

78 Mapping the Evolved Functional Organization of Mind and Brain

JOHN TOOBY AND LEDA COSMIDES

ABSTRACT The human brain is a biological system produced by the evolutionary process, and so cognitive neuroscience is itself a branch of modern evolutionary biology. Accordingly, cognitive neuroscientists can benefit by acquiring a professional knowledge of the recent technical advances made in evolutionary biology and by applying them to their research. Useful tools include the biologically rigorous concept of function that is appropriate to neural and cognitive systems; a growing list of the specialized functions the human brain evolved to perform; and criteria for distinguishing the narrowly functional aspects of the neural and cognitive architecture that are responsible for the brain's organization from the much larger set of properties that are by-products or noise. With such tools, researchers can construct biologically meaningful experimental tasks and stimuli. These are more likely to activate the large array of functionally dedicated mechanisms that constitute the core of human brain function, but which are at present largely unstudied.

Nothing in biology makes sense except in the light of evolution.
—T. Dobzhansky

It is the theory which decides what we can observe.
—Einstein

Seeing with new eyes: Toward an evolutionarily informed cognitive neuroscience

The task of cognitive neuroscience is to map the information-processing structure of the human mind, and to discover how this computational organization is implemented in the physical organization of the brain. The central impediment to progress is obvious: The human brain is, by many orders of magnitude, the most complex system humans have yet investigated.

JOHN TOOBY Department of Anthropology, and LEDA COSMIDES Department of Psychology, Center for Evolutionary Psychology, University of California, Santa Barbara, Calif.

Purely as a physical system, the vast intricacy of chemical and electrical interactions among roughly one hundred billion neurons defeats any straightforward attempt to build a comprehensive model, as one might attempt to do with particle collisions, geological processes, protein folding, or host-parasite interactions. At present, the underlying logic of the system seems lost among the torrent of observations that have been accumulated to date, and obscured by the inherent complexity of the system.

Historically, however, well-established theories from one discipline have functioned as organs of perception for others. They allow new relationships to be observed and make visible elegant systems of organization that had previously eluded detection. It seems worth exploring whether the same could be true for the brain sciences.

In fact, the brain is more than a physical system: It is both a computational system and an evolved biological system. Although cognitive neuroscience began with the recognition that studying the brain as a computational system would offer important new insights, the field has so far failed to take equal advantage of the fact that the brain is an evolved system as well. Indeed, the brain is a computational system that was organized and specifically designed to solve a narrowly identifiable set of biological information-processing problems. For this reason, evolutionary biology can supply a key missing element in the cognitive neuroscience research program: a list of the native information-processing functions that the human brain was built to execute. Our computational architecture evolved its distinctive sets of structured information-processing relationships as devices or modules to perform this particular targeted set of adaptive functions. In turn, our neural architecture evolved its distinctive physical configu-

ration because it brought these targeted sets of functional information-processing relationships into existence. By providing the functional engineering specifications to which human brains were built to conform, evolutionary biology can help researchers to isolate, identify, activate, and map the important functional aspects of the cognitive architecture, aspects that would otherwise be lost among the myriad irrelevant phenomena in which they are embedded. The resulting maps of the computational structure of each device will then allow researchers to isolate, identify and map the functional aspects of the neural architecture. The biologically implausible view that the brain is a general-purpose information-processing system provides little guidance for research in cognitive neuroscience. In contrast, an evolutionary approach allows cognitive neuroscientists to apply a sophisticated body of new knowledge to their problems. In short, because theories and principled systems of knowledge can function as organs of perception, the incorporation of a modern evolutionary framework into cognitive neuroscience may allow the community to detect ordered relationships in phenomena that otherwise seem too complex to be understood.

Over the last 30 years, evolutionary biology has made a number of important advances that have not yet diffused into allied disciplines such as the cognitive and neural sciences. These advances constitute a potent set of new principles relevant to dissecting and understanding the phenomena studied by cognitive neuroscientists (Tooby and Cosmides, 1992). Central to these advances is the modern technical theory of evolution. This consists of the logically derivable set of causal principles that necessarily govern the dynamics of reproducing systems. These principles account for the properties that reproducing systems cumulatively acquire over successive generations. The explicit identification of this core logic has allowed the biological community to develop an increasingly comprehensive set of principles about what kinds of features can and do become incorporated into the designs of reproducing systems down their chains of descent, and what kinds of features do not (Hamilton, 1964, 1972; Maynard Smith, 1964, 1982; Williams, 1966; Dawkins, 1976, 1982, 1986; Cosmides and Tooby 1981; Tooby, 1982). This set of principles has been tested, validated, and enriched through its integration with functional and comparative anatomy, biogeography, genetics,

immunology, embryology, behavioral ecology, and a number of other disciplines.

