

Three Controversial Hypotheses Concerning Computation in the Primate Cortex

Thomas Dean and Greg S. Corrado and Jonathon Shlens

Google Inc., Mountain View, CA
{tld, gcorrado, shlens}@google.com

Abstract

We consider three hypotheses concerning the primate neocortex which have influenced computational neuroscience in recent years. Is the mind modular in terms of its being profitably described as a collection of relatively independent functional units? Does the regular structure of the cortex imply a single algorithm at work, operating on many different inputs in parallel? Can the cognitive differences between humans and our closest primate relatives be explained in terms of a scalable cortical architecture? We bring to bear diverse sources of evidence to argue that the answers to each of these questions — with some judicious qualifications — are in the affirmative. In particular, we argue that while our higher cognitive functions may interact in a complicated fashion, many of the component functions operate through well-defined interfaces and, perhaps more important, are built on a neural substrate that scales easily under the control of a modular genetic architecture. Processing in the primary sensory cortices seem amenable to similar algorithmic principles, and, even for those cases where alternative principles are at play, the regular structure of cortex allows the same or greater advantages as the architecture scales. Similar genetic machinery to that used by nature to scale body plans has apparently been applied to scale cortical computations. The resulting replicated computing units can be used to build larger working memory and support deeper recursions needed to qualitatively improve our abilities to handle language, abstraction and social interaction.

Introduction

We consider three claims regarding the primate *neocortex*, referred to simply as the *cortex* in the following:

- *Modular Minds Hypothesis* — The mind is primarily composed of stable cortical circuits which encapsulate specific cognitive competences and exhibit a high degree of structural and informational modularity.
- *Single Algorithm Hypothesis* — There is one fundamental algorithm that underlies all or most cortical computations; it is implemented on a computationally homogeneous cortical substrate and runs simultaneously in multiple instances on different inputs.

- *Scalable Cortex Hypothesis* — Simply scaling the inherent computational architecture of our closest non-human cousins — perhaps by varying the scaling factor among the differing functional areas — can account for much if not all of the cognitive differences between homo sapiens and non-human primates.

These hypotheses do not assume that the cortex can perform without considerable input from and interaction with sub-cortical neural circuits. We merely choose to focus our attention on the cortex, given its large volume, evolutionarily recent expansion in humans, and tantalizingly regular structure.

Supporting evidence for theories in neuroscience can originate from a number of disciplines including the following sample, all of which we make use of in the sequel:

- Histology — dissect, slice, stain, inspect and render brain tissue — from Santiago Ramón y Cajal’s meticulous drawings of neurons to the Human Connectome Project;
- Neurology — observing patients with head injuries and brain disorders — from John Harlow’s 1848 observations of Phineas Gage to the clinical study of stroke patients;
- Neurosurgery — therapeutic opportunities to stimulate and probe brains — from Roger Sperry’s split-brain epileptics to non-invasive radiotherapies for tumors;
- Neurophysiology — expose, probe, stimulate and record from experimental animals — from Hubel and Wiesel’s work on macaque striate cortex to current optogenetics;
- Neuroimaging — non-invasive recording of neural activity in awake humans — from Caton and Berger’s development of EEG to the modern surge of interest in Functional Magnetic Resonance Imaging (fMRI);
- Evolutionary Neuroscience — study of the fossil record and sequencing of genomes of descendant species — from Charles Darwin to modern paleobiology and molecular phylogenetics;

Modular Minds Hypothesis

Franz Joseph Gall believed that human cognitive capabilities could be usefully divided into separate functions and those functions localized in the brain, or, according to Gall, the

surface of the skull. He developed the discipline of *phrenology* to pursue this idea. While the particular functions identified by adherents of phrenology and their localization relative to the skull surface merit barely a footnote in modern textbooks, the notions of functional decomposition and locality of function continue to influence our investigation of the brain to this day.

Korbinian Brodmann's division of the cortex into distinct regions, or *areas*, based on the distribution and density of cell types, or local *cytoarchitecture*, provides a more principled way of mapping the brain. Some degree of modularity in cortical computation based on Brodmann's areas is incontrovertible. For example, Brodmann area 4 is the unique and final common pathway for all volitional body movement. Any cortical processes intending to move the arm, for example, must do so by acting through the interface of a specific sub region of this Brodmann area.¹ Brodmann area 3 provides a complementary and similarly organized cortical interface for the bulk of sensory information related to touch, and other areas provide gateways for other senses.

