

Environmental Complexity and the Evolution of Cognition

Peter Godfrey-Smith
Department of Philosophy
Stanford University
Stanford CA, 94305-2155
pgs@csl.stanford.edu

**In R. Sternberg and J. Kaufman (eds.) The Evolution of Intelligence.
Lawrence Erlbaum Associates, 2001.**

- 1. Starting Simple**
- 2. The Environmental Complexity Thesis**
- 3. On Complexity**
- 4. The ECT as a Component in Many Evolutionary Scenarios**

1. Starting Simple

One problem faced in discussions of the evolution of intelligence is the need to get a precise fix on what is to be explained. Terms like "intelligence," "cognition" and "mind" do not have simple and agreed-upon meanings, and the differences between conceptions of intelligence have consequences for evolutionary explanation. I hope the papers in this volume will enable us to make progress on this problem. The present contribution is mostly focused on these basic and foundational issues, although the last section of the paper will look at some specific models and programs of empirical work.

Some people have a very demanding picture of what is required for intelligence, thinking that it always involves such sophisticated skills as planning, language-use, and perhaps even some sort of consciousness. To these people, intelligence is to be contrasted with instinct. Perhaps in this rich sense of the term, intelligence is even to be contrasted with the simpler types of learning, such as learning through reinforcement (operant conditioning). From this first point of view, the problem of explaining the evolution of intelligence is explaining why instinct and other simple behavioral capacities were not enough; why evolutionary processes took a few organisms so far beyond these basic behavioral skills.

Another approach uses terms like "intelligence" and "cognition" in much less demanding senses. On this second approach, intelligence is not restricted to a few exceptional cases in the biological world -- humans and perhaps some primates. Rather, intelligence exists to some degree in a huge range of living systems. Humans have a lot more of it than cockroaches do, but cockroaches do have some of it.

According to this second view, all the mechanisms that enable organisms to coordinate their behavior with conditions in the world involve some degree of intelligence. So there is no opposition between intelligence and what is often referred to as "instinct." An instinctive behavior can involve perception, and a good deal of processing and feedback to ensure the right match between behavior and circumstances. According to this second view, that is a low-level variety of intelligence.

Maybe it is difficult to motivate the second, less demanding approach with the term "intelligence" which has such definite common sense usages. Perhaps it is better to present this view with the aid of a less everyday term such as "cognition." So in this paper I will mostly discuss the "evolution of cognition" rather than the evolution of intelligence. The term "mind" is another broad one, but it has its own capacity to mislead, as for many people it is closely linked with consciousness and a sense of self.

Given this, we can describe the second approach by saying that a great range of living things have some cognitive capacities. In many cases these capacities are extremely limited. The capacities we habitually refer to as "intelligent" in humans, such as the capacities for planning and conscious reflection, comprise one type of cognition. But when a fish negotiates its way around a reef, or a rat finds its way back to a food source, the internal processes responsible for these behaviors are varieties of cognition as well.

Does it matter which of these two general approaches to the evolution of cognition we take? Certainly it does not much matter what we decide to refer to with the term "intelligent" or "cognitive." Either way, the problem remains of explaining how nervous systems evolved at all, and the problem remains of explaining how humans became so much smarter, in many respects, than other organisms. But I think it often does make a difference how we view and describe the continuities between human and non-human psychological capacities. A great deal depends on how much significance we place on the distinction between animals that use language and those that don't, for example. Views about non-human cognition often have ethical consequences, and consequences for a range of issues in the philosophy of mind.

In any case, my own approach is very much along the lines of the second option described above. I approach the problem with a broad and very undemanding concept of "cognition." My aim is to set out with a broad concept of cognition and ask: can we formulate a generalization about why these sorts of capacities will tend to evolve? Because I use "cognition" to refer to such a broad class of capacities, cognition is not a single evolutionary "discovery," restricted to a single lineage of organisms. Cognition of various

kinds has been discovered and rediscovered by evolution many times, just as eyes and wings have been discovered independently several times. And just as we find with eyes, cognitive machinery is very diverse. There are lots of ways to process information and control behavior; a central nervous system is one way, but not the only way. An important feature of this view, which will be discussed in more detail below, is that cognition "shades off" into other kinds of biological capacities and processes. In some cases it is hard to distinguish cognition from other control systems in the body, and hard to distinguish behavior from such things as growth, development and the regulation of metabolism.

Cognition is diverse, but it might be possible to find a common type of evolutionary story that applies in all or most of these diverse cases. With such a general framework in hand, we can then ask more specific questions about why certain types of cognition evolved. When is learning favored over less flexible strategies of dealing with the world? And if one is to learn, when is learning through individual trial-and-error better than learning by imitating a parent? What brings about the transition to a planning intelligence? When is what primatologists refer to as "theory of mind" (taking other individual organisms to have a mind) favored? And what on earth lay behind the explosion of mental capacities found in the evolution of humans?

It might turn out that all these explanations are so diverse that it is pointless to try to link them under a general principle. I do not deny that possibility. But my own approach here is to outline and cautiously defend one possible generalization about the adaptive value of cognition. The generalization is intended to be a fairly obvious one, something that has been expressed in partial or imperfect ways dozens of times before, dating back to the 19th century. My aim here is more to make a vague existing idea into a precise one, than to present a novel idea. I do think it could be of considerable help to discussions about the evolution of cognition if this underlying idea, and its possible rivals, were made explicit and precise in people's minds.

2. The Environmental Complexity Thesis

Here is my proposal for a general "first principle" about the evolution of cognition:

Environmental Complexity Thesis (ECT):

The function of cognition is to enable the agent to deal with environmental complexity.

Each of the key terms in the ECT requires a good deal of clarification. The term "function" is understood here in a strong sense. To ascribe a function in this sense is to offer an evolutionary hypothesis. The function of a trait or structure is the effect or capacity it has which has been responsible for its

success under a regime of natural selection. When we say that the function (in this strong sense) of the thorns on a plant is to deter herbivores from eating the plant, we are not just saying that the thorns help the plant by deterring herbivores. We are saying that thorns were selected for in evolutionary processes because they tended to have the effect of deterring herbivores. The ECT makes a similar claim about cognition.

The ECT is a broader and more abstract claim than the one about the thorns, but very abstract functional claims can certainly be made. Eyes have evolved many times, and they make use of various different types of mechanisms. But a general claim can be made about their evolutionary function: the function of eyes is to respond in discriminative ways to light, and hence to enable the organism to make use of information about the world that is carried in light (Gibson 1966, Dretske 1981). We can also formulate an even more general thesis about the function of perceptual mechanisms: they all respond, with some degree of sensitivity or discrimination, to some physical or chemical variables that impinge causally on the organism, in such a way as to enable the organism to make use of information about the world that is carried by these variables.

Some might say that when we formulate a broad generalization like this, we are saying something so obvious or empty as to make it not worth the effort of saying. I disagree. First, it is important to be able to embed specific functional claims made about biological structures within a more general picture of what the organism as a whole is doing. Second, it is actually quite hard to get these generalizations right, and lots of puzzling questions can get raised along the way. For example, in the claims about perception I made above, I used the concept of information. Although lots of people, including scientists, talk about information in a casual and unreflective way, it is a subtle and difficult concept. The philosopher Fred Dretske (1981) developed a detailed theory of where information is found in the physical world, and of how everyday talk about information is related to the technical discussions of "information theory" in engineering. Information, for Dretske, is found where there is contingency and correlation. Any variable in the world which has a range of possible states is a source of information. When the state of a source of information is correlated with the state of another variable, as a consequence of physical laws, the second variable carries information about the source. For Dretske, information is a resource that organisms use to make their way through the world; cognitive systems are information consuming, or information exploiting, systems.

