

On Language

Steven Pinker

Interview

■ Steven Pinker is a professor in the Department of Brain and Cognitive Sciences at MIT, and in 1994 will become director of its McDonnell-Pew Center for Cognitive Neuroscience. He received his B.A. from McGill University in 1976 and his Ph.D. from Harvard University in 1979, both in experimental psychology, and taught at Harvard and Stanford before joining the faculty of MIT in 1982. He has done research in visual cognition and the psychology of language, and is the author of *Language Learnability and Language Development* (1984) and *Learnability and Cognition* (1989) and the editor of *Visual Cognition*

(1985), *Connections and Symbols* (1988, with Jacques Mehler), and *Lexical and Conceptual Semantics* (1992, with Beth Levin). He was the recipient of the Early Career Award in 1984 and the Boyd McCandless Award in 1986 from the American Psychological Association, a Graduate Teaching Award from MIT in 1986, and the Troland Research Award from the National Academy of Sciences in 1993. His newest book, *The Language Instinct*, will be published by William Morrow & Company in January 1994. ■

JOCN: You have written a new book about the nature of human language for a general audience. Tell us what you want to communicate in *The Language Instinct*?

SP: The book is about all aspects of human language. One thing I wanted to do was to answer the kinds of questions I get asked when I tell people I study language, such as, are there really pockets of the Ozarks where people speak Elizabethan English? What language would a child speak if he was raised by wolves? What's going on when religious people "speak in tongues?" Why can't computers take dictation? Why does no one know the plural of "Walkman"? Why is English spelling so deranged? What's the scoop about the search for the mother of all languages?

But my main goal was to try to unify the study of language under a key idea: that language is an evolutionary adaptation, like echolocation in bats or the elephant's trunk. This may seem like a boring observation, but it buys a lot. It allows for a vertically integrated science of language, where everything from genes and neural networks to Orwell and dudespeak can be fit into a consistent framework. And a lot of controversies just disappear, such as whether syntactic form or semantic function is more important, or whether there would be some evolutionary paradox if humans turned out to be the only species with language. After all, no one gets upset at the idea that the elephant's trunk is both structured and useful, or that it is complex but found in only one species.

Also, treating language as a biological adaptation overturns many folk theories that pervade modern intellectual life. The books on language that you will find at your local B. Dalton all treat language as some obscure

body of lore, like how to use "hopefully" correctly or the correct term for a collection of larks, something that has to be carefully passed on by English teachers and "language mavens" (who are mostly quacks, by the way). I try to show that the complexity of language really comes from the minds of ordinary children and Joe Six-packs; the rules of the schoolmarm are just minor little decorations. Also, the general picture of the human mind that you find in books and magazines—basically the blank slate, together with the concession that of course heredity and environment are inseparably interconnectedly intertwiningly intermingled—turns out to be woefully lame. I think that with what we now know about language, we can do better.

JOCN: Well some might read you as saying that language is innate and Chomsky made that point in 1959 with his review of Skinner's "Verbal Behavior." Before going into details, how does your MIT view differ from other MIT views?

SP: Obviously, some of the key ideas in the book come from Chomsky—that there is an innate neural system dedicated to language; that his system uses a discrete combinatorial code, or grammar, to map between sound and meaning; that this code manipulates data structures that are dedicated to language and not reducible to perception, articulation, or concepts. But there are also some differences in style and substance. Chomsky's arguments for the innateness of language are based on technical analyses of word and sentence structure, together with some perfunctory remarks on universality and acquisition. I think converging evidence is crucial, and try to summarize the facts on children's development, cross-linguistic surveys, genetic language disorders, and so on.

In that sense the book is more in the tradition of George Miller and Eric Lenneberg than Chomsky. And there is one substantive difference: I argue that language is an adaptation, a product of natural selection, and hence has parts that are designed for specific functions involved in communication. Chomsky is agnostic-to-hostile about natural selection and its applicability to language. He suspects that the language faculty could have come about as an accidental product of the way the laws of physics act on the developing brain. He has even suggested that grammar appears to have been designed for beauty, not for usefulness.

JOCN: Is that a way of saying that Chomsky has not really thought about how language could have developed from a biological (evolutionary) point of view? It would seem he simply observed it must be established through genetic mechanisms and moved on to other matters.