The general trends of modularity and functional localization seemed to hold for more complex functions as well. Brodmann areas 44 and 45 of the dominant hemisphere — the left in 90% of us — were identified with language production by Piere Paul Broca and are now associated with his name. Area 22, also in the dominant hemisphere, was linked to the comprehension of language by Carl Wernicke. In both cases, identification was made by studying patients with brain lesions who exhibited specific language deficits.

Patients with damaged Wernicke's area but with intact Broca's area often produce elaborate, syntactically-correct sentences devoid of meaning. Conversely, in the case of patients with lesions in Broca's area but an intact Wernicke's area, the meaning of words is preserved but sentence structure is impoverished. Early brain imaging studies using Positron Emission Tomography (PET) seemed to strengthen the case, showing highly focused activity in awake subjects performing language tasks. Such observations were taken as evidence for the localization of specific language functions, and lent support to the idea that the brain as a whole is organized as a set of anatomically distinct functional modules. Claims of strong functional localization beyond the arena of basic sensory and motor processing, however, have tempered somewhat in recent years. For example, though Broca's Area and Wernicke's Area clearly play critical roles in human speech, recent studies using fMRI (more sensitive than PET) have implicated a wider network of brain as participating in language, many of which are known to subserve other functions as well (Dehaene 2009). It would be a gross oversimplification to imagine that there is a spatially distinct "language area" in the brain. Similar observations in a variety of systems have given credence to a *distributed circuits* view of modularity, in which a network of brain areas might participate in supporting one particular com-

¹The surgeon Wilder Penfield famously cataloged the organization and demonstrated the sufficiency of activity in area 4 to evoke movement by electrically simulating the cortex and moving the patient's limbs on the operating table.

putation, and then an overlapping network of areas might support another. This distributed circuits view is compatible with Jerry Fodor's influential conception of cognitive modules based on specificity of input and informational encapsulation (Fodor 1984) rather than simple locality.

Some evolutionary psychologists have rejoined that cognitive modules may be better defined in terms of the problems they solve rather than the inputs on which they operate. Evolutionary selection pressures might then act on circuits that encapsulate a pattern of behavior rather than tinker with information representations directly. Leda Cosmides and John Tooby have argued that this idea extends all the way up to behavioral primitives such as reciprocity, social dominance, hygiene, incest prevention and coalition alignment (Tooby and Cosmides 1992). At first blush these behaviors may seem quite complex, but this is misleading in that they are likely constructed from a set of "hard-coded" cues and an ancient repertoire of basic communicative actions and responses. Social primitives confer such enormous advantage to the organism that it seems plausible that they might end up deeply coded in our brains and genome.

Modular organization of basic sensory and motor function having been well established, and hints of modularity in social behavior and language, some effort has been made to identify modular functions in the sphere of more general cognitive capacities. The emerging view of postnatal developmental, however, argues for more integrative view of general cognition and awareness. Human development is now believed to extend well into adulthood, with a marked shift from relying primarily on local uncoordinated activity in early adolescence to relying on coordinated distributed activity spanning a large fraction of the cortex in early adulthood and beyond. This so-called local-to-distant functional change is at least evident on high-level cognitive tasks of the sort historically associated with tests of intelligence (Fair et al. 2009; Sepulcre et al. 2010).

Single Algorithm Hypothesis

In 1978, Vernon Mountcastle (1978) characterized the cortex as a collection of anatomically if not functionally distinct structures which he referred to as *cortical columns*. Mountcastle proposed that "all parts of the neocortex operate based on a common principle, with the cortical column being the unit of computation", a view which was held by other leaders in the field at the time (Creutzfeldt 1977). Each *hypercolumn* — a bundle of smaller columnar structures — consists of approximately 60,000 cells and a thousand times that many connections, most of which span no more than a couple of millimeters. The cells within columns are themselves organized in several layers and the hyper columns are grouped into larger functionally-related cortical areas. Columns in a given Brodmann area exhibit the cytoarchitectural differences originally identified by Brodmann, but the cortex as a whole is viewed as an homogeneous sheet of densely packed cortical columns each running the same basic algorithm, albeit with different inputs.