Without careful, explicit discussions like Dretske's, it would be unclear whether or not it is really justifiable to use the concept of information when making generalizations about the function of eyes and other perceptual systems. Whenever I use the term "information" in this paper, I have Dretske's sense in mind. Below I will discuss the concept of "environmental complexity" in some detail; this is another concept which it can be easy to throw around without having a clear idea of what is being said.

So far I have said that the ECT is an attempt to give a general functional explanation of cognition. Functional explanations have received a lot of attention in the philosophical literature, especially over the past 20 years or so. (For a collection of classic and recent articles, see Allen, Bekoff and Lauder 1998) In this paper I will not discuss the many issues that have arisen in these debates, as the present topic is not functional explanation in general but evolutionary hypotheses about cognition. All that is important for present purposes is the idea that functional explanations are attempts to describe, in a shorthand way, the processes of mutation and natural selection that were responsible for the origination and maintenance of biological structures. Functional explanations are attempts to isolate the effects or dispositions of a structure which were responsible for the natural selection of that structure. So functional claims are "teleological" only in this specific Darwinian sense.

In the first section I said I would be using a broad and undemanding sense of the term "cognition." But how broad? What exactly is cognition? What is the set of organisms in which it is found?

I understand cognition as a collection of capacities which, in combination, allow organisms to achieve certain kinds of coordination between their actions and the world. This collection typically includes the capacities for perception, internal representation of the world, memory, learning, decision-making and the production of behavior. This set of capacities, according to the ECT, has the function of making possible patterns of behavior which enable organisms to effectively deal with complex patterns and conditions in their environments.

So although the ECT is expressed as a connection between cognition and environmental complexity, it really embeds two separate claims. One is the claim that the immediate role of cognition is to control behavior. The other is the claim that the point of this control of behavior is to deal with environmental complexity.

Think of cognition here not as a single type of process, but as a biological "tool-kit" used to direct behavior. There is no single list of tools found across all the organisms with cognitive capacities; different organisms have different collections of tools, according to their circumstances and history. So when I listed "perception, internal representation of the world, memory, learning (etc.)" above, it should not be thought that this has to be a description of a set of recognizable and distinct "modules" common to all cognitive systems. Rather, I am referring to a set of capacities which are realized in very different ways in different organisms, a set of capacities which shade into each other and shade off into other, non-cognitive parts of the biological machinery. Also, the capacity for "internal representation of the world" is one about which there is enormous disagreement within philosophy and cognitive science (Stich and Warfield 1994). Here again, some people understand this ability in a very demanding way, perhaps requiring language, while others view a very wide range of mechanisms for registering external events as representational in some sense. Here I assume a very loose

and undemanding view of representation, but the issues surrounding that concept are too complex to go into here.

In the list of basic cognitive capacities given above, some are more fundamental than others. Perception is very fundamental; learning is somewhat less so. It is true that all macroscopic animals are thought to be able to "learn," in at least a minimal sense. Bees have been shown to have quite impressive ways of learning the location of food sources, and fruit flies have been conditioned to exhibit avoidance behaviors in some of the same sorts of ways found in more celebrated learners such as rats and pigeons. The neural basis for the most minimal kinds of learning is often studied using sea slugs. Still, learning is not in principle essential to cognition. A behavioral pattern which is completely insensitive to any refinement through learning, but which does involve the coordination of actions with perceived environmental conditions, does display a minimal type of cognition.

But if learning is not essential, where does cognition stop? Do plants have it? What about bacteria? By any normal standard, plants and bacteria do not have minds and do not exhibit cognition. But my suggestion is that cognition shades off into other kinds of biological processes. There is not much point in trying to draw an absolute line. Plants and bacteria do exhibit some capacities for flexible response to environmental conditions, using environmental cues to control development and metabolism. These are low-level cases of the same types of capacities that, in more elaborate cases, do constitute cognition.

Many bacteria can adjust in adaptive ways to changing circumstances around them. Dretske (1986) discusses aquatic bacteria which use little internal magnets to track the distinction between north and south, enabling them to move towards water with their required chemical properties. Bacteria also make use of external cues to adjust their metabolic activity. A famous case is the lac operon system in e. coli bacteria. These bacteria can respond to a change in local food type through processes in which the availability of a nutrient affects the regulation of genes which code for enzymes able to digest that nutrient.

Plants are able to direct a range of their activities with the aid of cues from the external world. Within "activities" here I include some cases which fall naturally under the headings of growth or individual development, and others which might be distinguished from development and considered genuine behavior. Silvertown and Gordon (1989) argue at length that plants can behave, but they use an extremely broad conception of behavior. I suggest a narrower construal; one rough way to distinguish plant behavior from plant development is to say that behavioral changes must be reasonably rapid and also reversible. Then there will be a large range of cases where plants adaptively control individual development with environmental cues, and a smaller range of cases where they control behavior. Genuine plant behaviors include the behaviors of Venus Fly Traps, and a large range of reversible responses to local light conditions.

Some of the ways in which plants use environmental cues are quite sophisticated. For example, many plants can determine not just that they are being shaded, but that they are being shaded by other plants. This is done by detecting the wavelength properties of reflected light. The plants respond to shading by growing longer shoots (Silvertown and Lovett Doust 1993 pp. 11-12). Lest I leave out the least glamorous biological kingdom, some soil fungi have a reflex which enables them to trap and digest tiny wandering worms.

Though it makes sense to distinguish control of plant behavior from control of plant development, for many theoretical purposes these can be seen as similar capacities. From the point of view of theoretical modeling, control of developmental processes and control of behavior have much in common (see section 4 below). And the temptation to use intentional and cognitivist terms when describing control of plant development in more informal ways can be strong. David Attenborough begins his book The Private Life of Plants as follows: "Plants can see. They can count and communicate with one another" (1995 p. 7). Most of the phenomena Attenborough is referring to here involve control of growth and development, such as tactile exploration by a young vine, looking for a tree to climb.

I said that plants use a lot of their "smarts" for controlling growth and development, rather than behavior. In fact this phenomenon is not restricted to plants, but is found in vertebrate animals as well. In certain fish, capacities for perception and information-processing are used in directly regulating central aspects of development. These fish determine whether they will develop as male or female via perception of their relative size within the population (Francis and Barlow 1993). So when I said earlier that the ECT links cognition first to behavior, and then to environmental complexity, this was a slight oversimplification. Even in animals, sometimes cognition controls things other than behavior.

I have spent some time discussing capacities for flexible response which do not constitute genuine cognition on any normal standard. My aim in discussing these cases is to suggest that cognition shades off into other kinds of biological processes; even though plants do not exhibit cognition and people do, there is no single scale between them and us with a threshold marking a transition to genuine cognition. Rather, all or practically all living organisms have some capacities for responding to environmental changes and conditions. Sometimes environmental cues are used to control metabolic processes or development, sometimes they are used to consciously choose where to plant crops. In ordinary talk and in theoretical discussion, we habitually pick out only some of these capacities as "intelligent" or "cognitive," and the decision to do so can be guided by a mixture of criteria. Complexity and flexibility play a role, but so does the time-scale at which responses occur. There is nothing wrong with that; my point is just that there are some fundamental similarities between real cognition and much simpler capacities for control of biological processes, and there is no reason to seek a sharp cut-off between the two classes.

As some terminology might be useful here, I will say that plants and bacteria have a number of "proto-cognitive" capacities. These are capacities for controlling individual growth, development, metabolism and behavior by means of adaptive response to environmental information. The term "development" refers here only to processes within an individual lifetime; evolutionary change is not classified as proto-cognitive. Complex multi-cellular organisms like ourselves also contain "subpersonal" systems with some of the proto-cognitive capacities of simpler whole organisms. The vertebrate endocrine and immune systems are examples. In this paper I will not discuss the very difficult questions raised by the attribution to proto-cognitive capacities to higher-level systems, such as ant colonies.

