Chapter 6

Towards a new orientation

The previous chapters have dealt with fundamental questions of what it means to exist as a human being, capable of thought and language. Our discourse concerning these questions grew out of seeing their direct relevance to our understanding of computers and the possibilities for the design of new computer technology. We do not have the pretension of creating a grand philosophical synthesis in which Maturana, Heidegger, Gadamer, Austin, Searle, and others all find a niche. The importance of their work lies in its potential for unconcealing the rationalistic tradition in which we are already immersed. Their unity lies in the elements of the tradition that they challenge, and thereby reveal.

As background to our study of computers and programming in Part II, this section summarizes the concerns raised in previous chapters, pointing out their areas of overlap and the role they play in our detailed examination of computer technology and design.

6.1 Cognition and being in the world

This book has the word 'cognition' in its title, and in the previous chapters we have given accounts of cognitive activity. But in using the term 'cognition' we fall into the danger of implicitly following the tradition that we are challenging. In labelling it as a distinct function like 'respiration' or 'locomotion,' we evoke an understanding that an activity of 'cognizing' can be separated from the rest of the activity of the organism. We need first to examine this understanding more carefully and to recognize its consequences for design.

In speaking of thinking as a kind of activity, we adopt a common preunderstanding that seems so obvious as to be unarguable. When you sit at your desk deciding where to go for lunch, it seems clear that you are engaged in 'thinking,' as opposed to other things you might be doing at the time. This activity can be characterized in terms of mental states and mental operations. An explanation of how it is carried out will be couched in terms of sentences and representations, concepts, and ideas. This kind of detached reflection is obviously a part of what people do. The blindness of the rationalistic tradition lies in assuming that it can serve as a basis for understanding the full range of what we might call 'cognition.' Each of the previous three chapters challenges this assumption.

One of the most fundamental aspects of Heidegger's discourse is his emphasis on the state of thrownness as a condition of being-in-the-world. We do at times engage in conscious reflection and systematic thought, but these are secondary to the pre-reflective experience of being thrown in a situation in which we are always already acting. We are always engaged in acting within a situation, without the opportunity to fully disengage ourselves and function as detached observers. Even what we call 'disengagement' occurs within thrownness: we do not escape our thrownness, but shift our domain of concern. Our acts always happen within thrownness and cannot be understood as the results of a process (conscious or non-conscious) of representing, planning, and reasoning.

Heidegger argues that our being-in-the-world is not a detached reflection on the external world as present-at-hand, but exists in the readiness-to-hand of the world as it is unconcealed in our actions. Maturana, through his examination of biological systems, arrived in a different way at a remarkably similar understanding. He states that our ability to function as observers is generated from our functioning as structure-determined systems, shaped by structural coupling. Every organism is engaged in a pattern of activity that is triggered by changes in its medium, and that has the potential to change the structure of the organism (and hence to change its future behavior).

Both authors recognize and analyze the phenomena that have generated our naive view of the connection between thinking and acting, and both argue that we must go beyond this view if we want to understand the nature of cognition—cognition viewed not as activity in some mental realm, but as a pattern of behavior that is relevant to the functioning of the person or organism in its world.

When we look at computer technology, this basic point guides our understanding in several ways. First, it is critical in our anticipation of the kinds of computer tools that will be useful. In a tradition that emphasizes thought as an independent activity, we will tend to design systems to work within that domain. In fact much of the current advertising rhetoric about

computers stresses the role they will play in 'applying knowledge' and 'making decisions.' If, on the other hand, we take action as primary, we will ask how computers can play a role in the kinds of actions that make up our lives—particularly the communicative acts that create requests and commitments and that serve to link us to others. The discussion of word processors in Chapter 1 (which pointed out the computer's role in a network of equipment and social interactions) illustrates how we can gain a new perspective on already existing systems and shape the direction of future ones.

We also want to better understand how people use computers. The rationalistic tradition emphasizes the role played by analytical understanding and reasoning in the process of interacting with our world, including our tools. Heidegger and Maturana, in their own ways, point to the importance of readiness-to-hand (structural coupling) and the ways in which objects and properties come into existence when there is an unreadiness or breakdown in that coupling. From this standpoint, the designer of a computer tool must work in the domain generated by the space of potential breakdowns. The current emphasis on creating 'user-friendly' computers is an expression of the implicit recognition that earlier systems were not designed with this domain sufficiently in mind. A good deal of wisdom has been gained through experience in the practical design of systems, and one of our goals is to provide a clearer theoretical foundation on which to base system design. We will come back to this issue in our discussion of design in Chapter 12.