SP: The vast majority of cognitive scientists and neuroscientists have not really thought about the evolution of the brain, but I don't think that is true of Chomsky; he has thought about evolution a great deal. But he is not familiar with mainstream evolutionary biology. His views are more in the tradition of the many physicists who suspect that the theory of natural selection is somehow vacuous and circular, or that the mechanism is too klugey and random to create interesting biological systems. Their esthetic is that biological structures should somehow follow deductively from general laws and principles, like the growth of crystals.

JOCN: Evolutionary theory does suggest the brain is a kluge, a collection of ad hoc systems that somehow get the job done. Does your account of language, driven as it is from an evolutionary perspective, document and/or identify various language operations as a collection of adaptations that has accumulated over time? Chomskyans talk about the language organ as an entity that suddenly appears on the scene (in the brain). Does your analysis suggest this first approximation is insufficient in light of what is now known about language organization?

SP: The point about biological systems being collections of kluges has been a bit overdone by psychologists. If you look at many biological systems you see astoundingly sophisticated engineering, the eye being the most famous example, but with a few kluges (like the retina being installed backward) that reveal their origin as the result of selection of random small changes from an ancestral form. I think in language we certainly see signs of engineering to carry out a function. Syntax and morphology are codes that map multidimensional semantic data structures onto strings of symbols that can be transmitted through a serial interface. Phonology allows a finite number of sound units to be rearranged to form an open-ended set of words, and phonetics compresses the units into a signal that transmits them at a rate that exceeds the resolving power of the ear.

And, as you would expect, there are oddities and quirks that suggest that language was not designed de-

liberately or from scratch. The descended larynx is the obvious physiological example (good for making speech sounds at the cost of making us more likely to choke on food). At the computational level, you find examples of seemingly needless redundancy, like the fact that certain information can be conveyed either by rote memorizing lexical items or by composing structures with grammatical rules, giving us the contrast between regular and irregular forms (*take-took* versus *bake-baked*) and the resulting grammatical mayhem. The verb and preposition system seem to have been designed to convey spatial and force-dynamic information, so we have to resort to motion metaphors when talking about abstract information such as state changes (e.g., *Sam went from being sick to being well*). And no one knows how to talk about something possessed by two people (*Rob and my mother? Rob and me's mother? Rob's and my mother?*).

Regarding language appearing all at once—I don't think Chomsky has made that claim, though Derek Bickerton has suggested that it appeared in two stages. It strikes me as unlikely, for standard evolutionary reasons. If language is a complex system involving many finely interacting parts that collectively do something interesting (as Chomsky himself has shown) then by the laws of probability you would not expect one random mutation to give some fortunate ancestor all of the necessary neural modifications in one thunderclap. I also think there is evidence from neuroscience and genetics that speaks against language emerging as an automatic physical by-product of some more global development such as a large brain (which is something Chomsky has conjectured as a possibility). Across normal variation and pathology you see big differences in brain size, shape, and global organization that can coexist with intact language. This suggests that it is a certain wiring of the microcircuitry that is essential. Also, in cases of genetically transmitted specific language impairment, you don't find language wiped out entirely, but different components affected to different extents. If a single lucky genetic change had given us language, it should be possible for a mutation in that gene to wipe language out entirely, but one never sees that.

JOCN: Hmm, sounds like a good topic for someone working on gene knockout to consider. But let's switch gears. Sitting majestically on the other side of the biologic interpretation of language acquisition are those that feel language is learned and built up through associations. Even though Chomsky seemed to have devoured Skinner's early analysis of language along these lines, there are modern reincarnations of the environmentalist views. Neurocomputationalists have any number of algorithms that they feel can learn and handle language learning. Ever since Rosenberg and Sejnowski demonstrated that a simple neural net could read, the field has blossomed with claims. How do you see it?

SP: I feel that a lot of the work in artificial "neural networks" is based more on eighteenth-century notions

of learning—mainly associative pairing and generalization by similarity—than on any systematic empirical study of what organisms' brains are computing. Randy Gallistel has made this point in reviewing the remarkable computational abilities of various animals in domains such as navigation and the perception of time and number, none of which has any need for the classical associative bond. I think his critique carries over to language (because I think psycholinguistics is a branch of ethology).

