

Published in final edited form as:

Top Cogn Sci. 2009 January 1; 1(1): 145–172. doi:10.1111/j.1756-8765.2008.01007.x.

How Does the Mind Work? Insights from Biology

Gary Marcus

Department of Psychology, New York University

Abstract

Cognitive scientists must understand not just what the mind does, but *how* it does what it does. In this paper, I consider four aspects of cognitive architecture: how the mind develops, the extent to which it is or is not modular, the extent to which it is or is not optimal, and the extent to which it should or should not be considered a symbol-manipulating device (as opposed to, say, an eliminative connectionist network). In each case, I argue that insights from developmental and evolutionary biology can lead to substantive and important compromises in historically vexed debates.

Keywords

Language; Connectionism; Domain-specificity; Cognitive development; Innate; Learning; Symbol manipulation; Syntactic representation; Syntactic trees; Linguistics; Cognitive architecture; Cognition; Optimality; Evolutionary psychology

1. Cognitive architecture

One mission of cognitive science is to figure out what the mind does: What is the capacity of short-term memory? Are there genuine gender differences, in say mathematics or navigation? Does the mind of a child fundamentally differ from that of an adult? Do nonhuman animals have a robust theory of mind?

Another mission is to figure out *how* the mind does whatever it does. Is the mind in some way like a computer? Does it compute in serial or parallel? Does it manipulate symbols? Is it a statistical inference engine?

In some sense, cognitive science was founded on “how” questions like these, not just questions about what people know about specific domains, but questions about *mechanism*, and how it is that the trick of biologically instantiated cognition is accomplished at all.

Ironically, one of the first explicit answers to that question came from behaviorism, which viewed the mind as a sort of black box that calculates contingencies between stimulus and response. As is now well known, this behaviorist notion that the mind consisted solely of stimulus–response associations ultimately fell short. In 1959, in an article that many see as pivotal to the formation of the very field of cognitive science (Gardner, 1985), Noam Chomsky issued a devastating critique of B. F. Skinner’s theory of verbal behavior (Chomsky, 1959). Chomsky argued that behaviorism said too little about how creatures like humans deal with novel stimuli, and that it seemed hopeless as a theory unless some sort of internal variables—what we would now recognize as mental representations—were incorporated.

Cognitive science formed, in no small part, as an effort to do better. Chomsky's devastating critique of Skinner came just as the computer revolution was beginning, and in those heady early days of cognitive science, a main goal of cognitive science was to figure out the nature of what Allen Newell later dubbed "cognitive architecture." If the mind does not just calculate stimulus–response contingencies, what *does* it do?

Interest in such questions was steady throughout the 1960s, 1970s, 1980s, and 1990s—yet, judging by the number of publications, has declined precipitously in the last several years. Discussions of whether the mind consists of "connections" or "symbols," for example, were wildly popular in the late 1980s and early 1990s yet receive very little attention in contemporary cognitive science. The question of whether the mind is a symbol processor or parallel distributed processing device (or something else altogether) remains foundational, but graduate students entering in the field now could easily be excused if they had only the dimmest notion of what that debate was about. No consensus was ever reached, but the issue has largely vanished from the discourse.

Nonetheless, questions of cognitive architecture—particularly the remarkable cognitive architecture of human beings—remains as profound as they ever were. Even as the power of synthetic minds—digital computers—has followed Moore's law and doubled, doubled and redoubled, again and again, many times over, no software engineer has devised a machine that can really match the overall power of the human mind. Although computers plainly exceed human beings in their raw processing speed and memory capacity, they still lack common sense, lag behind in their capacity to interpret the visual world, and fail spectacularly in the mission of understanding natural language. How do the networks of neurons in our 3 lb brains combine to yield the exquisitely powerful (if sometimes maddeningly limited) cognitive capacities of our species?

1.1. Four candidate answers

Broadly, cognitive scientists have put forward essentially four sorts of answers.

1. The human mind is a symbol-manipulation device (Fodor & Pylyshyn, 1988; Marcus, 2001; Newell, 1980) that consists in large part of a set of evolved modules that have been optimized to the needs of our Paleolithic ancestors (Cosmides & Tooby, 1992; Pinker, 1997).
2. The human mind consists of a set of heuristics and biases (Gigerenzer & Selten, 2002; Kahneman, Slovic, & Tversky, 1982; Tversky & Kahneman, 1974).
3. The human mind consists of a large network of neurons, initially connected in largely random ways, which become tuned through experience (Elman et al., 1996; McClelland, Rumelhart, & Hinton, 1986).
4. The human mind is a probabilistic inference engine, which may be near optimal in its capacity for core inference (Anderson, 1990; Chater, Tenenbaum, & Yuille, 2006; Kording & Wolpert, 2004).

Although these positions are not entirely mutually exclusive—for example, it is not obvious that options 2 and 3 are necessarily in contradiction—they offer a good starting point for discussion. Looking more closely, those four stances all revolve around four sets of core questions:

1. *Development*: To what extent is the fundamental structure of the mind organized in advance of (or independently of) experience?
2. *Modularity*: Does the mind consist, in whole, or in part, of informationally encapsulated computational units, à la Fodor (1983)?

3. *Optimality*: To what extent are the mind's operations optimal?
4. *Symbol manipulation*: Does the mind in fact manipulate symbols, in something like the fashion of a computer program?

1.2. The need for compromise

To the extent that researchers have waded into these controversies, they have often defended extreme positions. Either symbol-manipulation is the essence of human cognition (Marcus, 2001) or it is scarcely relevant (Rumelhart & McClelland, 1986b). Either the mind is "massively modular" (Carruthers, 2006), or modularization is, at best, an end product of development (Elman et al., 1996; Karmiloff-Smith, 1992), perhaps a consequence of the way the world itself is organized, rather than an intrinsic part of cognitive architecture (Elman et al., 1996). Either the mind, at least in its core computations, is an engine for deriving optimal statistical inference (Chater et al., 2006), or the mind is riddled with imperfections (Tversky & Kahneman, 1974). Either the basic architecture of the mind is specified genetically, largely in advance of experience (Pinker, 1997), or it's just an acquired response to the nature of the world (Elman et al., 1996; Watson, 1925).

Extremes often make for entertaining reading, yet they may also be a sign: To the extent that intelligent people differ on each of these debates, one must ask why. Could it be, for example, that bright people disagree because they focus on different elements of the evidence, a la classical debates about whether light is a wave or a particle? Is any sort of compromise possible?

Not all compromise, to be sure, is desirable. As Richard Dawkins once remarked, creationism and evolution cannot both be right, and bland pronouncements like "it's not nature or nurture, but both," while plausible, tell us little. If nature and nurture work together (as seems obvious), *how* do they work together? Compromise for compromise's sake is of no value.

At first glance, the possibility of genuine compromise might seem unrealistic; for decades and even centuries, the debates in question have been as recalcitrant as they are important. Innateness questions have been discussed since at least the time of Plato and Aristotle, and the rationality debate can probably be traced just as far. (Aristotle's remark that "man is a rational animal," for instance, might be juxtaposed with Bertrand Russell's sardonic retort: "It has been said that man is a rational animal. All my life I have been searching for evidence which could support this.") Modularity and connectionism are newer to the scene, but both have been around in some form or another for quite some time; the debate about modularity can be traced back to Gall's early 19th century discussions about faculty psychology, while the debate on connectionism goes back at least to the debate between Chomsky and Skinner.

In order to break out of these cycles, my strategy, over the last decade, has been to draw on the new field of "evo-devo" (e.g., Carroll, Grenier, & Weatherbee, 2001; Gerhart & Kirschner, 1997)—a synthesis of developmental and evolutionary—in order to bring new insights into the origins, nature, and development of human cognition. Developmental biology seeks to characterize transformation from genes to living, breathing organism, while evolutionary biology seeks to characterize the nature of evolution and how it has constrained biological form.

In what follows, I will try to argue that a series of core ideas borrowed from biology can cast significant light on the nature of cognitive architecture—and point the way toward precisely the sort of substantive compromises that we so desperately need.

2. Development

Few questions about cognitive architecture would be interesting were it not for the fact that there seem to exist a considerable number of human universals. If the mind were infinitely plastic as in John Watson's famous boast,¹ there might be nothing to say about how the mind works in general. Some humans might grow up to symbol-manipulators, others parallel distributed processors, some possessed of modular minds, others not, and so forth, all aspects of cognitive architecture open, determined sheerly as a function of experience, rather than prior constraint.

Instead, both within the species, between individuals, and within families of genetically related individuals, there is plenty of reason to think that development is constrained. On the one hand, despite considerable cultural variation, anthropologists have uncovered an abundance of cognitive universals (Brown, 1991; Pinker, 1997) and a reasonable degree of between-individual consistency in brain organization (Damoiseaux et al., 2006; Hasson, Nir, Levy, Fuhrmann, & Malach, 2004). On the other hand, such species-wide universals are complemented by the observation that—contra Watson's extreme empiricism—even when reared apart, identical twins resemble each other more than nonidentical twins in a vast array of psychological traits (Bouchard, 1994). On just about any behavioral trait that has ever been measured, there is some degree of heritability. Contra Watson's extreme view, development is not a blank check purely shaped by environment, but a highly constrained process in which experience and environment work hand and hand with prior genetic constraints. Mating behavior, for example, is constrained in part by culture, yet tiny genetic differences (e.g., in the alleles of the gene that govern the production of vasopressin receptors) appear to have considerable consequences, both in humans (Walum et al., 2008) and other animals (Pitkow et al., 2001).

Still, opponents of nativism have emphasized the fact that the brain is remarkably plastic (Merzenich et al., 1984), and the amount of information packed into the genome seems relatively small compared to the complexity of the brain itself (Bates et al., 1998; Ehrlich, 2000). This led scholars like Quartz and Sejnowski (1997) and Elman et al. (1996) to argue that innateness, at least in any strong form, is biologically implausible.

The core debate is, as it always has been, about nature and nurture. Is the mind's basic program "specified by our genetic program" (Pinker, 1997, p. 21) or the "endproduct of development, rather than its starting point" (Elman et al., 1996)?