In stressing that cognition shades off into other proto-cognitive biological processes, I am asserting a version of what is sometimes called a "continuity" assumption about cognition. The simplest biological capacities that we might consider proto-cognitive are cases of flexibility in behavior or development (etc.) controlled by a fixed response to a physically simple environmental cue, but where the nature of the response is not determined directly by the physical properties of the cue. (There has to be some "arbitrariness" in how the cue affects the system, to use a term due to Levins 1968.) As we add different types of flexibility of response, and different kinds of inner processing of the output of perceptual mechanisms, we reach clearer and clearer cases of cognition. But there is no single path that takes us from the simplest cases to the most elaborate. There are various ways of adding sophistication to the mechanisms of behavioral control, ways which will be useful to different organisms according to their circumstances. The ability to expand or contract the range of stimuli coupled to a given response is one important sophistication (Sterelny 1995). The ability to learn through reinforcement is another. Yet another is the ability to construct a "cognitive map" of spatial structure in the environment (this case will be discussed further below). It is an error to try to describe a single hierarchy of cognitive skills, from simplest to most complex. Here as elsewhere, there are many distinct kinds of complexity.

Compare two imaginary organisms which both have good spatial memory. One is "more sophisticated" than the other because it can remember more features of the environment, and can use its knowledge to find novel routes to where it wants to go. But this first organism can only acquire this spatial information by first-hand experience, by laboriously traveling and remembering the terrain. The second organism has a more limited capacity to remember features and to manipulate its internal model of the world, but it can acquire its knowledge in a richer variety of ways. It can infer spatial structure from the behavior of other organisms. In some respects the first organism is smarter, but in other respects the second is.

As we add sophistication to the tool-kit of behavior-guiding capacities, we eventually reach clear, unmistakable cases of cognition. And these clear cases do not all involve humans. Some birds, such as Clark's nutcrackers in the south-west of the US, hide stores of food when supplies are good, and

retrieve it in times of scarcity. To do this requires a sophisticated combination of perceptual abilities and spatial memory. In the sense in which I am using the term "cognition," there is nothing marginal about such cognitive capacities.

So I claim that we reach uncontroversial cases of cognition before we reach language use. And I have left out all mention of the "qualitative," first-person, "how it feels" side of mental life. My overall position is that we do have reasonably good evidence to posit rich qualitative states in non-human animals. But that is a separate point which does not matter to the task at hand.

To use a very broad sense of "cognition," as I do here, does not require postulation of fundamental similarity in the cognitive processes in all these diverse cases. Indeed, I have been stressing the opposite -- the diversity of ways in which cognitive and proto-cognitive capacities are realized. Many people have given general overarching theories of how all cognitive processes work. Examples include the general theories of learning that dominated psychology in the middle decades of this century. If one has an overarching theory of this kind, then one will want to have a broad term like "cognition" to capture what one is generalizing about. (Although many of the behaviorist psychologists who defended general theories of learning would not have liked the term "cognition.") I am skeptical about those overarching theories of cognitive processes and mechanisms, and I use a broad sense of "cognition" for what might be called more "ecological" reasons. There is a certain kind of job that the collection of processes I am calling "cognitive" performs. For one reason or another, organisms acquire capacities for behavior and machinery to control this behavior. The behavioral machinery acquired is diverse, and so are the mechanisms used to control this behavior. Our sophisticated human mental abilities are one instance of this evolutionary phenomenon, but the abilities of bees and jaguars are as well. Understanding the evolution of cognition is understanding this whole domain of evolution's products.

3. On Complexity

The ECT claims that the function of cognition is to enable organisms to deal with environmental complexity. But what exactly is environmental complexity? There is a good deal of unpacking to do here (see also Godfrey-Smith 1996.)

I suggest that the most useful concept of complexity here is a simple one. Complexity is heterogeneity. Complexity is variety, diversity, doing a lot of different things or having the capacity to occupy a lot of different states.

There are many different kinds of heterogeneity, hence many kinds of complexity. It is not just unnecessary, but positively mistaken, to try to devise a single scale to order all environments from the least to the most complex. Rather, any environment will be heterogeneous in some respects, and homogeneous in others. Environments can be heterogeneous in space and in

time, and spatial and temporal heterogeneity exists at many different scales. An environment with a large number of different possible states which come and go over time is a complex environment, in that respect. So is an environment which is a patchwork of different conditions across space. The heterogeneity property is not the same in these two cases, but in both cases heterogeneity can be opposed to homogeneity. A complex environment is in different states at different times, rather than the same state all the time; a complex environment is different in different places, rather than the same all over. Whether a particular type of complexity is relevant to an organism will depend on what the organism is like -- on the organism's size, physiology, needs and habits. The heterogeneity properties of environments are objective, organism-independent properties, but among the countless ways in which an environment is structured and patterned, only some will be relevant to any given organism. (See Levins 1968 for a classic discussion of some of these issues.)

In the ECT I said that cognition enables agents to "deal" with environmental complexity. That terminology suggests that environmental complexity is seen as posing problems for organisms. Often this is so, but I do not want to put too much weight on the concept of a "problem" here (Lewontin 1983). In some cases, environmental complexity provides what would normally be called an "opportunity" rather than a problem. A population might be located in a fairly benign set of circumstances, but one where tracking and adapting to environmental complexity makes it possible for some individuals gain a reproductive advantage over others. Natural selection works in a comparative way; the absolute level of hardship is in general not important in understanding evolutionary processes within a population. So while I will often write of the "problems" posed by environmental complexity in this paper, occasionally I will use the term "opportunities" as well. The distinction between the terms is mostly an everyday one which should not be taken too seriously in this context.

If we think of complexity just as heterogeneity, this concept of complexity can be applied to organisms as well as to environments. An organism is complex to the extent that it is heterogeneous. Here again, there are different kinds of heterogeneity; an organism can be heterogeneous in many different respects. (For different concepts of organismic complexity, see McShea 1991.)

Cognitive capacities themselves are complex, so the ECT can be seen as claiming that one kind of organic complexity has been produced by evolution to enable organisms to deal with environmental complexity. Dealing with complex problems by means of perception and action can be seen as a special case of a more general phenomenon: dealing with environmental complexity by means of flexibility.

This way of looking at the ECT is illustrated by the "proto-cognitive" capacities that were discussed in the previous section of this paper. When a plant has the ability to adaptively alter its development to suit its environment, this is a case of complexity in the plant's developmental

capacities which enables the plant to adapt to heterogeneity in its environment. Similarly, why do e. coli bacteria have their lac operon system of gene regulation? The preferred food of e. coli bacteria is glucose, but sometimes glucose is not available while other sugars are. The variability in the availability of different sugars is one type of environmental complexity faced by bacteria. Metabolic machinery is expensive, and e. coli have apparently been selected to economize in their production of enzymes. So the enzyme needed to digest lactose is not produced in the absence of lactose. Instead, the production of the enzyme is controlled by an environmental cue. Here the cue used is the presence of lactose itself (and also the amount of glucose available to the cell -- see Lodish et al. 1995 pp. 421-22). The system of gene regulation used by the bacteria here constitutes one kind of complexity in these organisms, and this complex mechanism has a functional explanation, of the strong type discussed earlier. The function of the lac operon system is to enable e. coli bacteria to deal effectively with one type of environmental complexity -- variation in the availability of different sugars.

As I said earlier, I do not claim that bacteria exhibit cognition; this is at most a case of proto-cognition. However, the ECT claims that the explanations for more complex and genuinely cognitive capacities tend to have a similar general shape as this explanation for a property of bacteria. The point of acquiring complex systems for behavioral control is to enable the organism to deal with variation in what the environment confronts the organism with, and variation in the opportunities the environment offers.