Just as the fields of electrical and mechanical engineering summarize our knowledge of principles that govern the design of human-built machines, the field of evolutionary biology summarizes our knowledge of the engineering principles that govern the design of organisms, which can be thought of as machines built by the evolutionary process (for overviews, see Dawkins, 1976, 1982, 1986; Daly and Wilson, 1984; Krebs and Davies, 1987). Modern evolutionary biology constitutes, in effect, a foundational organism design theory, whose principles can be used to fit together research findings into coherent models of specific cognitive and neural mechanisms.

First principles: Reproduction, feedback, and the entropic construction of organic design

Within an evolutionary framework, an organism is describable as a self-reproducing machine, and the defining property of life is the presence in a system of devices or organization that cause the system to construct new and similarly reproducing systems. From this defining property of self-reproduction, the entire deductive structure of modern Darwinism logically follows (Dawkins, 1976; Williams, 1985; Tooby and Cosmides, 1990b). Because the replication of the design of the parental machine is not always error-free, randomly modified designs (i.e., mutants) are introduced into populations of reproducers. Because such machines are highly organized so that they cause the otherwise improbable outcome of constructing offspring machines, the great majority of random modifications will interfere with the complex sequence of actions necessary for self-reproduction. Consequently, such modified designs will tend to remove themselves from the population—a case of negative feedback.

However, a small residual subset of design modifications will, by chance, happen to constitute improvements in the design's machinery for causing its own reproduction. Such improved designs (by definition) cause their own frequency to increase in the population—a case of positive feedback. This increase continues until (usually) such modified designs outreproduce and thereby replace all alternative designs in the population, leading to a new species-standard design. After such an event, the population of reproducing

machines is different from the ancestral population: The population- or species-standard design has taken a step "uphill" toward a greater degree of functional organization for reproduction. This spontaneous feedback process—natural selection—is the only known process by which functional organization emerges naturally in the world, without intelligent design and intervention. Hence, all naturally occurring functional organization in organisms must be ascribed to its operation and must be consistent with its principles.

Over the long run, down chains of descent, this feedback cycle pushes the design of a species stepwise "uphill" toward arrangements of elements that are increasingly improbably well organized to cause their own reproduction in the environment the species evolved in. Because the reproductive fates of the inherited traits that coexist in the same organism are linked together, traits will be selected to enhance each other's functionality (but see Cosmides and Tooby, 1981; Tooby and Cosmides, 1990b, for the relevant genetic analysis and qualifications). Consequently, accumulating design features will tend to fit themselves together sequentially into increasingly functionally elaborated machines for reproduction, composed of constituent mechanisms—called adaptations—that solve problems whose solutions either are necessary for reproduction or increase its likelihood (Darwin, 1859; Williams, 1966, 1985; Dawkins, 1986; Thornhill, 1990; Tooby and Cosmides, 1990b). Significantly, in species like humans, genetic processes insure that complex adaptations are virtually always species-typical (unlike non-functional aspects of the system)—so the functional aspects of the architecture will tend to be genetically universal (Tooby and Cosmides, 1990b).¹ This means that any complex device that cognitive neuroscientists find should be universal, at least at the genetic level.

Because design features are embodied in organisms, they can, generally speaking, propagate themselves in only two ways: by solving problems that will increase the probability that the organism they are situated in will produce offspring, or by solving problems that will increase the probability that the organism's kin will produce offspring (Williams and Williams, 1957; Hamilton, 1964; however, see Cosmides and Tooby, 1981, for intragenomic methods). An individual's relatives, by virtue of having descended from a common ancestor, have an increased likelihood of having the same design feature as compared to other conspecifics,

so their increased reproduction will tend to increase the frequency of the design feature. Accordingly, design features that promote both direct reproduction and kin reproduction, and that make efficient trade-offs between the two, will replace those that do not. To put this in standard biological terminology, design features are selected for to the extent that they promote their inclusive fitness (Hamilton, 1964).

Although largely unknown to cognitive neuroscientists, the promotion of inclusive fitness is the ultimate functional product of all evolved cognitive devices. That is, design changes were incorporated into the neural architecture of a species to the extent that they stably promoted inclusive fitness in the past, and were discarded to the extent that they did not. The human brain, to the extent that it is organized to do anything functional at all, will be organized to construct the information, make the decisions, and generate the behavior that would have tended to promote inclusive fitness in the ancestral environments and behavioral contexts of Pleistocene hunter-gatherers. (The preagricultural world of hunter-gatherers is the appropriate ancestral context, because natural selection operates far too slowly to have built complex information-processing adaptations to the post-hunter-gatherer world of the last few thousand years.) Although there are an infinite number of other standards of functionality one could sensibly have for various purposes (e.g., getting an individual to read English or to avoid shouting inappropriately), this biological standard is the only standard of functionality that is relevant to analyzing why the human brain architecture is organized in one fashion rather than another.