Finally, our orientation to cognition and action has a substantial impact on the way we understand computer programs that are characterized by their designers as 'thinking' and 'making decisions.' The fact that such labels can be applied seriously at all is a reflection of the rationalistic tradition. In Chapters 8 through 10, we will examine work in artificial intelligence, arguing that the current popular discourse on questions like "Can computers think?" needs to be reoriented.

6.2 Knowledge and representation

Our understanding of being is closely linked to our understanding of knowledge. The question of what it means to know is one of the oldest and most central issues of philosophy, and one that is at the heart of Heidegger's challenge. Chapter 2 described a 'naive realism' that is prominent within the rationalistic tradition. As we pointed out there, this is not a logical consequence of the tradition (and is not accepted by all philosophers within it), but it is part of the pervasive background that follows the tradition in our everyday understanding.

At its simplest, the rationalistic view accepts the existence of an objective reality, made up of things bearing properties and entering into relations. A cognitive being 'gathers information' about those things and builds up a 'mental model' which will be in some respects correct (a faithful representation of reality) and in other respects incorrect. Knowledge is a storehouse of representations, which can be called upon for use in reasoning and which can be translated into language. Thinking is a process of manipulating representations.

This naive ontology and epistemology is one of the central issues for both Maturana and Heidegger. Neither of them accepts the existence of 'things' that are the bearers of properties independently of interpretation. They argue that we can not talk coherently of an 'external' world, but are always concerned with interpretation. Maturana describes the nervous system as closed, and argues against the appropriateness of terms like 'perception' and 'information.' Heidegger begins with being-in-the-world, observing that present-at-hand objects emerge from a more fundamental state of being in which readiness-to-hand does not distinguish objects or properties.

For Heidegger, 'things' emerge in breakdown, when unreadiness-to-hand unconceals them as a matter of concern. Maturana sees the presence of objects and properties as relevant only in a domain of distinctions made by an observer. In the domain of biological mechanism they do not exist. Both authors recognize that we are situated in a world that is not of our own making. Their central insight is that this world, constituted as a world of objects and properties, arises only in the concernful activities of the person.

Maturana and Heidegger both oppose the assumption that cognition is based on the manipulation of mental models or representations of the world, although they do so on very different grounds. Maturana begins as a biologist, examining the workings of the nervous system. He argues that while there is a domain of description (the cognitive domain) in which it is appropriate to talk about the correspondence between effective behavior and the structure of the medium in which it takes place, we must not confuse this domain of description with the domain of structural (biological) mechanisms that operate to produce behavior. In saying that a representation is present in the nervous system, we are indulging in misplaced concreteness, and can easily be led into fruitless quests for the corresponding mechanisms. While the point is obvious in cases of reflex behavior like the frog and fly of his early research, Maturana sees it as central to our understanding of all behavior, including complex cognitive and linguistic activities.

Heidegger makes a more radical critique, questioning the distinction between a conscious, reflective, knowing 'subject' and a separable 'object.'

He sees representation as a derivative phenomenon, which occurs only when there is a breaking down of our concernful action. Knowledge lies in the being that situates us in the world, not in a reflective representation.

Chapter 2 described efforts being made to create a unified 'cognitive science,' concerned with cognition in people, animals, and machines. To the extent that there is intellectual unity in this quest, it centers around some form of the *representation hypothesis:* the assumption that cognition rests on the manipulation of symbolic representations that can be understood as referring to objects and properties in the world.[1]

When we turn to a careful examination of computer systems in Chapter 7, we will see that the corresponding representation hypothesis is not only true but is the key to understanding how such systems operate. The essence of computation lies in the correspondence between the manipulation of formal tokens and the attribution of a meaning to those tokens as representing elements in worlds of some kind. Explicit concern with representation is one of the criteria often used in distinguishing artificial intelligence from other areas of computer science.

The question of knowledge and representation is central to the design of computer-based devices intended as tools for 'knowledge amplification.' We may seek to devise means of manipulating knowledge, in the sense that a word processor allows us to manipulate text. We might attempt to build systems that 'apply knowledge' towards some desired end. In this effort, our choice of problems and solutions will be strongly affected by our overall understanding of what knowledge is and how it is used. Many of the expert systems being developed in 'knowledge engineering' research are based on a straightforward acceptance of the representation hypothesis. In Chapter 10 we will describe these efforts and their limitations, and characterize the kinds of *systematic domains* that can be successfully treated in representational terms.

