

1 INTRODUCTION

This book attempts to tackle the fundamental question of the linguistic sciences: *what is language?* It is a very old question, and the fact that the answers keep changing testifies to its complexity and importance, to the enormous intellectual effort invested in it through the ages—and, perhaps paradoxically, to the fact that the human condition is always in transition; it forces us to constantly re-examine and adjust our conceptions of ourselves, and at every point along the way, whenever we manage to take a look at ourselves from a new angle, we find that we also see language in a new light. There is not much choice here: the way we understand ourselves is inextricably and permanently entangled with the way we understand language.

The understanding of language that I wish to offer in these pages begins with a return to a long tradition of thought—a tradition that was expelled from the linguistic sciences by the *cognitive revolution*, fifty years ago, and is now being locally re-considered in certain quarters of the field: the human condition is deeply social, and language is a social entity. It is a property of the community, of the social network, the product of a collective process of invention and development. It resides *between* speakers, not *in* them, at a level of organization and complexity that transcends the individual mind—and cannot be reduced to it. The place to look at for the essence of language is not the mind-brain. It is social life.

This, however, is only the beginning, and as such, not very informative. The question has to be: what type of social entity is language? I will argue that it is essentially a *communication technology*. What this means, as a first approximation, is that language has to be ontologically classified together with the other communication technologies humans have invented, such as the book, fax, telephone, computer games, and Facebook—not together with social institutions (such as government or the family), or cognitive capacities (such as vision or rationality). Like other technologies, language has to be constructed before it can be used, and it constantly develops and changes as a result of usage. It spreads and propagates

like other technologies, and like them it has its experts, role models, innovators, guards, rebels, users, and abusers—active and passive to different degrees. Most importantly, its *modus operandi* is best analyzed in technological terms. Language is the first communication technology we ever invented. It has revolutionized human life and actually changed us as a biological species. It is still, today, the most powerful technology we use. To understand it, and its dialectic entanglement with the human condition, we have to adopt a *social-technological* approach to its study.

At the center of the social-technological approach to be developed in this book, stands the understanding that language cannot be a *general-purpose* communication technology. There is simply no such thing: had there been, we would never need to invent anything else. Every communication technology is functionally specific. Every technology employs a specific functional strategy, and the specificity of the strategy determines the technology's functional envelope: what it can do with high levels of efficiency, where its efficiency declines, where it collapses, and what it cannot do to begin with. To get to the bottom of language, to turn the acknowledgment of its social essence into a scientific theory, we have to ask: What is the *functional specificity* of language as a socially constructed communication technology?

The functional specificity of language, I will argue, lies in the very particular functional strategy it employs. It is dedicated to the systematic *instruction of imagination*: we use it to communicate directly with our interlocutors' imaginations. All the other systems of intentional communication, technological or not, used by humans and other species (with the possible exception of bee dances), work with what I will call the *experiential* strategy. They provide materials for the interlocutors to experience with their senses and thus allow for the actual *sharing* of experience. The experiential strategy is thus inherently limited: an experience can only be shared if it can be experienced. Language is the only system that goes beyond the sharing of experience. It allows speakers to intentionally and systematically instruct their interlocutors in the process of imagining the intended experience—*instead* of directly experiencing it. The speaker provides the receiver with a code, a plan, a skeletal list of the basic co-ordinates of the experience—which the receiver is then expected to use as a scaffold for experiential imagination. Following the code, the interlocutor raises past experiences from memory, and then reconstructs and recombines them to produce novel, imagined experiences. Language is thus the only system that allows for communication that actually *bridges the experiential gaps* between speakers. In doing that, it opens a venue for human sociality that would otherwise remain closed. This is the secret of its success, and it is also the way it actively participates in the construction of the human individual as a social being.