The easy way out, taken by the cultural critic Louis Menand, is to wish the problem away: "Every aspect of life has a biological foundation in exactly the same sense, which is that unless it was biologically possible it wouldn't exist. After that, it's up for grabs" (Menand, 2002). Any number of scholars have suggested that we scrap the nature–nurture distinction altogether, substituting catch phrases like "interactions all the way down" (Elman et al., 1996; Spencer et al., 2009).

But that easy way out tells us too little. Interactionism is surely true as far as it goes, but in the final analysis, such facile pronouncements bring us no further toward understanding how things work, what is likely and what is not, and so forth. Genes and the environment both clearly contribute to the development of the mind and brain, but their contributions are of very different sorts, shaped over vastly different time scales, with radically different consequences.

¹"Give me a dozen healthy infants, well-formed, and my own specified world to bring them up and I'll guarantee to take any one at random and train him to become any type of specialist I might select ... regardless of his talents, penchants, tendencies, [and] abilities ..." (Watson, 1925).

None of which can be understood, I would argue, without some basic grasp of ideas imported from developmental biology.

2.1. What the genome is—and is not

Perhaps the single most important thing that biology can teach cognitive scientists is this: Genetic does not equal “specified by a blueprint”—not withstanding many straw arguments against nativism. Ehrlich (2000), for example, has argued that the role of genes in brain development must be minimal because the number of neurons vastly outstrips the numbers of genes. Similarly, as part of larger antinativist argument, Bates et al. (1998) argued that “On mathematical grounds, it is difficult to understand how 10^{14} synaptic connections in the human brain could be controlled by a genome with approximately 10^6 genes.”

But, as I argued in Marcus (2004), the notion of prewired neural structure in no way rests on genomes being blueprints. There is no one-to-one mapping between genes and neurons, but there need not be: The role of individual genes is not so much to give pixelwise portraits of finished products as to provide something far more subtle: an *environmentally sensitive* set of instructions for constructing (and maintaining) organisms.

An appreciation of that point can go a long way toward resolving some of the ongoing conflict that has characterized developmental psychology. To begin with, as reviewed by Marcus (2004) but discovered by Jacob and Monod (1961), every individual gene has two functions. First, each gene serves as a template for building a particular protein. The insulin genes provide a template for insulin, the hemoglobin genes give templates for building hemoglobin, and so forth. Second, each gene contains what is called a *regulatory* sequence, a set of conditions that govern the circumstances in which that gene’s template gets converted into protein. Although every cell contains a complete copy of the genome, most of the genes in any given cell are silent. Your lung cells, for example, contain the recipe for insulin, but they do not produce any because in those cells the insulin gene is switched off (or “repressed”); each protein is produced only in the cells in which the relevant gene is switched on.

This basic logic applies as much to humans as to bacteria, and as much for the brain as for any other part of the body. Monod and Jacob aimed to understand how a bacterium (*Escherichia coli*) could switch almost instantaneously from a diet of glucose (its favorite) to a diet of lactose (an emergency backup food), but the same applies more generally. What they discovered was that the bacterium’s abrupt change in diet was accomplished by a process *that switched genes on and off*. To metabolize lactose, the bacterium needed to build a certain set of protein-based enzymes that for simplicity I will refer to collectively as lactase, the product of a cluster of lactase genes. Every *E. coli* had those lactase genes lying in wait, but they were only expressed—switched on—when a bit of lactose could bind (attach to) a certain spot of DNA that lay near them, and this in turn could happen only if there was no glucose around to get in the way. In essence, the simple bacterium had an IF-THEN—if lactose and not glucose, then build lactase—that is very much of a piece with the billions of IF-THENS that run the world’s computer software.

IF-THENS are as crucial and omnipresent in brain development as they are elsewhere. To take one recently worked-out example, rats, mice, and other rodents devote a particular region of the cerebral cortex known as barrel fields to the problem of analyzing the stimulation of their whiskers. The exact placement of those barrel fields appears to be driven by a gene or set of genes whose IF region is responsive to the quantity of a particular molecule, fibroblast growth factor 8 (FGF8). By altering the distribution of that molecule, and hence the input to the IFs of genes that are downstream of FGF8, researchers were able to alter barrel development. Increasing the concentration of FGF8 led to mice with barrel fields that were unusually far

anterior, while decreasing the concentration led to mice with barrel fields that were unusually far posterior.

Moreover, here is where one can appreciate not just the mere existence of gene–environment interactions but a key mechanism by which it can come to pass: The IFs of genes are responsive to the environment of the cells in which they are contained. Even a single environmental cue can radically reshape the course of development. In the African butterfly *Bicyclus anynana*, for example, hot temperatures during development (associated with rainy season in its native tropical climate) lead the butterfly to become brightly colored; cold temperatures (associated with a dry fall) lead the butterfly to become a dull brown. The growing butterfly does not *learn* (in the course of its development) how to blend in better—it will do the same thing in a lab where the temperature varies and the foliage is constant; instead, it is genetically programmed to develop into two drastically different ways in two different environments. Because genes incorporate IF, and not just THENs, processes of cellular development, such as cell division and cell migration, are dynamic, not static.

This in turn led me (Marcus, 2004) to the realization that the classic dichotomy between “hard-wired” nativism and the “plasticity” championed by anti-nativists was woefully off the mark. Historically, “Anti-nativists”—critics of the view that we might be born with significant mental structure prior to experience—have often attempted to downplay the significance of genes by appealing to neural plasticity, viz. the brain’s resilience to damage and its ability to modify itself in response to experience, while nativists often seem to think that their position rests on downplaying (or demonstrating limits on) plasticity.

In reality, nativism and plasticity are not in conflict; indeed, researchers such as Marler (1991) have often talked about innately guided learning. As I suggested in Marcus (2004), it may be more profitable to draw a distinction, between *prewiring* and *rewiring*—each of which can be had in abundance without precluding the other. On this view, innateness is about the extent to which the brain is prewired, plasticity about the extent to which it can be rewired. And these are simply two different sorts of processes, each underwritten by distinct (though not necessarily disjoint) sets of genes. Some organisms may be good at one but not the other; chimpanzees, for example, may have intricate innate wiring yet, in comparison to humans, relatively few mechanisms for rewiring their brains. Other organisms may be lacking in both: *Caenorhabditis elegans* worms have predictable but somewhat limited initial structure, and relatively little in the way of techniques for rewiring their nervous system on the basis of experience. And some organisms, such as humans, are well endowed in both respects, with enormously intricate initial architecture and fantastically powerful and flexible means for rewiring in the face of experience.

Crucially, an understanding of the dynamics of genes allows us to recognize the fact that rewiring, like prewiring, has its origins in the genome. Memory, for example, is the way of rewiring connections between neurons (or altering something inside individual neurons) that plainly depends on gene function. When one interferes with the process of synthesizing proteins from genes, one interferes with memory. In keeping with this view of genes as underwriting plasticity itself, closely related organisms (such as *Aplysia*, the sea slug that has been the central experimental animal for Nobel laureate Eric Kandel, and a closely related species, *Dolabrifera dolabrifera*) can differ significantly in their talents for learning. Apparently as a function of small differences in their genomes, both species are capable of associative learning, but only *Aplysia* appears capable of nonassociative learning (Wright, 1998). What an organism learns depends in no small part on what genes for learning it is born with.

Innateness and learning are thus not in complementary distribution, but rather distinct contributors with distinct but identifiable roles. Genes, shaped by vast expanses of time, supply

instructions for building proteins; learning is a process by which on-line, in-the-moment, in-the-lifetime-of-an-individual cues can undergird neural rewiring. There is no genuine conflict between the two, between prewiring and rewiring. If the human mind is neither rigid and fixed (like some caricature of nativism) nor arbitrarily malleable (as Watson argued), it is precisely because genes are conditional recipes rather than blueprints.

3. Modularity

Innateness is one thing; domain specificity (the extent to which brain structure is specialized for particular tasks) is another. But, sociologically speaking, much of the time—perhaps the vast majority of the time—when people talk about the former, they often really have in mind the latter. For example, debates about the existence of a language “faculty” or “instinct” really seem to boil down to the question of whether the machinery for acquiring language (clearly found in every normal human being) is *specific* to language? As only humans (and not chimpanzees, frogs, or fruit flies) can acquire language, it seems safe to assume that some fragment of the machinery that allows children to acquire language has to be innate. But does the capacity to acquire language depend on machinery that is specially adapted for language, or is whatever we come to know about language merely the output of some sort of general learning mechanism? A debate that is allegedly about innateness is really about domain specificity. Similar threads loom in the background of other debates about modularity, such as questions about whether the mind contains a specialized face-recognition faculty (Kanwisher, 2000; Tarr & Gauthier, 2000); for present purposes, I will focus on language as a case study.

The domain-specificity view itself is captured concisely by Pinker’s (1994) suggestion that “language is not just a solution thought up by a generally brainy species” (p. 45), as well as by his suggestion, later in the same book, that the mind can be carved neatly in distinct evolved mechanisms for tasks such as the tracking of individuals (“a mental rolodex”), social exchange, and intuitive physics. The opposing view, captured by Elman et al.’s (1996) position, quoted in part above, is that “domain-specific representations ... emerge from domain-general architectures and learning algorithms [with] modularization as the end product of development rather than its starting point” (p. 115).

Sociologically speaking, the canonical strong empiricist view is also strongly domain-generalist. The innate endowment is presumed to be absolutely minimal and to consist only of domain-general learning mechanisms, with no language-specific prior knowledge. Information about how to perceive speech, how syntax should be structured, and so forth would be acquired in exactly the same way as all other kinds of knowledge; learning about relative clauses would use very similar computations to learning about the social behavior of kin and con-specifics. Crucially, on this view, it follows that the only reason language is learned differently from any other material (if at all) is due to idiosyncrasies in either the input or (say) a transducer such as the ear.

The canonical nativist, in contrast, is also canonically an enthusiast for domain specificity—holding that the human mind is endowed with a significant amount of innate knowledge and/or structure that pertains to particular domains, such as language. Although in principle one could imagine a nativist position in which the only thing that was innate was an intricate yet domain-general learning mechanism, the debate has really been about whether there might be contributions that are not only innate but also domain-specific. To the extent that language acquisition depended largely on domain-specific mechanisms, one would expect the acquisition of language to be relatively independent of other cognitive abilities.