6.3 Pre-understanding and background

Chapter 3 emphasized that our openness to experience is grounded in a pre-understanding without which understanding itself would not be possible. An individual's pre-understanding is a result of experience within a tradition. Everything we say is said against the background of that experience and tradition, and makes sense only with respect to it. Language (as well as other meaningful actions) need express only what is not obvious, and can occur only between individuals who share to a large degree the same background. Knowledge is *always* the result of interpretation, which depends on the entire previous experience of the interpreter and on situatedness in a tradition. It is neither 'subjective' (particular to the individual) nor 'objective' (independent of the individual).

Maturana describes a closely related phenomenon in explaining how the previous structure of the system defines its domain of perturbations. The organism does not exist in an externally defined space. Its history of structural coupling generates a continually changing space of possible perturbations that will select among its states. Interacting systems engage in mutual structural coupling, in which the structure of each one plays a role in selecting among the perturbations (and hence the sequence of structures) of the others.

Our presentation of speech act theory has also emphasized the role of background and interpretation, while retaining a central focus on the commitment engendered by language acts. In this we move away from the individual-centered approach of looking at the mental state (intentions) of speaker and hearer, describing instead the patterns of interaction that occur within a shared background. As we will show in detail in Chapter 12, the pervasive importance of shared background has major consequences for the design of computer systems.

Artificial intelligence is an attempt to build a full account of human cognition into a formal system (a computer program). The computer operates with a background only to the extent that the background is articulated and embodied in its programs. But the articulation of the unspoken is a never-ending process. In order to describe our pre-understanding, we must do it in a language and a background that itself reflects a pre-understanding. The effort of articulation is important and useful, but it can never be complete.

This limitation on the possibility of articulation also affects more concrete issues in designing computer tools. If we begin with the implicit or explicit goal of producing an objective, background-free language for interacting with a computer system, then we must limit our domain to those areas in which the articulation can be complete (for the given purposes). This is possible, but not for the wide range of purposes to which computers are applied. Many of the problems that are popularly attributed to 'computerization' are the result of forcing our interactions into the narrow mold provided by a limited formalized domain.

At the other extreme lies the attempt to build systems that allow us to interact as though we were conversing with another person who shares our background. The result can easily be confusion and frustration, when breakdowns reveal the complex ways in which the computer fails to meet our unspoken assumptions about how we will be understood. The goal of creating computers that understand natural language must be reinterpreted (as we will argue in Chapter 9) in light of this. We must be

[1] This assumption, which has also been called the *physical symbol system hypothesis*, is discussed at length in Chapter 8.

especially careful in dealing with so-called 'expert systems.' The ideal of an objectively knowledgeable expert must be replaced with a recognition of the importance of background. This can lead to the design of tools that facilitate a dialog of evolving understanding among a knowledgeable community.

6.4 Language and action

Popular accounts of language often portray it as a means of communication by which information is passed from one person (or machine) to another. An important consequence of the critique presented in the first part of this book is that language cannot be understood as the transmission of information.

Language is a form of human social action, directed towards the creation of what Maturana calls 'mutual orientation.' This orientation is not grounded in a correspondence between language and the world, but exists as a consensual domain—as interlinked patterns of activity. The shift from language as description to language as action is the basis of speech act theory, which emphasizes the *act* of language rather than its representational role.

In our discussion of language we have particularly stressed that speech acts create commitment. In revealing commitment as the basis for language, we situate it in a social structure rather than in the mental activity of individuals. Our reason for this emphasis is to counteract the forgetfulness of commitment that pervades much of the discussion (both theoretical and commonplace) about language. The rationalistic tradition takes language as a representation—a carrier of information—and conceals its central social role. To be human is to be the kind of being that generates commitments, through speaking and listening. Without our ability to create and accept (or decline) commitments we are acting in a less than fully human way, and we are not fully using language.

This dimension is not explicitly developed in work on hermeneutics (including Heidegger) or in Maturana's account of linguistic domains. It is developed in speech act theory (especially in later work like that of Habermas) and is a crucial element in our analysis of the uses of computer technology. This key role develops from the recognition that computers are fundamentally tools for human action. Their power as tools for linguistic action derives from their ability to manipulate formal tokens of the kinds that constitute the structural elements of languages. But they are incapable of making commitments and cannot themselves enter into language.

