Conclusion

The most surprising aspect of the organization of language and thought is perhaps the fundamental unity of the cognitive operations that serve to construct the simple meanings of everyday life, the common sense reasoning of our daily existence, the more elaborate discussions and arguments that we engage in, and the superficially far more complex scientific theories and artistic and literary productions that entire cultures develop over the course of time. We have seen that the simplest meanings are in fact not simple at all. They rely on remarkable cognitive mapping capacities, immense arrays of intricately prestructured knowledge, and exceptional on-line creativity. They also rely on the impressive, and poorly understood, human ability to resolve massive underspecification at lightning speeds.

The cognitive operations focused on in the present work include cross-space mapping, access and spreading principles, space-tracking by shifting viewpoint and focus, matching, structure projection, and conceptual blending. We are not conscious of performing these operations when we speak, think, and listen. Everything takes place very fast, and only the words themselves and the global emergence of meaning are accessible to consciousness. At other levels of thought, such as science, poetry, or rhetoric, there may be more awareness of some of the operations. We may consciously perceive an

analogy, a metaphor, or a metonymy. Yet, typically, even then, most of the efficient cognitive processes are hidden from view, and their overall structure is seldom directly apprehended.

Language is not a code for such operations. Their complexity far exceeds the overt information that a language form could carry. Instead, language serves to prompt the cognitive constructions, by means of very partial, but contextually very efficient, clues and cues. Our subjective impression, as we speak and listen, is that when language occurs, meaning directly ensues, and therefore that meaning is straightforwardly contained in language. This fiction is harmless in many areas of everyday life - buying groceries or going fishing -, but may well be quite pernicious in others - trials, politics, and deeper social and human relationships.

As in other areas of science, knowledge is ignorance. Catching a glimpse of the elaborate mappings and blends that operate in meaning construction also reveals our vast ignorance of such processes. At this time, no one has a clear idea of how the brain carries out the high level operations in question. And there is no computational model that captures the essence of analogy, blends, or domain matching.[1] Clearly, symbolic rule production systems are inadequate. The consensus is that a connectionist neural network architecture is closest to what the brain seems to be doing. And yet, in the areas studied in this book, there is multiple binding from one space to another, from roles to values, and from multiple inputs to blends. Multiple binding of this sort, structure projection, and gist extraction[2] represent a formidable challenge for any modeling approach.

When we deal with a single language, the complexities of modeling meaning do not necessarily jump out, especially if the context has been artificially restricted to a micro-world, as in Winograd's celebrated SHRDLU.[3] It is perhaps in the domain of machine translation that some researchers have become most acutely aware of what Oettinger (1963) aptly calls the "very mysterious semantic processes that enable most reasonable people to interpret most reasonable sentences unequivocally most of the time."[4] The reason translation reveals some of the hidden complexity is that different languages have developed different ways to prompt the required cognitive constructions. In addition, of course, different cultures organize their background knowledge differently. Good translation, then, requires a quasi-total reconstruction of the cognitive configurations prompted by one language and a determination of how another language would set up a similar configuration with a radically different prompting system and prestructured background.[5] It was commonly assumed for a long time that this type of difficulty was a feature of literary translation, but that in everyday 'objective' domains, like commerce, science, or industry, the differences would be tractable. This has turned out not to be the case, which comes as no surprise in the context of the present book: the nature and the complexity of the cognitive operations are the same in all cases.