One source of evidence supporting the single-algorithm hypothesis in primary sensory cortex involves "cross-modal" experiments in which the cortex is surgically rewired

so that visual stimuli are routed to auditory cortex. Normally in mammals, the retinal input from the right visual field is routed to the lateral geniculate nucleus (LGN) of the thalamus in the left hemisphere and to the superior colliculus (SC). Primary visual cortex (V1) in adult mammals exhibit *retinotopic* maps with the property that if two neurons receive input from the same local region of the retina, then these neurons tend to be close to one another in V1. This property is achieved by a combination of chemical-gradient-guided axonal growth and highly-correlated spontaneous firing patterns produced by the retina prior to eye opening. Apparently the brain is plastic enough that these maps can be adjusted and fine-tuned postnatally (McLaughlin and O’Leary 2005).

In the experiments involving ferrets reported by von Melchner *et al* (2000), the connections to the SC are rerouted to the medial geniculate nucleus (MGN) of the thalamus which serves the role of the LGN for auditory stimuli. To extinguish any auditory stimuli from reaching the rewired cortex, the SC is ablated and the normal input to the MGN from the auditory tract via the inferior colliculus (IC) is severed. Not only is the ferret able to perform visual tasks using its auditory cortex, but the auditory cortex develops a retinotopic map similar to that found in a normally functioning ferret visual cortex. The neurons in the ferret primary auditory cortex (A1) develop visual response characteristics including orientation selectivity, direction selectivity and an orderly retinotopic map.

While similar to V1, the way in which orientation selectivity is mapped in the rewired A1 is less orderly than in V1, perhaps exhibiting a bias stemming from its original auditory function (Newton and Sur 2004). Métin and Frost (1989) provide related experimental evidence in their work rerouting visual stimuli to the primary somatosensory cortex (S1) in hamsters. They show that S1 neurons exhibit some of the same orientation- and direction-selectivity characteristic of V1 neurons.

Saxe *et al* (2011) point to these cortical rewiring experiments as evidence of a single-algorithm at work, and their analysis of several biologically-inspired learning algorithms provides a plausible list of properties that such an algorithm might possess. Saxe *et al* suggest that several existing algorithms — including independent components analysis, sparse autoencoder neural networks, restricted Boltzmann machines and sparse coding, all of which yield similar results with no single algorithm obviously superior — might provide a constructive realization of Mountcastle’s theory that a “qualitatively similar learning algorithm acts throughout the primary sensory cortices” (Saxe *et al*. 2011). The restriction to the primary sensory cortices represents (a possibly necessary) weakening of the general hypothesis.

Todorov (2009) argues the single-algorithm implements a form of Bayesian inference that combines both bottom-up top-down components and that additional, intermediate levels of representations can be stacked to form a hierarchical system much as was described in earlier work by Lee and Mumford (2003). He suggests that the same basic algorithm can be applied to the motor cortex using the notion of a *motor synergy* corresponding to an abstraction of the musculo-

skeletal state as an analogy for primary sensory features. Such a probabilistic algorithm would satisfy the properties of the class of algorithms examined by Saxe *et al* and behave similarly, at least in performing purely feed-forward inference.

If the inference performed by the motor, somatosensory and primary sensory cortices can be explained by a unifying algorithmic basis, then we have progressed some way in validating the single-algorithm hypothesis. It would be particularly interesting if primary motor cortex could be so explained given its extensive connections to the cerebellar cortex and its role in controlling the execution of movement which has parallels with some aspects of language generation. Charles Sherrington’s work (1906) on muscle reflex circuits underscored the importance of control, feedback, and prediction in neural systems. Indeed, some scientists believe that the algorithmic foundations of sensorimotor control are more fundamental given that our sensory apparatus would not be of much utility if we weren’t able to move and interact with the world (Franklin and Wolpert 2011). Granted, these areas do not cover those believed responsible for language generation and comprehension, high-level multi-modal associations, and the executive-control functions which govern abstract, stimulus independent thought.

Koechlin *et al* (2000) claim the same basic organizational principles apply to both motor and executive control. The latter term they reserve to mean the temporal and hierarchical organization of action or thought in relation to internal goals. Koechlin and Jubault (2006) suggest the posterior portion of the prefrontal cortex including Broca’s area (which, as we mentioned above, is critical to speech production) and its homologue in the right hemisphere are primarily responsible for implementing these organizational principles, citing their fMRI studies of humans as supporting evidence.

It is tempting to speculate that the ability to reason about temporal order and hierarchies of tasks described in the Koechlin papers was a precursor to and enabler of the language facility in modern humans. With respect to the single-algorithm hypothesis, assuming something like this conjectured cognitive capability is present in humans, does it imply a different sort of algorithm than those conjectured by Saxe *et al* (2011) or Todorov (2009)? It would certainly seem to require more sophisticated inference than either of the earlier candidates. It could also be the case — again assuming something like the Koechlin model is true — that some or all of the features that serve to differentiate the Koechlin algorithm are actually present in primary motor and sensory areas as well.

