1. The Self-Organization Revolution in Science

1.1 Self-organization: a new light on nature

Nature, especially inorganic nature, is full of fascinatingly organized forms and patterns. The silhouette of mountains is the same whether one views it at the scale of a rock, a summit, or a whole mountain range. Sand dunes often arrange themselves in long parallel stripes. Water crystallizes into symmetrical serrated flakes when the temperature is right. And when water flows in rivers and hurtles over cataracts, trumpet-shaped vortices appear and the bubbles collect together in structures which are sometimes polyhedral. Lightning flashes draw plant-like branches in the sky. Alternating freezing and thawing of the rocky ground of the tundra leaves polygonal impressions in the earth. The list of these forms rivals many human artefacts in complexity, as can be seen in Figure 1.1. And yet they have no designer, not even natural selection, Dawkins's (1982) 'blind watchmaker'. What, then, are the mysterious factors that explain their existence?

In fact, all these organized structures have a feature in common: they are the macroscopic outcomes of local interactions between the many components of the system from which they emerge. Their global organizational properties are not to be found at the local level. Indeed, the properties of the shape of a water molecule, as well as of its individual physico-chemical components, are qualitatively different from the properties of ice crystals (see Figure 1.2), whirlpools, or polyhedral bubbles. The polygonal impressions in the tundra do not correspond with the shape of the stones composing them, and have a spatial organization quite different from the temporal organization of freezing and thawing. This is the hallmark of a newly discovered phenomenon—self-organization.

This fundamental concept is the touchstone of the paradigm shift driven by the sciences of complexity in the twentieth century, developed by brilliant researchers such as William Ross Ashby, Heinz Von Foerster, Ilya Prigogine, Francesco Varela , René Thom, and Stuart Kauffman. Ever since Newton,

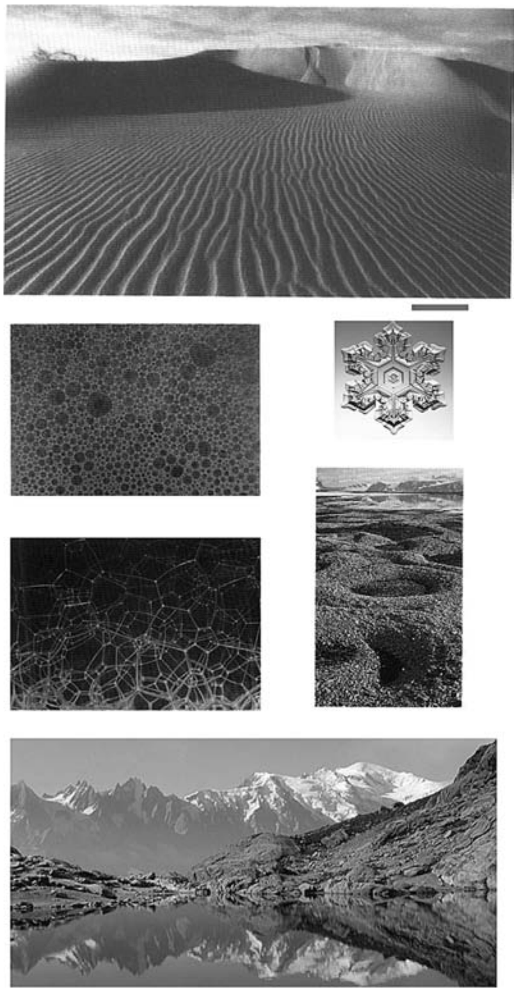

FIGURE 1.1. Nature is full of organized forms and patterns without there being anywhere any plans which might have served to build them; they are said to be self-organized. Here, parallel stripes running through sand dunes, water bubbles on the surface of liquid which has been stirred up and the polyhedral structures which are left when they dry out, an ice crystal, mountains whose shapes are the same whether one views them on the scale of a rock or a whole peak. (Photos: Nick Lancaster, Desert Research Institute, Nevada; Burkhard Prause, University of Notre Dame, Indiana; Bill Krantz, University of Colorado.)

