higher level approach manifests itself in rather different ways in the four chapters. (I shall save my comments about the value of the approach until the end.) In the first chapter, he presents what is perhaps his best illustration of it. His basic strategy is to derive conclusions about the functional importance of certain aspects of complex systems, including the brain, merely by observing regular changes (or failures to change) in these complex systems as they "scale up". In the case of the brain, this "scaling up" is the increase in number of neurons in the brain's network across animal types, as one moves from small to large animals, and from simple to complex animals.

The general method for deriving functional conclusions from scaling laws can be illustrated by a simple game. Player One is given a ball of wax,[1] and is instructed to shape it so as to maximize, minimize, or keep constant some particular quantity. You are Player Two, and you must divine the instructions given to Player One by presenting them with further pieces of wax, and observing what they produce. First, Player One hands you a very flat circular piece of wax. At this point, it is impossible to tell what their instructions were. To produce a cylinder of constant diameter d? Or to minimize surface area? To distinguish these hypotheses, you could present Player One with a piece of wax having a larger volume. If they hand it back again very flat (with a larger diameter), and continue to do so no matter what volume of wax you present, you could reasonably infer that they were told to maximize surface area (for a given volume). This hypothesis would fully explain how surface area scales with volume in their bits of wax. It turns out that many properties of mammalian neocortex scale in regular ways with one another (specifically, in accordance with power laws). The properties Changizi considers include volume of grey matter, number of areas, number of areas to which an area connects (on average), neuron number, neuron density, cortical thickness, surface area, soma radius, axon radius, and volume of white matter. Just as your hypothesis that the instruction to maximize surface area fully explains how surface area scales with volume in Player One's bits of wax, Changizi argues that a handful of simple hypotheses fully explain how these properties of neocortex scale with one another. Perhaps the most interesting of these hypotheses is:

> The principle of economical well-connectedness–The number of areas to which a cortical area connects, and the density of these connections, remains invariant as the cortex scales up (this is "well-connectedness"). Further, this is done using the minimum amount of tissue (i.e., wiring) necessary (thus *economical* well-connectedness) (see pp. 11–15).

Surprisingly, this hypothesis can *fully explain* why, as cortex gets bigger, it tends to have more areas. (Areas are individuated by their dense local connections within them; areas interconnect via long distance connections traveling in the white matter.) Changizi infers that "the number of cortical areas increases in bigger brains, then, not because of some kind of pressure to have more specialized areas, but because by not increasing the number of areas the network would become decreasingly well-connected, or would no longer be economically wired" (p. 17). The same principle can fully explain why cortex becomes more convoluted and has more synapses per neuron. These apparent signs of increased functional complexity are nothing of the sort, says Changizi. They are simple consequences of the selective pressure for a network to stay well-connected, cheaply (p. 26).

It seems, though, that Changizi slightly overstates his conclusion here. Regular scaling relations among properties does not reveal order of influence. Perhaps the reason that brains (or neocortices) get bigger is precisely because of the advantages in functional complexity conferred by having more areas and/or synapses per neuron. This is not to deny that economic well-connectedness is a driving force, Changizi is quite convincing on this point. He is just a bit hasty to conclude that it is the *only* driving force (or perhaps just a little sloppy in stating his conclusion).

However, he does independently address increase in brain complexity, understood as the number of types of things a system can do ("expression types" of the system). How can a system increase its number of expression types? Well, it could start *de novo* at each increase, building a brand new expression type from scratch. More economically, it could adopt some compositional method of expression construction, where expressions have (or are produced by something which has) components. Starting from a number of basic types of expression components, a system can increase its complexity by putting together those basic component types in longer, more complex ways (the "universal language" approach), or it can increase its complexity by increasing the number of basic component types (the "invariant length" approach), or some combination of these two approaches.

