# G R O U N D S  F O R  C O G N I T I O N

## How Goal-Guided Behavior Shapes the Mind

R a d u  J.  B o g d a n

Erlbaum, Hillsdale, N.J. 1994

## Table of Contents

## Chapter 5: Primitive Semantics



## Chapter 6: Intimations of Re-presentation: Vision



## Chapter 7: Mental Guidance



## Chapter 8: Social Guidance

## PART THREE: IMPLICATIONS

# Chapter 1

# THE PROJECT

---

## 1. THE ARGUMENT

**Questions and Answers.** Why do organisms cognize? That is, why do they process and store data, form concepts, or solve problems? To obtain behaviorally relevant information about their environments. Why would they need such information? Because organisms must guide themselves to their goals. Why guidance to goal? Because goals must be satisfied, and for that, organisms must locate and identify their goals. But why would organisms have goals, in the first place? Because they are material systems or complexities that are genetically programmed to maintain and replicate themselves by engaging their worlds in goal-directed ways. That is the trick of life. These, essentially, are the questions this essay asks, and the answers it proposes, as prerequisites for understanding what cognition is and how it works. The order of the questions suggests the order of grounding and explanation. Goal-directedness grounds and explains guidance to goal, and the latter in turn grounds and explains the design and operation of cognition.

**Methodological Stance.** This construal of cognition relies on a certain methodological stance. We are not viewing the cognitive mind merely as a complex physical system reacting mechanically to stimuli. Life and cognition are not solely particles or molecules bouncing around under causal pressures. Life and cognition also display functional properties whose rationale and work cannot be reductively explained by their physical composition. The cognitive mind is a tool with a job to

do. The question is what is that job. The assumption behind this question is that a description of the job would reveal the nature and operation of the cognitive mind.

This is the methodological stance of evolutionary biology. It recommends, first, figuring out the program or modus operandi of an organ from a reconstruction of the task or job for which the organ was naturally selected, and then explaining the functional mechanism or architecture of the organ in terms of its program. The explanatory sequence, thus, goes from task to program to mechanism. (This, as we shall see, is also the explanatory policy adopted recently in the foundations of cognitive science.) Evolutionary biologists use this approach for every organ and capability. When they ask, for example, why skunks smell bad, the immediate causal answer is that skunks secrete chemical substances whose molecular structures result in bad smells. This answer bears on the program and mechanism involved. But when the question turns to why skunks smell bad in the first place, the inquiry must probe the deeper and historically more distant reason for having smelling programs and mechanisms at all. That reason defines the task those programs and mechanisms had evolved to handle. The answer is that smelling bad is a good defense, and natural selection favored those with the worst smells, which survived to produce the most babies. The specific molecular structure of the bad-smelling mechanisms is an accident; any other chemical structure, or any other defense mechanism, programmed in some other way to spread bad smells could have had the same results by handling the required defense task.

This is the familiar distinction between proximate and ultimate explanation. Proximate explanations are about mechanisms, their programs, and the surrounding external conditions (inputs, contexts) of operation. I lump all these factors under the notion of proximate arrangement. Proximate

explanation, then, is about proximate arrangements, as it proceeds from functional mechanisms and programs (as explanantia), and their contexts of operation as boundary conditions, to particular organ performances (as explananda).

Ultimate explanations are about the evolutionary shapers that had configured the proximate causes (functional mechanisms and their programs) in the first place. The direction of the ultimate explanation is from evolutionary shapers (genetic variations, natural selection, guidance to goal) to the tasks or jobs to be accomplished and then to the programs carrying out the tasks and the functional mechanisms running the programs in specific proximate arrangements.

Proximate explanations are causally or functionally subsumptive, ultimate explanations typically reconstructive. Our first explanation of the skunk's smell was proximate, the second ultimate. In the psychological domain, the explanation of an inference or behavior in terms of some program is proximate, whereas an evolutionary explanation of why the program was first naturally selected, for what task, is ultimate.

The Genetic Grounding of Goal-directedness. Our teleological approach takes goal-directedness as the ultimate evolutionary reason for guidance to goal, and the latter as the ultimate evolutionary shaper of the cognitive mind. Guidance to goal defines the information tasks that cognitive programs and their mechanisms had evolved to handle. This is the overall direction of our project. For the project to take off, it must first establish that goal-directedness itself, the ultimate explainer, is a respectable and scientifically intelligible material phenomenon in the world, a property of biomatter, indeed, the key property of life. This is the aim of chapter 2. There I argue that goal-directedness characterizes the generic form in which the genes program the design and behavior of living beings.