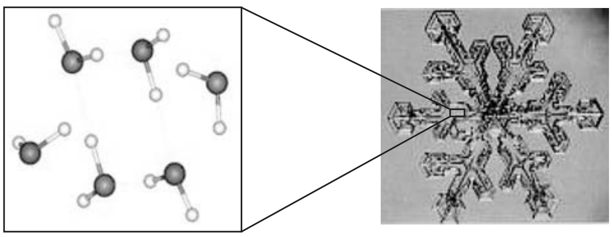

FIGURE 1.2. The self-organization of ice crystals.

good science has been supposed to be reductionist, and has consisted in decomposing natural systems into simpler subsystems. For example, to understand the functioning of the human body, it was appropriate to study the respective parts, such as the heart, the nervous system, or the limbic system. It did not stop there: study of the nervous system, for example, was subdivided into study of the cortex, of the thalamus, or of the peripheral motor innervations, and each of these subparts was studied by hyper-specialists in separated, dedicated university departments. This method has obviously enabled us to accumulate an impressive bank of knowledge.

But the prophets of complexity have broken up this paradigm. Their credo is: 'The sum of the parts is greater than the parts taken independently.' This is because nature is composed of complex systems with many interacting subsystems, and complex systems have a very strong tendency to self-organize. This includes even systems in biology, in which the ascendancy of natural selection is not total but must work alongside self-organization (Kauffman, 1996; Ball, 2001).

This is why it now seems that many natural systems cannot simply be explained by a reductionist study of their parts. One of the most emblematic examples is that of the collectively built artefacts of insect societies (Camazine et al., 2001). For example, termites make immense nests, rising several metres above the ground and with an architecture which is reminiscent of human structures, as Figure 1.3 shows. To try to explain how these structures are built, the study of individual termites, for example the precise study of all their neural wiring, is absolutely not sufficient. One could know everything about the anatomy and the brain of a termite without ever understanding how their

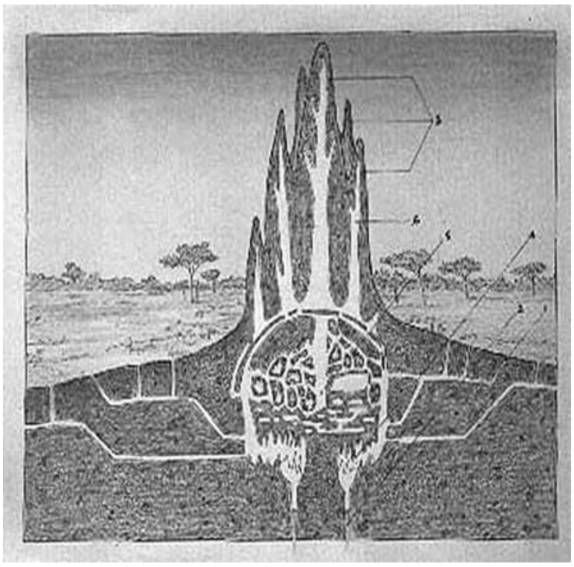

FIGURE 1.3. The architecture of termite nests is the self-organizd result of local interactions among thousands of individuals.

nests are built; because no termite has the equivalent of a plan, even partial, of the superstructure. The knowhow that they possess is infinitely more basic; it is of the type 'if I come across a lump of earth, pick it up and place it where the pheromone signal is strongest'. The superstructure is rather the result of the dynamic interactions in the environment of thousands of termites, in the same way as the symmetrical structure of ice crystals is the result of the interactions of water molecules, and not a projection to the macroscopic level of structures already present at the microscopic level.

Such concrete examples of the use of the concept of self-organization, and of explanations of natural forms in terms of systemic properties, are now abundant, and are at the heart of the most advanced research of more and more physicists and biologists. Further examples include hunting or foraging patterns of bees and ants, the dynamic shapes of shoals of fish or flocks of birds, symmetrical patterns on butterfly wings, the regular spots on a leopard's skin, the stripes of fish and shellfish, the magnetization of magnets, the formation of whirlpools in rivers, the birth of galaxies, demographic oscillations in predator–prey ecology, the formation of patterns in bacterial cultures and chemical reaction–diffusion systems, crystallization, lasers, superconductivity, the distribution of the sizes of avalanches, auto-catalytic chemical systems, the formation of lipid membranes, and the dynamics of traffic jams on freeways (see Ball, 2001).