Scalable Cortex Hypothesis

Charles Darwin (1859) believed that the “difference in mind between man and the higher animals, great as it is, is one of degree not of kind.” Today, there are many who would disagree, but Robert Sapolsky is not one of them. Sapolsky (February 2011) in answering a question following his talk “Are Humans Just Another Primate?” said that the biggest difference between chimp and human brains is that we have “three times as many neurons.” According to Sapol-

sky the genes responsible for this difference govern the number of rounds of cell division during fetal brain development.

Clearly Sapolsky is not saying that it is just the quantity of neurons but, rather, he is depending on all that follows from the interplay of the complicated genomic programs that initiate the additional stages of neurogenesis.² Understanding what makes us uniquely human from the perspective of our neural architecture is certainly more complicated than it was at the time when Darwin made his sweeping statement. We have a wealth of new information coming from sources that Darwin could not have anticipated. Unfortunately, the new data is incomplete and open to interpretation.

In comparing humans and non-human primates, we are most interested in chimpanzees and bonobos with whom our last common ancestor was 6–8 million years ago, and macaques and baboons with whom our last common ancestor was about 25 million years ago. We have some detailed comparative analyses of cortical cytoarchitecture, but the data sparsely samples the space of cortical areas, subject ages, and species we care about. Comparative histological studies are also hampered by the difficulty of reliably registering areal landmarks across species (Amunts, Schleicher, and Zilles 2007), and, in cases where neuroanatomical differences are uncontested, the differences don't necessarily imply variation in function (Semendeferi et al. 2001). With such caveats and given the data, we are more similar than different from our great ape cousins. Even when we go looking for differences in, for example, language areas we find intriguing parallels, e.g., homologues of Broca's and Wernicke's areas have been identified in non-human primates.

But there are also differences that hint at deeper divisions, e.g., both humans and most apes possess a distinctive class of so-called *spindle cells*, but these neurons are larger and more numerous in humans (Nimchinsky et al. 1999). Constantin von Economo (1929) is credited with discovering these large, uniquely-shaped neurons. They are found in multiple cortical areas, including anterior-cingulate cortex, fronto-insular cortex, and, most recently dorsolateral-prefrontal cortex. It is suspected these cells play a role in higher-level cognitive functions concerning decision making, value judgements, error recognition, appetite regulation and conscious awareness — hence their relevance in comparing human and non-human primates (Allman et al. 2010).

There are also conjectures concerning the possible role of spindle neurons in accelerating communication between distant cortical regions, thus suggesting a possible evolutionary advantage to explain their presence in large-brained mammals. Action potentials can propagate along an axon at anywhere from 1 to 120 meters per second depending on the size of the fiber and whether it is myelinated. The fact that spindle cells are quite large suggests they may increase impulse transmission speed, but there is more mystery than light to be had in applying these observations and conjectures to ex-

²Michael Gazzaniga (2008) is fond of pointing out that split-brain patients — whose corpus callosum is severed and the two hemispheres isolated from one another — effectively have half of the cortical neurons available to normal people, and yet they exhibit no loss of performance on standard intelligence tests.

plain how humans differ from apes.

The sequencing and comparative analysis of chimpanzee chromosome 22 with its human ortholog chromosome 21 was enormously exciting (Watanabe 2004), and we can expect this trove of data to be mined for some time to come.³ The FOXP2 gene is to be one of the genetic differences distinguishing humans from other animals and is of particular interest because it is implicated in language disorders (Enard et al. 2002). While Sapolsky drew upon some details of this analysis in making his claim, there are numerous other genomic differences that he chose not to emphasize, but which other researchers have found significant, e.g., reports of novel genes differentially expressed in the cortex during late neurogenesis (Mühlfriedel et al. 2007). Genomic tools are also at work in the use of microarrays to compare gene expression levels in chimpanzee and human (Preuss et al. 2004). Cáceres *et al* (2003) have identified more than one hundred genes that are more highly expressed in human than in chimpanzee cortex, in contrast to, say, the heart and liver where the number of genes more highly expressed in humans is approximately the same as the number of genes expressed at lower levels in chimps.

Humans appear to be more lateralized than non-human primates and some of the well studied asymmetries in humans, including those associated with language, are found in apes. Despite predictions of humans possessing more areas than apes the latest evidence suggests that humans did not require new cortical areas (Preuss 2009).