Environmental complexity figures in evolutionary processes that give rise to cognition. But where does environmental complexity itself come from? And what should we make of cases where environmental complexity is itself the product of organisms and their activities?

Environmental complexity itself has many sources. For the purposes at hand I will make a loose distinction between two main categories. One source is the class of physical processes which are more or less independent of the activities of the organisms under consideration. Seasonal cycles provide an obvious example. And many resources that are relevant to an organism's well-being will be scattered through space in a way that is largely independent of the organism's own actions and properties.

When some type of organism acquires, through evolutionary processes, a way of tracking and dealing with environmental complexity of this first kind, the explanatory pattern described by the ECT has a straightforward causal directionality. But there are other cases in which the situation is more complicated. These are cases where the environmental complexity that organisms must deal with is either a causal product of, or is constituted by, the activities of other organisms within the same population. Then we have a situation that can exhibit feedback, or a "coupling" of organism and environment. (Lewontin 1983, Odling-Smee 1988).

The most graphic examples are probably those that involve competitive interactions between animals. If the only way for you to obtain and hold a resource is by winning contests with other individuals in the same

population, then these other organisms constitute a key part of your environment. Their behavioral complexity constitutes part of the environmental complexity you must deal with, so the behavioral capacities of organisms similar to yourself are the source of a crucial kind of complexity in your own environment.

In behavioral ecology, contests of this kind are modeled with game theory (Krebs and Davies 1987). Most mathematical game theory models only remain simple enough to be comprehensible when many idealizations are made. Surrounding a few well-understood cases explicitly modeled with game theory, there is now a great deal of informal verbal "modeling" (in scare quotes) and computer simulation of these sorts of interactions going on. In fact, some have claimed that feedback processes of this kind are the key to understanding the evolutionary transition to genuine human intelligence. Those suggestions will be discussed in my final section below.

My present point is that these phenomena are not incompatible with the ECT. The ECT need not be understood in a way in which the processes generating environmental complexity are casually autonomous, or independent of the activities of the evolving population in question. The ECT is compatible with the view that a centrally important aspect of environmental complexity for many organisms is complexity that is made up by, or caused by, the activities of other organisms of the same species. In those cases, the ECT describes one part of a larger causal "cycle" -- the part in which environmental complexity puts selective pressure on organisms' cognitive capacities. The other part of the "cycle" is where the behaviors of organisms influence or determine the relevant patterns of environmental complexity.

So some environmental complexity for a given organism is made up by the activities of other organisms in that population. What about organisms from other species, which constitute sources of food, or sources of danger, for the organisms we are concerned with? I stress again that my two-way distinction here is rough and ready. To the extent that the relevant activities of other species are causally influenced by the properties and activities of the population under consideration, we have a case of the second "coupled" type. Predator-prey interactions are a classic example. In general descriptions of ecological relationships, people often stress that every species is connected to virtually ever other, through direct or indirect causal chains. Clearly however this is a matter of degree. It is an error to over-generalize about the richness of inter-specific connections, just as it is an error to treat organisms as if all they ever have to deal with is an independent, causally autonomous, physical environment.

I have been contrasting relatively simple cases in which organisms are responding to environmental complexity that is causally independent of them, and more complicated cases where the environmental complexity itself depends on the organisms in significant ways. But even in the "simpler" class of cases, it should not be thought that I am suggesting that the evolutionary processes themselves are simple and predictable. Much environmental complexity is not relevant to any given organism, and the factors that

contribute to some aspect of complexity posing a problem are diverse and subtle. Suppose you move through the world like a monkey, swinging from tree branches. Then the relevance of diversity in the size and strength of these branches depends a great deal on your size. If you are small, most branches will support you and in any case a fall is unlikely to lead to serious harm. If you are larger, paths through the forest must be chosen with care and pose a significant information-processing problem because of the causal role of your own weight (Povinelli and Cant 1995). Some branches will break, leading to a dangerous fall, while others will bend, in ways that affect your possible next moves. Heterogeneity in the properties of tree branches is thus relevant in different ways, and in different degrees, to differently sized organisms.

The mere presence of environmental complexity that is relevant for a given type of organism does not automatically generate cognition, or even natural selection for it. The consequences of relevant environmental complexity also depend on many other features of the organism and its ecology. The fact that the ECT is expressed as a simple generalization should not be taken to downplay the role for "architectural constraints" in explaining why evolution takes the course it takes in a particular case. (These constraints are famously discussed in Gould and Lewontin 1979.) For some organisms, getting smart is not really an evolutionary option, as a consequence of their basic biological lay-out, their characteristic developmental sequence, or their overall ecology. Even for those that could, in principle, start to respond to environmental variation by tracking and behaviorally adapting to it, the appropriate genetic variation has to arise, and there will be costs associated with the machinery required to take a smart approach.

It is also well understood that some kinds of environmental complexity can be effectively dealt with by buffering it or blocking it out. One can respond to a threat by being smart, but also by becoming impervious to it, via a strong shell or via sheer size. Some organisms, including many insects, deal with certain kinds of environmental complexity with an "r-selected" strategy for reproduction, in which there is massive reproduction in good times, and little activity in bad times. To take this strategy it is necessary to be able to produce huge numbers of quickly-maturing offspring when times are good. All this makes for a lack of cognitive machinery in r-selected organisms.

So whether evolution takes a lineage of organisms down a path towards increased cognitive capacity is contingent on a great range of factors, many of them having to do with the "raw materials" that evolution has to work with in that particular case. But this fact does not make the ECT false or naive. The ECT, when it applies to some particular case, is one part of a more complicated and detailed explanation.

4. The ECT as a Component in Many Evolutionary Scenarios

So far this paper has discussed the ECT in extremely general terms. In this final section I will look at some specific models and programs of empirical

work. I will discuss four examples, each focused on understanding a specific type of cognition (or proto-cognition). I suggest that the ECT is one component in many diverse scenarios that have been discussed in connection with the evolution of cognition.

(i) Phenotypic plasticity

I have said several times in this paper that cognition, understood in my broad way, shades off into other biological processes, especially those that use signals (from the environment or elsewhere) to control adaptive responses. One important class of cases in the category I have been calling "proto-cognitive" is the phenomenon of phenotypic plasticity, especially in plants.

In the paradigm cases, phenotypic plasticity is a phenomenon in which a single plant genotype can produce variety of forms (phenotypes) or can take a variety of developmental paths, where the "choice" is determined by an environmental cue transduced by the plant (Bradshaw 1965, Sultan 1987, Schlichting and Pigliucci 1998). A plant might have a wet-environment and a dry-environment phenotype, for example, or might alter its form according to altitude and accompanying climatic conditions, as in the classic experiments of Clausen, Keck and Hiesey (1948). So in these cases the plant has some mechanism for transducing an environmental cue, and of controlling development as a function of the state of the cue. The cue, the plant phenotype, and the environmental variables that are being adapted to might be discrete or continuous. (When the organism's response is a discrete choice this is sometimes called "polyphenism," but I will not make that terminological distinction here.)

No nervous system is involved in these cases, and in general it is growth and development, rather than behavior, that is being controlled. But this type of phenomenon is a useful "zero order" case for discussions of models of adaptive response to environmental conditions. There are formal similarities between these capacities and cases of real cognition. Indeed, in the 1990s two mathematically identical evolutionary models were published independently (Moran 1992, Sober 1994). One was presented as a model of the advantages of learning (Sober), while the other was presented as a model of the advantages of plastic control of development (Moran).