There is, however, a second family of evolutionary processes, in addition to selection, by which mutations can become incorporated into species-typical designs: chance processes. For example, the sheer impact of many random accidents may cumulatively propel a useless mutation upward in frequency until it crowds out all alternative design features from the population. Clearly, the presence of such a trait in the architecture is not explained by the (nonexistent) functional consequences it had over many generations on the design's reproduction; as a result, chance-injected traits will not tend to be coordinated with the rest of the organism's architecture in a functional way.

Although such chance events play a restricted role in evolution and explain the existence and distribution of

many simple and trivial properties, organisms are not primarily chance agglomerations of stray properties. In the first place, reproduction is a highly improbable outcome in the absence of functional machinery designed to bring it about, and only designs that retain all of the necessary functional organization avoid being selected out. Secondly, to the extent that a mutation has a significant systematic impact on the functional organization leading to reproduction, selection will act on it. For this reason, the significant and consequential aspects of organismic architectures are organized primarily by natural selection. Reciprocally, those modifications that are so minor that their consequences are negligible on reproduction are invisible to selection and are therefore not organized by it. Thus, chance properties do indeed drift through the standard designs of species in a random way, but they are unable to account for the complex, organized design in organisms and are, correspondingly, usually peripheralized into those aspects of the system that do not make a significant impact on its functional operation (Tooby and Cosmides, 1990a,b, 1992). Random walks do not systematically build intricate and improbably functional arrangements such as the visual system, the language faculty, face-recognition competences, emotion-recognition modules, food-aversion circuits, cheater-detection devices, or motor control.

Brains are composed primarily of adaptive problem-solving devices

The foregoing leads to the most important point for cognitive neuroscientists to abstract from modern evolutionary biology: Although not everything in the design of organisms is the product of selection, all complex functional organization is. This is because the only known cause of and explanation for complex functional design in organic systems is natural selection: It is the single natural hill-climbing process that pushes designs through state space toward increasingly well organized—and otherwise improbable—functional arrangements (Williams, 1966, 1985; Dawkins, 1986; Tooby and Cosmides, 1990a,b, 1992; Pinker and Bloom, 1992). Specifically, this means that all of the functional organization present in the human brain was built by natural selection during our evolutionary history. Indeed, selection can account only for functionality of a very narrow kind: approximately, those design features organized to promote the reproduction

of an individual and his or her relatives (Williams, 1966; Dawkins, 1986). Fortunately for the modern theory of evolution, the only naturally occurring, complex functionality that has ever been documented in undomesticated plants, animals, or other organisms is functionality of just this kind, along with its derivatives and by-products. This has several important implications for cognitive neuroscientists.

First, whenever one finds functional organization built into our cognitive or neural architecture, one is looking at adaptations—devices that acquired their distinctive organization from natural selection acting on our hunter-gatherer or more distant primate ancestors. Reciprocally, when one is searching for intelligible functional organization underlying a set of cognitive or neural phenomena, one is far more likely to discover it by using an adaptationist framework for organizing observations, because adaptive organization is the only kind of *functional* organization that is there to be found.

Second, because the reliably developing mechanisms (i.e., modules, circuits, functionally isolable units, mental organs, or computational devices) that cognitive neuroscientists study are evolved adaptations, all of the biological principles that apply to adaptations apply to cognitive devices. Thus cognitive neuroscience and evolutionary biology are connected in the most direct possible way. This conclusion should be a welcome one, because it is the logical doorway through which a very extensive body of new expertise and principles can be applied to cognitive neuroscience, stringently constraining the range of valid hypotheses about the functions and structures of cognitive mechanisms. Because cognitive neuroscientists are usually studying adaptations and their effects, they can supplement their present research methods with carefully derived adaptationist analytic tools (e.g., Shepard, 1981, 1984, 1987a, 1987b; Marr, 1982; Freyd, 1987; Leslie, 1987, 1988; Sherry and Schacter, 1987; Cosmides, 1989; Gallistel, 1990; Ramachadran, 1990; Pinker, 1991; Cosmides and Tooby, 1992; Jackendoff, 1992; Pinker and Bloom, 1992; Baron-Cohen, 1994).