I don't think the interesting issue is whether such-and-such a class of model is capable of learning X given a suitable training regimen. That's a mathematical point, not a scientific one, and there is a consensus that most of the commonly discussed artificial neural networks can be designed either with Turing power or as universal function approximators. The interesting question is empirical: how, in fact, does the neural circuitry underlying language (or any other mental ability) work? To answer that you can't just wire up any old model, train the daylights out of it, and declare victory; you have to check if the wiring diagram really corresponds to the plausible innate organization of the creature, and whether the training schedule plausibly corresponds to its experience. In many of the connectionist models of language, neither is done. Just take the various reading-aloud models, where the designer wires together a network designed to map from visual symbols to phonology, and the network has to learn the exact mapping. Taken literally, this is a claim that we are innately designed to read—the very ability where we are 100% sure that there can be no innate faculty! Similarly, neural net modelers have no compunctions about building in innate wiring to perform artificial tasks of the late twentieth-century experimental psychologist, like lexical decision. The issue is not whether there is learning or innate wiring—obviously there's both. The issue is what in fact *is* the innate wiring and learning experience.

This is especially clear to me in my own empirical work, on the linguistic computation we do when creating past tense forms such as *faxed* or *broke*. In 1986 Rumelhart and McClelland published a brilliant study of how a simple feedforward network learns these mappings. Alan Prince and I noted a number of ways in which the model behaved systematically differently from people. Many involved linguistic quirks such as the fact that people can easily inflect weird-sounding verbs such as *to out-Gorbachev*, that *uring* and *ring* are homophonous but have different past tense forms (and so the input to the past tense system cannot be sound alone), and that verbs formed from nouns and adjectives, such as *flied out to center field* and *ringed the city*, always take regular *-ed* even if homophonous with an irregular verb.

The response of many connectionists to our critique (not Rumelhart and McClelland themselves) was to “improve” the model by adding hacks designed to handle each one of these quirks and train the improved model with the crucial examples—amounting to the bizarre

claim that the brain is specifically wired, and children are specifically taught, in such a way as to make the quirks come about! This was all meant to show that a connectionist model can, in principle, handle the phenomena, but that was never our dispute. Our point was that these quirks are by-products of some fundamental ways in which the language system is organized, and that any model of how language is implemented in the brain will have to reflect that organization. In particular, we showed that the past tense computation requires at least three things: a division into subsystems (most fundamentally, the mental dictionary and the mental grammar); some way of representing the identity of entities as distinct individuals, independent of their phonological and semantic content; and a computational operation that can concatenate variables, not just analogize. These are not particularly extravagant claims, and one can imagine all kinds of neural networks that can implement them. But the standard model of a single associative network has become such doctrine that people will go to any lengths to maintain it, even if it involves innately wiring in peculiarities of English grammar and sticking exotic cases into the training set.

JOCN: By yanking language learning out of the field of learning mechanisms and marking it down as one of our instincts, don't you sort of also trumpet the end of its study? After all, sex is an instinct and while people study the physiological basis for aspects of sex such as arousal, there is not much more to say about it. In short, once the descriptive work is done, and the rules are written for biologically based grammar, can't you go fishing?

SP: No, I would disagree with all those assumptions, starting with sex. There is plenty to say about the cognitive psychology of sex, as shown in the work of evolutionary psychologists such as Don Symons, David Buss, Margo Wilson, and Martin Daly. The sexual “instinct” surely involves many complex information-processing mechanisms—the psychophysics of sexual attractiveness, short-term and long-term strategies for courtship and manipulation and for evaluating and resisting such tactics, and decision rules for commitment versus desertion. And many of these mechanisms are surely specific to the domain of sex, not social relations in general—as Fran Lebowitz said, you would never choose someone as a close friend because he had a really cute nose.

One of the achievements of linguistics is to show that even if a language instinct is innate, that does not mean that we announce that language belongs to the physiologists and leave it at that. There is a huge body of ongoing research showing how detailed facts of English fall out of the computational organization of the mental grammar for English, and how the mental grammar for English falls out of the universal grammar underlying all languages. But this is technical, somewhat difficult work, and many psychologists find their eyes glazing over when, say, Chomsky starts going on about *John is too stubborn for anyone to talk to*. They read the first and

last 10 pages of his books, with all the exciting rhetoric about innateness, skip the actual content, and then dimly remember the claims in a bumper-sticker form (“language is innate”) that seems to leave no room for anyone but the physiologists.