The sciences of complexity have thus demonstrated the fundamental usefulness of the concept of self-organization for the explanation of natural phenomena involving physical structures and certain biological structures characterizing the morphology or behaviour of simple animals like insects. We are now at the dawn of a new and decisive phase in this scientific revolution: researchers in complexity theory are beginning to tackle the understanding of mankind itself, using these new tools. The understanding of the vital functions of the human body was the first to be transformed by the emerging wave of what is called "integrative" or "systemic" biology (Chauvet, 1995; Kitano, 2002). Rather than concentrating on each organ in isolation, there is now an attempt to understand their complex interactions in an organism considered as a whole in which each element is integrated with the others. This has opened up new theoretical vistas, for example the understanding of cancers (Kitano, 2004) or of morphogenesis (Kupiec and Sonigo, 2000), which, according to the authors, is seen not as the serial execution of a genetic programme but as the self-organized dynamic of the whole ecosystem formed by the cells competing for nourishment.

The advocates of self-organization do not stop there: the human brain, and thus the phenomena of sensation and thought, are also under the strong influence of features of spontaneous organization in their structure. Indeed, the brain, composed of billions of neurons dynamically interacting among themselves and with the outside world, is the prototype of a complex system. For example, as I will show in this book, self-organization could be at the heart of the capacity of our brains to categorize the perceived world, that is, to organize the continuous flux of perceptions into atomic psychological objects.

But the main subject of this book goes beyond speculation about the brain as a self-organized system. We are today on the brink of a major advance in science: that of a naturalized understanding of what makes humans so exceptional, their culture and language. Indeed, if culture and language have been the subjects of investigation by social sciences for centuries, scientists have not yet succeeded in anchoring an understanding of them in terms of their material biological substance, i.e. the set of human brains in ongoing complex dynamic interactions. Now, the tools of complexity begin to make it possible. This is what I will illustrate in this book, concentrating on a quite precise example, that of the origin and shaping of one of the pillars of human language—speech, the outward form and vehicle of language—seen as systems of shared and combinatorial sounds particular to each language community. To understand the ongoing revolution on this particular question, I will first outline the main trends in its history.

1.2 Language origins

It is an obvious fact, only matched by the mystery that follows from it: humans speak. It is their main activity, an activity which, moreover, sets them off from the rest of the animal kingdom. Human language is a communication medium of unequalled complexity. It is a conventionalized code which lets one individual share his ideas and emotions with others, talk of the colours in the sky and also of distant landscapes, of past events, even of how he imagines the future, of mathematical theorems, of invisible properties of matter, and of language itself. Besides that, each language defines a system which is peculiar to its speakers, an original way of organizing sounds, syllables, words, and sentences, and of spelling out the relationship between these sentences and the concepts which they convey. Today there are thousands of languages spoken in human communities. Over time, some languages die and others are born. The number of languages which have existed is estimated at over half a million. It is hard to imagine humanity without language. And yet, a long time in the past, humans did not speak.

This raises one of the most difficult questions in science: how did humans come to talk? A further question follows naturally: how do languages evolve?

These two questions, concerning the origin of the language faculty and the evolution of languages, have been focused on by many thinkers in centuries gone by, particularly in the nineteenth century. They are prominent in Darwin's speculations (Darwin, 1999[1859]). Many such theories were developed without the benefit of any empirical or experimental constraint. They were equally devoid of reasoned arguments and scientific method, to the point where the Linguistic Society of Paris ruled that such questions should be raised no more in the context of scientific discussion. This ruling initiated a century of almost total lack of progress in research in this domain.