Nothing in the social-technological approach to be developed here denies that language raises crucial questions about our cognitions. On the contrary: language is a technology that requires much of its users, a whole array of cognitive (and emotional) capacities, and these deserve our full attention. What the approach insists on is that in order to understand the capacities (in acquisition, in actual communication)—and in order to understand how they evolved—we must first of all position the technology itself at the right ontological level, where it belongs. We have learned an enormous amount in the last fifty years about language-related cognition, and we now know enough to finally abandon the *computer metaphor*—the idea that language, like other cognitive capacities, should be investigated as if it were a piece of software in the computer that we call the mind-brain—and replace it with the updated *metaphor of the Web*: language is indeed a piece of software—a *communication software*. It resides on the Net. Human individuals are *end-users*: they download *copies* of the software into their mind-brains, and these allow them to communicate with one another (to the extent that the copies are similar enough). What we understand today about the processes and capacities involved in the downloading and use of the software by the end-users should teach us exactly this: that the essence of the software itself does not lie there. The essence of the software lies in the fact that it resides between the end-users; that it facilitates something that is not within the individual capacity of any of them. As Humboldt (1936 [1999], p. 42) puts it, “the existence of *languages* proves that there are also mental creations which in no way whatever pass out from a single individual to the remainder, but can only emanate from the simultaneous self-activity of all.” Language cannot be explained in terms of the cognitive dynamics taking place within the individual mind-brain, for the exact same reason that social networks on the Internet cannot be explained on the basis of whatever is happening in the personal computers of the networks' users. As I will try to show, the fact that the cognitively oriented theories of language currently under discussion in the literature are so baroquely complex testifies to exactly this: they are attempts to achieve an impossible goal. Language is a much simpler, much more reasonable entity, but to see that, we have to see it for what it is.

My hypothesis, then, is that everything in need of theoretical explanation in linguistics and around it follows, directly or indirectly, from the essence of language as a socially constructed, imagination-instructing communication technology: the architecture of language, the way it is socially constructed and the way it works in actual communication; the properties of words and their meanings; the inextricable relationship between language and truth; the dialectical relationship between language and the ways we experience the world; the patterns of its grammatical structures; its dual life as a universally diverse

phenomenon; the inherent difficulty and fragility of linguistic communication, its successes and failures; the dynamics of language change; the intricacies of language-based identity; the political power of language; the nature of language acquisition; and the evolutionary dynamics that brought language about in the first place: pre-linguistic human collectivities gradually distanced themselves from ape sociality, and developed a multi-faceted and highly sophisticated culture, still without words—a culture that brought the strategy of experience-sharing to the limits of its functional envelope. This was, at the very same time, the foundation upon which language was built—and the cause for its emergence. As human existence came to depend more and more on the overall quality of the collective, the need for a way of communicating that could go beyond the limits of the envelope gradually turned into a necessity—as always, the mother of invention. As the explorations into the new realm of communication stabilized to form the first working prototypes, human evolution entered a new era. Language continued to evolve at the cultural level, and it dragged human societies and human individuals—their behaviors and identities, cognitions and emotions, physiologies and genetic makeups—into a fascinating web of co-evolutionary spirals. First we invented language. Then language changed us.

It goes without saying that I have no intention of actually explaining everything in need of explanation about language—in this book or in general. What I will try to do is present a general theory of language explicit and detailed enough to provide a new conceptual foundation for the scientific inquiry into everything in and around language. Philosopher of science Carl Hempel's classical definition of scientific theory captures what I aim to achieve with great precision and elegance. It deserves to be quoted in its entirety:

A scientific theory might therefore be likened to a complex spatial network. Its terms are represented by the knots, while the threads connecting the latter correspond, in part, to the definitions and, in part, to the fundamental and derivative hypotheses included in the theory. The whole system floats, as it were, above the plane of observation and is anchored to it by the rules of interpretation. These might be viewed as strings which are not part of the network but link certain points of the latter with specific places in the plane of observation. By virtue of these interpretive connections, the network can function as a scientific theory: from certain observational data, we may ascend, via an interpretive string, to some point in the theoretical network, thence proceed, via definitions and hypotheses, to other points, from which another interpretive string permits a descent to the plane of observation. (Hempel 1952, p. 36).