Although the canonical views, empiricist and nativist, are strikingly different and would appear to make very different predictions, there has been a notorious lack of consensus, even after several decades of research (Cosmides & Tooby, 1994; Crain, 1991; Elman et al., 1996;

Karmiloff-Smith, 2000; van der Lely, 2005; Pinker, 1997, 2002). The problem is not that the two theories do not actually differ (as some cynics have suggested) or that either side lacks evidence. Rather, both sides have an embarrassment of riches: Language dissociates with cognition, in keeping with modularist views, but it also overlaps with cognition, hinting at something domain-general. How can both be true? Before I suggest a way out of this puzzle, it is necessary to examine in some detail both sorts of evidence, that which points to dissociation, and that which points to overlap or co morbidity (the cooccurrence of disorders).

3.1. Evidence for comorbidity and dissociation

There are at least three reasons to believe that there is some sort of dissociation between language and cognition. First, there is evidence from studies of developmental disorders, such as in contrast between specific language impairment (SLI) and Williams syndrome. People with Williams syndrome show significantly impaired and unusual cognition (Bellugi, Lichtenberger, Jones, Lai, & St. George, 2000; Mervis, Morris, Bertrand, & Robinson, 1999). For instance, they exhibit very low performance on tasks requiring them to reason spatially (such as copying a drawing; Bellugi et al., 2000), frequently failing to put the components of a model in the correct spatial relation. Likewise, they have difficulty in reasoning about concepts such as “living thing” and “species kind” that underlie everyday folk biology (Johnson & Carey, 1998). In contrast, their language is relatively intact. Intriguingly, these dissociations extend to a remarkably fine grain. For instance, despite their poor reasoning about spatial relations, people with Williams syndrome display relative mastery of prepositions (Landau & Hoffman, 2005), and their knowledge of syntax is far greater than is seen in other disorders (such as Down syndrome) with similar levels of mental retardation.

Children with SLI, in contrast, have by definition normal to nearly normal cognitive capacities, yet display significant impairments in at least one element of language production or comprehension (The SLI Consortium, 2002). Whereas Williams syndrome leads to (comparatively) intact language in the face of large-scale cognitive disabilities, SLI appears to have the opposite pattern. A particularly striking example of dissociation has come from van der Lely’s work on “G-SLI” (van der Lely, 2005; van der Lely, Rosen, & McClelland, 1998), a form of language impairment that appears specifically targeted at the development of grammar. Affected individuals exhibit an extreme deficit in the comprehension and production of grammatical relations in sentences, but otherwise normal cognitive and meta linguistic skills (such as pragmatics). For example, a 10-year-old child, AZ, demonstrated auditory processing, analogical, and logical reasoning indistinguishable from age-matched controls (van der Lely et al., 1998). At the same time, his knowledge and use of syntax were poor. He frequently omitted agreement morphemes (e.g., the plural—s) and could not use or understand complex sentence structures without the aid of context. AZ was at chance when deciding on the referents of *him* or *himself* in sentences like *Mowgli says Baloo is tickling him* or *Mowgli says Balooo is tickling himself*, which are completely transparent to normal readers. At the same time, he could easily use the available context to understand sentences such as *Grandpa says Granny is tickling him* or *Grandpa says Granny is tickling herself*. The strong dissociation between linguistic and nonlinguistic skills led van der Lely and colleagues to conclude, “The case of AZ provides evidence supporting the existence of a genetically determined, specialized mechanism that is necessary for the normal development of human language” (p. 1253).

A second line of evidence suggests that language dissociates from cognition even in the pattern of normal development. Whereas language is acquired quickly and robustly across a broad range of cultural conditions, before children start formal schooling and without any specific instruction (Lenneberg, 1967; Pinker, 1994), the acquisition of systems, like mathematics or logical reasoning takes many years of instruction, is often tied to literacy and is only found in some cultures (Gordon, 2004; Luria, 1979; Pica, Lemer, Izard, & Dehaene, 2004). Recently

observed cases of children developing their own languages (Goldin-Meadow, 2003; Goldin-Meadow & Mylander, 1998; Senghas & Coppola, 2001) do not appear to find a parallel in mathematical or formal reasoning ability. (It is likewise striking that language seems to be somewhat self-contained, in that children can acquire it with relatively little real-world background knowledge, in contrast to, say, rules of social interaction or global politics.)

Third, language can be acquired—and indeed is best acquired—early in life, even before many other cognitive abilities have matured, and thus appears temporally dissociated from a general maturation of cognitive skills (Johnson & Newport, 1989; Newport, 1990). The fluency with which both first and second languages are acquired decreases with increasing age, even as general cognitive skills improve. This does not mean that language is unique in having a “critical” or sensitive period (cf., e.g., musical ability; Schlaug, 2001), but it does further the notion that language (but not, for example, rules for card games) can be acquired in a way that dissociates from a number of other cognitive systems.

Such findings are inconsistent with a strong empiricist learning theory yet fit naturally with the notion of an innately constrained theory of domain-specificity. Were this recounting sufficient to exhaust the evidence, one might wonder why any controversy remains. However, an entirely different set of facts points in the opposite direction:

First, across the population as a whole, disorders in language are correlated with disorders in cognition and motor control. Although strong dissociations are possible, comorbidity is actually the more typical situation. Across SLI children taken as a whole, cases such as van der Lely’s G-SLI are comparatively rare. Disorders of language are rarely isolated. Instead, language impairments frequent co-occur with other impairments, for example, in motor control (Hill, 2001); similarly, language abilities tend to correlate with general intelligence (Colledge et al., 2002). Furthermore, language disorders with strong dissociations are rare relative to disorders that impair both language and general cognitive ability. Whereas Williams syndrome occurs once in every 20,000 births (Morris & Mervis, 2000), Down syndrome occurs once in every 800 births (Nadel, 1999). Conversely, verbal and nonverbal skills, as measured, for example, by SAT verbal and SAT math, are significantly correlated, across the normal population (Frey & Detterman, 2004).

Second, evidence from neuropsychology and brain imaging suggests that many of the neural substrates that contribute to language also contribute to other aspects of cognition. Whereas the textbook view attributes linguistic function largely to two areas that were once thought to be essentially unique to language, Broca’s and Wernicke’s areas, it is now clear that regions such as the cerebellum and basal ganglia (once thought to be of little significance for language) also play important roles. Meanwhile, those same previously unique areas are now implicated in numerous nonlinguistic processes, such as music perception (Maess, Koelsch, Gunter, & Friederici, 2001) and motoric imitation (Iacoboni et al., 1999).

Genetic evidence also hints at a common substrate for language and cognition. The human genome differs by only a small percentage (<1.5% measured by nucleotides) from the chimpanzee genome, which suggests cognitive differences between the two species must be limited. Similarly, multivariate genetic research—the analysis of covariation between traits—consistently points to links between genetic influences on different domains (Kovas & Plomin, 2006). That is, tasks such as reading and mathematics, or different tests of general cognitive ability, show highly heritable covariation with language, suggesting a common genetic basis (Plomin & Kovas, 2005).

3.2. The evolution of a Language Acquisition Device—and a way out of the apparent paradox

The key to resolving this apparent paradox—the juxtaposition of comorbidity and dissociation—may again come from biology, and in particular evolutionary biology, but not in the ways in which that theory has been typically applied to psychology.

Cognitive scientists often assume (habitually, but not necessarily explicitly) that if two neural or cognitive mechanisms subserve different systems, they are separate not only in their current function but in their evolutionary history. For example, evolutionary psychologists (Cosmides & Tooby, 1994) argue that “the human mind can be expected to include a number of functionally distinct cognitive adaptive specializations ... Both empirically and theoretically, there is no more reason to expect any two cognitive mechanisms to be alike than to expect the eye and the spleen, or the pancreas and the pituitary to be alike” (p. 92).

Yet even where two different neural systems are dedicated (or become specialized) in two different ways, they may well share evolutionary history. Whereas the eye and the spleen diverged roughly 500 million years ago, language evolved quite recently (perhaps in the last several 100,000 years). As the Nobel Laureate François Jacob (1977) put it, evolution is a tinkerer, who “often without knowing what he is going to produce ... uses what ever he finds around him, old cardboards, pieces of strings, fragments of wood or metal, to make some kind of workable object ... [The result is] a patchwork of odd sets pieced together when and where the opportunity arose” (p. 1163–1166). Marcus (2006) has suggested that the rapid evolution of language suggests it should be seen in similar terms, as more of a “tinkering” with preexisting systems, rather than a wholesale innovation from new cloth.

Two things follow from this. First, phylogeny does not necessarily map transparently onto ontogeny. The hand and the foot, for example, are, in contemporary organisms, functionally and anatomically distinct (“modules,” if you will) yet transparently evolved from a common source. Second, contemporary systems that are physically (or behaviorally) separate may derive from common ancestry. The hand and the foot subserve different functions but depend in part (although not exclusively) on a large number of overlapping genes. What we are suggesting is that “cognitive modules” or “linguistic modules” be viewed in a similar fashion. Comorbidity follows from common ancestry; dissociation follows from a divergence during that portion of evolutionary history that separates systems that derive from a once common origin. In Darwin’s terminology, we are describing the consequence of the process called “descent with modification.”

Fig. 1 depicts three possible relationships between the substrates (neural, cognitive, or genetic) of language and the substrates of cognition. Panel A depicts a strong modularist account in which language is almost wholly separate from cognition; Panel B depicts a purely domain general account. Panel C illustrates the predictions one might derive from the principle of descent with modification. The area of each oval represents a set of mechanisms; in Fig. 1A these are distributed between mechanisms specified for language and those specialized for other forms of cognition. In Fig. 1B, language represents a subset of general cognitive mechanisms. In Fig. 1C, language is taken to be an orchestration of components, some distinctively associated with language, but most shared with other cognitive systems.

Fig. 1B is essentially the position Elizabeth Bates defended in her (Bates, 2004) suggestion that language is “a new machine that Nature has constructed out of old parts” (p. 250). In essence, Bates suggested that relevant evolutionary changes were likely to have been (only) quantitative and domain-general, essentially expanding the cognitive circle in Fig. 1B—while leaving language as a proper subset of that circle; no allowance was made for humanly unique, linguistically unique adaptation. Bates often likened the evolution of language to the evolution of the giraffe’s neck: “[the] uniquely human capacity for language, culture and technology may

have been acquired across the course of evolution by a similar process—*quantitative changes* in primate abilities that bring about and insure a qualitative leap in cognition and communication” (p. 250, emphasis added).

