

This PDF includes a chapter from the following book:

Linguistics for the Age of AI

© 2021 Marjorie McShane and Sergei Nirenburg

License Terms:

Made available under a Creative Commons
Attribution-NonCommercial-NoDerivatives 4.0 International Public License

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

OA Funding Provided By:

The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.

The title-level DOI for this work is:

[doi:10.7551/mitpress/13618.001.0001](https://doi.org/10.7551/mitpress/13618.001.0001)

Epilogue

This book presented an approach to operationalizing natural language understanding capabilities. The approach is rooted in the hypothesis that it is both scientifically fruitful and practically expedient to model intelligent agents after the humans they are intended to emulate. This involves endowing agents with a host of interconnected capabilities that represent our best functional approximation of *modeling what people do, how they do it*. This *how* is not at the level of brain or biology; it is at the level of folk-psychological explanation, introspection, and commonsense human reasoning. The multifaceted interconnectedness of cognitive processes does not lend itself to a reassuringly small and targeted program of R&D—and that is the rationale behind the large-scale program of R&D that we have dubbed Linguistics for the Age of AI.

Generations of linguists have acknowledged that the comprehensive analysis of language use must invoke world knowledge, general reasoning, linguistic reasoning, mindreading of the interlocutor, and the interpretation of nonlinguistic features of the real world. But, having acknowledged this, the vast majority have chosen to work on quite narrowly defined linguistic subproblems. This is understandable on two counts: first, modularization fosters the development of certain kinds of theories; and, second, linguists are not necessarily drawn to modeling all of the nonlinguistic capabilities that interact with linguistic ones.

In this book we have attempted to explain why holistically addressing a broad range of cognitive capabilities is the only realistic path toward cracking the problem of natural language understanding and artificial intelligence overall. It is not that we are *choosing* to solve dozens of problems rather than select a more manageable handful; instead, we are acknowledging the inevitability of this course of action. For example, dynamically tracking the plans and goals of one's interlocutors is a well-known basic challenge of AI, and one that linguists might not consider a first priority, but it is the key to fully interpreting elliptical and fragmented dialog turns. That is, any attempt to interpret incomplete utterances without understanding their function in the discourse would have nothing in common with what people do and, at best, it would be no more than a temporary stopgap to give a system the veneer of intelligence.

Let us return to the notion of *the knowledge bottleneck*, which has been a key reason why knowledge-based methods have been at the periphery of AI for decades. In order for agents to become truly humanlike, they must acquire knowledge about language and the world through a process of lifelong learning. From preschool to their college years and beyond, people learn largely by reading and interacting—in natural language!—with other people. Our goal should be to impart this capability to artificial agents. To bootstrap this process, agents must be endowed with an initial knowledge base of sufficient size and quality. The ontology/lexicon knowledge base that we developed in our lab is a good candidate. People will have to help the system learn by providing explanations, answering questions, and checking the quality of the system’s output. Importantly, the labor requirements for this kind of project will not be unusual for NLP/NLU. In fact, the machine learning community has devoted extensive resources to the manual preparation of datasets over the past several decades, and it continues to do so with no signs of letting up. If those resources had, instead, gone toward building the kinds of knowledge bases we describe in this book, it is entirely possible that we would already be benefiting from highly functional LEIAs who would already be overcoming the knowledge bottleneck with an effective level of lifelong learning.

Knowledge-based and statistical methods are not in competition; they offer different methods of achieving different types of analysis results that can serve as input to intelligent systems. The most plausible path to human-level AI is to integrate the results of both of these approaches into hybrid environments. In fact, combining the best that each approach has to offer has been at the center of attention of AI researchers for quite some time. However, because of the high profile of machine learning, most of the thinking has focused on how to improve the results of current applications by adding a sprinkling of stored, human-acquired knowledge. We believe that this is the exact opposite of the most promising long-term direction. We hypothesize that the most fruitful path of integration will be for statistical methods both to support an agent’s lifelong learning and to supply knowledge-based systems with high-quality modules for subtasks that lend themselves well to knowledge-lean approaches—such as named-entity recognition and syntactic parsing.

While writing this epilogue, we fortuitously attended a lecture by Jeffrey Siskind, who, like us, is working toward integrating multiple cognitive capabilities in AI systems. When asked to project into the future, he explained that he thinks about the future only to the level of his “great-grandstudents”: anything beyond that is over the horizon. This resonated with us. The field of computational linguistics—with *linguistics* as a central topic of study—has been woefully underexplored over the past generation. This means that a tremendous amount remains to be done, and there is no telling what the field might look like in forty years’ time. We hope that the program of research that we call Linguistics for the Age of AI will serve as a shot in the arm to the linguistics community, renewing excitement in the challenge of endowing agents with human-level language capabilities.