Hempel's characterization of scientific theory as a spatial network of terms, definitions, and hypotheses, floating above the plane of observation as an interpretative map, makes it immediately clear that the issue is not just the explanation of this set of data or the other, or the formulation of hypotheses and predictions, or the development of tools for empirical investigation. What is at stake is the establishment of a complex web of relationships (of many different types) between ideas, arguments, observations, and interpretations that eventually produces a general sense of clarity, coherence, and unity and allows for that freedom of interpretative movement, up and down, between the plane of theory and the plane of observation, so beautifully described above. The challenge of theory is of the architectural type.

What I will do, then, is begin with the careful construction of the theoretical network—with its terms, definitions, and hypotheses—and then show how it allows for a major re-thinking of the plane of observation, on three interrelated levels. First, I will select a representative set of problem-clusters from the different sub-domains of the linguistic sciences and show how the theory offers to deal with them—not in order to provide them with full accounts, but in order to demonstrate their susceptibility to the type of treatment offered by the theory. The goal here will be to show that the theory offers a fresh outlook on the various problem-clusters; suggests new answers to foundational questions; resolves persistent difficulties; allows for serious re-arrangements of the data; invests major findings and discoveries in the different fields with new meaning; releases the investigation from deadlocked controversies and unwarranted stipulations; simplifies analyses; opens new venues for empirical research; and so on. As the theory attempts to re-position language in the social domain, the most important challenge here will be to demonstrate that it does not lose sight of the enormous achievements of linguistics as a cognitive science. There is no reason to make *that* mistake again.

The technical discussions of the various problem-clusters will thus be as explicit and detailed as the task of demonstration requires, not more than that. They should be read as attempts to show, in a rather informal manner, how the theory allows for new types of principled explanations—and what more serious work on the empirical issues should look like. Most of the examples, moreover, are going to be in *English*, but nothing is implied by this choice—definitely not that there is something universal about the conventions of English. It is exactly because language is universally a conventional system, that any demonstration will have to make use of the particular conventions of this language or the other. The fact that you and I are at the moment using English for communication suggests that a demonstration in English would probably be the most reasonable choice.

At the second level, I will try to demonstrate that the particular analyses of the various problem-clusters converge to produce a new and *unified* outlook on the general picture of human language: that the theory presents enough evidence to support the claim that all the different problem-clusters, currently subjected in the literature to highly specialized and incompatible explanatory apparatuses, are in fact susceptible to the same type of treatment; that the theory re-arranges the problem-clusters with respect to each other in fruitful ways, and raises new explanatory bridges between them; that the entire move shows the way toward the reassembly of the puzzle of human language as a unified phenomenon.

At the third level, I will be interested in the impact of the theory on the relationships between theoretical linguistics and all the other disciplines around it which have a vested interest in the linguistic plane of observation—all the way from psychology and biology to sociology, communication studies, and critical studies. I will try to show that the theory creates new harmonies between theoretical linguistics and its neighbors; moves them closer together and helps rationalize their relationships; opens new venues for mutual influence, inspiration, and co-operation; and thus in effect helps re-position language in its rightful place within the overall picture of human life. Here, again, the most important challenge will be to demonstrate how the theory re-connects linguistics with the socially oriented traditions (including the postmodern ones)—without severing the ties it has worked hard to establish with its cognitively oriented neighbors.

1.1 Setting the Theoretical Stage: Where Are We Today?

Fifty years ago, a proposal for a general theory of language was put forward by Noam Chomsky. Not everybody agreed, of course, but no one offered a comprehensive alternative. Chomsky's theoretical strategy was clear, bold, and very exciting, and it remained virtually untouched through the entire developmental history of his theory, from *the logical structure of linguistic theory* (1955) and *syntactic structures* (1957), through the *aspects* model (1965), *government and binding* (1981) and *barriers* (1986), all the way to the current version of the theory, *the minimalist program* (1995). At the center of his spatial network, Chomsky positioned a set of definitional statements (often characterized, quite misleadingly, as working hypotheses):

- (a) Language is a *generative* system: it allows for the production of an infinite number of sentences from a finite arsenal of primitives. The constitutive principles of the generative system are abstract and formal. They are the essence of language.