In framing things in this way, Bates’s interpretation of the evolutionary record excludes, without argument, the possibility of true qualitative change, and hence the small—but, in my view, essential—nonoverlapping region of language in Fig. 1C (Fisher & Marcus, 2006; Marcus, 2006). Essentially, Bates argued for descent but ignored the possibility of language-specific modification.

Given the recency of language’s evolution, language probably relies on a set of mechanisms mainly shared across multiple domains, but given its unique position in the animal world, it seems likely that there are also novel specializations (small white region in Fig. 1C). One can see exactly this sort of thing in the evolution of the hand and foot, where the bulk of genetic material is quite ancient, but there is still a small amount of distinction specialization for the hand versus the foot.

The perspective represented in Fig. 1C, informed by ideas borrowed from developmental and evolutionary biology, can cast immediate light on the apparently paradoxical conclusion that language is both domain-general and -specific, and that it is both dissociable from other cognitive abilities and comorbid with them. Any distinct cognitive system must have evolved from a prior structure. The genes (and neural/cognitive circuitry) that underlie linguistic ability are descendents, presumably with modification, of genes (and neural/cognitive circuitry) that contributed to other, evolutionarily prior, abilities. Comorbidity comes from descent, from those substrates of language that are shared with or descended from other cognitive systems. Dissociation comes from divergence, from the ways in which language’s substrates have been modified as it diverged and developed into its current unique form.

More broadly speaking, advocates of modularity have enumerated a great number of possible modules, such as (putatively) specialized mechanisms for face recognition, intuitive physics, mate selection, and the tracking of social exchanges. In each case, we may find that even where systems are more or less separable in current organisms, there may be some benefit to considering how and whether those faculties could have descended with modification from faculties possessed by ancestors that lack such specializations.

4. Optimality

The third fundamental question, scarcely less contentious than the first two, concerns the question of whether human beings are rational. Researchers such as Kahneman and Tversky (Kahneman, 2003; Tversky & Kahneman, 1974) have documented numerous apparent deviations from rationality, ranging from the well-known conjunction fallacy² to anchoring effects (in which judgments are modulated by arbitrary and irrelevant information), yet in fields ranging from reasoning to linguistics, the idea of humans as perfect, rational, optimal creatures is making a comeback. Noted evolutionary psychologists (Tooby & Cosmides, 1995) have argued that the mind consists of an “accumulation of superlatively well-engineered designs,” while Bayesian cognitive scientists (Chater et al., 2006) have argued that “it seems increasingly plausible that human cognition may be explicable in rational probabilistic terms and that, in core domains, human cognition approaches an optimal level of performance.” Ultimately, these

²The *conjunction fallacy* is a cognitive error in which people find conjunctions more probable than their constituent parts. In its canonical form, subjects are told that a person named Linda “is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice and also participated in anti-nuclear demonstrations” and asked to evaluate a series of statements. Most people appear to think that the conclusion that “Linda is a feminist bankteller” is more likely than the conclusion that “Linda is a bankteller,” a mathematical impossibility.

two camps have yielded very different pictures of how cognition works: Kahneman and Tversky implied a kind of eclectic mind that draws a grab bag of different mechanisms; Chater and Tenenbaum and colleagues have at times pointed to a very different picture, in which the essence of cognition might consist of a form of ideal-observer-like Bayesian inference.

In keeping with this latter, rather optimistic view, a growing body of evidence (see Chater et al., 2006, for a review) shows circumstances in which human inference approaches theoretical limits. For example, Griffiths and Tenenbaum (2006) recently conducted a series of studies in which people answered questions, such as “If someone has thus far served in the Senate for 12 years, how many more might you expect them to serve?” People are remarkably accurate, extrapolating such facts with stunning accuracy—leading to authors’ suggestion that “everyday cognitive judgments follow ... optimal statistical principles.”

Implicit, I believe, in much of the effort to defend human rationality is the intuition that natural selection is so powerful it is unlikely that human beings would be anything *other* than rational. One colleague, for example, said to me (via E-mail), “I’m a serious believer in the power of evolution to optimize ... I would be very surprised if you can find even one convincing case where the brain does something that you can prove to be non-optimal!”

Behind this confidence in evolution is presumably the assumption that, as Tooby and Cosmides (1995) put it, “superlatively well-engineered functional designs” tend to accumulate

because natural selection is a hill-climbing process that tends to choose the best of the variant designs that actually appear, and because of the immense numbers of alternatives that appear over the vast expanse of evolutionary time.

In reality, pockets of rationality (e.g., Griffiths and Tenenbaum’s demonstrations of near-normative capacities for extrapolation) are often juxtaposed with striking demonstrations of *irrationality*, for example, anchoring effects, in which people blithely extrapolate from manifestly *irrelevant* information reaching erroneous conclusions.³ As Shafir and LeBoeuf (2002) have argued, some residue of irrationality remains, even after all reasonable objections have been considered.

Of course, it is probably too much to expect absolute rationality; as Cherniak (1986) and others have pointed out, inferences requires resources and resources are finite. No finite agent could make every possible inference or follow every possible deductive chain to infinite limits; memory and time are limited, and so forth. Anderson’s (1990) book was an explicit attempt to derive generalizations about what an *optimal* creature might do given limited resources, and Gigerenzer and Goldstein (1996) have written extensively about the idea of “bounded cognition,” a further attempt to outline the sorts of inferences (and accompanying biases) a resource-bounded creature might make. It seems highly likely that at least some deviations from rationality and a kind of idealized, resource-independent optimality can be captured in this way.

But there is a further reason to expect that real, biologically instantiated creatures might deviate from even an idealization that is downscaled to take into account resource limitations, what one might call evolutionary inertia.

4.1. Evolutionary inertia

Evolutionary inertia (Marcus, 2008) is the notion that evolution’s solutions are inevitably constrained by previous history. Because evolution works like a blind tinkerer, not like a

³A canonical example, again from Kahneman and Tversky, asked subjects to estimate the percentage of African countries in the United Nations—after first spinning a wheel of fortune; people’s estimates tended to be correlated with their (irrelevant) spins.

designer with foresight, it lacks privileged access into what would be an optimal (or even optimal-relative-to-bounded-resources) solution. It may tend on average to make things better, but there is no guarantee. Rather, as a hill-climbing process that essentially follows the algorithm of “take small steps, go up, but never down,” it is highly vulnerable to the problem of local maxima, occasionally landing on good but less than ideal solutions—with no wherewithal to go back down a short peak in order to climb some distant but higher mountain. Thus, evolution is not an optimizer (as implied by a naïve reading of the popular phrase “survival of the fittest”) but a “meliorizer” (Dawkins, 1982).

A quick look at the awkward design of the injury-prone human spinal column is enough to remind us that this is so; the vertical human spine takes the form that it does because it is a ready modification of the horizontal spine that we inherited from our ancestors. The design made sense for our four-legged ancestors—a horizontal spine distributes a creature’s weight in a balanced fashion—but rotated 90° it becomes an awkward solution, more like a flagpole than a table. (Better might have been a system that distributed our weight across four equal and cross-braced columns.) We are stuck with a clumsy kluge of a solution not because it is the best conceivable way to support the weight of a biped, but because the spine’s structure evolved from the (more sensible) spine of four-legged ancestors; evolution tends to *tinker* (Jacob, 1977) with what was already there rather than starting from scratch. Just as an object that is in motion tends to stay in motion, evolution, once headed in a particular direction, tends to continue in that direction. And this, too, may lead to mechanisms that are less than optimal (even once resource limitations are factored in).

4.1.1. Poorly integrated subsystems—For example, evolutionary inertia may lead to what the neuroscientist Allman (1999) has called the “progressive overlay of technologies”; the notion here is that when evolution adds something new, it generally adds it on top of earlier mechanisms without fully eliminating older mechanisms, even if those older mechanisms are less efficient.

In this way, evolutionary inertia may help to explain a common observation—dating back at least to Freud, if not earlier—that the mind often seems torn between two courses of action, one short-term, one long-term. Freud’s id and ego have counterparts in modern dichotomies such as rational–emotional, reflexive–reflective, and Kahneman’s (2003) less committal Systems I and II. However, one parses the distinction—I prefer the terms “deliberative” and “reflexive” (see Marcus, 2008)—the tension seems undeniable. It is perhaps most salient in everyday tensions like the everyday battles for self-control (“Should I dip into that tempting fondue even though it runs counter to my New Year’s resolution?”), but it is also present in a variety of more subtle ways, as when our judgments of other people’s competence (e.g., for jobs) is clouded by extraneous factors such as physical appearance. Todorov, Mandisodza, Goren, and Hall (2005), for example, found that in nearly three out of four elections for U.S. Congress, the winner was the candidate that independent raters judged to be more competent-appearing.

More broadly speaking, even though deliberative reasoning mechanisms might in some ways be considered more sophisticated, they are the first to break down under time pressure, exhaustion, or cognitive load (Ferreira, Garcia-Marques, Sherman, & Sherman, 2006); when push comes to shove, our older reflexive systems still seem to hold the cards.⁴

⁴The phenomena of motivated reasoning (Kunda, 1990), whereby we apply greater scrutiny to threatening conclusions than to conclusions that we hope are true, may similarly stem from an accident of evolutionary inertia. When evolution overlaid the technology of deliberative reasoning, it did so in a clumsy fashion that left the *intensity* of deliberative reasoning largely up to the whims of older emotional systems (see Marcus, 2008, for further discussion).

One of the saddest ironies of human nature is that we are simultaneously clever enough to make thoughtful, long-term plans yet foolish enough to abandon those plans in the face of temptation—while still being intelligent enough to feel remorse about it. Such tensions may not persist because they are intrinsically adaptive but simply because evolution, lacking forethought, could not figure out how to do better.

4.1.2. Mismatched parts—A subtler but even more striking illustration of the power of evolutionary inertia comes from the nature of vertebrate “context-dependent” memory, an otherwise exquisite adaptation that is ill-suited to many of the requirements of human cognition.

The term *context-dependent* refers to the fact that the memory system found in humans, like the memory of all known biological creatures, appears to be accessed primarily by context or cues. Scuba-divers, for example, are better able to remember facts that they learn underwater when then they are *tested* underwater (Godden & Baddeley, 1975), and rats perform better in a maze if the lighting is the same at test it was during a learning phase (Carr, 1917). (For a recent review, see Smith, 2006.)