How do the models and theoretical discussions look? Let us be very abstract. Assume that an organism confronts an environment which has a range of alternative possible states. The organism itself has a range of possible developmental options. The alternative environmental states have consequences for the organism's chances of surviving and reproducing, and the best developmental option for one environmental state is not the best choice for another. The organism receives imperfect information about the actual state of the environment, as a consequence of correlations between environmental conditions which matter to it and environmental conditions which directly affect its periphery. There are several ways to respond to the problem. One way is to be able to buffer out the environmental variation -- perhaps by being big, or strong, or restricting exposure in some way. Another

way is to adapt to the most common or the most critically important environmental state. But yet another way is to use a flexible strategy -- to use environmental information to determine the organism's phenotype in accordance with how the environment is perceived to be.

For example, Drew Harvell (1986) investigates defences against predators produced by colonial marine invertebrate animals called "bryozoans," or sea moss. The bryozoans Harvell studies are able to detect the presence of predatory sea slugs, making use of a water-borne chemical cue. When sea slugs are around, the bryozoans produce spines. The spines have been shown to effectively reduce predation, but also to incur a significant cost in terms of growth, so they are detrimental when sea slugs are not around. Here we have a rudimentary form of perception. The bryozoans show sensitivity to an environmental cue which is not itself practically important, but which carries information about a more important state of the world, the presence of predators. The organism use a cue to produce an adaptive response to a more important "distal" environmental state.

In cases like these, a complete explanation for the organism evolving a "smart" or proto-cognitive capacity includes a description of the problem posed by environmental complexity -- the fact that predators are sometimes, but not always, present. But the explanation includes much more as well. We need to also know the reasons for a proto-cognitive response being favored over buffering, adaptation to the most common condition, or some other "dumb" strategy.

When will the proto-cognitive strategy be favored? There is no simple answer, but many models developed by biologists and others can be pieced together to give a partial answer (Godfrey-Smith 1996, chapters 7-9). Some parts of the story are intuitive. To use a proto-cognitive strategy, the organism needs a suitable signal from the environment. If there is no way of tracking the relevant states of the environment, it is better to produce a single "cover-all" phenotype, or to adapt to the single most common or important environmental state. Parts of information theory and signal detection theory can be used to describe exactly what sorts of properties an environmental cue must have, in order for it to be worth using. When do you want to choose a flexible strategy over an inflexible one? Only when your environmental cue is good enough so you don't make too many of the wrong kinds of errors. Even you can track the world with some reliability, if one type of wrong decision is sufficiently disastrous, it may be best never to behave in a way which risks this error. Principles like those are close to common sense. But these models also have a number of more subtle features. For example, it can matter a great deal how payoffs from individual encounters or "trials" are related to each other in their effects on overall fitness -- whether payoffs are summed or multiplied (Levins 1968, Seger and Brockman 1987). The principles discussed by these models of plasticity cast light on both proto-cognition and genuine cognition as well. The models describe the first step towards cognition -- opting for a flexible response to a heterogeneous environment.

(ii) The evolution of associative learning.

The second example I will discuss is a computer simulation of the evolution of associative learning, due to Todd and Miller (1991). This simulation explores the evolution of the architectural properties of simple networks of neurons, using what is known as the "genetic algorithm." The aim is to see when evolution will select for organisms that exhibit one of the simplest kinds of learning -- classical conditioning.

The neural networks can usefully be imagined as embodied in simple marine animals which are born in the open sea, but which settle down to an immobile life feeding on passing food particles. Once an individual has settled, its only problem is to decide whether to feed or not feed, when presented with each item of possible food. The environment contains both food and also inedible or poisonous particles, in equal proportions. When food is eaten the organism gains an energetic benefit, and when poison is eaten the organism pays a cost, though the error is not fatal.

Particles of possible food have two sorts of properties that the organisms can perceive -- color and smell. Food smells sweet and poison smells sour, but in this turbulent environment smells can mislead. The probability of a sweet smell, given the presence of food, is 0.75. The probability of a sour smell, given poison, is also 0.75.

The color of food is not affected by turbulence, but color is unpredictable in a different respect. In half of the population's environment food is red and poison is green, but in the other half the colors are reversed. Within each of these two micro-environments, color is 100% reliable.

Each generation contains a large number of individuals of different types, which settle at random in the two different micro-environments. At the end of a fixed period they reproduce (sexually) according to their accumulated fitness, with the possibility of mutation and recombination of genes. The new generation then floats about and settles in the environment at random and the cycle begins again.

The neural networks placed in this scenario are constrained to have three "units," or nodes, only. These nodes are like idealized nerve cells. Despite these limited resources there are lots of possibilities for the networks' architectures, and these architectures evolve in the model by natural selection across generations. Units can function as input devices of various kinds (red-detectors, green-detectors, sweet-detectors or sour-detectors). There is just one type of output unit (eating), and a "hidden" unit, which mediates between a detector and a motor unit, is also a possibility. The range of units an individual has, and which ones are connected to which, are determined by its genetic make-up.

Connections between units can be hard-wired with either an excitatory or inhibitory one-way connection, or they can be plastic and altered by the individual's experience. If the genotype specifies a plastic connection, then the connection is shaped over time by what is known as a "Hebbian" learning rule. If those two units tend to fire at the same time in the individual's

experience, they acquire a positive connection between them -- one unit comes to have a (one-way) excitatory connection to the other. If they do not tend to fire together, the connection becomes negative or inhibitory. The question the model is intended to address is: when and in what ways will individuals with the ability to learn evolve in the population? The question is interesting because Hebbian learning is discussed a good deal by neuroscientists, and they hunt for Hebbian learning in the synapses of the brain. But from the evolutionary point of view, it is often not obvious what use Hebbian learning has. If two neurons tend to fire together, what is the point of also making one excite the other?

At the start of a "run" of the Todd and Miller simulation, the population consists of randomly configured individuals, most of which do not fare well. For example, some will not have a motor unit and will never eat, or will have a motor unit connected to an input unit which has the wrong setting -- it might tell the organism to always eat when the present food particle smells sour. Another type of miswiring might be called "the academic." An individual can have two input units and a motor unit, but only learnable connections between all three units, connections which are initially set at zero. Suppose such a creature lands in a patch where food is red. Then it will learn the statistical association between redness and a sweet smell -- the red-color input unit will tend to be on at the same time as the sweet-smell unit. But nothing is inducing the individual to eat. The motor unit will never be turned on, and its knowledge of the world will not do the individual any good, as far as nutrition is concerned.

Two kinds of wiring do work well for the organism though. One has a fixed positive connection between a sweet-smell sensor and a motor unit, and nothing else which influences behavior. This organism will generally eat when there is food present -- in the present case, it will make the right decision 75% of the time. After a short period, these individuals tend to proliferate in the population.

The best possible wiring is a variant on this one, which has a fixed connection between a sweetness sensor and a motor unit (as above), but also a learnable connection between a color sensor and the motor unit. From the start this individual will tend to eat when there is food, as the smell sensor is controlling the motor unit. But in addition there will be a correlation between eating and some particular state of the color sensor. If the micro-environment is a red-food one, then when the organism eats it will also tend to be seeing red. This correlation establishes a connection between the color sensor and the motor unit, and (given the right initial settings) this connection will eventually be strong enough to control the motor unit by itself. Then the eating behavior will be controlled by a 100% reliable cue for the remainder of the individual's life. Typical runs of the simulation begin with the fairly rapid evolution of the simple, hard-wired smell-guided networks, and some time afterwards learners appear and take over.