So is the human cortex substantially different from that of our closest cousins? It depends on whom you ask, what function you're interested in, and what you mean by substantial. Our guess is that there are some differences in the areas associated with the production and processing of language, visual perception, and the prefrontal cortex that can't be accounted for by a few extra rounds of cell division during fetal brain development. Surely these are likely to provide advantages, but keep in mind that as far as we know these differences primarily enhance capabilities that were already available to our closest relatives to some degree. Consider the computational advantages humans might get just from the extra rounds of cell division and how these advantages could improve existing capabilities in chimps.

Transistor count is sometimes given as a proxy for the performance of a new processor chip, but every computer scientist knows it is not the number of transistors or even the number of logic gates that matter, but how those components are organized. The reason transistor count is at all interesting is that processor architectures are modular and highly scalable. Registers, caches, processor cores and SIMD lanes all scale — more is generally better, though we have yet to figure out how to write code or build compilers that make the most of multi-core hardware.

In terms of morphology, natural selection is particularly adept at designing robust modular systems. Hox genes gov-

³However, as new data becomes available, there is a growing concern that we have substantially underestimated the genetic differences between human and chimp, which could be as high of 4% in the full DNA sequence by some estimates (Cohen 2007).

erning the basic body plans in most animals illustrate this modular design. The PAX-6 gene has the capability that if expressed in a fruit fly it builds a fruit-fly eye and if expressed in a mouse it builds a mouse eye (Callaerts, Halder, and Gehring 1997). Genomic regulatory programs control the number, type and location of body segments and appendages. Cell differentiation and migration in the developing embryo control the type of neuron, its placement, and the distinctive cytoarchitecture of different areas in the cortex. Much of this machinery is common to all vertebrates.

Scaling is fundamental in development, e.g., capillary networks have to be scaled to handle wide variation in tissue type to bring oxygenated, nutrient-rich blood to within reach of every cell in the body as well as carry away waste products. The developing embryonic brain employs diverse signaling mechanisms, e.g., following chemical gradients, to route axonal processes to their final or near final locations as well as induce topographic mappings that preserve local invariants, e.g., adjacent points of the visual field are represented in adjacent regions of the striate cortex.

Cortical neurogenesis can be divided into an early and a late period. The length of time and number of cell cycles spent in the early period of cell division determine the number of cortical columns. The length of time and the number of cell cycles spent in the later period is believed to determine the number of individual neurons within a cortical column (Gazzaniga 2008).

While humans have pretty much the same basic functional areas as chimps, there are differences in the morphology and connectivity of some areas. There is evidence that the prefrontal cortex has more neurons and some layers are more densely packed (Semendeferi et al. 2001). While there is no difference in the average number connections per neuron, humans have somewhat more intra-regional connections and somewhat faster inter-regional connections. The corpus callosum is somewhat thinner in humans, which may indicate some reduction in the communication needs between the hemispheres.

While more cortical columns and more densely packed neurons in layers could help to accelerate some computations, the biggest potential gains would likely come from an increase in the depth of combinatorial circuits that can be constructed from the neural substrate. The human brain can't implement stacks or recursion as we commonly do on von Neumann machines. Instead it must replicate structures and maintain information on the stack, perhaps using some form of what O'Reilly calls "limited variable binding" (O'Reilly 2006).

Granger (2006a) believes the algorithmic foundations of the cortex have their basis in *thalamo-cortico-striatal* (TCS) loops which can be combined to implement hierarchical inference, iterative algorithms and recursive embedding in language. Such loops are evident in most cortical areas including prefrontal cortex and are augmented with sub-cortical structures including the thalamus and basal ganglia. Granger conjectures human language "arises in the brain as a function of the number of TCS loops".

Deeper recursive embedding enables more complicated linguistic structures, and richer epistemological and emo-

tional theories of other minds. The ability to handle deeper layers of abstraction and richer compositional models would facilitate managing with larger cliques and more complicated social arrangements.

As far as higher-level and executive cognitive functions are concerned, the prefrontal cortex (PFC) is most often implicated. O'Reilly *et al* (2010) cite converging evidence "to suggest a hierarchy organized along the posterior-anterior anatomical axis of the PFC, with more anterior areas providing higher level control representations."