Context-dependent memory is, in many respects, a sensible way for a finite creature to organize memory (Anderson, 1990); our brains tend to make most accessible information that is frequent, information that is recent, and information that is contextually relevant. Given the intrinsically slow nature of neurons (relative to the processing levels achievable by contemporary microprocessors), the organization of human memory is remarkably efficient, capable of quickly providing accurate information through parallel search.

But while context-dependent memory is speedy (relative to the comparatively slow speed of neurons), it also has a number of limitations; information that is used infrequently, for example, is difficult to retrieve, and, overall, context-dependent memory is better tuned toward the retrieval of general tendencies (gist) rather than specific memories.⁵ Retrieving specific bits of information—where I put my keys yesterday, for example—can often be remarkably challenging. It is par for the human course to lose keys, wallets, and cell phones, to confuse what we did today with what we did yesterday (was it waffles today, eggs yesterday, or the other way around?), and to drive directly home from work (“on autopilot”) when one meant to pick up groceries or dry-cleaning first.

The remarkable thing about these minor glitches of memory is not that they can lead to serious problems (pilots forget to pull up their landing gear, and eyewitnesses saddled with unreliable and distortion-prone memory sometimes unwittingly given false testimony) but that in principle they seem so easy to avoid: Memory could have been organized in a different way, mapped by internal location (as in digital computers). In so-called location-addressable memory, any bit of information can be retrieved with high reliability, whether it is frequent, recent, or contextually relevant. In a computer, keeping track of (say) the last known location of a set of keys would be trivial, a matter of routinely updating a specific *buffer* or memory location. Because each memory is allocated to a particular spot, there is no risk of distortion, no problem in updating memories, no risk of interference with earlier memories, and, more generally, none of the fragility of human memory.

Humans (and perhaps other creatures) might have been far better off had evolution stumbled on the trick of location-addressable memory (ideally layering context-dependent memory on top, a la Google), but the pressures that might have led to such a system appeared too late; by

⁵Could context-dependent memory reduce overall storage requirements, as one anonymous reviewer suggested? Perhaps, but it is not obvious why that should be the case. Information is information, whether it is stored in context-dependent or location-addressable information. The difference is in how information is *accessed*, not the amount of information stored.

the time humans began to emerge, context-dependent memory had a history of several hundred million years. Context-dependent memory is found in species ranging from spiders to mice (see Marcus, 2008, for a review), and the rest of cognition, for better and for worse, appears to have built on that substrate.

A great deal of human irrationality may stem from this fact. Take, for example, a phenomenon known as the focusing illusion, in which people's current thoughts tend to be overly influenced by whatever they most recently were thinking about, as in an elegant but telling two-question survey run by Strack, Martin, and Schwarz (1988). Undergraduates were either asked, "1. How happy are you with your life in general?" followed by "2. How many dates have you had in the last month?" or "1. How many dates have you had in the last month?" followed by "2. How happy are you with your life in general?" Answers were uncorrelated when questions were presented in the first order, but they were highly correlated when presented in the second order.⁶ A perspective of cognition shaped by evolutionary inertia helps explain these results: Our judgments (even of intimate personal details, such as how happy we are) are inevitably mediated by what we can retrieve from memory, and our memories are heavily biased toward recent information.

Own-contribution bias (by which people overestimate their own contributions to joint enterprises such as housework or scientific collaboration), *confirmation bias* (by which people overestimate the importance of data that matches their intuitive theories relative to nonmatching, hence harder to remember, information), and *framing effects* (by which the statement of a problem influences our reasoning about a given problem) and a host of other facts that Wilson and Brekke (1994) dubbed "mental contamination" can all be seen, similarly, as reflexes of a deliberative system that cannot compensate for inherited biases in our underlying memory systems. When we consider purchasing meat that is 80% lean, we may rationally evaluate whatever associations we have with leanness, but fail to bring to mind our associations with 20% fat, and therefore fail to rationally evaluate that body of evidence. Because we cannot systematically search our memory structures, we are largely at their mercy, and some degree of irrationality may inevitably follow (Marcus, 2008).

4.2. Recap

If one thinks of evolution as a process of hill climbing, the fundamental limit on evolution is the nature of the landscape. If there is but a single peak, a "policy" of always climbing up, never down, will eventually lead to the top. But in a landscape with multiple peaks, a system that lacks foresight can easily wind up stuck on a peak that is higher than its neighbors but well short of the highest peak, what computer scientists call a *local maximum*. The conservative nature of evolution—whereby new species tend to have genomes that are closely related to those of their ancestors—ensures that local maxima will be common.

While it is true that any given moment in evolution we would expect a better reasoning creature to outperform a creature that reasons less well, natural selection can only choose among available alternatives—and the alternatives that are available at any given moment tend to be heavily constrained by what already is in place. In a creature that (in some respects) differs importantly from its predecessors, as humans do relative to our last nonspeaking ancestors, some degree of evolutionary friction—suboptimalities due to inertia—may be inevitable.

To the extent that modern humans endeavor to solve problems in novel ways, our solutions may be effective but far from optimal. A cognitive mechanism that is found only in modern humans is likely to be close to optimal only to the extent that mechanisms are compatible with

⁶It is not just undergraduates: older adults show similar results when the "dates" question is replaced by one asking about their health or quality of their marriages.

earlier cognitive machinery we inherited; when new solutions demand significant repurposing of old cognitive machinery, suboptimality may be a common result.

5. Symbol manipulation

As a final foundational question, let us consider the “connections and symbols” debate that begin in the mid-1980s, with publication of Rumelhart and McClelland’s two-volume collection on parallel distributed processing (McClelland, Rumelhart, and The PDP Research Group, 1986; Rumelhart, McClelland, and The PDP Research Group, 1986). Up until that point, most (though not certainly all) cognitive scientists presumed that symbols were the primary currency of mental computation. Newell and Simon (1975), for example, wrote about the human mind as a “physical symbol system,” in which much of cognition was built upon the storage, comparison, and manipulation of symbols.

Rumelhart and McClelland (1986a) challenged this widespread presumption by showing that a system that ostensibly lacked rules could apparently capture a phenomenon—children’s overregularization errors—that heretofore had been the signal example of rule learning in language development. On traditional accounts, overregularizations (e.g., *singed* rather *sang*) were seen as the product of mentally represented rule (e.g., *past tense = stem + -ed*). In Rumelhart and McClelland’s model, overregularizations emerged not through the application of an explicit rule, but through the collaborative efforts of hundreds of individual units that represented individual sequence of phonetic features that were distributed across a large network, with a structure akin to that in Fig. 2.

A flurry of critiques soon followed (Fodor & Pylyshyn, 1988; Lachter & Bever, 1988; Pinker & Prince, 1988), and the subsequent years were characterized by literally dozens of papers on the development of the English past tense, both empirical (e.g., Kim, Marcus, Pinker, Hollander, & Coppola, 1994; Kim, Pinker, Prince, & Prasada, 1991; Marcus, Brinkmann, Clahsen, Wiese, & Pinker, 1995; Marcusa et al., 1992; Pinker, 1991; Prasada & Pinker, 1993) and computational (e.g., Ling & Marinov, 1993; Plunkett & Marchman, 1991, 1993; Taatgen & Anderson, 2002).

In the late 1990s, I began to take a step back from the empirical details of particular models—which after all were highly malleable—to try to understand something general about how the models worked and what their strengths and limitations were (Marcus, 1998a,b, 2001). In rough outline, the argument was that the class of connectionist models that were then popular was inadequate, and that without significant modification they would never be able to capture a broad range of empirical phenomena.

In 2001, in a full-length monograph on the topic (Marcus, 2001), I suggested that—consistent with what had often been assumed—that the mind had very much the same symbolic capacities as the pioneering computer programming language Lisp, articulated in terms of the following seven claims.

1. The mind has a neurally realized way of representing “symbols.”
2. The mind has a neurally realized way of representing “variables.”
3. The mind has a neurally realized way of representing “operations over variables,” to form the progressive form of a verb, take the stem and add *-ing*.
4. The mind has a neurally realized way of representing “distinguishing types from tokens,” such as one particular coffee mug as opposed to mugs in general.
5. The mind has a neurally realized way of representing “ordered pairs ($AB \neq BA$)”; *man bites dog* is not equivalent to *dog bites man*.

6. The mind has a neurally realized way of representing “structured units” (element C is composed of elements A and B, and distinct from A and B on their own).
7. The mind has a neurally realized way of representing “arbitrary trees,” such as the syntactic trees commonly found in linguistics.

Seven years later, I see no reason to doubt any of the first six claims; no serious critique of *The Algebraic Mind* was ever published, and to my knowledge there has been no serious PDP attempt in recent years to capture the phenomena highlighted therein (e.g., the human facility with distinguishing types from tokens). Instead, recent theoretical works such as Rogers and McClelland (2004) continue to rely on architectures that were introduced over a decade ago, such as the PDP model of Rumelhart and Todd (1993), and remain vulnerable to the same criticisms as their predecessors (Marcus & Keil, 2008).

Yet if I remain confident in the accuracy of the first six conjectures, I now believe I was quite wrong about the seventh claim—in a way that may cast considerable light on the whole debate. The problem with the seventh claim is, to put it bluntly, people do not behave as if they really *can* represent full trees (Marcus & Wagers, 2009). We humans have trouble in remembering sentences verbatim (Jarvella, 1971; Lombardi & Potter, 1992); we have enormous difficulty in properly parsing center-embedded sentences (*the cat the rat the mouse chased bit died*) (Miller & Chomsky, 1963) and we can easily be seduced by sentences that are globally incoherent (Tabor, Galantucci, & Richardson, 2004), provided that individuals chunk of the sentence seem sufficiently coherent at a local level (e.g., *More people have been to Russia than I have*). Although tree representations are feasible in principle—computers use them routinely—there is no direct evidence that *humans* can actually use them as a form of mental representation. Although humans may have the abstract competence to use (or at least discuss trees), actual performance seems to consistently fall short.