This simulation illustrates the advantages associated with two kinds of behavioral complexity. Consider first the contest between individuals which

always eat every particle that drifts by, and individuals which use smell as a cue. If everything in the environment was food, there would be little point in controlling behavior with perception, especially with an only partially reliable environmental cue. But the environment used by Todd and Miller is one where food and poison drift by with equal frequencies. This is one type of environmental complexity, and it has the consequence that a permanently-eating architecture will have low fitness when compared to an individual that uses smell as a cue. A different type of environmental complexity, and a different reliability relationship between a cue and the world, explains the evolution of learning. The total environment in which these organisms live is spatially heterogeneous -- in half the environment food is red and in the other half food is green. If food was red in the whole environment, it would not be worth taking the time to learn that food is red, and a network with a red-detector hard-wired to the output unit would be optimal. But Todd and Miller use an environment which is heterogeneous in this respect as well. And although the color of food is not predictable in advance, the past experience of an individual is a good guide to the future. That is what is needed for learning to be more useful than an inflexible behavioral program.

(iii) Spatial memory and cognitive maps

For a mobile animal, one very important kind of environmental heterogeneity is heterogeneity in the distribution of resources, dangers and other factors in space. Spatial structure plays a very different role for a plant, of course, or for an animal like a clam which does not move around the world. But once an organism is on the move, as most terrestrial animals must sometimes be, spatial structure in the environment is of prime importance.

Here as in general, some of this environmental heterogeneity can be dealt with by various forms of buffering. But evolutionary responses to the problem of dealing with space have produced some impressive and complicated forms of cognition, even in small and otherwise behaviorally simple animals. The mechanisms associated with bee dances, which direct workers from the hive to sources of food, are one famous example, but it has turned out that bees as individuals also show good spatial skills. They can learn to reliably associate a source of food with either single landmarks or with geometrical structure in a set of landmarks. In recent years a lot of attention has been directed on spatial memory in food storing birds, such as the Clark's nutcracker in the south-western US, and marsh tits in England. Clark's nutcrackers hide thousands of pine seeds as a food source for the winter. It appears that these birds have specialized spatial memory abilities which do not extrapolate (as far as has been determined) to superiority over other birds in non-spatial memory tasks. (See Roberts 1998 chapter 7 for the bee and bird examples in this paragraph.)

Back in 1948, E. C. Tolman suggested that both animals and humans make their way through space by using "cognitive maps," or rich internal representations of spatial structure in the environment. After initial

controversy and some decades of neglect during the heyday of strict forms of behaviorism, the concept of a cognitive map is again being used by ethologists and comparative psychologists (Tolman 1948, Thinus-Blanc 1988, Roberts 1998). The concept of a cognitive map is controversial in a number of respects. First, there is a good deal of vagueness and ambiguity in how it is applied by different researchers (Bennett 1996). Some use the term to refer specifically to postulated internal structures which work in psychological processing in ways reminiscent of ordinary, external maps. In this narrow sense, the hypothesis that an animal has a "cognitive map" requires, at a minimum, a capacity to devise novel detours and shortcuts in response to obstruction of more familiar paths. But the term is sometimes used more broadly, to refer to almost any kind of spatial memory. Tolman, for example, distinguished "strip maps" and "comprehensive maps" within the more general category of cognitive maps. Strip maps represent only a path to a goal; they are dependent on the starting point of the animal. Comprehensive maps are richer representations of the overall spatial structure in some domain, so they can be used despite variation in starting points, new obstacles and so on. But it can be argued that the sort of behavior associated with "strip maps" is easily explained without talking of inner "maps" at all; the animal is just executing a sequence of behaviors in response to a set of cues or landmarks. Some sophistication in memory is clearly involved, but there is no need to postulate an inner map-like structure.

However, there are experimental results which do justify a richer interpretation of some animals' inner processing of spatial information. A simple and striking case is found in an experiment by Tolman and Honzik (1930). Rats were trained in a maze that has three different paths to a single supply of food. Path 1 is shorter than path 2, and path 2 is shorter than path 3. The rats were easily able to learn to prefer the best available path. After a few days of training, path 1 was almost always chosen first. If path 1 was blocked (at an early point), they would go back and take path 2. If path 2 was also blocked, they would settle for path 3. So far, this only shows a fairly routine (but very useful) type of reinforced learning. The impressive behavior resulted when path 1 was blocked in a novel way for the first time. Path 1 has its final section in common with path 2, but path 3 reaches the food independently of this common section. So when this final part of path 1 is blocked, that has the effect of making path 2 useless as well. What will the rats do when path 1 is blocked in this novel way? Their history of conditioning has taught them that when path 1 is blocked, path 2 is the next choice. But if the rats are smart enough to realize the consequences of this novel way of blocking path 1, they should choose path 3 directly and not waste time on path 2. In Tolman and Honzik's experiment, a large majority of rats, on encountering the novel obstruction on the later part of path 1 for the first time, returned to the junction point of the three paths and immediately chose path 3. They did not follow their failure on path 1 with an attempt at path 2; they had somehow been able to represent the new obstacle as rendering path 2 useless as well.

Tolman and Honzik interpreted this as showing "insight" on the part of the rats, following the gestalt psychologist Köhler, who had found similar results with chimps. Tolman did not, for some reason, use this experiment in the famous 1948 discussion which introduced the concept of a "cognitive map." Instead he used results which, to my mind, were a good deal less convincing than his 1930 "insight" experiment. (Thinus-Blanc 1988 erroneously reports Tolman 1948 as actually discussing the "insight" experiment in support of the cognitive map concept, an interesting case of wishful thinking, or post-hoc improvement of Tolman's paper!) Perhaps Tolman did not think of the "insight" experiment as showing specifically spatial cognitive skills, but a more general capacity to draw conclusions and reason beyond the immediate lessons of conditioning. However the "cognitive map" concept is used, this 1930 experiment does appear to show a capacity to construct and manipulate some sort of internal model of spatial structure in the environment. In qualification of this, I should note that Tolman and Honzik found it quite tricky to devise a maze in which most of their rats would consistently show this spatial "insight." For example, insight was consistently shown only when the maze was made of elevated tracks, not tunnels.

Another kind of experiment designed to investigate innovative behavior based on representation of spatial structure studies the use of short cuts. Both dogs and chimps can be shown hidden pieces of food in an environment, with the order of their exposure to the food corresponding to an inefficient path from food item to food item. Once released, the animals in both cases are able to find the food items and move from item to item using a path that is new and more direct and efficient than the one they were trained on. (In effect, the chimps do a reasonable job at what mathematicians refer to as a "traveling salesman problem.") Here again, the behavior produced does not correspond to any motor routine or action pattern that the animal was trained on, and these experiments also control for such possibilities as olfactory detection of the food (Roberts 1998, Menzel 1997).

All of these experiments are associated with some controversy. For example, it can be argued that to the extent that animals are simply moving from one memorized landmark (or point specified by its relation to a set of landmarks) to another, there is justification for attributing memory to the animal but not a cognitive map (Bennett 1996). Menzel (1997) found deviations from "traveling salesman" optimality when macaques were tested on whether they chose the optimal path from item to item, or simply went towards the closest piece of food at each decision point. But it is noteworthy that the best of this work, such as the Tolman and Honzik "insight" experiment, does involve behavior which appears to justify the postulation of internal representation of the world, and fairly complex use of the representations in guiding behavior, without the animal having a capacity for public language. There is a tradition in philosophy of denying that any animal which lacks language can properly be said to "think" or "represent." In recent years, Davidson (1975) has been the most influential defender of this

view. It was held in a different form by Dewey (1929) and is often associated with Wittgenstein (1953) and his followers. Such views struggle to make sense of the skills in animal path-choice discussed above. Whether or not these animals are able to "represent" the world in the richest, most philosophically loaded sense of "represent," they do seem to be doing some kind of representation or mapping of the spatial structure of their environment (see also Allen and Bekoff 1997). Problems and opportunities associated with spatial structure in environments have apparently generated a range of sophisticated cognitive skills.