Organisms are genetically primed to pursue goals.

Although natural selection eventually decides which goal-pursuing policies work, as adaptations, it is goal-directedness that ultimately explains the most general and systematic pattern that all adaptations, including the cognitive ones, display in their relations to environments. As Lynn Margulis is reported to have said, so very well, natural selection is the editor, not the author. The author is the genetic code. Chapter 2 shows that, deep down, the genetic writing is teleological. For, when we set aside the richness of variations spawned by evolution, goal-directedness appears as the fundamental property of life, the systematically recurrent theme that underlies all the variations.

From Teleology to Cognition. The next step in our project is to bring this genetically grounded teleology to bear on our understanding of cognition. Its aim is to connect goal-directedness (the evolutionary reason) to cognition (the family of programs) by way of guidance to goal (the family of tasks the programs had evolved to execute). We need, therefore, an analysis that effects the transition from the ultimate to the proximate explanation of cognition. The basic insight behind the analysis is the following. We noted that goal-directedness captures the general patterns of all organisms-environments interactions. Guidance to goal is one family of such patterns, those characterizing the informational transactions between organisms and environments. Cognition is the overall adaptation (or family of adaptations) evolved to instantiate such patterns of informational transactions.

To understand the programs of cognition, and their proximate operation, is to understand that the information tasks facing these programs eventuate in guidance to goal. For guidance to goal describes the systematically recurrent pattern that all cognitive adaptations display in their informational relations to their environments. The reason for this simple,

fundamental, but often overlooked truth is that to pursue goals, organisms evolve ways to guide themselves to goals, and to do that, they come to exploit systematically pervasive and recurrent patterns of information relations available to them. The exploitable information patterns (as tasks) in turn select for appropriate (adaptive) cognitive programs and functional mechanisms. How this selection works and what are its real life outcomes is the business of evolutionary psychology. Although our discussion will often handwaive at cognitive adaptations, its direction is more ultimate and concerns guidance to goal and its information tasks.

To effect the proposed transition from goal-directedness to cognition by way of guidance to goal and the information relations that secure it, we need a suitable methodological framework. This is the object of chapter 3. Guidance to goal defines the sort of knowledge that goal-directed systems or organisms need to manage in their worlds. Guidance to goal is secured by patterns of information relations that must be in place among an organism, its goals, its behaviors, and its ecology. Finding the goal-guiding patterns of information relations constitutes the information tasks facing an organism's cognition. The environmental opportunities exploited by the organism and the cognitive programs that do the exploiting can be regarded as the executors of the information tasks. Inside the organism, the program execution is implemented by appropriate functional mechanisms (the operating system). The grounding and explanatory order thus goes, top-down, from goal-directedness to guidance to goal to information tasks to cognitive programs (and ecological opportunities) to functional mechanisms. We can call it the KICM method of analysis: from Knowledge to Information to Cognition to Mechanism.

Forms of Guidance. We have noted a paragraph ago that our focus will be on the first two, and most ultimate,

components of the KICM framework, knowledge as guidance to goal, and the information patterns that make it possible. Given the account of knowledge and information offered in chapter 3, the next questions we want to ask is how many and what sort of forms of guidance are possible, and what it takes informationally (i.e., what tasks must be executed in what configurations) to effect guidance in each of these forms. Or, put more plainly, the question is how many ways are there to spot and track a goal. These are not bioevolutionary questions about specific cognitive adaptations (programs and mechanisms) that had evolved in specific species. Our questions are more fundamental and speculative. And their answers, suggested in chapters 4 through 8, sketch the teleoevolutionary profiles of possible forms of guidance. These profiles offer theoretical guidelines and heuristic hints to the more proximately and empirically oriented cognitive sciences whose job is to confirm (or not) which of the suggested forms of guidance and types of information tasks got embodied in the real biological (or computer) world, in what ways, and why.