Advances in neuroscience, cognitive science, and genetics towards the end of the twentieth century have put these questions back into the centre of the scientific arena. On the one hand, modern neuroscience and cognitive science have made enormous progress in understanding the general functioning of the brain, and especially the way in which language is acquired and processed. These developments have allowed the study of language to relate to the natural sciences, that is, to ground the abstract systems which linguists describe in the biological matter of which humans and their environment are composed. In short, natural sciences have taken over questions previously in the domain of the social sciences. This new light on the workings of language and the brain have provided researchers with the constraints whose absence undermined the speculations about the origins of language of the nineteenth century.

On the other hand, progress in genetics has turned the spotlight on neo-Darwinian theories of evolution, both confirming some of its foundations (with the discovery of genes, for example, along with their mechanisms of variation) and allowing its predictions to be tested, often successfully, thanks to the sequencing of the genomes of different species of animals so that we can reconstruct their phylogenetic trees and trace their evolutionary history. In particular, the sequencing of the human genome, along with that of other animals such as chimpanzees and monkeys, makes it possible to specify the relationships between humans and their ancestors. Thus, driven by vigorous evolutionary biology, which simultaneously provides an impressive body of observations and a solid explanatory framework, the question of human origins has become a central theme in science. And, quite naturally, the origins of language (this being one of the distinctive features of modern humans) has become, as in the nineteenth century, a beacon for research.

1.2.1 Interdisciplinarity

There is an emerging consensus among researchers who are today getting down to questions of the origin of the human language faculty and the evolution of languages: this research must be interdisciplinary. In fact it poses a puzzle with immense ramifications which go beyond the competence of each individual discipline. First, the two big questions must be decomposed into subquestions which are themselves already quite complex. What, in fact, is the language faculty? What is a language? How are sounds, words, sentences, and representations of meaning related to each other? How does the brain represent and process these sounds and sentences and the concepts which they convey? How do we learn to speak? What are the respective roles of nature and nurture? What is language for? What is its role in a community? How does a language form and change in the course of successive generations of speakers? What do we know of the history of each particular language? Why are the language faculty and languages the way they are? Why do we see universal tendencies and at the same time great diversity in languages? How does language influence the way we perceive and understand the world? What do we know of the history of the human capacity for speech? Is it mainly the result of genetic evolution, like the evolution of the eyes, or a cultural invention, like writing? Is language an adaptation to a changing environment? An internal change in an individual which increased its chances of reproduction? Is it an exaptation, a side effect of changes which were not at first tied to communicative behaviour? What are the evolutionary prerequisites which paved the way for the capacity of speech? And how did these prerequisites appear? Independently? Genetically? Culturally?

Ranged against the diversity of these questions is an even greater diversity of research disciplines and methods. Linguists, even though they may continue to provide crucial data on the history of languages, are no longer the main actors. Developmental and cognitive psychology and neuropsychology carry out behavioural studies of language acquisition and language pathology, and these often reveal cognitive mechanisms involved in language processing. Neuroscience—especially with equipment for brain imaging allowing us to see which brain regions are active for given tasks—attempts to find neural correlates of verbal behaviour, to discover its organization in the brain. Some researchers also study the physiology of the vocal tract, to try to understand how we produce speech sounds. The physiology of the ear, the essential receptor in the speech-decoding chain (or vision, in the case of signed languages), is also a focus of research. Archaeologists examine fossils and artefacts left by the first hominids, and try on the one hand to deduce our anatomical evolution (especially of the larynx) and on the other to get an idea of what activities they were engaged in (what tools did they make? and how did they use them? and what can these tools tell us about the degree of cognitive development?). Anthropologists do fieldwork on isolated peoples and report on cultural differences, especially those related to languages and the meanings they convey. Primatologists try to report on the communicative capacities of some of our ancestors and to compare them with our own. Geneticists on the one hand sequence the human genome and that of potential ancestral species when it is possible to specify their phylogenetic relatedness, and on the other hand use genetic information from different people across the planet to help in reconstructing the history of languages, which is often correlated with the genetic history of their speakers.