More generally, it’s wrong to equate “instinct” and “physiology” (or “innateness” and “wiring”). Neural tissue, at the level at which a physiologist studies it, is not going to do anything that corresponds directly to interesting psychology. That is, you might discover that the geometry of a dendrite or the strength of a synapse functions in some way that corresponds to a logic gate or a paired association, but it’s not going to correspond to the image of a sexy mate or the constituents of a prepositional phrase. I believe in the fairly standard view that cognitive abilities consist of some sequence or network of more elementary information processes, and that the elementary information processes are the kinds of things that neurons and simple neural nets can do—that is the level where a cognitive scientist might hand the baton to the physiologist and go fishing. But most important, this hierarchical explanation is needed whether an ability is learned or innate, just as the same sequence of instructions can either be programmed into a computer’s memory by the user or burned into a ROM at the factory (a crude analogy, I know).

JOCN: In your book you constantly argue from concrete, almost home-spun everyday examples from ordinary language use to make strong arguments for the genetic basis of language. Could you speak for a moment about the study of linguistics and how it approaches a problem empirically so as to allow the use of those examples?

SP: Actually, both in the book and in my day-to-day research, I try to get data of very different kinds to converge before concluding anything. The most commonly used data in linguistics are judgments about whether some word or sentence sounds natural to a speaker of the language, and what the speaker takes it to mean. It’s a kind of psychophysics done on oneself and one’s readers, a lot like a demonstration that a Necker cube flips or that isoluminant pictures lack depth. Sometimes, to get higher-precision data on squishy cases I get numerical ratings of grammaticality or meaning from sophomores, but it’s the same kind of data, and the judgments and ratings always coincide. (In fact, the *F*-ratios are often in the 600s, so these studies can get published in the psychology journals where they would otherwise reject anything that seemed too “linguistic.”) Generally the linguistic judgments are the most information-rich data, but it’s also important to bolster any conclusion by other means. That is because some pattern in a person’s judgments may not have been caused by a rule of grammar implemented in his brain but by some set of individual cases that fossilized in the language centuries ago and have been memorized individually.

Take the claim I am currently working on, that regular past tense and plural forms are usually assembled on-

line by a mental rule, but irregular ones are retrieved or analogized from memory. You begin with the simple observation that new verbs entering the language automatically get regular forms—*fax*, *faxed*, *faxing*, *faxes*. This suggests some rule that adds *-ed* or *-s*, but that is just a beginning, the level of detail that connectionist memory models can also handle.

From linguistics, you can add a couple of more subtle phenomena. Verbs based on noun roots can’t have past tenses listed with their roots in memory (nouns inherently don’t have past tenses), and they turn out to be always regular, even when their sound pattern would seem to call out for an irregular form—hence *flied out* and *ringed the city*, not *flew out* and *rang the city*. This confirms that the process creating regular forms is a default operation that applies whenever memory doesn’t supply a form. Second, people accept *mice-infested* and *men-bashing* but not *rats-infested* and *guys-bashing*. Say the compounding operation takes two words from memory and glues them together. *Mice* and *men* are in memory, because they are unpredictable, whereas *rats* and *guys* are not, because they can be generated by rule when needed. Therefore the compounding operation finds *mice* in lexical memory and can glue it to *infested*, but can’t find *rats* and has to use *rat*.

Now go to the laboratory, and you can show that both these effects can be replicated in ratings by sophomores, and in experiments where you elicit new words from preschool children. Also, sophomores give low ratings and slow reaction times to low-frequency irregular forms such as *strive-strove* and unfamiliar-sounding ones such as *nist-nust* (because of their weak memory traces) but give high ratings and quick reaction times to low-frequency regular forms such as *stint-stinted* and unfamiliar-sounding ones such as *ploampb-ploampbed* (because they don’t have to be retrieved from memory, so the weakness of any memory trace is irrelevant). Now look at naturalistic speech errors in children: kids overapply the regular rule, saying things such as *buyed*, and the errors are not correlated either in time or over words with how many forms such as *tied* or *fried* their parents use, suggesting that the errors are rule products, not analogies from memory.