In a proximal sense, the rate-limiting step may be the human mind’s difficulties with dealing with rapidly creating large numbers of bindings. For example, by rough count, the sentence *the man bit the dog* demands the stable encoding of at least a dozen bindings, on the reasonable assumption that each connection between a node and its daughters (e.g., S(entence) and N(oun) P(hrase)) requires at least one distinct binding; on some accounts that sentence alone might require as many as 42 (if each node bore three pointers, one for its own identity, and one for each of two daughters).

Although numbers of between 12 and 42 (more in more complex sentences) might at first blush seem feasible, they far exceed the amount of short-term information-binding bandwidth seen in other domains of cognition (Treisman & Gelade, 1980). Miller (1956) famously put the number of elements a person could remember at 7 ± 2 , but more recent work suggested that Miller significantly overestimated; realistic estimates are closer to 4 or even fewer (Cowan, 2001; McElree, 2001). Similarly low limits on binding seem to hold in the domain of visual object tracking (Pylyshyn & Storm, 1988). Although it is certainly possible that language affords a far greater degree of binding than in other domains, the general facts about human memory capacity clearly raise questions.

In a more distal sense, the rate-limiting step may have been the conservative nature of evolution, and in particular the very mechanism of context-dependent memory that was mentioned earlier. As mentioned, computers succeed in readily representing trees because their underlying memory structures are organized by location (“or address”); someone’s cell phone number might be stored in location 43,212, their work number in location 43,213, and so forth. Human memory, in contrast—and indeed probably all of biological memory—appears to be organized around a different principle, known as content addressability, meaning that memories are retrieved by content or context, rather than location (Anderson, 1983). Given the short

evolutionary history of language (Marcus, 2008), and the fundamentally conservative nature of evolution (Darwin, 1859; Marcus, 2008), context-dependent memory seems likely to be the *only* memory substrate that is available to the neural systems that support language. Although content addressability affords rapid retrieval, by itself it does not suffice to yield tree-geometric traversability. With content addressability, one can retrieve elements from memory based on their properties (e.g., animate, nominative, plural, etc.), but not (absent location-addressable memory) their location.

Constrained in this way, we may thus be forced to rely on a sort of cobbled-together substitute for trees, in which linguistic structure can only be represented in approximate fashion, by mean of sets of subtrees (“treelets”) that are bound together in transitory and incomplete fashion. (See Marcus & Wagers, 2009, for further details.) Our brains may thus be able to afford some sort of approximate reconstruction but not with the degree of reliability and precision that veridically represented trees would demand. Though space precludes making that argument in full detail here, Wagers and I have argued that a host of facts, ranging from the human difficulty with parsing center-embedding sentences to our difficulties in remembering verbatim syntactic structure, this are consistent with this notion.

If this conjecture is correct, the upshot would be that human beings possess most—but critically, *not all*—of the symbol-manipulating apparatus that we associate with digital computers, viz. the first six elements listed above, but not the seventh, leaving us with vastly more symbolic power than simple neural networks, but still notably less than machines. Which seems, in hindsight, to be quite plausible. We are not nearly as good as machines at representing arbitrary taxonomies (what normal human, for example, can instantaneously acquire the complete family tree of the British Monarchy?), yet we are fully capable of symbolic, rule-based generalizations, and still far outstrip PDP networks in the capacity to think about individuals versus kinds (Marcus, 2001; Marcus & Keil, 2008). The conjecture that humans possess most but not quite all of the classic tools of symbol manipulation would capture some of the spirit (though not the letter) of Smolensky’s (1988) suggestion that human representations might be approximations to their digital counterparts. In retrospect, advocates of PDP probably underestimated our symbolic capacities but may have been right that humans lack the full complement of symbol-manipulating faculties seen in machines. Connectionism’s common denial of symbols still seems excessive and undermotivated, but there may well be plenty of important work to be performed in finding instances in which human capacities deviate from full-scale symbol manipulators.

6. Coda

The overall picture presented here is one in which there is considerable room for innate structure yet substantial room for learning and plasticity. Genes underwrite a complex cognition architecture with a good degree of specialization, yet even the most humanly unique aspects of cognitive architecture can trace many of their properties to a bedrock shaped long before language and deliberate reasoning emerged. Significantly, extant cognitive architecture likely represents a sort of compromise between that which would be optimal and that which was actually evolvable. The apparent lack of tree structures is a case in point. An ideally designed linguistic creature would almost certainly have recourse to tree structures as a representational format, but we appear to make due with a somewhat clumsy substitute, fashioned from context-dependent rather than location-addressable substrates, yielding a kind of approximation to full-scale symbol manipulation that is less than ideal.

Whether the details of the current view are correct, it seems vital that cognitive science return some of its focus to questions about cognitive architecture, especially as the prospect of a genuinely integrated cognitive neuroscience looms closer. Cognitive neuroscience is

ultimately, at least in principle, supposed to be a bridging science, one that shows us the mechanistic nature of the connection between cognitive constructs and neural constructs.

At the moment, the gap between cognition and neuroscience remains enormous. Although contemporary cognitive neuroscience has made considerable strides in mapping out the geography of which parts of the brain participate in particular skills, it has still yielded little in the way of concrete information about how individual neural circuits do their job. If we are to make real progress—if we are to genuinely bridge cognition with neuroscience—we will need to understand the cognitive architecture every bit as well as we understand diffusion-tensor imaging and Hodgkin–Huxley synapses.

Without that, we are lost. Mere mappings—recognizing that vision takes place in the occipital cortex, for example—tell us too little. They do not tell us why the brain is organized as it is, and they do not tell us how various bits of cortex do what they do. To move further, we need to know not just the general sorts of things the brain does (e.g., processing semantics or recognizing faces) but the processes and representational formats by which it achieves those computations.

As it stands, these questions have received comparatively little attention, and there are real risks that the cognitive entities that we seek may not even exist. For example, linguistic theories have historically been couched in terms of syntactic trees, postulated on the basis of idealized linguistic data, but if the arguments of the previous section are correct, such trees may turn out to be a sort of idealized unicorn never actually observable in empirical practice. If so, there might be little chance of properly bringing the ontology of cognitive science into alignment with the mechanisms of neurophysiology. A proper unification would explain *how* the brain achieves syntactic representation, but it will do us no good to seek the neural realization of ontological entities that do not actually exist as such.

Because of scenarios like these, it seems imperative that we as a field return forthwith to the tough questions that preoccupied an earlier generation of cognitive scientists: an understanding of the nuts and bolts of cognition, and how those elements combine, in development, and in the adult form, to produce the wondrous yet imperfect creatures we know as human beings.

Acknowledgments

For helpful discussion and comments on an earlier draft, I thank Wayne Gray, Peter Todd, and athena Vouloumanos. This article was supported by NIH Grant HD 048733.

References

- Allman, JM. *Evolving brains*. New York: Scientific American Library, distributed by W.H. Freeman; 1999.
- Anderson, JR. *The architecture of cognition*. Cambridge, MA: Harvard University Press; 1983.
- Anderson, JR. *The adaptive character of thought*. Hillsdale, NJ: Erlbaum Associates; 1990.
- Bates EA. Explaining and interpreting deficits in language development across clinical groups: Where do we go from here? *Brain and Language* 2004;88(2):248–253. [PubMed: 14965545]
- Bates, E.; Elman, J.; Johnson, M.; Karmiloff-Smith, A.; Parisi, D.; Plunkett, K. Innateness and emergentism. In: Bechtel, W.; Graham, G., editors. *A companion to cognitive science*. Malden, MA: Blackwell; 1998. p. xvii-791.
- Bellugi U, Lichtenberger L, Jones W, Lai Z, St George M. I. The neurocognitive profile of Williams syndrome: A complex pattern of strengths and weaknesses. *Journal of Cognitive Neuroscience* 2000;12 (Suppl 1):7–29. [PubMed: 10953231]
- Bouchard TJ Jr. Genes, environment, and personality. *Science* 1994;264(5166):1700–1701. [PubMed: 8209250]

- Brown, DE. Human universals. New York: McGraw-Hill; 1991.
- Carr H. Maze studies with the white rat. *Journal of Animal Behavior* 1917;7(4):259–275.
- Carroll, SB.; Grenier, JK.; Weatherbee, SD. From DNA to diversity: Molecular genetics and the evolution of animal design. Oxford; Malden, MA: Blackwell Science; 2001.
- Carruthers, P. The architecture of the mind: Massive modularity and the flexibility of thought. Oxford; New York: Clarendon Press; Oxford University Press; 2006.
- Chater N, Tenenbaum JB, Yuille A. Probabilistic models of cognition: conceptual foundations. *Trends in Cognitive Sciences* 2006;10(7):287–291. [PubMed: 16807064]
- Cherniak, C. Minimal rationality. Cambridge: MIT Press; 1986.
- Chomsky NA. Review of verbal behavior. *Language* 1959;35:26–58.
- Colledge E, Bishop DV, Koeppen-Schomerus G, Price TS, Happe FG, Eley TC, Dale PS, Plomin R. The structure of language abilities at 4 years: A twin study. *Developmental Psychology* 2002;38(5):749–757. [PubMed: 12220052]
- Cosmides, L.; Tooby, J. Cognitive adaptations for social exchange. In: Barkow, J.; Tooby, J.; Cosmides, L., editors. The adapted mind: Evolutionary psychology and the generation of culture. Oxford, England: Oxford University Press; 1992. p. 163–228.
- Cosmides, L.; Tooby, J. Origins of domain specificity: The evolution of functional organization. In: Hirschfeld, LA.; Gelman, SA., editors. Mapping the mind: Domain specificity in cognition and culture. Cambridge; New York: Cambridge University Press; 1994. p. 85–116.
- Cowan N. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences* 2001;24(1):87–114. [PubMed: 11515286]discussion 114–185
- Crain S. Language acquisition in the absence of experience. *Behavioral and Brain Sciences* 1991;14:597–650.
- Damoiseaux JS, Rombouts SA, Barkhof F, Scheltens P, Stam CJ, Smith SM, Beckmann CF. Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences of the United States of America* 2006;103(37):13848–13853. [PubMed: 16945915]
- Darwin, C. On the origin of species. London: Murray; 1859.
- Dawkins, R. The extended phenotype: The gene as the unit of selection. Oxford: San Francisco: W.H. Freeman; 1982.
- Ehrlich, PR. Human natures: Genes, cultures, and the human prospect. Washington, DC: Island Press/Shearwater Books; 2000.
- Elman, JL.; Bates, E.; Johnson, MH.; Karmiloff-Smith, A.; Parisi, D.; Plunkett, K. Rethinking innateness: A connectionist perspective on development. Cambridge, MA: MIT Press; 1996.
- Ferreira MB, Garcia-Marques L, Sherman SJ, Sherman JW. Automatic and controlled components of judgment and decision making. *Journal of Personality and Social Psychology* 2006;91(5):797–813. [PubMed: 17059302]
- Fisher SE, Marcus GF. The eloquent ape: Genes, brains and the evolution of language. *Nature Reviews Genetics* 2006;7(1):9–20.
- Fodor, JA. Modularity of mind. Cambridge, MA: MIT Press; 1983.
- Fodor JA, Pylyshyn Z. Connectionism and cognitive architecture: A critical analysis. *Cognition* 1988;28:3–71. [PubMed: 2450716]
- Frey MC, Detterman DK. Scholastic assessment or g? The relationship between the Scholastic Assessment Test and general cognitive ability. *Psychological Science* 2004;15(6):373–378. [PubMed: 15147489]
- Gardner, H. The mind's new science: A history of the cognitive revolution. New York: Basic Books; 1985.
- Gerhart, J.; Kirschner, M. Cells, embryos, and evolution. Cambridge, MA: Blackwell Science; 1997.
- Gigerenzer G, Goldstein DG. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 1996;103(4):650–669. [PubMed: 8888650]
- Gigerenzer, G.; Selten, R. Bounded rationality: The adaptive toolbox (1st MIT Press paperback ed.). Cambridge, MA: MIT Press; 2002.
- Godden DR, Baddeley AD. Context-dependent memory in two natural environments: On land and underwater. *British Journal of Psychology* 1975;66(3):325–331.