Tai-Sing Lee and David Mumford (2003) and Dileep George and Jeff Hawkins (2009) associate the layers of their hierarchical models with specific cortical areas and hence their models of the primate visual ventral pathway have layers corresponding to V1, V2, V4, IT, etc. Earlier papers by Mumford (1991; 1992) on cortico-thalamic loops make a similar assumption. This is fine except that it doesn't help the argument *vis a vis* why humans are cognitively more capable than non-human primates, since we have more or less the same number and arrangement of cortical areas as chimps and even macaques. The work by Granger (2006b; 2004) may not provide solid evidence for deeper combinatorial circuits, but it does provide some complementary hypotheses and interesting conjectures about how such circuits might be implemented in the cortical substrate — and sub-cortical structures in the case of the thalamic-cortico-striatal loops. .

Where does this leave us? There are significant limitations imposed on gathering information on primate brains. Evidence on the neural correlates of higher cognitive function in humans and apes is less direct — primarily based on noninvasive techniques such as fMRI, the observation of patients with stroke or tumor, and postmortem examination, and often less compelling than we would generally consider in the case of hypotheses concerning simpler organisms. Genomic information has been immensely useful in suggesting places to look, and postmortem studies provide insights into cytoarchitectural differences that might yield functional differences. But the questions are too intriguing to simply postpone until better brain-imaging and cell-recording technologies become available.

By some accounts, there is four-fold difference in brain size between humans and our closest primate relations and a "far greater size difference if just the anterior cortical areas underlying language abilities are concerned" (Granger 2006a). How could such differences confer computational advantages that might account for the observed cognitive differences between the species? Certainly a larger working memory and support for representing more complicated relationships might be at play, but we suggest here that the key is the ability to realize deeper combinatorial circuits which would enable us to handle longer chains of inference, deeper recursive embedding, and nested representational structures.

Discussion

There is substantial evidence for each of the three hypotheses: cortical modularity, a single common algorithm, and intrinsic scalability. In no case, however, is the question decided. Some degree of modularity is incontrovertible given

the predictable relationship between the location of brain injuries and their behavioral impacts, yet the best measurements of our time show that most cognitive functions entail the coordinated activity of a network of brain areas. The notion of a single cortical learning algorithm is suggested by the repetitive anatomy of the cortex and bolstered by direct experiments in the primary sensory cortices, but lacks a mature theory of what that single algorithm might be. Yet where these first two hypotheses seem most vague, in the abstract computing capability of the frontal cortex, the third hypothesis seems strongest. The cortical expansion of the human frontal cortex is our best explanation of our powerful cognitive abilities, and that conjecture of scalability would seem to rest, at least partially, on the validity of cortical modularity and the presence of a single cortical learning algorithm.

In general, a deeper stack, whether this be realized in software or by replicating cortical structures, allows for deeper procedural nesting and richer representations. The larger your memory and the more powerful the ability to make fine distinctions and infer complex relationships, the larger your social circle and the more subtle the laws that govern its behavior. Better means of transferring knowledge by acting to encourage mimicry and communicating to convey abstract ideas augment a society's ability to govern itself and create and share technology. In short, more computation can propel a species to civilization if it starts out with a sufficiently broad repertoire of basic physical and social behaviors.

The mystery of homo sapiens' dominance might also be resolved by appeal to our strong social instinct. Noting that apes have the capacity for abstract thinking and evidence localizing such function in the prefrontal cortex, O'Reilly (2006) suggests the possibility that the critical difference may be due not to the hardware, despite its being quantitatively superior, but to "the motivations that drive us to spend so much time learning and communicating what we have learned to others." And recent evidence (Shultz, Opie, and Atkinson 2011) supports the hypothesis that social behavior is deeply rooted in genetics and thus a "species has to operate with whatever social structure it inherits."

A skeptic might view the three hypotheses featured in this paper as belonging to the same class as Ernst Haeckle's (1883) recapitulation theory or Paul Maclean's (1990) hypothesis that the primate brain is comprised of a reptilian component (basal ganglia), a mammalian component (limbic system), and a primate component (neocortex) — conjectures that made sense until technology provided us the means to delve deeper into the brain's mysteries. But these three hypotheses are not based on superficial appearances which — as in the case of Haeckle and Maclean — will be rendered obsolete by the next generation of recording and imaging techniques. Structural (morphological) modularity, algorithmic parsimony, and scalable architectures are powerful computational principles which one might imagine natural selection would likely stumble on eventually. It might even be argued that these principles were necessary to produce these remarkably efficient computing devices which have enabled our species to survive long enough to realize the benefit of shouldering such a punishing metabolic load.