Third, our cognitive architectures are designed to incorporate only the precise, narrow, and strange kinds of functional organization that natural selection spontaneously builds, rather than any other kind of functional organization. What this means is that the problems our cognitive devices are designed to solve do not reflect the problems our modern life experiences lead

us to see as normal, such as reading, driving cars, working for large organizations, reading insurance forms, learning the oboe, or playing Go. Instead, they are the ancient and seemingly esoteric problems that our hunter-gatherer ancestors encountered generation after generation over hominid evolution. These include such problems as foraging, kin recognition, "mind reading" (i.e., inferring the motives, intentions, and knowledge of others based on their situation, history, and behavior), engaging in social exchange, avoiding incest, choosing mates, interpreting threats, recognizing emotions, caring for children, regulating immune function, and so on, as well as the already well-known problems involved in perception, language acquisition, and motor control. For biological reasons discussed elsewhere, such devices should be far more numerous and far more content-specialized than is usually appreciated even by cognitive scientists sympathetic to modular approaches (for a review of the issues, see Cosmides and Tooby, 1987, 1994, this volume; Tooby and Cosmides, 1992). From an evolutionary perspective, the human cognitive architecture is far more likely to resemble a confederation of hundreds or thousands of functionally dedicated computers, designed to solve problems endemic to the Pleistocene, than it is to resemble a single general-purpose computer equipped with a small number of general-purpose procedures such as association formation, categorization, or production-rule formation (Cosmides and Tooby, 1987; Tooby and Cosmides, 1992; see also Gallistel, 1990, this volume). Although our architectures may be capable of other kinds of functionality or activities (e.g., weaving, playing pianos), these are incidental by-products of selection for our Pleistocene competences—just as a machine built to be a hair dryer can, incidentally, dehydrate fruit or electrocute. But it will be difficult to make sense of our cognitive mechanisms if one attempts to interpret them as devices designed to perform functions that were not selectively important for our hunter-gatherer ancestors, or if one fails to consider the functions they were designed to accomplish (for the importance of functional analysis, see chapter 79).

Fourth, evolutionary biology gives the concept of function a specific and rigorous content that is otherwise lacking, and imposes strict rules on its use. Although many cognitive scientists are unaware of it, every time the function of a computational device is discussed, this automatically invokes a biological concept with a very specific and narrow technical meaning

(Williams, 1966; Dawkins, 1986; Tooby and Cosmides, 1990a, 1992). As discussed, it is only the narrow biological meaning that is relevant in explaining why a system is structured as it is—that is, what specific consequences of a design feature caused it to be propagated and made species-standard within ancestrally structured environments. So, not only are the problems that our devices were designed to solve esoteric to modern sensibilities, but the standards that define what counts as functional solutions to these problems are evolutionary standards, and hence odd and nonintuitive as well. Cognitive neuroscientists need to recognize that in explaining or exploring the reliably developing organization of a cognitive device, the "function" of a design refers solely to how it systematically caused its own propagation in ancestral environments. It does not validly refer to any of the various intuitive or folk definitions of function such as "contributing to the attainment of the individual's goals," "contributing to one's well-being," "contributing to society," or even "making a valid inference." These other kinds of usefulness may or may not exist as side effects of a given evolved design, but they can play no role in explaining how such designs came into existence or why they have the organization that they do. The only kind of functional organization that has been built into our cognitive architectures—and hence that researchers should spend their time looking for—is the kind that matches this peculiar biological standard of functionality: enhancing propagation in ancestral environments. The fact that sexual jealousy, for example, has no truth-value, and may not contribute to any individual's well-being or to any positive social good, is irrelevant to why the cognitive mechanisms that reliably produce it under certain limited conditions became part of our species-typical computational architecture (Daly, Wilson, and Weghorst, 1982). In short, the technical criteria that define what solutions our cognitive devices are designed to produce—that is, what counts as functional design and successful processing—are evolutionary in nature, and usually cannot be supplied by simply consulting common sense.

Fifth, the standard of parsimony imported from physics, from traditional philosophy of science, or from habits of economical programming is inappropriate and misleading in biology and, hence, in neuroscience and cognitive science, which study biological systems. The evolutionary process never starts with a clean workboard, has no foresight, and incorporates new fea-

tures solely on the basis of whether they lead to systematically enhanced propagation. Enhanced functionality in a complexly structured series of environments is the only criterion for organizing designs. Indeed, when one examines the brain, one sees an amazingly heterogeneous physical structure. A correct theory of evolved cognitive functions should be no less complex and heterogeneous than the evolved physical structure itself, and should map onto the heterogeneous set of recurring adaptive tasks faced by hominid foragers over evolutionary time. Indeed, analyses of the adaptive problems that humans and other animals must regularly have solved over evolutionary time to remain in the world suggests that the mind contains a far greater number of functional specializations than has traditionally been supposed (for discussion, see Cosmides and Tooby, 1987, 1993; Symons, 1987; Tooby and Cosmides, 1992).