- Goldin-Meadow, S. The resilience of language: What gesture creation in deaf children can tell us about how all children learn language. New York: Psychology Press; 2003.
- Goldin-Meadow S, Mylander C. Spontaneous sign systems created by deaf children in two cultures. *Nature* 1998;391(6664):279–281. [PubMed: 9440690]
- Gordon P. Numerical cognition without words: Evidence from Amazonia. *Science* 2004;306(5695):496–499. [PubMed: 15319490]
- Griffiths TL, Tenenbaum JB. Optimal predictions in everyday cognition. *Psychological Science* 2006;17(9):767–773. [PubMed: 16984293]
- Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R. Intersubject synchronization of cortical activity during natural vision. *Science* 2004;303(5664):1634–1640. [PubMed: 15016991]
- Hill EL. Non-specific nature of specific language impairment: A review of the literature with regard to concomitant motor impairments. *International Journal of Language and Communication Disorders* 2001;36(2):149–171. [PubMed: 11344592]
- Iacoboni M, Woods RP, Brass M, Bekkering H, Mazziotta JC, Rizzolatti G. Cortical mechanisms of human imitation. *Science* 1999;286(5449):2526–2528. [PubMed: 10617472]
- Jacob F. Evolution and tinkering. *Science* 1977;196:1161–1166. [PubMed: 860134]
- Jacob F, Monod J. On the regulation of gene activity. *Cold Spring Harbor Symposium on Quantitative Biology* 1961;26:193–211.
- Jarvella RJ. Syntactic processing of connected speech. *Journal of Verbal Learning and Verbal Behavior* 1971;10:409–416.
- Johnson SC, Carey S. Knowledge enrichment and conceptual change in folkbiology: Evidence from Williams syndrome. *Cognitive Psychology* 1998;37(2):156–200. [PubMed: 9878105]
- Johnson JS, Newport EL. Critical period effects in second language learning: The influence of maturational state on the acquisition of English as second language. *Cognitive Psychology* 1989;21:60–99. [PubMed: 2920538]
- Kahneman D. A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist* 2003;58(9):697–720. [PubMed: 14584987]
- Kahneman, D.; Slovic, P.; Tversky, A. *Judgment under uncertainty: Heuristics and biases*. Cambridge; New York: Cambridge University Press; 1982.
- Kanwisher N. Domain specificity in face perception. *Nature Neuroscience* 2000;3(8):759–763.
- Karmiloff-Smith, A. *Beyond modularity*. Cambridge, MA: MIT Press; 1992.
- Karmiloff-Smith, A. Why babies' brains are not Swiss army knives. In: Rose, H.; Rose, SPR., editors. *Alas, poor Darwin: Arguments against evolutionary psychology*. London: Jonathan Cape; 2000. p. 144–156.
- Kim JJ, Marcus GF, Pinker S, Hollander M, Coppola M. Sensitivity of children's inflection to grammatical structure. *Journal of Child Language* 1994;21:173–209. [PubMed: 8006091]
- Kim JJ, Pinker S, Prince A, Prasada S. Why no mere mortal has ever flown out to center field. *Cognitive Science* 1991;15:173–218.
- Kording KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature* 2004;427(6971):244–247. [PubMed: 14724638]
- Kovas Y, Plomin R. Generalist genes: Implications for the cognitive sciences. *Trends in Cognitive Sciences* 2006;10(5):198–203. [PubMed: 16580870]
- Kunda Z. The case for motivated reasoning. *Psychological Bulletin* 1990;108(3):480–498. [PubMed: 2270237]
- Lachter J, Bever TG. The relation between linguistic structure and associative theories of language learning: A constructive critique of some connectionist learning models. *Cognition* 1988;28:195–247. [PubMed: 3349761]
- Landau B, Hoffman JE. Parallels between spatial cognition and spatial language: Evidence from Williams syndrome. *Journal of Memory and Language* 2005;53:163–185.
- van der Lely HK. Domain-specific cognitive systems: Insight from Grammatical-SLI. *Trends in Cognitive Sciences* 2005;9(2):53–59. [PubMed: 15668097]
- van der Lely HK, Rosen S, McClelland A. Evidence for a grammar-specific deficit in children. *Current Biology* 1998;8(23):1253–1258. [PubMed: 9822577]

- Lenneberg, EH. Biological foundations of language. New York: Wiley; 1967.
- Ling CX, Marinov M. Answering the connectionist challenge: A symbolic model of learning the past tense of English verbs. *Cognition* 1993;49:235–290. [PubMed: 8131377]
- Lombardi L, Potter M. The regeneration of syntax in short term memory. *Journal of Memory and Language* 1992;31:713–733.
- Luria, A. The making of mind. Cambridge, MA: Harvard University Press; 1979.
- Maess B, Koelsch S, Gunter TC, Friederici AD. Musical syntax is processed in Broca's area: an MEG study. *Nature Neuroscience* 2001;4(5):540–545.
- Marcus GF. Can connectionism save constructivism? *Cognition* 1998a;66:153–182. [PubMed: 9677762]
- Marcus GF. Rethinking eliminative connectionism. *Cognitive Psychology* 1998b;37(3):243–282. [PubMed: 9892549]
- Marcus, GF. The algebraic mind: Integrating connectionism and cognitive science. Cambridge, MA: MIT Press; 2001.
- Marcus, GF. The birth of the mind: How a tiny number of genes creates the complexities of human thought. New York: Basic Books; 2004.
- Marcus GF. Cognitive architecture and descent with modification. *Cognition* 2006;101:443–465. [PubMed: 16750185]
- Marcus, G. Kluge: The haphazard construction of the human mind. Boston: Houghton-Mifflin; 2008.
- Marcus GF, Brinkmann U, Clahsen H, Wiese R, Pinker S. German inflection: The exception that proves the rule. *Cognitive Psychology* 1995;29:186–256.
- Marcus GF, Keil FC. Concepts, correlations, and some challenges for connectionist cognition. *Behavioral and Brain Sciences*. 2008Forthcoming
- Marcus GF, Pinker S, Ullman M, Hollander JM, Rosen TJ, Xu F. Overregularization in language acquisition. *Monographs of the Society for Research in Child Development* 1992;57(4 Serial No 228)
- Marcus, GF. Tree structure and the representation of sentences: a reappraisal. In: Bolhuis, J.; Everaert, M., editors. *Birdsong, Speech & Language*. Cambridge, MA: MIT Press; 2009.
- Marler, P. The instinct to learn. In: Carey, S.; Gelman, R., editors. *The epigenesis of mind: Essays on biology and cognition*. Hillsdale, NJ: Erlbaum; 1991. p. 37-66.
- McClelland, JL.; Rumelhart, DE.; Hinton, GE. The appeal of parallel distributed processing. In: McClelland, JL.; Rumelhart, DE.; The PDP Research Group. , editors. *Parallel distributed processing: Explorations in the microstructure of cognition Volume 1: Foundations*. Cambridge, MA: MIT Press; 1986a. p. 365-422.
- McClelland, JL.; Rumelhart, DE.; The PDP Research Group. *Parallel distributed processing: Explorations in the microstructure of cognition Volume 1: Foundations*. Cambridge, MA: MIT Press; 1986b.
- McElree B. Working memory and focal attention. *Journal of Experimental Psychology. Learning, Memory, and Cognition* 2001;27(3):817–835.
- Menand L. What comes naturally: Does evolution explain who we are? *The New Yorker* 2002 November 25;;96–101.
- Mervis, CB.; Morris, CA.; Bertrand, J.; Robinson, BF. Williams syndrome: Findings from an integrated program of research. In: Tager-Flusberg, H., editor. *Neurodevelopmental disorders*. Cambridge, MA: MIT Press; 1999. p. 65-110.
- Merzenich MM, Nelson RJ, Stryker MP, Cynader MS, Schoppmann A, Zook JM. Somatosensory cortical map changes following digit amputation in adult monkeys. *Journal of Comparative Neurology* 1984;224(4):591–605. [PubMed: 6725633]
- Miller GA. The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review* 1956;63(2):81–97. [PubMed: 13310704]
- Miller, G.; Chomsky, NA. Finitary models of language users. In: Luce, RD.; Bush, RR.; Galanter, E., editors. *Handbook of mathematical psychology. Vol. II*. New York: Wiley; 1963. p. 419-492.
- Morris CA, Mervis CB. Williams syndrome and related disorders. *Annual Review of Genomics and Human Genetics* 2000;1:461–484.