Acknowledgments

The authors would like to thank Fernando Pereira, Marc'Aurelio Ranzato, and Andrew Saxe for their helpful comments and discussions.

References

- Allman, J.; Tetreault, N.; Hakeem, A.; Manaye, K.; Semendeferi, K.; Erwin, J.; Park, S.; Goubert, V.; and Hof, P. 2010. The von Economo neurons in fronto-insular and anterior cingulate cortex in great apes and humans. *Brain Structure and Function* 214(5-6):495–517.
- Amunts, K.; Schleicher, A.; and Zilles, K. 2007. Cytoarchitecture of the cerebral cortex — more than localization. *NeuroImage* 37(4):1061–1065.
- Burgess, P. W.; Simons, J. S.; Dumontheil, I.; and Gilbert, S. J. 2005. The gateway hypothesis of rostral prefrontal cortex (area 10) function. In Duncan, J.; Phillips, L.; and McLeod, P., eds., *Measuring the Mind: Speed, Control, and Age*. Oxford, UK: Oxford University Press. 217–248.
- Cáceres, M.; Lachuer, J.; Zapala, M. A.; Redmond, J. C.; Kudo, L.; Geschwind, D. H.; Lockhart, D. J.; Preuss, T. M.; and Barlow, C. 2003. Elevated gene expression levels distinguish human from non-human primate brains. *Proceedings of the National Academy of Sciences of the United States of America* 100(22):13030–13035.
- Callaerts, P.; Halder, G.; and Gehring, W. J. 1997. PAX-6 in development and evolution. *Annual Review of Neuroscience* 20:483–532.
- Cohen, J. 2007. Relative differences: The myth of 1%. *Science* 316(5833):1836.
- Creutzfeldt, O. D. 1977. Generality of the functional structure of the neocortex. *Naturwissenschaften* 64:507–517.
- Darwin, C. 1859. *The Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*. London: John Murray.
- Dehaene, S. 2009. *Reading in the Brain: The Science and Evolution of a Human Invention*. Viking Press.
- Enard, W.; Przeworski, M.; Fisher, S. E.; Lai, C. S.; Wiebe, V.; Kitano, T.; Monaco, A. P.; and Paabo, S. 2002. Molecular evolution of FOXP2, a gene involved in speech and language. *Nature* 418(6900):869–872.
- Fair, D.; Cohen, A.; Power, J.; Dosenbach, N.; and Church, J. 2009. Functional brain networks develop from a local to distributed organization. *PLoS Computational Biology* 5(5).
- Fodor, J. 1984. *Modularity of Mind*. Cambridge, Massachusetts: MIT Press.
- Franklin, D., and Wolpert, D. 2011. Computational mechanisms of sensorimotor control. *Neuron* 72(3):425–442.
- Gazzaniga, M. S. 2008. *Human*. New York: Harper Collins.
- George, D., and Hawkins, J. 2009. Towards a mathematical theory of cortical micro-circuits. *PLoS Computational Biology* 5(10).
- Granger, R. 2006a. Engines of the brain: the computational instruction set of human cognition. *AI Magazine* 27:15–32.