With that empirical confirmation, the information tasks that passed the test of natural selection can then be mapped into causal/functional analyses of the successful cognitive programs and their mechanisms. The latter, in turn, can proximately explain particular program applications. Our project does not go so far and proximately. How the forms of guidance and their information tasks are empirically fleshed out by natural selection theories and specific cognitive sciences is a matter well covered in the recent literature, and about which I will have little to say here.

The following are the principled forms of guidance that we are going to examine. When goals are gross, diffuse and normally proximal, and the behaviors directed at them simple, reflex, and indiscriminate, the guidance is managed by simple information tasks whose execution relies heavily, rigidly, and

almost exclusively on laws of nature and ecological cooperation. Guidance, in this case, is parasitic on the ways of nature, with little functional initiative on the part of the organism's cognition. This is the teleonomic guidance of chapter 4. Finely individualized goals, revealed by specific and distal properties, and satisfied by complex and delicate behaviors, normally require a systematic targeting by means of discriminating information tasks able to triangulate the goal-revealing properties from the internal encodings left by the sensory inputs. This is teleosemantic guidance. When the triangulation in question is primitive, in the sense that the semantic information tasks range over simple "signs" or "presentations" of external correlations, and are not executed in terms of simpler semantic information tasks, we have guidance by primitive semantics (chapter 5). A semantics that secures guidance thorough a flexible and combinatorially versatile triangulation of goal-revealing properties from internal signs or presentations of other such signs or presentations, with the effect that complex semantic tasks are executed in terms of simpler such tasks, is re-presentational. Human vision appears to be an instance of the transition from a primitive to a re-presentational semantics, as chapter 6 argues.

Together with other program developments, re-presentation in turn facilitates a massive internalization of guidance through conceptualization and the mental ability to model goal situations. So argues chapter 7. Besides being goal-directed, re-presentational systems that model goal situations are also goal-directors. They create goals by means of desires and plans, and anticipate their guidance to these goals mentally. Chapter 8 brings in, briefly, the social form of guidance. Mental modelers of goal situations, like ourselves, can spontaneously conceive of each other as goal-directed agents, pursue their goals by knowing those of others, and use each other as sources of information in guidance and as tools

in the pursuit of goals. In so doing, such goal directors and modelers develop a conceptual framework in which they think of each other and figure each other out. That is our commonsense psychology.

Reasons for Speculation at the Top. Although in the discussion and particularly the illustrations of each form of guidance, we will say a few (tentative) words about the real life programs and mechanisms likely to instantiate them, the emphasis is consistently on the two top levels of analysis, guidance and the information relations that make it possible.

There are two good reasons for this division of labor, with guidance and information, on one side, and cognitive programs and mechanisms, on the other side, and for the speculative exercise of profiling possible forms of guidance solely in terms of (theoretically) plausible configurations of information relations. The first reason, already noted, has to do with the ultimate-proximate distinction and the resulting top-down analysis now prominent both in evolutionary biology and cognitive science. It is almost a logical point that if the proximate or the down is ultimately explained by the ultimate or the top, then, given sufficiently realistic though not necessarily biological constraints, the possible scenarios of guidance through information patterns envisaged at the top levels of analysis must have a bearing on the particular, often deviant and messy program and mechanism instantiations in the real world. The converse perhaps makes the point better. If the speculation on possible forms of guidance had no bearing on understanding real cognitive adaptations, then it would be hard to see the point and usefulness of the ultimate-proximate or top-down analyses. Indeed, one would expect those who deny the antecedent to also deny the consequent.

The second reason for playing our theoretical game at the speculative heights of guidance and information, far from their

biological instances, is that the cognitive (and all other biological) adaptations installed by natural selection are provincial relative to how I construe goal-directedness and guidance to goal. Biological life is not the only form of life, and biocognition not the only form of cognition. This is true in principle and may become an empirical truth in a few decades.

Although I take the genetics of goal-directedness very seriously, I regard it merely as a biological expression of the more general phenomenon of goal-directed self-organization and self-reproduction. Although our essay does not explore this more general phenomenon, I think that its scientific analysis would one day ground biogenetic goal-directedness, as one of many forms of goal-directedness displayed by self-organizing and self-reproducing systems. So, while taking goal-directedness seriously, I am liberal about its worldly and specifically biological incarnations. This is why my teleoevolutionary stance is generic, does not entail biology, and is necessarily hybrid by combining bioevolution with abstract speculations about alternative forms of goal-directedness.