Sixth, understanding the neural organization of the brain depends on understanding the functional organization of its cognitive devices. The brain originally came into existence, and accumulated its particular set of design features, only because these features functionally contributed to the organism's propagation. This contribution—that is, the evolutionary function of the brain—is obviously the adaptive regulation of behavior and physiology on the basis of information derived from the body and from the environment. The brain performs no significant mechanical, metabolic, or chemical service for the organism; its function is purely informational, computational, and regulatory in nature. Because the function of the brain is informational in nature, its precise functional organization can be described accurately only in a language that is capable of expressing its informational functions—that is, in cognitive terms, rather than in cellular, anatomical, or chemical terms. Cognitive investigations are not some soft, optional activity that goes on only until the “real” neural analysis can be performed. Instead, the mapping of the computational adaptations of the brain is an unavoidable and indispensable step in the neuroscience research enterprise; it must proceed in tandem with neural investigations, and indeed will provide one of the primary frameworks necessary for organizing the body of neuroscience results.

The reasons why are straightforward. Natural selection operating on hominids in complexly structured ancestral environments posed adaptive information-

processing problems, such as effective foraging, object recognition, motivational allocation, contagion avoidance, and so on. These recurrent problems selected for specialized devices that could solve these information-processing problems—that is, cognitive or computational devices. That in turn selected for those precise physical arrangements of cells (and modifications of the internal organization of cells) that could embody the particular computational relationships that reliably solved those adaptive problems. Natural selection retained neural structures on the basis of their ability to create adaptively organized relationships between information and behavior (e.g., the sight of a predator activates inference procedures that cause the organism to hide or flee) or between information and physiology (e.g., the sight of a predator increases the organism's heart rate in preparation for flight). Thus, it is the information-processing structure of the human psychological architecture that has been functionally organized by natural selection, and the neural structures and processes have been organized insofar as they physically realize this cognitive organization. Brains exist and have the structure that they do because of the computational requirements imposed by selection on our ancestors. The adaptive structure of our computational devices provides a skeleton around which a modern understanding of our neural architecture will be constructed.

That is why cognitive neuroscience is pivotal to the progress of the brain sciences. There are an astronomical number of physical interactions and relationships in the brain, and blind empiricism rapidly drowns itself among the deluge of manic and enigmatic measurements. The same is true at the cognitive level: The blind empiricist will drown in the sea of irrelevant phenomena that our computational devices can generate—everything from writing theology or dancing the limbo to calling for the restoration of the Plantagenets to the throne of England. Fortunately, however, evolutionary biology, behavioral ecology, and hunter-gatherer studies can be used to identify and supply descriptions of the recurrent adaptive problems humans faced during their evolution. Supplemented with this knowledge, cognitive research techniques can abstract out of the welter of human cognitive performance a series of maps of the functional information-processing relationships that constitute our computational devices and that evolved to solve this particular

set of problems: our cognitive architecture. With these computational maps in hand, we can navigate the ocean of physical relationships in the brain, abstracting out that exact and minute subset that implement those information-processing relationships. It is only those relationships that explain the existence and functional organization of the system. The immense number of other physical relationships in the brain are incidental by-products of those narrow aspects that implement the functional computational architecture. Consequently, an adaptationist inventory and functional mapping of our cognitive devices can provide the essential theoretical guidance for neuroscientists, allowing them to home in on these narrow but meaningful aspects of neural organization, and to distinguish them from the sea of irrelevant neural phenomena.

Brain architectures consist of adaptations, by-products, and random effects

To understand the human computational or neural architecture (or that of any living species) is a problem in reverse engineering: We have working exemplars of the design in front of us, but we must organize our observations of these exemplars into a systematic functional and causal description of the design. One can describe and decompose brains into properties according to any of an infinite set of alternative systems, and hence there are an indefinitely large number of cognitive and neural phenomena that could potentially be defined and measured. However, describing and investigating the architecture in terms of its adaptations is a useful place to begin, because (1) the adaptations are the cause of the system's organization (the reason for the system's existence); (2) organisms, properly described, consist largely of collections of adaptations (evolved problem-solvers); (3) all of the complex, functionally organized subsystems in the architecture are adaptations; (4) an adaptationist frame of reference allows cognitive neuroscientists to apply to their research problems the formidable array of knowledge that evolutionary biologists have accumulated about adaptations; and (5) such a frame of reference permits the construction of economical and principled models of the important features of the system, models in which the wealth of varied phenomena fall into intelligible, functional, and predictable patterns. As Ernst Mayr put it, summarizing the historical record, "the

adaptationist question, 'What is the function of a given structure or organ?' has been for centuries the basis for every advance in physiology" (1983, 32). It should prove no less productive for cognitive neuroscientists.