- (b) The fact that we acquire language and use it means that we know the generative system. We have linguistic *competence*. The essential question about language is thus a question about *individual cognition*: how does the speaker know the generative system? How does competence come about?
- (c) Knowledge does not emerge from experience: the constitutive principles of language cannot be learned from the input offered to the language-acquiring child by the surrounding community of speakers. The input is “meager and degenerate”; the principles are too abstract. The child must come to the word already equipped with the essential knowledge of language. Competence must be *innate*.
- (d) Children are capable of acquiring whichever language they find around them, which means that the constitutive principles are the same in all the languages of the world. The observed differences between languages are surface phenomena. Where it counts, all languages are identical (or almost identical, as in the program of principles-and-parameters, Chomsky and Lasnik 1993): they are all founded on a Universal Grammar (UG), coded in our genes.
- (e) The scientific goal of Linguistics is to uncover the principles of UG, and thus to figure out the essence of the innate and autonomous *language organ* (an organ of the *mind*, irreducible to the physical properties of the *brain*). Linguistics is a branch of psychology (itself a branch of biology). “The tasks of the psychologist,” Chomsky writes, “divide into several sub-tasks. The first is to discover the innate schema that characterizes the class of potential languages—that defines the ‘essence’ of human language. This sub-task falls to that branch of human psychology known as linguistics; it is the problem of traditional universal grammar, of contemporary linguistic theory. The second sub-task is the detailed study of the actual character of the stimulation and the organism-environment interaction that sets the innate cognitive mechanism into operation . . . A third task is that of determining just what it means for a hypothesis about the generative grammar of a language to be ‘consistent’ with the data of sense.” (Chomsky 1998)

From this heavy set of definitional statements emerges a complex and highly particular strategic attitude toward the task of interpreting the plane of observation. To begin with, empirical work should concentrate on the observed patterns of generativity, especially as they manifest themselves in the syntactic structures of languages. Syntax is where the essence of human language lies. The patterns of syntactic generativity should be figured out on the basis of native speakers’ grammaticality judgments, not on the basis of their actual usage of language: linguistic *performance* is full of mistakes, hesitations, repetitions, and repairs (and

this is why it cannot provide the child with sufficient input); competence is only reflected in judgment. Through the analysis of speakers' judgments, the linguist discovers the principles that govern them—those innate principles that make up the human capacity for language. These principles are abstract and formal, which means that the analysis cannot, and therefore should not, attempt to relate them to the meanings of the sentences, or to the communicative intents of their speakers, or to the context of conversation. In this challenge of distillation, everything that does not directly relate to the foundations of generativity is moved aside. Other issues may be interesting, even important, but they are not essential: performance, meaning, communication, context, and also general cognition, and social learning, and language change (the linguist's task is that of *synchronic*, not *diachronic* analysis), and society, culture, semiotics, rhetoric, literature—virtually everything that the non-linguist might consider relevant for linguistic research. Generativity deserves to be investigated in isolation.

This was a radical strategy indeed, and it brought about a huge revolution in our understanding of language. There was, however, something deeply paradoxical in the way it did that: the theory itself, as it continued to be developed by Chomsky and his colleagues of the generative camp, developed into an exceedingly esoteric discourse, baroquely complex and deeply abstract, often interested more in the internal relationships between its terms, definitions, and hypotheses than in the strings of interpretation that were supposed to anchor the whole theoretical machinery to the plane of observation. With time, even the judgments lost much of their importance. In a way, this was inevitable. The third task of the psychologist-linguist, as Chomsky defined it, was “that of determining just what it means for a hypothesis about the generative grammar of a language to be ‘consistent’ with the data of sense.” If the child could not trust the input, there was no a priori reason to assume that the linguist should. Empirical science does not grow very well on rationalist ground.

This, however, was only one side of the coin, actually the less interesting. Chomsky's general perspective sent a shock wave through the linguistic community and beyond it, across the scientific world. His picture of language invested everything around it with new meaning and new energy. The issue was no longer the investigation of languages. This was *passé*. It was not even that hidden software of language inside our minds, not as such. The issue was the essence of being human. More than anybody else in the second half of the twentieth century, Chomsky was the one who asked the formative question: as a present-day Plato, he formulated the question as one about the computational foundations of the human mind. It was no accident that computers, computer science, and artificial intelligence emerged at the very same time: to program machines that do what we do, we need to discover the programs that we carry within.