The reason for my position is simple. There is an objective sense in which goal-directedness obtains whenever certain teleological constraints (specified in chapter 2, section 5) are met by some form of matter or another. Goal-directedness need not be intrinsic to any particular form of matter, and is not necessarily biochemical. If other forms of matter (mineral, plasmatic, robotic, whatever) in functional configurations other than biological (e.g., synthesized on a computer, embodied in a robot, or manufactured in some other artificial manner) run means/ends programs that maintain and replicate an internal structure in some environment, then, according to the argument of chapter 2, these forms display goal-directedness. This means that the bioevolutionary constraints on DNA-run life on earth need not be the same as the teleoevolutionary constraints operating on other forms of goal-directed life.

This angle on goal-directedness explains why the forms of guidance and the information tasks examined in this work are not intended to match real biological taxonomies and at times may even fail to have any biological counterparts. This is all right, as long as we recall that the biological taxa are but one version of goal-directedness among the multiplicity of versions abstractly contemplated by our story. This approach is not new. Decades ago, general system theorists and cyberneticians made similar proposals, and some (Wiener 1948; 1950) grounded teleology in a comparative analysis of brains and servomechanisms. Without the teleology, I find the same spirit in Braintenberg's (1984) splendid essay on synthetic psychology, and I expect the von Neumann-inspired views of artificial life to rehabilitate and further develop these early and good insights (Levy 1992).

This generic and hybrid stance on goal-directedness should also allay the misgivings of those who are skeptical of evolutionary accounts of cognition and of top-down reconstructions that derive biological features, including real cognitive programs, from rational models (Lewontin 1990; also the discussion in Pinker and Bloom 1990). These skepticisms draw on the idea that bioevolution is too messy and opportunistic, biocognition too complex, and its rational models too rarefied and abstract to allow for a neat and enlightening fit. I agree. This essay does not offer an evolutionary account of biocognition, nor does it reconstruct biocognition from rational models. Although my project is inspired by genetics and bioevolution, its focus is on guidance and its possible forms and implications, not on biological embodiments.

Evolutionary teleology, I said, need not be only biological, and evolution need not reduce to bioevolution. To the extent to which cognitive science, generally, and psychology, in particular, concern themselves with the information tasks and

cognitive programs involved in guidance to goal, it follows that neither can be reduced to, or exclusively grounded in, biology.

## 2. CRITICAL TARGETS

Most of this essay is dedicated to the constructive argument just previewed. Occasional criticisms and exegetical analyses intrude (as in chapter 6) solely to support the ongoing argument. Yet throughout this work, and notably in its third part, there are two critical and related targets that are pursued and shot at systematically. Dominant in current cognitive science and philosophy of mind, these targets are major obstacles to a teleological understanding of cognition. Let me introduce the first target, and the most relevant to our project, by contrasting the standard and the teleological styles of top-down analysis of cognition.

**Psychosemanticism.** Knowledge, almost everybody agrees, is information about the world on which organisms act. An ancient, popular, and implausible view that has become standard sees knowledge in two dimensions. One dimension is the hook-up relation to the world, the other the causal or functional impact of this relation, once cognized, upon behavior. To simplify, I will label the first dimension in terms of its most prominent version, the semantic form of hook-up, although the force of my critique extends to any form of hook-up, whether (strictly) semantic or nonsemantic. So construed, then, the semantic dimension takes the job of cognition to be about, or indicate, or refer to, or reflect, or picture, or represent, or covary with, aspects of the world of interest to the organism and its behavior; the causal dimension indicates that, once encoded and processed, the information is causally/functionally efficacious in virtue of its hook-up relation

to the world, of what it represents. This is to say that the forms in which information is encoded and processed, as data, have causal or functional efficacy in virtue of how the data hook up with the world.

Going top-down now: to yield semantic knowledge, cognition must provide the organism with information that causes behavior in virtue of what it represents. Cognition causally converts semantic information into behavior; the semantics runs the psyche. The form and function of the cognitive programs are psychosemantic since they reflect this semantics-to-causation conversion. The operating system running such programs must be a semantic mind. Evolution, therefore, must have installed a cognitive mind whose rationale and role is to map semantic information into behaviors. The view committed to this top-down analysis and its implications is psychosemanticism. It is the main and explicit villain of our story.