Indeed, all of the inherited design features of organisms can be partitioned into three adaptively defined categories: adaptations (often, though not always, complex), the by-products or concomitants of adaptations, and random effects. Chance and selection, the two components of the evolutionary process, explain different types of design properties in organisms, and all aspects of design must be attributed to one of these two forces. The conspicuously distinctive cumulative impacts of chance and selection allow the development of rigorous standards of evidence for recognizing and establishing the existence of adaptations and distinguishing them from the nonadaptive aspects of organisms caused by the nonselectionist mechanisms of evolutionary change (Williams, 1966, 1985; Tooby and Cosmides, 1990b; Thornhill, 1991; Pinker and Bloom, 1992; Symons, 1992).

Adaptations are systems of properties (called mechanisms) crafted by natural selection to solve the specific problems posed by the regularities of the physical, chemical, developmental, ecological, demographic, social, and informational environments encountered by ancestral populations during the course of a species' or population's evolution (figure 78.1; for other discussions of adaptation, see Williams, 1966, 1985; Dawkins, 1986; Symons, 1989, 1992; Thornhill, 1990; Tooby and Cosmides, 1990a, 1992; Pinker and Bloom, 1992). Adaptations are recognizable by "evidence of special design" (Williams, 1966); that is, certain features of the evolved species-typical design of an organism are recognized "as components of some special problem-solving machinery" (Williams, 1985, 1). Moreover, they are so well organized and represent such good engineering solutions to adaptive problems that a chance coordination between problem and solution is effectively ruled out as a counterhypothesis. Standards for recognizing special design include whether the problem solved by the structure is an evolutionarily long-standing adaptive problem, and such factors as economy, efficiency, complexity, precision, specialization, and reliability, which, like a key fitting a lock, render the design too good a solution to a defined adaptive problem to be coincidence (Williams, 1966). Like most other methods of empirical hypothesis-

The formal properties of an adaptation

An adaptation is:

- (1) A cross-generationally recurring set of characteristics of the phenotype,
- (2) that is reliably manufactured over the developmental life history of the organism
- (3) according to instructions contained in its genetic specification,
- (4) in interaction with stable and recurring features of the environment (i.e., it reliably develops normally when exposed to normal ontogenetic environments),
- (5) whose genetic basis became established and organized in the species (or population) over evolutionary time, because
- (6) the set of characteristics systematically interacted with stable and recurring features of the ancestral environment (the "adaptive problem")
- (7) in a way that systematically promoted the propagation of the genetic basis of the set of characteristics better than the alternative designs existing in the population during the period of selection. This promotion virtually always takes place through enhancing the reproduction of the individual bearing the set of characteristics, or the reproduction of the relatives of that individual.

FIGURE 78.1 The most fundamental analytic tool for organizing observations about a species' functional architecture is the definition of an adaptation. In order to function, adaptations evolve such that their causal properties rely on and exploit these stable and enduring statistical and structural regularities in the world, and in other parts of the organism. Things worth noticing include the fact that an adaptation (such as teeth or breasts) can develop at any time during the life cycle, and need not be present at birth; that an adaptation can express itself differently in different environments (e.g., speaking English, speaking Tagalog); that an adaptation is not just any individually beneficial trait, but one built over evolutionary time and expressed in many individuals; that it may not be producing functional outcomes now (e.g., agoraphobia), but was needed to function well in ancestral environments; and finally, that an adaptation is the product of gene-environment interaction, like every other aspect of the phenotype; however, unlike many other phenotypic properties, it is the result of the interaction of the species-standard set of genes with those aspects of the environment that were present and relevant during the evolution of the species. (For a more extensive definition of the concept of adaptation, see Tooby and Cosmides, 1990a, 1992.)

testing, the demonstration that something is an adaptation is always, at core, a probability assessment concerning how likely a set of events is to have arisen by chance. Such assessments are made by investigating

whether there is a highly nonrandom coordination between the recurring properties of the phenotype and the structured properties of the adaptive problem, in a way that meshed to promote fitness (genetic propagation) in ancestral environments (Tooby and Cosmides, 1990a, 1992). For example, the lens, pupil, iris, retina, visual cortex, and so on are too well coordinated both with each other and with features of the world, such as the properties of light, optics, geometry, and the reflectant properties of surfaces, to have co-occurred by chance. In short, like the functional aspects of any other engineered system (e.g., the electron gun in a television), they are recognizable by their organized and functional relationships to the rest of the design and to the structure of the world.