Chomsky's question was a major engine behind the stellar appearance of cognitive science in the world of science, and language came to be permanently positioned at the very center of the drama. A huge number of new research programs on language and human cognition appeared, and virtually all of them were strategically designed to deal with different facets of Chomsky's question. Many corners of the plane of observation were visited for the first time, and many were re-examined in new ways. New specialties developed, and sub-disciplines went from first programmatic papers to yearly international conferences. From the very beginning, the booming field acquired a certain energy, highly argumentative and polemic: researchers were either coming back from their travels on the observational plane with the ultimate proof that Chomsky was perfectly right, or with the strongest demonstration possible that he was utterly wrong. In this sense, Chomsky's program directed the development of his rivals as much as it directed that of his supporters. They played against him on *his* field.

The last fifty years, then, have seen enormous developments in language research, and it seems fair to state that most of the work has been dedicated to the attempt to bring back into the puzzle those pieces that Chomsky deliberately left out. Meaning and communicative interaction were brought back with the development of functional and cognitive linguistics (Fauconnier 1994, Fillmore 1968, Lakoff 1987, Langacker 1987), and then frame semantics (Fillmore 1982, Fillmore and Baker 2009) and the theory of constructionism (Croft 2001, Goldberg 1995, 2006, Östman and Fried 2005); with the development of pragmatic theory (Austin 1962, Searle 1969, Grice 1975, Levinson 2000, Sperber and Wilson 1986); with the emergence of a variety of semantic theories (Jackendoff 1983, 1990, Montague 1973, Rosch 1975, Wierzbicka 1996); and with the advances made in the understanding of lexical semantics and its relation to syntax (Jackendoff 1983, 1990, Cruse 1986, Dowty 1979, Kiparsky 1997, Levin 1993, Levin and Rappaport Hovav 1995).¹ Structure and meaning are now understood to be connected in bewilderingly complex webs of interrelations.

Society and culture came back into the puzzle with the sociolinguistic quest for socially determined variation (Eckert 2000, 2005, Labov 1966, Trudgill 2011); the anthropological linguistic search for the relationships between language, culture and mind (Duranti 1997, Gumperz and Levinson, 1996); and the growing understanding of the intersubjective nature of language (Meltzoff and Brooks 2007, Tomasello 2008, Tomasello et al. 2005, Zlatev et al. 2008). All these have contributed to the current positioning of language somewhere between human cognition and human society: for many in the field, the implicit working assumption is that language is grounded *both* in cognition and in society.

All of the above, together with what we now know from conversational analysis (Enfield and Stivers 2007, Sacks 1995, Shegloff 2007) and interactional linguistics (Ochs et al. 1996, Couper-Kuhlen and Selting 2001), has challenged the distinction between the original notion of competence and the attested patterns of linguistic performance. All this, together with the results accumulated in linguistic typology (Comrie 1989, Croft 2002)—and the simple fact that more and more languages came to be deeply researched—has changed our views on the universality of language. What we see today is a web of similarities and differences between languages—many restricted, implicational universals; not too many absolute ones: nothing, definitely, that looks like a set of properties worthy of the name Universal Grammar (Evans and Levinson 2009, Levinson and Evans 2010).