The spirit of psychosemanticism is aptly encapsulated by Ramsey's notorious metaphor of belief (cognition, generally) as a map by which we steer. This metaphor begins to tell us what is right and wrong with psychosemanticism. It is right that maps have the job of causally steering travel in virtue of what they represent, just as cognitive programs have the job of causally steering behavior in virtue of the information they make available to the organism. Right but incomplete. Steering is not by maps only. Nor does a map steer merely in virtue of what it represents. A map must be used to yield information, and the context and purpose of its use make a vast difference to the steering. The map user has access to information that the map does not provide (starting points, destinations, difficulty of access, various other values, and so much more).

The use of a map relies on and exploits but does not represent the contexts, assumptions, limitations, and choices associated with its use. The same, I will argue, is true of

guidance and cognition. The guiding knowledge is the information that gets us from here (current condition) to there (goal) by means and under assumptions that are not represented (cognized) at all, while what is represented, or generally linked up with, often has little if anything to do with goals pursued. The information cognized is only part of the total information that guides to goals, and the semantic information itself (when available) often is only part of the information cognized. A semantic mind would be a very narrow and unadapted mind.

My criticisms of psychosemanticism surface often as we proceed with our argument but are brought together in a focused and integrated form in chapter 9. The limitations of psychosemanticism are symptomatic of too narrow a view of the tasks of cognition, and of the resulting failure to see the wider arrangements in which, and the deeper reasons for which, organisms have knowledge. This narrow view is encouraged not only by the standard account of knowledge.

Psychosemanticism is a symptom of a deeper conceptual ailment. Which brings us to our second critical target, which I call psychological Newtonianism. Unlike its psychosemantic version, which bears the brunt of our critique, psychological Newtonianism is mostly in the background, the implicit villain. Whereas psychosemanticism obscures the phenomenon of guidance to goal and thus the teleological determination of cognition, deeper down in the order of theorizing, psychological Newtonianism is inimical to the more fundamental phenomenon of goal-directedness itself.

Psychological Newtonianism. This is the view that psychology ought to be a sort of physics functionalized. It is a view animated and held together by several doctrines: that the cognitive mind is a complex body subject to external and internal causal interactions (physicalism) that had been

functionalized (functionalism) either by accident or deliberate design or natural selection (historicism). Physicalism assigns the mind to the physical order of nature, whereas functionalism redescribes the causal work of the cognitive programs as patterns of input-internal states-output functions. In the domain of cognition, psychosemanticism legitimizes these functions as mappings of semantic relations into cognitive and behavioral causation. Historicism is a recent and not always welcome addition to psychological Newtonianism. When adopted, typically in the form of natural selection, historicism tells us which task-specific psychosemantic functions had prevailed over time and why. The explanatory policy of psychological Newtonianism, both in the ultimate and proximate domains, is that of causal/functional explanation: ultimately, natural selection causally shapes the mind's programs; proximally, their psychosemantic applications can be explained functionally, in terms of the rules of the programs and the causal push of the sensory inputs.

To follow the contours of our critique of psychosemanticism and psychonewtonianism, we should mark the areas where teleology does not make a difference to our view of cognition, and so where nonteleological accounts are legitimate and useful. One such area is that of the proximate analyses of dedicated cognitive programs (vision, language, sensory-motor) which handle domain-specific tasks, often psychosemantic nature. Teleology is also unimportant in the evolutionary psychology that explains more ultimately, by natural selection, such tasks and programs. In neither of these areas does teleology make a difference, not because goal-directedness is not present (it is), but because what is examined are its partial embodiments that can be explained ultimately and causally by natural selection, and proximately by the causal/functional work of dedicated programs. It will be a constant tactics of my argument (i) to note the mutually

reinforcing links between psychosemanticism, domain-specific tasks, and dedicated programs, on the one hand, and their ultimate yet causal explanation by natural selection, on the other hand; (ii) to acknowledge the validity of these links within definite (normally, modular) limits; but (iii) challenge the links outside those limits, where teleology makes a difference, both ultimately and proximately.