In contrast, concomitants or by-products of adaptations are those properties of the phenotype that do not contribute to functional design *per se*, but that happen to be coupled to properties that are. They were, consequently, dragged along into the species-typical architecture because of selection for the functional design features to which they are linked. For example, bones are adaptations but the fact that they are white is an incidental by-product. Bones were selected to include calcium because it conferred hardness and rigidity to the structure (and was dietarily available), and it simply happens that alkaline earth metals appear white in many compounds, including the insoluble calcium salts that are a constituent of bone. From the viewpoint of functional design, by-products are the result of chance, in the sense that the process that led to their incorporation into the design was blind to their consequences (assuming the consequences were not negative). They are distinguishable from adaptations by the absence of complexly arranged functional consequences (e.g., the whiteness of bone does nothing for the vertebrate). In general, by-products will be far less informative as a focus of study than adaptations, because they are consequences and not causes of the organization of the system. (Hence they are functionally arbitrary and unregulated, and may, for example, vary capriciously between individuals.) Unfortunately, unless researchers actively seek to study organisms in terms of their adaptations, they will usually end up measuring and investigating arbitrary and random admixtures of functional and functionless aspects of organisms—and this hampers the discovery of the underlying organization of a biological system. We do not yet know, for example, which exact aspects of the neuron are relevant

to its function and which are by-products, thus many computational neuroscientists are stuck using a model of the neuron that is inaccurate and outdated by decades.

Finally, of course, entropic effects of many types are always acting to introduce functional disorder into the design of organisms. Traits introduced by accident or by evolutionary random walks are recognizable by the lack of coordination they produce within the architecture, or between the architecture and the environment, as well as by the fact that they frequently cause uncalibrated variation between individuals. Examples of such entropic processes include genetic mutation, recent change in ancestrally stable environmental features, and developmentally anomalous circumstances.

How well engineered are adaptations?

The design of our cognitive and neural mechanisms should reflect the structure of the adaptive problems our ancestors faced only to the extent that natural selection is an effective process. Is it one? How well or poorly engineered are adaptations? Some researchers have argued that evolution produces mostly inept designs, because selection does not produce perfect optimality (Gould and Lewontin, 1979). In fact, evolutionary biologists since Darwin have been well aware that selection does not produce perfect designs (Darwin, 1859; Williams, 1966; Dawkins, 1976, 1982, 1986; for a recent convert from the position that organisms are optimally designed to the more traditional adaptationist position, see Lewontin, 1967, 1978; see Dawkins, 1982, for an extensive discussion of the many processes that prevent selection from reaching perfect optimality). Still, 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, natural selection tends to cause the accumulation of superlatively well engineered functional designs.

Empirical confirmation can be gained by comparing how well evolved devices and human-engineered devices perform on evolutionarily recurrent adaptive problems (as opposed to arbitrary, artificial modern tasks, such as chess). For example, the claim that language competence is a simple and poorly engineered adaptation cannot be taken seriously, given the total amount of time, engineering, and genius that has gone

into the still unsuccessful effort to produce artificial systems that can remotely approach—let alone equal—human speech perception, comprehension, acquisition, and production (Pinker and Bloom, 1992).

Even more strikingly, the visual system is composed of collections of cognitive adaptations that are well-engineered products of the evolutionary process, and while they may not be “perfect” or “optimal”—however these somewhat vague concepts may be interpreted—they are far better at vision than any human-engineered system yet developed. Wherever the standard of biological functionality can be clearly defined—from semantic induction to capturing solar energy, to object recognition, to color constancy, to echolocation, to relevant problem-solving generalization, to chemical recognition (olfaction), to mimicry, to scene analysis, to chemical synthesis—evolved adaptations are at least as good as, and usually strikingly better than, human-engineered systems, in those rare situations when humans can build systems that can accomplish these tasks at all. It seems reasonable to insist that before a system is criticized as being poorly designed, the critic ought to be able to construct a better alternative—a requirement, it need hardly be pointed out, that has never been met by anyone who has argued that adaptations are poorly designed. Thus, while adaptations are certainly suboptimal in some ultimate sense, it is an empirically demonstrable fact that the short-run constraints on selective optimization do not prevent the emergence of superlatively organized computational adaptations in brains. Indeed, aside from the exotic nature of the problems the brain is designed to solve, it is exactly this sheer functional intricacy that makes our architecture so difficult to reverse-engineer and to understand.