Finally, go to the neuropsychology clinic. A postdoc in my lab, Michael Ullman, has shown that patients with memory disorders and unimpaired grammar, such as Alzheimer’s, are fine with regular verbs and with nonsense verbs such as *wugged* but often make errors on irregulars, such as *swimmed*. Patients with impaired grammar but less impaired memory retrieval, such as Broca’s aphasics, have trouble with regular verbs and nonsense verbs, but less trouble with irregulars. The data of the linguist from everyday speech and the other kinds of data, at least in this case, fit together almost perfectly.

More generally, Chomsky’s argument for the innateness of the language system is based on the discovery that there is information in people’s judgments of words

and sentences that is not in the input they heard as children. To make the claim more precise and concrete, the psychologist Peter Gordon ran an experiment showing that preschoolers say *mice-eater* but never *rats-eater*, even though one can estimate that they have never heard their parents use either kind of construction and thus could not have learned the distinction from the input; it must have come from the way their lexicons are inherently organized vis-à-vis their grammars. And there are other kinds of data that line up with the claim from linguistic examples: children's precocity at mastering fine points of grammar, even the useless ones; various syndromes in which severely retarded people have intact language; language universals that cannot be attributed to mere utility; the uniform grammatical sophistication across cultures and subcultures despite vast differences in other measures of cultural sophistication; and, perhaps most interesting, cases where children create a grammatically more sophisticated language than the one they hear from their parents.

JOCN: And needless to say, I assume these kind of analysis work for all languages. They are not flukes of English.

SP: Most immediately, we find the family of phenomena related to regular and irregular morphology in other languages, such as French and Arabic. The best comparison is German, because its statistics are completely different from English. Its version of the plural *-s* and participle *-ed* applies to a *minority* of words, not a majority; the majority of words are irregular. But the suffixes show almost the entire set of effects we find in the English versions, quirks and all—the frequency and similarity effects, the *mice-infested/rats-infested* effect, the *flied out* effect, about a dozen in all. This shows that the hallmarks of a mental rule are not an epiphenomenon of the supposedly “rule-governed” inflection being the majority of cases in the child's experience. They involve a qualitative difference in the way that the brain computes a rule and the way it looks up items in memory.

There's no doubt that languages differ a lot. But I think the evidence is that the same kinds of computational machinery are used in all of them—the division into components, the kinds of data structures used in each component such as nouns and verbs and lexical and phrasal heads, and so on. Different languages use each of these gadgets to different extents, giving the appearance of radical differences. For example, in many native American languages you can build an entire sentence out of a verb by sticking strings of prefixes and suffixes onto it that specify key properties of the verb's arguments, without having to select and place a bunch of noun phrases. It seems completely different from English. But then we have this silly little agreement rule—*the girl eats* versus *the girls eat*—that is essentially the same mechanism. The isn't doing much in English; if it disappeared, no one would miss it. It's computationally costly to use and even harder to learn, but English children use it

correctly more than 90% of the time by the time they turn 4. This suggests that the mental algorithms necessary for supposedly radically different languages are available to all humans.

JOCN: Well, if language is to be understood in this biological sense, in how our species actually operates as opposed to how nineteenth-century grammarians would like us to speak and write, do you see the language cultists such as William Safire as Johnny-come-lately technologists? I mean all of us like sex in the biological sense, but some of us are artists.

SP: Yes, quite right. The guidelines for good style, standardized “proper” grammar, and so on, are at their best technological add-ons that help us use language for purposes that it was not designed for, basically, putting esoteric thoughts on paper for the benefit of strangers. It's an important technology but it is quite different (and far less interesting) than the basic unconscious grammar that we all use to put words together in ordinary conversation. An analogy might be the rules for an illegal defense in basketball compared to the motor control programs for bipedal locomotion. And at its worst, “proper” grammar is just plain dumb, like the screwball rules against *hopefully*, split infinitives, *Everybody returned to their seats*, and so on.

Actually, the so-called “language mavens” such as Safire are more like witchdoctors than technologists. Like many linguists, I am always astounded at how ignorant they are about language. It's not just that they don't know this week's version of formal Chomskyan theory; they can't work through freshman problems of grammatical analysis (such as telling a verb from an adjective), and they have no knowledge of the basic facts of English—what kinds of idioms and constructions there are, and how they are used and pronounced. It comes from a general condescension about the speech of the common person, which they consistently underestimate, and ultimately from a nonscientific, uncurious attitude—like many people, they are blasé about unconscious mental processes that work most of the time.