The Teleological Difference. As the preview of our argument has indicated, there are two areas of explanation where teleology does make a substantial difference, and where, therefore, the conflict with psychological Newtonianism and psychosemanticism is bound to be the sharpest. These areas form the object of this essay. One concerns the grounding and ultimate explanation of cognition in terms of guidance to goal, and the grounding and ultimate explanation of the latter in terms of goal-directedness. In this area psychological Newtonianism fails to recognize the teleological patterns that the genes (or other shapers of life) cast around to insert their products (organs, functional capabilities, behaviors) and thus pursue their maintenance and replication policies. This is why psychological Newtonianism looks like a reform Cartesianism that views the cognitive mind as a functionally autonomous unit whose design and operation owes nothing essential to how it engages the world and why. The historical appeal to natural selection (when attempted) probes only the pedigree of individual and task-specific functions, and never inquires into the larger picture that gives reason and rhyme to the overall interplay of these functions.

The second and proximate area of explanation where teleology matters (literally) and where psychological Newtonianism and its psychosemantic cousin again fail is thinking. In simpler forms of guidance, the cognitive component is typically a rigid cog in a larger arrangement, and instantiates

evolutionarily preassigned and fixed functional paths that lawfully connect with inputs, ecological accidents and natural regularities, to provide guidance to goal. By contrast, in mental guidance, the cognitive component itself, in the form of thinking, is intrinsically teleological since its design has evolved to posit and flexibly script its own goal models before action is initiated. The intrinsic teleology of thinking simply eludes the causal/functional style of explanation, so dear to psychological Newtonianism and psychosemanticism, because, like any goal-directed activity, from metabolism to escape behavior, mental goal scripting is systematically and irremediably underdetermined by its causal and functional implementations. So argues chapter 10.

## 3. RESPECTS AND DEMARCATIONS

This book approaches the cognitive mind in terms of its natural properties and evolutionary history. We can call such an approach mind naturalization. To naturalize the mind is to determine its origin, place and role in nature in the metaphysical terms of materialism (everything exists in space and time, as some form of matter, at some level of complexity), the logical terms of noncircularity (the analysis should not employ the notions to be naturalized), the methodological terms of intelligibility and explanatoriness (the analysis should make intelligible and explain the properties of its object), and the ideological terms of scientific respectability (the analysis must be compatible with science).

In recent years a good number of works have significantly advanced the cause of mind naturalization: Churchland (1989), Cummins (1983; 1989), Dretske (1981; 1986; 1988), Fodor (1975; 1983; 1987; 1990a), Pylyshyn (1984), Quine (1960), Searle (1983), Stich (1983), as well as influential papers by

Davidson, Putnam, and many others. Many of these works are committed to a proximate and often psychosemantic mind naturalization. A more ultimate and natural selection-based mind naturalization, with faint teleological echoes, has been attempted most notably by Dennett (1969; 1978; 1987; 1991a) and Millikan (1984; 1986;1989), and anticipated in earlier and important works of evolutionary epistemologists like Piaget, Popper, Campbell, and of cyberneticians like Wiener, Ross Ashby, Somerhoff, and others. I have learned very much from these works and developed my views in constant and fruitful dialogue with them.

It has been widely assumed that mind naturalization is incompatible with teleology. This essay shows not only that there is no such incompatibility but also that the mind naturalization project cannot be completed without teleology. In pushing for this line of argument, I have benefited considerably from a number of works on mind naturalization that have shown a lively and systematic interest in teleology. Among them I would list those of Bennett (1976; 1991), Collins (1984; 1987), Lycan (1981; 1989), Matthen and Levi (1984), Matthen (1988), McGinn (1989), and Papineau (1984; 1987). My project is intended to fortify and extend this common enterprise. Yet there are differences, often important. I find some of these teleological analyses too indebted to the psychosemanticist agenda. Instead of being primarily puzzled (as most teleoevolutionary psychosemanticists are) by the narrower question of how cognitive programs link up with the world in ways that steer behavior, I am more puzzled, antecedently, by the deeper question of how self-organizing systems such as organisms enter in complex, dynamic, and evolving arrangements with their natural and social ecologies, and manage to settle into ecological-cognitive patterns of informational transactions that instantiate guidance to goal. Also, unlike most of my teleological colleagues, I take genetics

very seriously and use it to ground goal-directedness, while at the same time remaining open to forms of goal-directedness which are not biological.