(iv) Social intelligence models

It has usually been thought not too hard to explain the evolution of such capacities as learning (Example (ii) above) and the ability to represent spatial structure in an environment (Example (iii)). But it is a different matter to explain the highly developed and distinctive cognitive capacities of human beings and non-human primates. Over 20 years ago, Nicholas Humphrey suggested that primates seem to have too much brain-power to be explained by the demands of such activities as foraging for food. He suggested that the problems primates are using this intelligence to deal with stem from the social complexity of their environments (Humphrey 1976). In particular, much primate life is concerned with the formation and maintenance of alliances, the policing of dominance hierarchies, and a variety of other social tasks that involve a mixture of competition and cooperation. Humphrey's suggestion (which had been partially anticipated by others) was that high primate intelligence evolved in response to the problem of dealing with this kind of complexity. As each individual primate comprises part of the environment for the others in a population, we have the ingredients here for a process of feedback, in which each increase in intelligence produced by evolution adds to the complexity of the social environment that individuals face.

This idea has come to be known as the "Machiavellian intelligence hypothesis" (Byrne and Whiten 1988, Whiten and Byrne 1997). Here I use the term "social intelligence" rather than "Machiavellian intelligence." As Whiten and Byrne say, the term "Machiavellian" should not be taken to suggest that all the behaviors involved in the hypothesis are manipulative and deceitful. They mean to include "cunning cooperation" which contributes to individual reproductive success (Whiten and Byrne 1997 pp. 12-13). Whiten and Byrne do mean to exclude, however, any suggestion of natural selection operating at the level of groups rather than individuals -- they mean to exclude the idea that intelligent cooperation could evolve "for the good of the group." As I do think the term "Machiavellian intelligence" continues to mislead, suggesting the darker side of social behavior, and I also do not want to rule out some role for group selection in these processes, I prefer to use the term "social intelligence hypothesis." Like Gigerenzer (1997), I reserve "Machiavellian" as a narrower term, specifically for behaviors involving exploitation rather than cooperation.

The social intelligence hypothesis is compatible with the ECT; it is a specific instance of the ECT which involves a special kind of environment. In the previous section I argued that the ECT need not be understood in a way that requires the environment in question to be independent of the population that is evolving. In any social animal, a key part of an individual's environment is made up of the other members of the social group, with all their behavioral capacities. The same is true to a lesser extent with many non-social animals. In cases like these, the ECT describes one particular explanatory arrow within a larger explanatory structure, a structure which links cognitive capacities and environmental complexity in a "coupled" way. At any given time, the individuals in a social population face environmental complexity in the form of the behavioral patterns of the other local members of the population. This environmental complexity may (or may not) give more intelligent individuals an advantage over less intelligent individuals, by some specific measure of intelligence relevant to the situation. If intelligence is favored, and if this kind of intelligence is inherited, then over time intelligence will increase in the population. If the intelligent individuals themselves display more complex patterns of behavior than others in the population, then this increase in intelligence will in turn entail an increase in the complexity of the social environment faced by later individuals. This process may or may not have a "runaway," positive-feedback character. One should not assume that a runaway process is the only outcome. For example, other constraints and costs might start to assume a larger role once the individuals reach a certain level of intelligence.

The social intelligence hypothesis has been developed in a number of different specific versions. Stronger versions claim that social complexity has been the key factor in producing the high levels of primate intelligence; weaker versions see this as one explanatory factor which might work in conjunction with others. For example, some suggest a role for special problems associated with foraging for the ripe fruits favored by primates -- it might be that primates require especially sophisticated cognitive maps of their (non-social) environments. (Whiten and Byrne 1997 contains discussions of a range of alternatives to the social intelligence hypothesis.) Different versions of the hypothesis also stress different aspects of social living -- direct competition between males for mates, dealing with dominance hierarchies, cooperative foraging and so on.

As Gigerenzer (1997) notes, the social intelligence hypothesis is sometimes associated with the suggestion that (i) the overall degrees of complexity in social and non-social environments (or social and non-social aspects of an environment) can be compared, and (ii) social environments are more complex. As I said back in section 3 of this paper, I am a skeptic about the project of giving overall measures of complexity across environments; any environment is complex in some respects and simple in others. Gigerenzer is similarly skeptical about these complexity measures. But where Gigerenzer appears to think that this problem makes it pointless to generalize about the role of complexity in the evolution of cognition, I think the

environmental complexity thesis is a useful general principle despite the absence of a unitary scale of environmental complexity.

As I understand them, most but not all of the specific hypotheses discussed under the general category of "social intelligence" can be seen as applications of the ECT. The versions that do fall under the ECT are those that stress the role of cognition in dealing with the behavioral complexity of other individuals within a social group. An example of a hypothesis in this general area that does not fall under the ECT is one version of the "protean behavior" hypothesis, discussed by Geoffrey Miller (1997). Miller suggests that animals like primates have been selected to be able to produce genuinely unpredictable behavior. Unpredictable behavior has fairly obvious advantages when an animal is escaping from predators, and more subtle advantages in situations involving conflict and bluffing, as discussed in game theory. In those cases, producing unpredictable behavior itself requires no special cognitive sophistication. But Miller also suggests that capacities for novel, creative behavior will help individuals attract mates, especially in situations of female choice, and these behavioral capacities are facilitated by a large brain. There are several ways in which this scenario might work, but consider one case. Suppose there has been selection for behavioral novelty in males, and suppose these novel mating displays succeed by taking advantage of a general feature of perceptual and cognitive mechanisms -- the fact that novelty attracts attention. Then any evolved increase in cognitive sophistication due to this process cannot be seen as an application of the ECT. A complex social environment might be created by male behaviors in this case, but cognition is not being selected as a way of dealing with complexity. On the other hand, if females are being strongly selected for the capacity to see through all this behavioral noise and make adaptive choices, there will be selection, compatible with the ECT, for cognitive sophistication in females.

All these "sex-specific" hypotheses about the advantages of cognition in primates have a problem stemming from the basic similarity between male and female brains in the most intelligent primates (as Richard Francis stressed to me). If selection only favors elaborate cognition in one sex, then the other sex will have much of its brain treated by the theory as a mere byproduct. Big brains are too expensive to be treated like male nipples; big brains as byproducts would be analogous to peacock tails on peahens. If both sexes are being selected to be smart, but in very different ways, then the problem posed by the lack of obvious sexual dimorphism is more subtle and hard to assess. In any case, I introduce these speculations about unpredictable behavior not to endorse them, but to illustrate the fact that while the ECT is very broad, it does not trivially encompass any possible explanation for the evolution of cognition. The ECT only covers cases where environmental complexity (whether social or nonsocial) creates a problem (or an opportunity) for some type of organism, and the problem leads to natural selection favoring individuals with an ability to use cognition to coordinate behavior with the state of the environment.

In earlier discussions of social intelligence, there was sometimes the appearance of a sharp "either-or" characteristic to the debates about social and non-social complexity; either primates became smart for social reasons, or they became smart for reasons having to do with non-social aspects of their ecology. But as Byrne (1997) argues, there is plenty of room for mixed explanations, in which a number of factors have a role. Byrne himself posits three distinct evolutionary transitions in the evolution of intelligence in monkeys, apes and humans. In this scenario, the evolutionary branch containing monkeys and apes (haplorhines) became smarter than its relatives because of selection for social intelligence. But the "Great Apes" (chimps, bonobos, humans, gorillas, orangutans) branched off from this group and became smarter because of selection for "technical intelligence," which involves planning and sophisticated tool-use in activities such as foraging. And then the branch that led to modern humans was perhaps again the subject of selection for social intelligence, in part because of larger group size. This scenario is very speculative, as Byrne stresses, but it provides a good example of a "mixed" story about the evolution of human cognition. It would be a mistake to only pursue "pure" social intelligence hypotheses, out of an overly strong attachment to explanatory simplicity.

In closing, I will try to give a more "fleshed out" summary of what the ECT claims.