JOCN: Do you see language as a system separate and distinct from those brain processes that allow for complex thought?

SP: To a large extent, yes. For one thing, the algorithms for grammar cut across the logic of conceptual categories, as in the *mice-eaters* who eat mice differing from the *rat-eaters* who eat rats. Grammar is a communications protocol, not our knowledge database or our inference system. Moreover it is a protocol that has to interface the ear, the mouth, and the mind. So it's no surprise that it doesn't reflect any one of them directly, but has a logic of its own. Also, you find language and general intelligence dissociating in many populations—beginning with toddlers, who are grammatical geniuses but incompetent at just about everything else. In Williams Syndrome, hydrocephalus, and Alzheimer's Disease, you can find intact

grammar despite severe deficits in general intelligence, and in some cases of aphasia and genetically transmitted language impairment you find the opposite.

Of course they cannot be completely separate. Language has to interface with the conceptual world, so there has to be a level of semantic representation that is built out of the same kinds of primitives as concepts. When we use language we are engaging in a kind of interpersonal interaction that must involve social reasoning processes. And at the lower neural levels I surely wouldn't expect language to use green neurons and complex thought to use red neurons or anything like that.

JOCN: Many scientists are now trying to understand the neural basis of language. There are many new imaging techniques (PET, functional MR, ERPS) and laboratories all over the world are studying language processes and looking for patterns of cerebral activation. What is your opinion of this work? Are they asking the right question? Or is perhaps the answer to how the brain enables language to be found more in considering the properties of local circuits?

SP: Ultimately the answer is in the local circuitry that actually does the computing. Methods such as aphasiology and neuroimaging are a bit like using bomb craters and blurred satellite photos to understand the long-distance telephone network. But of course the neural basis of language has to be studied at many scales, and the neuroimaging methods are very exciting and important. The work to date is very intriguing and a good start, but when you look closely at the whole literature you see some problems. David Poeppel, a graduate student in our program, reviewed the three published studies that claimed to have found the areas involved in phonological processing. He found that the overlap in the three teams' lists of "phonology areas" was zero! Poeppel found the same thing in his review of studies that claimed to have found the areas involved in semantic processing—three completely nonoverlapping lists of "semantic areas." Even more depressing, for both the semantics and the phonology areas, is that all of the teams managed to cite studies in the aphasiology literature that they claimed were consistent with their PET finding. Obviously there are some bugs to be worked out.

One of the problems is that none of the teams studying "language processing" has included a linguist or psycholinguist. Their models of language processing just seem to be made up on the spur of the moment. So they'll have some task such as judging whether two syllables end in the same consonant as an example of "phonological processing." Now, any circuitry for phonology is going to be doing much finer-grained analyses than called for in that task—we did not evolve a brain area to press buttons indicating whether a nonsense syllable

ends in a consonant. The task surely involves a whole slew of linguistic and cognitive processes other than phonology, such as parsing and remembering words, perhaps orthographic recoding, and the overhead of remembering the task and generating the appropriate button-press. So it's not surprising that the areas that light up after you subtract passive listening are not at all specific to phonology, and could show no overlap with the areas involved in the "phonology" task in some other lab, which might string together some other arbitrary collection of procedures. Likewise in looking for semantic areas, a task such as generating a verb that goes with a noun just doesn't correspond to any cohesive cognitive process, let alone being a test of language.

More generally, I wonder whether PET research so far has taken the methods of experimental psychology too seriously. In standard psychology we need to have the subject do some task with an externalizable yes-or-no answer so that we have some reaction times and error rates to analyze—those are our only data. But with neuroimaging you're looking at the brain directly so you literally don't need the button-press or the overt blurring. I wonder whether we can be more clever in figuring out how to get subjects to think certain kinds of thoughts silently, without forcing them to do some arbitrary classification task as well. I suspect that when you have people do some artificial task and look at their brains, the strongest activity you'll see is in the parts of the brain that are responsible for doing artificial tasks. Still, it's an intriguing beginning and like most cognitive scientists I'm following it eagerly.