The enormous advances we have made in the understanding of language acquisition (Bates et al. 1995, Berman 2004, Bloom 1970, Bruner 1983, Clark 2009, Elman et al. 1996, Slobin 1997, Tomasello 2003), actual linguistic processing (Levelt 1989, Harley 2008), and brain activity related to language (Ahlsén 2006) now position language somewhere between its original autonomy and the realm of general human cognition—between innate constraints (much softer than originally suggested) on language acquisition and usage and the general human capacity for learning, especially for social learning. All this has renewed our interest in the dynamics of language change (Aitchison 2001, Croft 2000), especially the process of grammaticalization (Deutscher 2005, Lehmann 2002, Traugott and Heine 1991, Traugott and Dasher 2002), and has positioned language, again, somewhere in between the synchronic and the diachronic. A rich and lively discourse, unimaginable fifty years ago, now attempts to tackle the most difficult question of all (which Chomsky himself refused to deal with for four decades): the question of the evolution of language (Arbib 2003, Bickerton 1990, Botha 2003, Botha and Knight 2009, Christiansen and Kirby 2003a, b, Corballis 2003, Deacon 1997, Donald 1991, Dor and Jablonka 2000, 2010, Dor, Knight, and Lewis 2014, Dunbar 1998, Fitch 2010, Hurford 2007, 2011, Larson, Deprez and Yamakido 2010, Lieberman 2007, Pinker and Bloom 1990, Richerson and Boyd 2005, Steels 2001, 2014, Tomasello 1999, 2008, and many more).

We have moved quite a long way from Chomsky's original picture, and we know much more than we ever did. But the accomplishments did not come without a price: language, the entire thing, has disappeared on the way. Today, we do not have a general theory of language. We have many pieces of the puzzle, but the puzzle itself does not assemble. The plane of observation has been parcelized, and different areas came to be governed by different rules. The science of language has developed into an extremely fragmented field, in which different explanatory

apparatuses, incompatible with each other, serve the theoretical needs of highly specialized sub-domains. As a consequence, when we look at our fragmented pieces of language, situated as they are away from Chomsky's original picture, we still look at them, as if by default, *from the point of view* of that very same picture. We have found new answers, but they are still answers to Chomsky's question. Our partial theories of language are still about human cognition; they are still answers to the question of knowledge. We have very different grammars, but we still think about them in the same way, as computational characterizations of levels of cognitive representation. We struggle to find a place for diversity within a framework that is still universalistic. We look at language as a social entity from the point of view of social *psychology*.

At the most foundational level, then, what I intend to suggest in this book amounts to the claim that the pieces of the puzzle, the partial answers, would only fall into place if we agree to look at them as partial answers to a very different *question*, if we manage to re-interpret them, as Hempel puts it, under the light of a new network of terms, definitions, and hypotheses. To find language again, to assemble the pieces, we need a new organizing principle, a new theoretical characterization of the *essence* of language. I will suggest that we need to look at language as a socially constructed communication technology, not as a cognitive capacity.

1.2 Language as a Communication Technology

The general idea that language is a socially constructed tool of communication is far from new. It has featured prominently in the linguistic discourse of the first half of the twentieth century (Saussure 1916, 1966, Sapir 1921, Meillet 1921, Gardiner 1932, and others); it has informed much of the semiotic literature on language (Itkonen 2003, 2008, Zlatev 2009); and it has been debated in the philosophy of language (Wittgenstein 1953, Lewis 1969, Davidson 1984, Dummett 1996). Most of the other disciplines interested in language—communication studies (Peters 2000), sociology, anthropology, literature, critical studies, and so on—have always worked with a general and often implicit view of language as an institution of society.

In the wake of the Chomskian revolution in linguistics, the idea of language as a tool was explicitly marginalized. Reddy's (1979) *toolmakers' paradigm* was a notable exception: its influence on my theory is enormous, and I will get back to it in chapter 3. In the last three decades, however, the idea has gradually come to be accepted, implicitly or explicitly, in many quarters of the linguistic sciences: if the capacities involved in language acquisition and usage are general human

capacities, rather than innate pieces of linguistic knowledge, then language itself, the object of acquisition, must be out there—available for the child in the social-cultural environment. “This perspective on human communication and language,” writes Tomasello (2008, p. 10–11), “thus basically turns the Chomskian proposal on its head, as the most fundamental aspects of human communication are seen as biological adaptations for cooperation and social interaction in general, whereas the more purely linguistic, including grammatical, dimensions of language are culturally constructed and passed along by individual linguistic communities.” For Croft (2000, p. 26), “a language is the population of utterances in a speech community,” whereas “a grammar is the cognitive structure in a speaker’s mind that contains her knowledge of her language.” Based on his lifelong inquiry into the lives of the *Pirahã* in Brazil, Everett (2012, p. 6) says that “languages are tools. Tools to solve the twin problems of communication and social cohesion. Tools shaped by the distinctive pressures of their cultural niches—pressures that include cultural values and history and which account in many cases for the similarities and differences between languages.”