Appealing to a large and varied class of systems (both naturally and socially produced), Changizi shows that as selectionally evolving systems scale up, they tend to follow the invariant length approach. Again, this can be determined by observing scaling regularities. The number of expression types *E*, is proportional to the number of component types *C*, raised to the power of how many component "slots", *d*, there are in an expression:

$$E \sim C^d$$

If a network uses the invariant length approach, simultaneous increases in the number of expression types and component types will be describable by the above equation, where *d*, the "combinatorial degree", will be a constant (and greater than 1).

The English language is an example of a system that has increased its expressive power over time at least in part through the creation of new component types (namely words). It turns out that English obeys the power law above, and its constant combinatorial degree is around 5 (or a bit higher, taking word extinction into account, p. 32). This means that in an average English sentence, there are about five slots for basic expression components. (The slots, or degrees of freedom, therefore do not correspond to words–Changizi provides some evidence that they correspond to *content* words, as opposed to *function* words like prepositions, conjunctions, auxiliary verbs, etc.) By contrast, birdsong has a combinatorial degree not significantly different from 1, i.e., it is not combinatorial at all. (This fundamental disanalogy, Changizi points out, calls into question attempts to draw lessons about language from birdsong–see, e.g., Doupe and Kuhl 1999.)

Taking neurons to be the equivalent of words in a language, Changizi shows that the neocortex most likely uses the invariant length approach in increasing complexity, and that its combinatorial degree is about 5. This is an exciting result, which could serve as an important test for mechanistic proposals of the cortical alphabet, as well as a guide to the discovery of such proposals. (Changizi speculates that these 5 degrees of freedom in the construction of functional cortical behaviors correspond to the 5-cell-rich layers in cortex.)

In Chapter 2, Changizi defends a non-standard Bayesian approach to perception. In the standard Bayesian approach, it is assumed that the perception generated by a stimulus represents the most probable cause of that stimulus. Changizi points out that, due to transmission and computation delays, this would result in our perceiving the *past*. What the perceptual system ought to do instead is use the stimulus to predict the *present*, i.e., what the scene is likely to be after transmission and computation is finished. (Indeed, why stop at the present? No doubt it would often be useful for a perceptual system to anticipate the future.) On the basis of this idea that our visual systems perform "latency correction", he is able to explain a large class of visual illusions (mostly the classical geometrical illusions like the Müller–Lyer and the Poggendorff), and, in ideal philosophy of science form, even predict some novel illusions (pp. 130, 146). In each case, the stimulus contains cues that the visual system reasonably construes as indicating the perceiver is moving forward, and infers that the scene will change in the next moment in accord with that forward motion. (There are alternative accounts of these illusions, which Changizi criticizes, most effectively on p. 85, where he shows that they fail to generalize.)

Just as in Chapter 1, Chapter 2 illustrates the higher level approach by demonstrating the fruitfulness of hypotheses of optimality. Given the physical impossibility of instantaneous computation, the best way for a perceptual system to operate is for it to "predict the present" (in accordance with good Bayesian principles). Supposing that our brains are in fact ideal in this way can explain facts about perception, and generate testable predictions. It could possibly even help guide the neuroscientific search for perceptual mechanisms.

In Chapter 3, Changizi defends "paradigm theory", which he describes as a "best we can hope for" solution to the riddle of induction (p. 231). The riddle of induction says that there is no rational way to choose among inductive policies, and thus no rational way to decide what we ought to believe given our evidence. Bayes' theorem decomposes this variability in inductive policy into a fixed component and a variable component. The fixed component is an *a priori* rule for modifying your beliefs about the world, given the evidence plus your prior beliefs about the world. Variability in inductive method is due to variability in these prior beliefs (prior probabilities). Thus the riddle of induction becomes: there is no rational way to choose among these sets of prior probabilities.