- Nadel, L. Down syndrome in cognitive neuroscience perspective. In: Tager-Flusberg, H., editor. *Neurodevelopmental disorders*. Cambridge, MA: MIT Press; 1999. p. 197-222.
- Newell A. Physical symbol systems. *Cognitive Science* 1980;4:135-183.
- Newell A, Simon HA. Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery* 1975;19:113-136.
- Newport EL. Maturational constraints on language learning. *Cognitive Science* 1990;14:11-28.
- Pica P, Lemer C, Izard V, Dehaene S. Exact and approximate arithmetic in an Amazonian indigene group. *Science* 2004;306(5695):499-503. [PubMed: 15486303]
- Pinker S. Rules of language. *Science* 1991;253:530-535. [PubMed: 1857983]
- Pinker, S. *The language instinct*. New York: Morrow; 1994.
- Pinker, S. *How the mind works*. New York: Norton; 1997.
- Pinker, S. *The Blank Slate*. New York: Viking Penguin; 2002.
- Pinker S, Prince A. On language and connectionism: analysis of a parallel distributed processing model of language acquisition. *Cognition* 1988;28:73-193. [PubMed: 2450717]
- Pitkow LJ, Sharer CA, Ren X, Insel TR, Terwilliger EF, Young LJ. Facilitation of affiliation and pair-bond formation by vasopressin receptor gene transfer into the ventral forebrain of a monogamous vole. *Journal of Neuroscience* 2001;21(18):7392-7396. [PubMed: 11549749]
- Plomin R, Kovas Y. Generalist genes and learning disabilities. *Psychological Bulletin* 2005;131(4):592-617. [PubMed: 16060804]
- Plunkett K, Marchman V. U-shaped learning and frequency effects in a multi-layered perceptron: implications for child language acquisition. *Cognition* 1991;38(1):43-102. [PubMed: 2015756]
- Plunkett K, Marchman V. From rote learning to system building: acquiring verb morphology in children and connectionist nets. *Cognition* 1993;48(1):21-69. [PubMed: 8403834]
- Prasada S, Pinker S. Similarity-based and rule-based generalizations in inflectional morphology. *Language and Cognitive Processes* 1993;8:1-56.
- Pylyshyn ZW, Storm RW. Tracking multiple independent targets: Evidence for a parallel tracking system. *Spatial Vision* 1988;3:179-197. [PubMed: 3153671]
- Quartz SR, Sejnowski TJ. The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences* 1997;20:537-556. [PubMed: 10097006]discussion 556-596
- Rogers, TT.; McClelland, JL. *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press; 2004.
- Rumelhart, DE.; McClelland, JL. On learning the past tenses of English verbs. In: McClelland, JL.; Rumelhart, DE.; The PDP Research Group. , editors. *Parallel distributed processing: Explorations in the microstructure of cognition Volume 2: Psychological and biological models*. Cambridge, MA: MIT Press; 1986a. p. 216-271.
- Rumelhart, DE.; McClelland, JL. PDP models and general issues in cognitive science. In: Rumelhart, DE.; McClelland, JL.; The PDP Research Group. , editors. *Parallel distributed processing: Explorations in the microstructure of cognition Volume 1: Foundations*. Cambridge, MA: MIT Press; 1986b. p. 110-146.
- Rumelhart, DE.; McClelland, JL.; The PDP Research Group. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models*. Cambridge, MA: MIT Press; 1986.
- Rumelhart, DE.; Todd, PM. Learning and connectionist representations. In: Meyer, DE.; Kornblum, S., editors. *Attention and Performance XIV*. Cambridge, MA: MIT Press; 1993. p. 3-30.
- Schlaug G. The brain of musicians. A model for functional and structural adaptation. *Annals of the New York Academy of Sciences* 2001;930:281-299. [PubMed: 11458836]
- Senghas A, Coppola M. Children creating language: How Nicaraguan sign language acquired a spatial grammar. *Psychological Science* 2001;12(4):323-328. [PubMed: 11476100]
- Shafir E, LeBoeuf RA. Rationality. *Annual Review of Psychology* 2002;53:491-517.
- Smith, SM. Context and human memory. In: Roediger, HL.; Dudai, Y.; Fitzpatrick, S., editors. *Science of memory: Concepts*. New York: Oxford University Press; 2006. p. 111-114.
- Smolensky P. On the proper treatment of connectionism. *Behavioral and Brain Sciences* 1988;11:1-74.

- Spencer JP, Blumberg MS, McMurray R, Robinson SR, Samuelson LK, Tomblin JB. Short arms and talking eggs: Why we should no longer abide the nativist-empiricist debate. *Child Development Perspectives*. 2009forthcoming
- Strack F, Martin LL, Schwarz N. Priming and communication: Social determinants of information use in judgments of life satisfaction. *European Journal of Social Psychology* 1988;18(5):429–442.
- Taatgen NA, Anderson JR. Why do children learn to say “Broke”? A model of learning the past tense without feedback. *Cognition* 2002;86(2):123–155. [PubMed: 12435534]
- Tabor W, Galantucci B, Richardson D. Effects of merely local syntactic coherence on sentence processing. *Journal of Memory and Language* 2004;50(4):355–370.
- Tarr MJ, Gauthier I. FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience* 2000;3(8):764–769.
- The SLI Consortium. A genomewide scan identifies two novel loci involved in specific language impairment. *American Journal of Human Genetics* 2002;70(2):384–398. [PubMed: 11791209]
- Todorov A, Mandisodza AN, Goren A, Hall CC. Inferences of competence from faces predict election outcomes. *Science* 2005;308(5728):1623–1626. [PubMed: 15947187]
- Tooby, J.; Cosmides, L. Mapping the evolved functional organization of mind and brain. In: Gazzaniga, MS., editor. *The cognitive neurosciences*. Cambridge, MA: MIT Press; 1995. p. 1185–1197.
- Treisman AM, Gelade G. A feature-integration theory of attention. *Cognitive Psychology* 1980;12(1): 97–136. [PubMed: 7351125]
- Tversky A, Kahneman D. Judgment under uncertainty: Heuristics and biases. *Science* 1974;185(4157): 1124–1131. [PubMed: 17835457]
- Walum H, Westberg L, Henningsson S, Neiderhiser JM, Reiss D, Igl W, Ganiban JM, Spotts EL, Pedersen NL, Eriksson E, Lichtenstein P. Genetic variation in the vasopressin receptor 1a gene (AVPR1A) associates with pair-bonding behavior in humans. *Proceedings of the National Academy of Sciences of the United States of America* 2008;105(37):14153–14156. [PubMed: 18765804]
- Watson, JB. *Behaviorism* [microform]. New York: W.W. Norton; 1925.
- Wilson TD, Brekke N. Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin* 1994;116(1):117–142. [PubMed: 8078969]
- Wright WG. Evolution of nonassociative learning: Behavioral analysis of a phylogenetic lesion. *Neurobiology of Learning and Memory* 1998;69(3):326–337. [PubMed: 9707494]




Fig. 1. The relationship between language and other aspects of cognition (e.g., reasoning and motor control) under three evolutionary scenarios. Panel A: Language as separate from cognition. Panel B: Language as a subset of cognition. Panel C: Language as descended with modification from cognition.

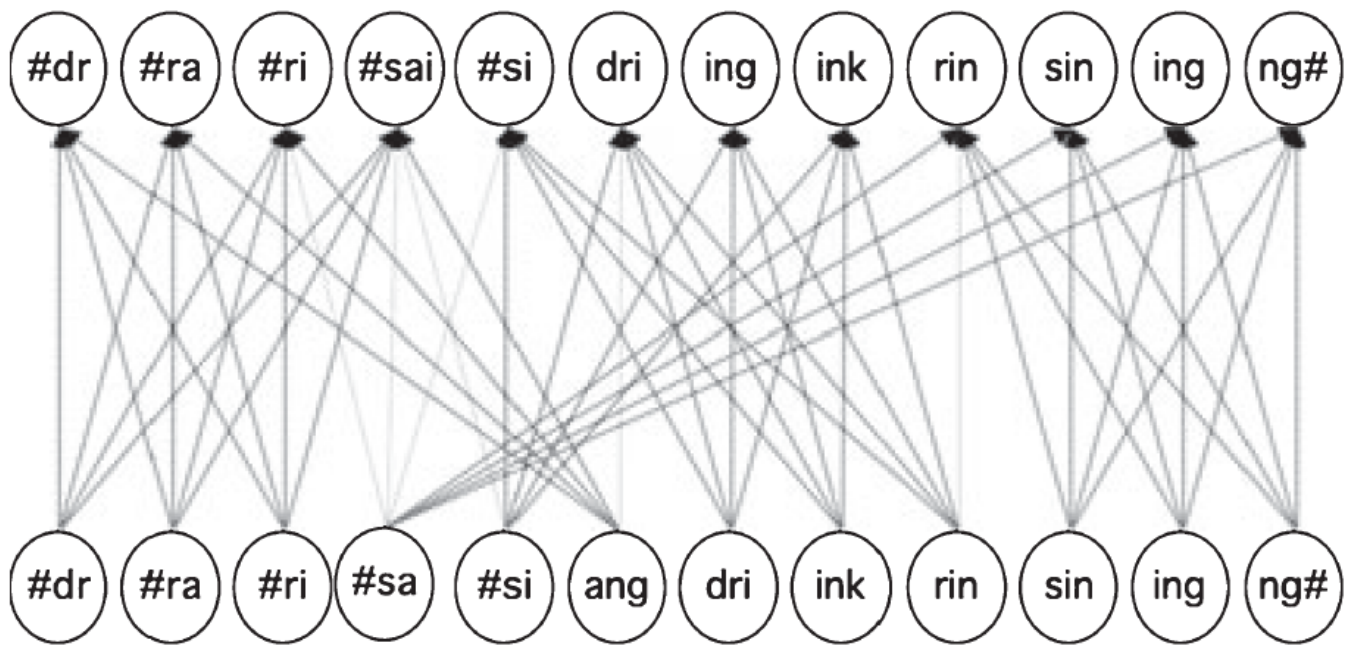


Fig. 2.
Simplified version of Rumelhart and McClelland's model of the acquisition of the English past tense. Input nodes (bottom) represent a verb's stem, in terms of a distributed set of nodes that denote sequences of phonological elements. Output nodes (top) represent that verb's predicted past tense form, again in terms of distributed sets of phonological elements. Lines between nodes represent adjustable connections, tuned through experience.