Thus language involves a multitude of components interacting in complex ways in parallel on several time-scales: the ontogenetic time-scale, characterizing the growth of an individual person; the glossogenetic or cultural time-scale, which characterizes the evolution of cultures; and the phylogenetic time-scale, which characterizes the evolution of species. Moreover, not only is it essential to study each of these components independently, to reduce the complexity of the problem; it is also necessary to study their interactions. In fact, as I stressed in the first part of this chapter, the sciences of complexity have taught us that in many of the complex systems in nature, there are global phenomena that are the irreducible result of local interactions between components whose individual study would not allow us to see the global properties of the whole combined system. Thus, a growing number of researchers think that many properties of language are not directly encoded by any of the components involved, but are the self-organized outcomes of the interactions of the components. Yet these self-organizational phenomena are often complicated to understand or to foresee intuitively, and to formulate in words.

1.2.2 Computer modelling

It is for this reason that we find, in addition to the scientific activities already mentioned, and in the framework of language origins research, work by researchers in artificial intelligence, mathematicians, and theoretical biologists, who construct operational models of these interactions between the components involved in language. An operational model is one which defines the set of its assumptions explicitly and above all shows how to calculate their consequences, that is, to prove that they lead to a certain set of conclusions. There are two main types of operational model. The first, used by mathematicians and some theoretical biologists, consists in abstracting from the phenomenon of language a certain number of variables, along with the rules of their evolution in the form of mathematical equations. Most often this resembles systems of coupled differential equations, and benefits from the framework of dynamic systems theory. The second type, which allows for modelling of more complex phenomena than the first, is that used by researchers in artificial intelligence:

it consists in the construction of artificial systems implemented in computers or in robots. These artificial systems are made of programs which most often take the form of artificial software or robotic agents, endowed with artificial brains and bodies. These are then put into interaction with an artificial environment (or a real environment in the case of robots), and their dynamics can be studied.

Building artificial systems in the context of research into language origins and the evolution of languages is enjoying growing popularity in the scientific community, exactly because it is a crucial tool for studying the phenomena of language in relation to the complex interactions of its components. These systems are put to two main types of use: (1) they serve to evaluate the internal coherence of verbally expressed theories already proposed by clarifying all their hypotheses and verifying that they do indeed lead to the proposed conclusions (and quite often one discovers errors in the assumptions as well as in the conclusions, which need to be revised); (2) they serve to explore and generate new theories, which themselves often appear when one simply tries to build an artificial system reproducing the verbal behaviour of humans. A number of definitive results have already been obtained, and have opened the way for the resolution of previously unanswered questions: the decentralized generation of lexical and semantic conventions in populations of agents (e.g. Steels, 1997; Kaplan, 2001); the formation of shared inventories of vowels or syllables in groups of agents, with features of structural regularities greatly resembling those of human languages (e.g. de Boer, 2001; Oudeyer, 2001*b*; 2001*c*; 2001*d*); the formation of conventionalized syntactic structures (e.g. Batali, 1998); the conditions under which combinatoriality, the property of sytematic reuse, can be selected (e.g. Kirby, 1998).[1]

The work to be presented in this book belongs in this methodological tradition of building artificial models. It will concentrate on the origin of one particular aspect of language: speech sounds. Sounds, as I will explain in detail in the next chapter, constitute a conventional code providing each language with a repertoire of forms for conveying its messages. This code, which has an acoustic and an articulatory side, organizes sounds into categories which are special to each linguistic community, and regulates the ways in which they can be combined (these rules of sound-syntax are also cultural conventions). The system is thus discrete and combinatorial. Without such a code, which could also be implemented in a manual modality for signed languages, there

[1] This list is in no way exhaustive, and more examples can be found in Cangelosi and Parisi (2002).

could be no form, no content, and hence no linguistic communication. How could such a code have arisen? In particular, how were the first codes able to form themselves before there was any conventional linguistic communication, of which they are prerequisites? Why are the sound patterns of human languages the way they are? Such are the questions which drive the work in this book. They are simultaneously very ambitious, as they include many complex aspects, both individual and social, and very modest in relation to the overall programme of research into the origin of language and the evolution of languages. In fact, they only concern the origin of one prerequisite of language among many others (such as the capacity to form symbolic representations, or the pragmatic capacity to infer the intentions of others by means of behavioural signals).