Changizi's goal is to push the variability in inductive method back a step further, so that the variability is accounted for, not by variability in undefended assumptions that are, like prior probabilities, *about the world*, but undefended assumptions that are *not* about the world. Sets of such undefended assumptions that lack empirical content he calls a "paradigm", and he presents some *a priori* rules, analogous to Bayes' theorem, for generating prior probabilities given any paradigm. The benefit is supposed to be in having an *a priori* method for deciding what we ought to believe, given our evidence, that does not rely upon substantive assumptions about the world. We need only start with one of the possible paradigms. (The riddle of induction survives, though, since there is no rational method for choosing a paradigm.)

What are these paradigms, such that they lack empirical content? Changizi compares them to conceptual frameworks; he also says that a paradigm is or determines the kinds of similarities and differences one allows among hypotheses. As far as I can tell, a paradigm is something that divides objects and states of affairs into kinds; alternatively, it describes what properties one is allowing into one's ontology. Only an extreme nominalist or conceptualist about properties would claim that this does not amount to a substantive assumption about the world. So whatever the value of the rules he proposes as analogous to Bayes' theorem, I don't think he's made the progress on the riddle of induction that he claims to.

But suppose that Changizi is right to maintain that paradigms are not substantive assumptions about the world. What does paradigm theory have to do with the brain? Changizi makes a rather weak connection here to *innateness*. It opens up the possibility, he suggests, that what is innate to a brain is a paradigm, rather than a set of substantive assumptions about the world (p. 169). (He regards the latter possibility as "a little preposterous" [p. 231], though it is not clear why. After all, most computational theories of perception take it to be obviously true.) He leaves it entirely open whether brains actually *do* have innate paradigms (plus a suite of principles of rationality). Besides, the notion of a paradigm is so abstract, it is not clear how it could carry *any* implications about the mechanism. This is a much less impressive demonstration of the higher-level approach than we saw in Chapter 1, where Changizi derived a mechanistically relevant hypothesis about the brain that was not merely possible, but probably actual.

In Chapter 4, Changizi tries to show that any thinking machine is bound to have vague lexical concepts. Concepts are construed as decision procedures, or programs, for classifying inputs, and learning the semantics of a language is construed as assigning such programs to substantive terms. Changizi assumes that the language learner is capable, at least in principle, of choosing from any of the programs that halts, as well as from ones that don't. Since there is no higher-level program for choosing only classificatory programs that are guaranteed to halt (i.e.,the halting problem is undecidable), the language learner is doomed to risk choosing classificatory programs that don't halt. (In fact, the language learner is vastly more likely to choose such programs.) Changizi takes the borderline regions of vague predicates to be inputs where the corresponding classificatory program does not halt. Higher-order vagueness arises because it is not in general possible to determine whether a program will halt on some given input, so it is not in general possible to determine the boundaries of these borderline regions. This chapter is interesting, but its relevance to the brain is hard to discern.

What, in the end, is the value of the higher-level approach? There are at least two ways in which Changizi takes it to prove its value. One way makes its first appearance in an extended analogy with which the book begins. In a post-apocalyptic future, some cavemen happen upon an entire working city, whose residents have disappeared. While they might learn the low-level causal patterns characterizing the various complex objects they find, this would leave them with a very poor understanding of those objects: traffic lights, cars, houses with their climate control and gadgets, and computers. True understanding of these things requires the understanding that the designers had – from physics, mathematics, and theoretical computer science, to specialized branches of engineering and principals of architecture (p.xiii).[2] Similarly for understanding the brain.

However, it's far from clear that the analogy carries over. Engineering design and natural design (by evolution) tend to be quite different. First, nature designs by trial and error, whereas the engineer designs through planning. As a result, there is a sense in which nature need not rely upon general principles as much as the engineer; rather, she can make use of incidental properties of the materials available. We can see this when human engineers adopt nature's approach, and design real robots using genetic algorithms (Clark 2001). Navigation robots "evolutionarily engineered" from real electronic circuits can rely on the transient dynamics or subtle output delays of these components, "incidental" features that a human engineer would never make use of. Indeed,

> it can be expected that all of the detailed physics of the hardware will be brought to bear on the problem at hand: time delays, parasitic capacitances, cross-talk, metastability constraints and other low-level characteristics might all be used in generating the evolved behaviour. (Thompson et al. 1996)

Understanding how these robots work requires little appreciation for higher-level general principles, and much knowledge of the mechanistic details, unlike Changizi's vision of the brain. Since the brain is naturally designed rather than engineered, we might expect it to be similarly opaque to higher-level conceptual tools.