JOCN: Well, it is an exciting story and a superb and fascinating book. I would like to conclude by asking your thoughts about how your studies of language offer insights into mind design.

SP: If language, the quintessential higher cognitive process, is an instinct, maybe the rest of cognition is a bunch of instincts too—complex circuits designed by natural selection, each dedicated to solving a particular family of computational problems posed by the ways of life we adopted millions of years ago. Aside from language, these might include systems for intuitive physics, biology, and psychology, mental maps, habitat, kinship, mating, danger, food, disease, justice, friendship, and self-monitoring. This is very different from the standard conception of some nondescript but all-powerful "culture" from the social sciences, "information processing" from cognitive psychology, or "association cortex" from neuroscience, concepts that I suspect will go the way of "protoplasm" in biology. But admittedly this is a big leap from irregular German participles.

JOCN: Thank you.

STATEMENT OF OWNERSHIP, MANAGEMENT AND CIRCULATION (Required by 39 U.S.C. 3685)

1. A. Title of Publication: Journal of Cognitive Neuroscience B. Publication No.: 0898929X 2. Date of Filing: 10/1/93 3. Frequency of Issue: Quarterly; Winter, Spring, Summer, Fall A. No. of issues published annually: Four B. Annual subscription price: \$58/Ind \$145/Inst 4. Location of known office of publication (not printer): MIT Press, 55 Hayward Street, Cambridge, Middlesex, MA 02142-1399 5. Location of Headquarters of General Business office of the publisher (not printer): same as item 4 6. Names and Complete Addresses of Publisher, Editor, and Managing Editor: Publisher: MIT Press, 55 Hayward Street, Cambridge, MA 02142-1399 Editor: Michael Gazzaniga, Center for Neurobiology, Univ. of California Davis, CA 95616 Managing Editor: Charlotte Smylie, JOURNAL OF COGNITIVE NEUROSCIENCE, Center for Neurobiology, Univ. of California, Davis, CA 95616 7. Owner (if owned by a corporation, its name and address must be stated and also immediately thereunder the names and addresses of stock-holders owning or holding 1 percent or more of total amount of stock. If not owned by a corporation, the names and addresses of the individual owners must be given. If owned by a partnership or other unincorporated firm, its name and address, as well as that of each individual, must be given. If the publication is published by a nonprofit organization, its name and address must be stated.): MIT Press, 55 Hayward Street, Cambridge, MA 021428. Known bondholders, mortgagees, and other security holders owning or holding 1 percent or more of total amount of bonds, mortgages, or other securities: None 9. For completion by nonprofit organizations authorized to mail at special rates: Has not changed during preceding 12 months 10. Extent and nature of circulation: A. Total number of copies printed (not press run): Average number of copies each issue during preceding 12 months: 1416 Actual number of copies of single issue published nearest to filing date: 1523 B. Paid and/or requested circulation: 1. Sales through dealers and carriers, street vendors and counter sales: Average number of copies each issue during preceding 12 months: 0 Actual number of copies of single issue published nearest to filing date: 0 2. Mail Subscriptions: Average number of copies each issue during preceding 12 months: 779 Actual number of copies of single issue published nearest to filing date: 976 C. Total paid circulation: Average number of copies each issue during preceding 12 months: 779 Actual number of copies of single issue published nearest to filing date: 976 D. Free distribution by mail, carrier or other means - samples, complimentary and other free copies: Average number of copies each issue during preceding 12 months: 26 Actual number of copies of single issue published nearest to filing date: 58 E. Total distribution (Sum of C and D): Average number of copies each issue during preceding 12 months: 805 Actual number of copies of single issue published nearest to filing date: 1034 F. Copies not distributed: 1. Office use, left-over, unaccounted, spoiled after printing: Average number of copies each issue during preceding 12 months: 611 Actual number of copies of single issue published nearest to filing date: 489. Returns from news agents: Average number of copies each issue during preceding 12 months: 0 Actual number of copies of single issue published nearest to filing date: 0 G. Total (Sum of E and F-should equal net press run shown in A): Average number of copies each issue during preceding 12 months: 1416 Actual number of copies of single issue published nearest to filing date: 1523 I certify that the statements made above by me are correct and complete. Dorothy Devereaux, Circulation Manager