The following chapters introduce discussions of the possibilities for 'intelligent computers,' 'computer language understanding,' 'expert systems,' and 'computer decision making.' In each case there is a pervasive misunderstanding based on the failure to recognize the role of commitment in language. For example, a computer program is not an expert, although it can be a highly sophisticated medium for communication among experts, or between an expert and someone needing help in a specialized domain. This understanding leads us to re-evaluate current research directions and suggest alternatives.

One possibility we will describe at some length in Chapter 11 is the design of tools that facilitate human communication through explicit application of speech act theory. As we pointed out in the introduction, computers are linguistic tools. On the basis of our understanding of commitment, we can create devices whose form of readiness-to-hand leads to more effective communication. We discuss a particular family of devices called 'coordination systems' that help us to recognize and create the commitment structures in our linguistic acts. In using such tools, people will be directed into a greater awareness of the social dimensions of their language and of its role in effective action.

6.5 Breakdown and the ontology of design

The preceding sections have discussed background and commitment. The third major discussion in the preceding chapters was about 'breakdown,' which is especially relevant to the question of design.

In designing new artifacts, tools, organizational structures, managerial practices, and so forth, a standard approach is to talk about 'problems' and 'problem solving.' A great deal of literature has been devoted to this topic, in a variety of disciplines. The difficulty with such an approach, which has been deeply influenced by the rationalistic tradition, is that it tends to grant problems some kind of objective existence, failing to take account of the blindness inherent in the way problems are formulated.

A 'problem' always arises for human beings in situations where they live—in other words, it arises in relation to a background. Different interpreters will see and talk about different problems requiring different tools, potential actions, and design solutions. In some cases, what is a problem for one person won't be a problem at all for someone else.

Here, as elsewhere, we want to break with the rationalistic tradition, proposing a different language for situations in which 'problems' arise. Following Heidegger, we prefer to talk about 'breakdowns.' By this we mean the interrupted moment of our habitual, standard, comfortable 'being-in-the-world .' Breakdowns serve an extremely important cognitive function, revealing to us the nature of our practices and equipment, making them

'present-to-hand' to us, perhaps for the first time. In this sense they function in a positive rather than a negative way.

New design can be created and implemented only in the space that emerges in the recurrent structure of breakdown. A design constitutes an interpretation of breakdown and a committed attempt to anticipate future breakdowns. In Chapter 10 we will discuss breakdowns in relation to the design of expert systems, and in Chapter 11 their role in management and decision making.

Most important, though, is the fundamental role of breakdown in creating the space of what can be said, and the role of language in creating our world. The key to much of what we have been saying in the preceding chapters lies in recognizing the fundamental importance of the shift from an individual-centered conception of understanding to one that is socially based. Knowledge and understanding (in both the cognitive and linguistic senses) do not result from formal operations on mental representations of an objectively existing world. Rather, they arise from the individual's committed participation in mutually oriented patterns of behavior that are embedded in a socially shared background of concerns, actions, and beliefs. This shift from an individual to a social perspective—from mental representation to patterned interaction—permits language and cognition to merge. Because of what Heidegger calls our 'thrownness,' we are largely forgetful of the social dimension of understanding and the commitment it entails. It is only when a breakdown occurs that we become aware of the fact that 'things' in our world exist not as the result of individual acts of cognition but through our active participation in a domain of discourse and mutual concern.

In this view, language—the public manifestation in speech and writing of this mutual orientation—is no longer merely a reflective but rather a constitutive medium. We create and give meaning to the world we live in and share with others. To put the point in a more radical form, we design ourselves (and the social and technological networks in which our lives have meaning) in language.

Computers do not exist, in the sense of things possessing objective features and functions, outside of language. They are created in the conversations human beings engage in when they cope with and anticipate breakdown. Our central claim in this book is that the current theoretical discourse about computers is based on a misinterpretation of the nature of human cognition and language. Computers designed on the basis of this misconception provide only impoverished possibilities for modelling and enlarging the scope of human understanding. They are restricted to representing knowledge as the acquisition and manipulation of facts, and communication as the transferring of information. As a result, we are now witnessing a major breakdown in the design of computer technology—a breakdown that reveals the rationalistically oriented background of discourse in which our current understanding is embedded.

The question we now have to deal with is how to design computers on the basis of the new discourse about language and thought that we have been elaborating. Computers are not only designed in language but are themselves equipment for language. They will not just reflect our understanding of language, but will at the same time create new possibilities for the speaking and listening that we do—for creating ourselves in language.