That said, Changizi has made it clear that, at least sometimes, a higher-level understanding is crucial for understanding some aspects of the mechanism. Without a higher level understanding of economic well-connectedness, for instance, we would be in the dark about the significance of the number of cortical areas in different sized brains. By both drawing our attention to principles like economic well-connectedness, and showing clearly how a higher-level methodology can provide evidence for their truth, Changizi has done us a great service. No matter if he is a little over-enthusiastic about how *much* the higher-level approach is likely to reveal.

A second respect in which it seems Changizi takes the higher-level approach to exhibit its value is as a source of fruitful hypotheses, in particular optimality hypotheses. Disanalogies between engineering and natural design ought to temper our enthusiasm here as well. Like MacGyver,[3] nature is very limited in the materials she has available. She must make do with what she is given: for example, the lung was built by tinkering with the swim bladder of a fish (Jacob 1977). By contrast, the engineer is much less limited in this respect. The result is that nature tends to cobble together multiple subsystems which perform multiple tasks relatively poorly, and overcome their imperfections only through redundancy, where available. Engineering solutions tend to be streamlined and elegant by comparison (though not necessarily more effective, in the end).

So optimality hypotheses will often fail in the face of naturally designed systems, especially as guides to mechanistic discovery (since even if a natural system performs ideally, it will likely be composed of multiple, poorly performing redundant subsystems). Again, this is not to say that Changizi's project is fundamentally flawed. On the contrary, in the first two chapters at least, he provides substantial evidence that the brain *is* ideal in the respects he considers. He is just a bit over-optimistic about how often the idealization hypothesis is likely to pan out. (He is less convincing about applicability in Chapters 3 and 4, but I repeat, this may be more of a problem with the title than the book. Officially, his strategy is to use his conclusions about the limits of all possible brains to understand *our* brain, but in practice he often doesn't seem to care much about our brain except qua possible brain.)

This is a very rich and exceedingly well-written book. I learned a lot. While I hesitate to recommend anything beyond the first chapter to neuroscientists, most philosophers with a tolerance for technical material will read it with pleasure.

NOTES

[1] Note the Cartesian symbolism.

[2] It also requires an understanding of the *function* of each object and its parts. Changizi calls this a "higher-level" approach, but if so, it is a ubiquitous one. Anyone engaged in a biological science is interested in function, whether they work on the lowest mechanistic level (e.g., the function of the parts of the glutamate receptor), or higher levels like Changizi. So the value of this so-called "higher-level" approach is not news.

[3] MacGyver was the eponymous hero of a 1980s television show who was constantly making serviceable gadgets from everyday items, like a defibrillator from candles and an electric cord, or a fuel line from a ballpoint pen case.

REFERENCES

Clark, A.: 2001, *Mindware*, Oxford University Press, Oxford.

Doupe, A. J. and Kuhl, P. K.: 1999, 'Birdsong and Human Speech: Common Themes and Mechanisms', *Annual Review of Neuroscience* **22**, 567–631.

Jacob, F.: 1977, 'Evolution and tinkering', *Science* **196** (4295), 1161–66.

Thompson, A., Harvey, I. and Husbands, P.: 1996, 'Unconstrained Evolution and Hard Consequences', *University of Sussex Cognitive Science Research Report*, 397.

Dan Ryder
Department of Philosophy
University of Connecticut
344, Mansfield Road
Storrs, CT 06269-2054
U.S.A.
E-mail: dan@danryder.com