Environmental Complexity Thesis (more detailed version):

The basic pattern found in the evolution of cognition is a pattern in which individual organisms derive an advantage from cognitive capacities in their attempts to deal with problems and opportunities posed by environmental complexity of various kinds. Cognitive capacities confer this advantage by enabling organisms to coordinate their behavior with the state of the environment. Cognition itself should be thought of as a diverse "tool-kit" of capacities for behavioral control, including capacities for perception, internal representation of the world, memory, learning, and decision-making. These capacities vary across different types of organism and are not sharply distinguished from other biological capacities, some of which have a "proto-cognitive" character. The "environment" referred to in the ECT includes the social environment, and there are some reasons to believe that problems posed by social complexity have been very important in the evolution of primate and human intelligence. Many specific evolutionary scenarios that have been discussed as possible explanations of particular cognitive capacities are instances of the ECT, or have the ECT as a part.

* * * * *

Acknowledgment

Thanks to Richard Francis, Lori Gruen and Kim Sterelny for discussions and correspondence on these issues.

References

- Attenborough, D. (1995). The Private Life of Plants. Princeton: Princeton University Press.
- Allen, C. Bekoff, M. and G. Lauder (eds.) (1998). Nature's Purposes: Analyses of Function and Design in Biology. Cambridge MA: MIT Press.
- Allen, C. and M. Bekoff (1997). Species of Mind. Cambridge: MIT Press.
- Bennett, A. T. D. (1996). "Do Animals Have Cognitive Maps?" Journal of Experimental Biology 199: 219-224.
- Bonner, J. (1988). The Evolution of Complexity. Princeton; Princeton University Press.
- Bradshaw, A. D. (1965). "Evolutionary Significance of Phenotypic Plasticity in Plants," Advances in Genetics 13: 115-55.
- Byrne, R. W. and A. Whiten (eds.) (1988). Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans. Oxford: Clarendon Press.
- Byrne, R. W. (1997). "The Technical Intelligence Hypothesis: An Additional Evolutionary Stimulus to Intelligence," in Whiten and Byrne (1997) pp. 289-311.
- Clausen, J., D. Keck, and W. M. Hiesey (1948). Experimental Studies on the Nature of Plant Species III. Environmental Responses of Climatic Races of Achillea. Carnegie Institution of Washington. Publication 581. Washington DC.
- Davidson, D. (1975). "Thought and Talk," reprinted in Essays on Truth and Interpretation. Oxford: Oxford University Press, 1984, pp. 155-170.
- Dewey, J. (1929). Experience and Nature (revised edition). New York: Dover, 1958.
- Dretske, F. (1981). Knowledge and the Flow of Information. Cambridge MA: MIT Press.
- Dretske, F. (1986). "Misrepresentation," reprinted in Stich and Warfield (1994) pp. 157-73.
- Francis, R. and G. W. Barlow (1993). "Social Control of Primary Sex Determination in the Midas Cichlid," Proceedings of the National Academy of Sciences, USA 90: 10673-10675.
- Gibson, J. J. (1966). The Senses Considered as Perceptual Systems. Boston: Houghton Mifflin.
- Gigerenzer, G. (1997). "The Modularity of Social Intelligence," in Whiten and Byrne (1997) pp. 264-288.
- Godfrey-Smith, P. (1996). Complexity and the Function of Mind in Nature. Cambridge: Cambridge University Press.
- Gould, S. J. and R. C. Lewontin (1979). "The Spandrels of San Marco and the Panglossian Paradigm: a Critique of the Adaptationist Program," Proceedings of the Royal Society, London 205: 581-598.

- Harvell, D. (1986). "The Ecology and Evolution of Inducible Defences in a Marine Bryozoan: Cues, Costs, and Consequences", American Naturalist 128: 810-23.
- Humphrey, N. (1976). "The Social Function of Intellect," Reprinted in Byrne and Whiten (1988).
- Krebs J. and N. Davies (1987). An Introduction to Behavioural Ecology. 2nd edition. Oxford: Blackwell.
- Levins, R. (1968). Evolution in Changing Environments. Princeton: Princeton University Press.
- Lewontin, R. C. (1983). "The Organism as the Subject and Object of Evolution," reprinted in R. Levins and R. C. Lewontin, The Dialectical Biologist. Cambridge MA: Harvard University Press, 1985, pp. 85-106.
- Lodish, H., D. Baltimore, A. Berk, S. L. Zipursky, P. Matsudaira and J. Darnell (1995). Molecular Cell Biology, 3rd edition. New York: Freeman.
- McShea, D. (1991). "Complexity and Evolution: What Everybody Knows," Biology and Philosophy 6: 303-24.
- Menzel, C. R. (1997). "Primates' Knowledge of their Natural Habitat," in Whiten and Byrne (1997) pp. 207-239.
- Miller, G. (1997). "Protean Primates; The Evolution of Adaptive Unpredictability in Competition and Courtship," in Whiten and Byrne (1997) pp. 312-340.
- Moran, N. (1992). "The Evolutionary Maintenance of Alternative Phenotypes," American Naturalist 139: 971-89.
- Odling-Smee, F. J. (1988). "Niche-Constructing Phenotypes," in H. C. Plotkin, (ed.) (1988). The Role of Behavior in Evolution. Cambridge MA: MIT Press, pp. 73-132.
- Povinelli, D. J. and J. G. Cant (1995). "Arboreal Climbing and the Evolution of Self-Conception," Quarterly Review of Biology 70: 393-421.
- Roberts, W. A. (1998). Principles of Animal Cognition. Boston: McGraw Hill.
- Schlichting, C. and M. Pigliucci (1998). Phenotypic Evolution: A Reaction Norm Perspective. Sunderland MA: Sinauer.
- Seger, J. and J. Brockman (1987). "What is Bet-Hedging?" Oxford Surveys in Evolutionary Biology 4: 181-211.
- Silvertown, J. and D. Gordon (1989). "A Framework for Plant Behavior," Annual Review of Ecology and Systematics 20: 349-366.
- Silvertown, J. and J. Lovett Doust (1993). Introduction to Plant Population Biology. Oxford: Blackwell.
- Sober, E. (1994). "The Adaptive Advantage of Learning Versus A Priori Prejudice," in From a Biological Point of View. Cambridge: Cambridge University Press, pp. 50-70.
- Sterelny, K. (1995). "Basic Minds," Philosophical Perspectives 9: 251-270.
- Stich, S. P. and T. A. Warfield (eds.) (1994). Mental Representation: A Reader. Oxford: Blackwell.
- Sultan, S. (1987). "Evolutionary Implications of Phenotypic Plasticity in Plants," in Evolutionary Biology, Volume 21 (M. Hecht, B. Wallace and G.T. Prance, Eds.) New York: Plenum, pp. 127-78.

- Thinus-Blanc, C. (1988) "Animal Spatial Cognition," in L. Weiskrantz (ed.) Thought Without Language. Oxford: Clarendon Press, pp. 371-395.
- Todd, P. M. and G. F. Miller (1991). "Exploring Adaptive Agency II: Simulating the Evolution of Associative Learning," in J.-A. Meyer and S. W. Wilson (eds.), From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior. Cambridge MA: MIT Press, 1991, pp. 306-15.
- Tolman, E. C. (1948). "Cognitive Maps in Rats and Men," Psychological Review 55: 189-208.
- Tolman, E. C. and T. H. Honzik (1930). "'Insight' in Rats," University of California Publications in Psychology 4: 257-275.
- Whiten, A. and R. W. Byrne (eds.) (1997). Machiavellian Intelligence II: Extensions and Evaluations. Cambridge; Cambridge University Press.
- Wittgenstein, L. (1953). Philosophical Investigations. Translated by G. Anscombe. New York: Macmillan.