- Granger, R. 2006b. The evolution of computation in brain circuitry. *Behavioral and Brain Sciences* 17–18.
- Haeckel, E. 1883. *The History of Creation*. London: Kegan Paul, Trench, Trübner & Co.
- Koechlin, E., and Jubault, T. 2006. Broca's area and the hierarchical organization of human behavior. *Neuron* 50(6):963–974.
- Koechlin, E.; Corrado, G.; Pietrini, P.; and Grafman, J. 2000. Dissociating the role of the medial and lateral anterior prefrontal cortex in human planning. *Proceedings of the National Academy of Sciences* 97(13):7651–7656.
- Lee, T. S., and Mumford, D. 2003. Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America* 2(7):1434–1448.
- MacLean, P. D. 1990. *The Triune Brain in Evolution: Role in Paleocerebral Functions*. Berlin: Springer.
- McLaughlin, T., and O'Leary, D. D. 2005. Molecular gradients and development of retinotopic maps. *Annual Review of Neuroscience* 28(1):327–355.
- Métin, C., and Frost, D. 1989. Visual responses of neurons in somatosensory cortex of hamsters with experimentally induced retinal projections to somatosensory thalamus. *Proceedings of the National Academy of Sciences* 86(1):357–361.
- Mountcastle, V. 1978. An organizing principle for cerebral function: the unit model and the distributed system. In Edelman, G., and Mountcastle, V., eds., *The Mindful Brain*. Cambridge, MA: MIT Press. 7–50.
- Mühlfriedel, S.; Kirsch, F.; Gruss, P.; Chowdhury, K.; and Stoykova, A. 2007. Novel genes differentially expressed in cortical regions during late neurogenesis. *European Journal of Neuroscience* 26(1):33–50.
- Mumford, D. 1991. On the computational architecture of the neocortex I: The role of the thalamo-cortical loop. *Biological Cybernetics* 65:135–145.
- Mumford, D. 1992. On the computational architecture of the neocortex II: The role of cortico-cortical loops. *Biological Cybernetics* 66:241–251.
- Newton, J. R., and Sur, M. 2004. Rewiring cortex: Functional plasticity of the auditory cortex during development. In Syka, J., and Merzenich, M. M., eds., *Plasticity of the central auditory system and processing of complex acoustic signals*. Springer Verlag. 127–138.
- Nimchinsky, E. A.; Gilissen, E.; Allman, J. M.; Perl, D. P.; Erwin, J. M.; and Hof, P. R. 1999. A neuronal morphologic type unique to humans and great apes. *Proceedings of the National Academy of Sciences* 96(9):5268–5273.
- O'Reilly, R. C.; Herd, S. A.; and Pauli, W. M. 2010. Computational models of cognitive control. *Current Opinion in Neurobiology* 20(2):257–261.
- O'Reilly, R. C. 2006. Biologically based computational models of high-level cognition. *Science* 314(5796):91–94.
- Preuss, T. M.; Cáceres, M.; Oldham, M. C.; and Geschwind, D. H. 2004. Human brain evolution: Insights from microarrays. *Nature Review Genetics* 5(11):850–60.
- Preuss, T. M. 2009. The cognitive neuroscience of human uniqueness. In Gazzaniga, M., ed., *The Cognitive Neurosciences, 4th Edition*. Cambridge, MA: MIT Press. 49–66.
- Rodriguez, A.; Whitson, J.; and Granger, R. 2004. Derivation and analysis of basic computational operations of thalamocortical circuits. *Journal of Cognitive Neuroscience* 16:856–877.
- Sapolsky, R. February, 2011. Are humans just another primate? <http://fora.tv/2011/02/15/Robert.Sapolsky.Are.Humans.Just.Another.Primate>, California Academy of Sciences.
- Saxe, A.; Bhand, M.; Mudur, R.; Suresh, B.; and Ng, A. 2011. Unsupervised learning models of primary cortical receptive fields and receptive field plasticity. In Shawe-Taylor, J.; Zemel, R.; Bartlett, P.; Pereira, F.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems 24*. MIT Press.
- Semendeferi, K.; Armstrong, E.; Schleicher, A.; Zilles, K.; and Hoesen, G. W. V. 2001. Prefrontal cortex in humans and apes: A comparative study of area 10. *American Journal of Physical Anthropology* 114(3):224–241.
- Sepulcre, J.; Liu, H.; Talukdar, T.; Martincorena, I.; and Yeo, B. 2010. The organization of local and distant functional connectivity in the human brain. *PLoS Computational Biology* 6(6).
- Sherrington, C. 1906. *The integrative action of the nervous system*. New Haven, CT: Yale University Press.
- Shultz, S.; Opie, C.; and Atkinson, Q. D. 2011. Stepwise evolution of stable sociality in primates. *Nature* 479(7372):219–222.
- Todorov, E. 2009. Parallels between sensory and motor information processing. In Gazzaniga, M., ed., *The Cognitive Neurosciences, 4th Edition*. Cambridge, MA: MIT Press. 613–624.
- Tooby, J., and Cosmides, L. 1992. The psychological foundations of culture. In Barkow, J.; Cosmides, L.; and Tooby, J., eds., *The Adapted Mind: Evolutionary psychology and the generation of culture*. New York: Oxford University Press.
- von Economo, C., and Koskinas, G. N. 1929. *The cytoarchitectonics of the human cerebral cortex*. Oxford Medical Publications. London: Oxford University Press.
- von Melchner, L.; Pallas, S. L.; and Sur, M. 2000. Visual behaviour mediated by retinal projections directed to the auditory pathway. *Nature* 404(6780):871–876.
- Watanabe, H. *etal.* 2004. DNA sequence and comparative analysis of chimpanzee chromosome 22. *Nature* 429(6990):382–388.