I do not claim to suggest direct or definitive answers. This book is rather of the type of exploratory philosophy which drives the construction of artificial systems. In this spirit, I construct a population of agents, endowed with well-defined mental, vocal, and perceptual capacities, based more or less exactly on their real human counterparts, which will enable us to establish the sufficient conditions for the formation of speech codes like those of humans. It will be shown that these sufficient conditions are interesting when their generality and simplicity is compared with the structures that they generate, namely speech codes. It is the phenomenon of self-organization which enables this linking between qualitatively different properties of the same system both on a local and on a global scale. This enables us to to define and suggest a possible and original kind of mechanism for answering the questions posed above, and to show its internal coherence. It will also be explained why these conditions are interesting in that they not only support a classical adaptationist, neo-Darwinian scenario of the origins of speech, but also open new perspectives, and in particular make possible a scenario in which the speech code, or at least some of its properties, might be an exaptation[2].

I will stress that the relationship between the artificial system and the human system is not a relationship of identity or close modelling; it is a relationship of analogy. I will not attempt to show that these assumptions correlate with precise features of the real world,[3] but rather that they are useful for defining the dimensions of the space of possible explanations which have already been proposed, and indeed for generating new types of explanation.

[2] We will explain in detail the differences between adaptations and exaptations in Ch. 3.

[3] This is largely infeasible today, due to the modest level of our knowledge of the phenomena of speech, particularly its neural correlates.

To be precise, the goal of this book is not to propose direct answers, but rather to engage in the structuration of the theoretical thinking involved in the research on the origins of speech. For this reason, the main evaluation criterion for this work is the impact it will have on the thinking of researchers in this domain. In summary, this book should be read while asking, not whether what is written is true or false, but whether it is useful or not.

Chapter 2 will introduce the way the speech code works, and will specify the questions to be asked about its origins. Chapter 3 locates the problems of the origin of speech in the general framework of the origins of form in biology: I will explain the phenomenon of self-organization along with that of natural selection, both of which are features of the mechanisms of creation of forms in the living world. In particular, I will present a linkage between the concepts of self-organization and natural selection, which will lead us to discuss how one should structure the arguments explaining the origin of living forms. Chapter 4 makes a detailed survey of the literature, to review what answers have already been proposed, and will support the broad outline of the approach at the centre of this work. Chapter 5 details my methodology, that is, the building of artificial systems, along with the goals and scientific philosophy which drive it. Chapter 6 gives a formal description of a first version of the artificial system, and illustrates its dynamics, which involves the formation of a discrete speech code shared by a population of agents who initially only pronounce unstructured and holistic vocalizations, and do not follow any rules of coordinated interaction. I will discuss, in particular, the role that morphological constraints on the vocal and perceptual apparatus may or may not play in the formation of speech codes. Chapter 7 presents a variant of the artificial system in which, by contrast with Chapter 6, it will not be assumed that the agents are capable from the outset of retrieving articulatory representations of the sounds which they hear; this capacity will be learned, thanks to a quite generic neural architecture. I will also use a model of the human vocal tract for producing vowels, which will enable us to specify the analogy between artificial and human systems: it will be shown that the statistical regularities which characterize the vowel systems of populations of artificial agents are very similar to those of the vowel systems of human languages.

Chapter 8 presents an extension of the artificial system of Chapter 6, and shows how non-trivial combinatoriality and rules of sound-syntax, that is, phonotactic rules, can arise. Chapter 9 discusses the results obtained, and in particular the contribution that the system makes to research on the origins of speech. I show that the generality and simplicity of the assumptions built into the artificial system, and the fact that they suffice for a speech code to self-organize, allow us to build up a convincing adaptationist scenario in which the neural structures have been chosen by natural selection driven by pressure for linguistic communication. Moreover, I show that the generality and the non-specificity of these assumptions suggest several alternative scenarios in which the speech code could have arisen independently from linguistic communication. Chapter 10 winds matters up.