When meaning construction is taken into account, the fundamental cognitive issues of learning and evolution appear in a different light. Clearly, what children learn is not language structure in the abstract. They acquire entire systems of mappings, blends, and framing, along with their concomitant language manifestations. If anything, the poverty of stimulus argument[6] is

even stronger for the learning of such elaborate systems than for syntax alone. There is little doubt that children come into the world innately endowed with powerful capacities to develop such systems in the appropriate environments. But it is also true that the required cognitive operations operate generally in thought and action. They are not specific to language, and the poverty of stimulus argument is not an argument for an autonomous language faculty. In fact, as we have seen repeatedly throughout the present work, the cognitive operations that play a central role in the construction of everyday meaning are also the operations that apply to reasoning, thinking, and understanding quite generally. There might still be, of course, purely structural language universals that hold independently of the system of meaning that they help to deploy. But we cannot know this by studying word combinations in isolation, any more than we can understand the structural properties of addition or multiplication routines across cultures without reference to their mathematical function. For a child, to know a grammar is not primarily to know which strings are well-formed or ill-formed; it is to know how to apply partial grammatical instructions in context to produce appropriate cognitive configurations. For a cognitive scientist, linguist, or philosopher of language, to understand the nature of grammar, and to analyze specific grammars, is to provide explicit accounts of how grammatical constructions contribute in context to the elaboration of cognitive configurations.[7]

Most of the data in this book comes from language in the spoken modality. But there is now impressive evidence for the key role of mental spaces and mappings in signed languages (Liddell (1995a, 1995b), Van Hoek

(1996), Poulin (1996)).[8] Because the modality is different, overt linguistic manifestations of mental space organization can be quite different and revealing in signed languages. Scott Liddell, in the work cited, has developed a powerful notion of grounded mental space. He shows how grounded conceptual spaces are part of our general thinking capacities, and how they are put to specific use, and grammatically signaled, in the anaphoric system of ASL. These important results fit in with our general theme that meaning at the most basic levels is supported and driven by general, not specifically linguistic, cognitive operations.

The other major cognitive issue is the question of how language appeared in the course of biological evolution. We seem to lack the continuity often found for development of organs and physical capacities. There is no record of species with rudimentary language abilities that would provide missing links between the absence of language and its full blown instantiation in homo sapiens. However, if we consider the bigger picture outlined in this book, there is perhaps more continuity than we had thought. The essence of language, under the view I have presented, is the meaning construction system - mappings, frames, and spaces. The words and sentences are a surface manifestation of this complex activity. And there is every reason to think that the general cognitive processes we have considered are not restricted to humans. Other biological organisms (mammals in particular) engage in building frames, projecting structure, and making analogies, and many species have elaborate social structures and cultural models, aspects of which are internalized by individual members.

Footnotes

[1] Douglas Hofstadter, in his review of the book *Mental Leaps*, by Holyoak and Thagard, points to serious problems in contemporary modeling of analogy. Hofstadter's view is that the hardest part of analogical reasoning, gist extraction, is already done covertly by the analysts themselves in the way that they set up the inputs to their models. The modelers and those who evaluate their models are thus victims of a 'giant ELIZA effect'. See Hofstadter (1995b)

[2] As we saw in Chapters IV and VI, only very partial structures get exploited for the purpose of analogy formation, metaphor, and blending. Extracting the right partial structures - the gist - from the considerable amount of structure and organization available for any real domain is beyond the capabilities of current models, which skirt the issue by focusing on the analogy process after suitable gist-extraction.

[3] Winograd (1972).

[4] Cited in Dreyfus (1979). Dreyfus discusses very cogently some of the considerable problems posed by natural language understanding for explicit computation or formalization. The phenomena we have studied in the present book suggest that the obstacles to modeling are even greater than Dreyfus envisione

[5] Mandelblit (1995a).

[6] Noam Chomsky has often argued that the evidence available to a child learning a language in a particular environment radically underdetermines the syntactic system actually learned by the child. See, for example, Chomsky (1972).

[7] This very ancient tradition of language study - trying to explain what grammar does, not just what grammar looks like - has been revived in cognitive-functional linguistics (e.g. Langacker (1987, 1991), Deane, Croft, Fillmore and Kay, Fauconnier & Sweetser (1996), ...). It goes far beyond the narrow characterization of syntactic well-formedness privileged by structuralist and generative approaches.

[8] This line of research was first undertaken, in unpublished work, by Richard Lacy in the late nineteen seventies.