These formulations and others are very important developments. They do not, however, go beyond the level of declaration. As Botha (1992, p. 237) puts it, “to say that something—for example an entity, structure, practice or whole realm of reality—is ‘social’ is not to characterize it ontologically in a fundamental way.” The authors position language in the social domain and then move on to more pressing issues: Tomasello to the social psychology of communication, Croft to language change, Everett to the relationship between language and culture. Under some readings, this theoretical dynamic has also characterized the writings of de Saussure and Sapir: according to Koerner (1982), de Saussure was interested in the development of a general theory of signs—a semiological, or semiotic theory—and he “appears to have claimed that language is a social fact just because of its semiological character” (p. 57). Sapir’s work often seems to deal more with the psychological significance of language than with its social nature. The same is true of the literature in sociolinguistics and anthropological linguistics: the social-cultural essence of language is stated informally and then taken for granted, and the discourse concentrates on the web of relationships between language and other social-cultural facts—identity, class, power, values, gender, and so on. This is what Coulmas (1998, p. 3) calls “the theoretical deficit of sociolinguistics.” In the philosophy of language, the debates revolve around the question of meaning and its determination vis-à-vis the questions of convention, usage, and truth. With the exception of Wittgenstein, there is relatively little that attempts to investigate the ontological foundations of the social reality of language. There is a regular pattern here: the declarative endorsement of the social-cultural nature of language has allowed researchers, past and present, to

delve deeply into different aspects of the context within which language operates, with extremely significant results that will play a central role in the construction of my theory. Language itself, however, the social-cultural entity as such, has somehow been neglected on the way: it is still severely under-theorized.

What are we looking for, then? Well, basically what we need is a much better understanding of what language *does*—a theoretical characterization of its *function* as a socially constructed tool of communication. Crucially, we should not confuse the question of what language does with the question of what we, as communicators, *do with it*. These are two separate levels of functional analysis. To see the difference, consider the following example. Suppose I wish to inform you that I'm arriving on Tuesday, and ask you to wait for me at the train station at 7:15 p.m. And assume that I can do this in three different ways: send you a fax, an email, or a regular letter by mail. So, two questions: (a) What would I do with each of these technologies? And (b) What would each of the technologies do if I chose to use it? As far as the first question is concerned, it seems I'll be doing the same two things in all three cases: telling you something and asking you to do something. In Jakobson's (1960) terms, I will be performing two functions: the referential and the conative. In Searle's (1969) terms, I will be performing two speech acts: the assertive and the directive. If we wish to break the speech acts down to their locutionary, illocutionary, and perlocutionary components (Austin 1962), we'll find that nothing there hinges on my choice of technology either. As far as the second question is concerned, however, each of the technologies would give us a radically different answer. The fax would look at my written note as visual information, convert it into a bitmap, and transmit it directly through the phone system to your fax. My email software would read my note as text, convert it into an electronic signal, and send it to my email server, from where it would be re-directed through the Internet to your server, and then to your computer. The mail system would require an envelope and a stamp, and it would physically carry my written note, through a series of processing centers, all the way to your door. The essence of the three technologies does not lie in what I do with them, but in the unique and specific functional strategies that they employ as technologies. This, then, is the first point: to understand language as a technology, we have to understand what it does, not what we do with it. We have to figure out its functional strategy as such.

Positioning the question of functional specificity at the center of our theoretical quest, we may finally break away from the deadlocked debate over the functionality of language as it developed on Chomsky's field, under Chomsky's question, in the last fifty years. The debate revolved around a much more fundamental question: Is language functional or not in the first place? The generativists claimed that it simply was not; the functionalists insisted that it was a

(cognitively based) general-purpose system of communication (for a thorough discussion of the debate, see Newmeyer 1998). If we agree to think of language as a technology, we have to immediately reject both positions. On the one hand, the essence of a functioning technology simply cannot lie in the non-functional specificity of its architecture: non-functionality may explain *useless* technology, but it is not an explanatory option with such an enormously successful technology as language. On the other hand, no technology is ever general-purpose. Just like any other technology, the essence of language must lie in the fact that its function is specific. This is where we should start.