Cognitive adaptations reflect the structure of the adaptive problem and the ancestral world

Looking at known adaptations not only can tell us about the overall engineering quality of evolved computational devices, but also can inform us about the general character of cognitive adaptations and adaptive problems. For example, hundreds of vision researchers, working over decades, have been mapping the exquisitely structured set of information-processing adaptations involved in vision. As Marr (1982) put it, the evolutionary function of vision is scene analysis—the reconstruction of models of real-world conditions

from a two-dimensional visual array. As more and more functional subcomponents are explored, and as artificial intelligence researchers try to duplicate vision in computational systems attached to electronic cameras, four things have become clear (Marr, 1982; Poggio, Torre, and Koch, 1985). The first is that the magnitude of the computational problem posed by scene analysis is immensely greater than anyone had suspected before trying to duplicate it. Even something so seemingly simple as perceiving the same object as having the same color at different times of day turns out to require intensely specialized and complex computational machinery, because the spectral distribution of light reflected by the object changes widely with changes in natural illumination (see, e.g., Shepard, 1992). The second conclusion is that, as discussed, our visual system is a very well engineered set of cognitive adaptations, capable of recovering far more sophisticated information from two-dimensional light arrays than the best of the artificially engineered systems developed so far. The third is that successful vision requires specialized neural circuits or computational machinery designed particularly for solving the adaptive problem of scene analysis (Marr, 1982). And the fourth is that scene analysis is an unsolvable computational problem unless the design features of this specialized machinery “assume” that objects and events in the world manifest many specific regularities—that is, unless the cognitive procedures embody a complementary structure that reflects the problem-relevant parts of the world (Shepard, 1981, 1984, 1987a; Marr, 1982, Poggio, Torre, and Koch, 1985).

These four lessons—complexity of the adaptive information-processing problem, well-engineered problem-solving machinery as the evolved solution, specialization of the problem-solving machinery to fit the particular nature of the problem, and the requirement that the machinery embody substantial and contentful innate knowledge about the adaptive problem—recur throughout the study of the evolved computational subcomponents of our cognitive architecture (Cosmides and Tooby, 1987, 1992; Tooby and Cosmides, 1990a, 1990b; on language, see Chomsky, 1975, and Pinker, 1989, 1991; on vision, see Marr, 1982, and Poggio, Torre, and Koch, 1985). Well-studied adaptations overwhelmingly achieve their functional outcomes because they display an intricately engineered coordination between their specialized design features and the detailed structure of the task and the task

environment. Like a code that has been torn in two and given to separate couriers, the two halves (the structure of the mechanism and the structure of the task) must be put together to be understood. In order to function, adaptations evolve such that their causal properties rely on and exploit these stable and enduring statistical and structural regularities in the world. Thus, to map the structures of our cognitive devices, we need to understand the structures of the problems they solve, and the problem-relevant parts of the hunter-gatherer world. If studying face recognition mechanisms, one must study the recurrent structure of faces. If studying social cognition, one must study the recurrent structure of hunter-gatherer social life. For vision the problems are not so very different for a modern scientist and a Pleistocene hunter-gatherer, so the folk notions of function that perception researchers use are not a serious problem. But the more one strays away from low-level perception, the more one needs to know about human behavioral ecology and the structure of the ancestral world.

Experimenting with ancestrally valid tasks and stimuli

Although bringing cognitive neuroscience current with modern evolutionary biology offers many new research tools (see, e.g., Preuss, this volume), we have out of necessity limited discussion to only one: an evolutionary functionalist research strategy (see chapter 79 and Tooby and Cosmides, 1992, for a description; for examples, see Barkow, Cosmides, and Tooby, 1992, and chapters 80, 83, and 84). The adoption of such an approach will modify research practice in many ways. Perhaps most significantly, researchers will no longer have to operate purely by intuition or guesswork in deciding what kinds of tasks and stimuli to expose subjects to. Using knowledge from evolutionary biology, behavioral ecology, animal behavior, and hunter-gatherer studies, they can construct ancestrally or adaptively valid stimuli and tasks. These are stimuli that would have had adaptive significance in ancestral environments, and tasks that resemble (at least in some ways) the adaptive problems our ancestors would have been selected for the ability to solve.

The currently widespread practice of using only arbitrary stimuli of no adaptive significance (e.g., lists of random words, colored geometric shapes), or abstract experimental tasks of unknown relevance to Pleisto-