Finally, I hope that the novel methodological approach to goal-directedness and guidance taken in this essay can bring teleology in line and in communication with the top-down analyses practiced in evolutionary biology, the emerging theory of self-organized and self-replicating systems, and cognitive science. If I were to capture the methodological spirit of my enterprise in a simple formula, I will choose (with some trepidation) the inelegant notion of transcendental naturalism -- in a modest and unusual form. Kant's notorious transcendental question was, given that we have knowledge, how is it possible? To simplify enormously, Kant's view was that knowledge results from, and encompasses, the application of concepts to sensory experiences. So, when he asks his transcendental question about knowledge, his interest is in what the human mind must be like in order to know what it does.

I share neither the narrow (psychological) scope of Kant's question nor the idealist and apriori answer to it. For my notion of knowledge is so much different from his. I think of knowledge in terms of guidance to goal by way of information. Kant's notion covers only the cognitive leg of the guidance journey; and Kant's notion is exclusively human, in most respects, while mine isn't. I would go even further, and say that Kant and those who followed him have not been radical enough in their transcendentalism. The most radical question is not what the human mind must be like, in order for it to cognize as it does and know what it does. For me, a deeper foundational question, the question of this essay, is how knowledge and cognition are possible at all, and why, in the larger scheme of matter and evolutionary history, so that simple and then

complex minds evolve to embody various forms of cognition and have various forms of knowledge.

Since I propose a more comprehensive story of knowledge, I also have a wider framework in which to examine its possibility. That examination will be modestly transcendental, as opposed to empirical or historical. For I think that the notions of guidance to goal and information tasks (not just cognitive tasks, for that would be more like Kant's angle) explicate the general conditions in which knowledge and cognition are evolutionarily possible, and even suggest some general constraints on, and forms of, their possibility. Therefore, prior to the empirical details of cognitive science and the historical details of evolutionary biology, the notions of guidance to goal and information tasks, which will be the key notions of our analysis, characterize the logic of the cognitive situation in general -- indeed, the very conditions for the possibility of cognition.

A good deal of work in cognitive science, artificial intelligence, artificial life, evolutionary biology and the design of self-maintaining and self-replicating systems is modestly transcendental, in my sense. People in these fields ask what a system (biological, robotic, whatever) must be like to interact successfully with the world; what its interaction tasks must be like, to account for this success; and what the executing programs must be like, to carry out the tasks. The transcendental stance is compatible with, and can be very useful to, scientific naturalism in its many forms. I think that stance needs a place in mind naturalization. It is also great fun.

The language of this work is going to be neither ordinary nor elegant. It is largely technical. Its basic vocabulary is made of notions such as guidance, information task, program, assumptions, goal situations, and others. Most of our analysis is in such terms. To help keep track of these notions and the places where they are first explicated, the reader may consult a

short glossary at the end of the book. This terminological decision was not taken lightly. There are several reasons for it.

One is that we are going to cross many disciplinary boundaries, with their own vocabularies and habits of thinking and explaining. A single frame of terminological and conceptual reference can only help. Another reason is that in philosophy, a discipline with a very long and unforgiving memory, such concepts as knowledge, belief, goal, intentionality, representation, and the like, carry associations and implications that one has to track and fight for, or against, constantly. This is not what I want to do here. As a result, I have either abandoned these standard philosophical notions or else (as in the case of knowledge or goal-directedness) re-explicated them in terms of my technical vocabulary.

The final reason has to do with commonsense psychology, which implicitly defines our understanding of such widely used notions as belief, desire, goal, thinking, and the like. In recent years, grand doctrines and research programs analyze the commonsense wisdom for what it can tell us about the mind, and enlist it in the service of scientific theorizing, or else try to prove it false of the mind and hence detrimental to cognitive science. I happen to think (and argue a little in chapter 8 and other works) that these are the wrong ways to look at commonsense psychology, and that, in particular, its notions have little to say about the matters discussed in this essay.

Pending a clarification of what the commonsense notions are all about, their uncritical use could only obscure our inquiry here. So, again, I have tried either to avoid using the commonsense notions or, as in the case of some, such as goal, desire or belief, redefine them in our technical vocabulary. I can only hope that what is lost in agreeableness and familiarity is compensated by some uniform precision and focus.