As I will show throughout the book, this new perspective does much more than simply reject the two rivaling positions. It actually allows us to accept many of the *positively* formulated *observations* of both camps, while rejecting their *negatively* formulated *explanations*. No, language is not non-functional or non-specific. As a technology, it cannot be any of those. Yes, language is functional *and* specific. As a technology, it must be both. From our new perspective, we will find again and again that the two camps looked at two sides of the same coin, concentrating on the observations of their own side, and working hard to discredit the idea that there is another. Released from the debate, we will be able to re-evaluate the relevant observations and assign them a new meaning. Having rejected Chomsky's question, we will even find, from time to time, that we should adopt some of his answers.

1.3 The Structure of the Argument

The next two chapters provide a first systematic look at the spatial network of the theory—its terms, definitions, and hypotheses. In chapter 2, I define the specific functional strategy of language, and show how it is different from the strategies employed by all the other communication systems, technological or not, that are used by us and other biological species. In chapter 3, I present a technical description of language the way I see it—its constitutive parts, their social construction, the way they fit, how they function together to allow for the production and interpretation of complex sets of instructions for imagination, and the parameters that determine the overall quality of the instruction.

Following this definitional discussion, chapters 4 and 5 descend to the plane of observation and deal with a set of issues having to do with *meaning*. In chapter 4, I show how the theory re-frames some of the foundational questions of *word* meaning: how do words mean? How do their meanings connect them to the world of experience and to other words? Why do they manifest the very particular behavioral patterns that they do? In chapter 5, I present a major

re-interpretation of the question of *linguistic relativity*, as a question about the dialectic influence of a technology on its users—a re-interpretation that positions the question in its rightful place at the very center of linguistic theory, assigns the accumulated results in the field with new significance, and opens the way toward new ways of research.

Chapter 6 takes us back up to the spatial network of the theory, with a more detailed definitional discussion of the processes involved in the production and comprehension of linguistic utterances. It then descends back to the plane of observation, and shows how the definitional description captures some of the most important recent findings in psycholinguistics and how the stormy relationship between psycholinguistics and general linguistics can be rationalized. In chapter 7, I use all this to deconstruct the question of syntax, demonstrate the very different way the theory handles syntactic complexity, and claim that Chomsky was actually right in his insistence on the autonomy of syntax from general, individual cognition—but for the wrong reason: syntactic complexity is not a matter of individual cognition; it is socially constructed, prescriptive, and specifically suited for the instruction of imagination.

Chapters 8 and 9 are all about *variability*. In chapter 8, I show how the theory captures the attested patterns of linguistic diversity around the world, and how it re-conceptualizes the universality of language as a foundationally social fact—not a cognitive one. This carries important implications for research. In chapter 9, I claim that language-related cognitive capacities are also variably spread within communities of speakers. The universal mind of the human speaker is a myth, and so is the universal mind of the language-acquiring child. This move allows for a new hypothesis about the collective essence of the process of language acquisition, which also sheds new light on ongoing construction of sign languages around the world, in Nicaragua and elsewhere.

Chapter 10 connects all these issues together in a new hypothetical explanation of the evolution of language as a collectively constructed communication technology. Based on my work with evolutionary biologist Eva Jablonka, and on a wealth of insights from current evolutionary biology, the evolution of the human species in general, the evolution of technology, and the evolution of language, I claim that the specific function of the instruction of imagination was collectively invented, on the basis of everything that was achieved by pre-linguistic societies, before individual speakers were specifically adapted to it. Our variable, innate dispositions toward the participation in the collective activities of language emerged *for* a technology that was already there in the social domain. In the conclusion, I show how the theory, with its derivative explanations of the different areas on the plane of observation, begins to allow for the systematic reassembly of the puzzle of language.