

Contents

Acknowledgments xi

Setting the Stage xiii

1 Our Vision of Linguistics for the Age of AI 1

1.1 What Is Linguistics for the Age of AI? 1

1.2 What Is So Hard about Language? 2

1.3 Relevant Aspects of the History of Natural Language Processing 4

1.4 The Four Pillars of Linguistics for the Age of AI 8

1.4.1 Pillar 1: Language Processing Capabilities Are Developed within an Integrated, Comprehensive Agent Architecture 9

1.4.2 Pillar 2: Modeling Is Human-Inspired in Service of Explanatory AI and Actionability 13

1.4.3 Pillar 3: Insights Are Gleaned from Linguistic Scholarship and, in Turn, Contribute to That Scholarship 15

1.4.3.1 Theoretical syntax 16

1.4.3.2 Psycholinguistics 17

1.4.3.3 Semantics 17

1.4.3.4 Pragmatics 20

1.4.3.5 Cognitive linguistics 22

1.4.3.6 Language evolution 23

1.4.4 Pillar 4: All Available Heuristic Evidence Is Incorporated When Extracting and Representing the Meaning of Language Inputs 25

1.4.4.1 Handcrafted knowledge bases for NLP 25

1.4.4.2 Using results from empirical NLP 27

1.5 The Goals of This Book 29

1.6 Deep Dives 31

1.6.1 The Phenomenological Stance 31

1.6.2 Learning 33

1.6.3 NLP and NLU: It's Not Either-Or 33

1.6.4 Cognitive Systems: A Bird's-Eye View 36

1.6.5 Explanation in AI 38

1.6.6 Incrementality in the History of NLP 40

1.6.7 Why Machine-Readable, Human-Oriented Resources Are Not Enough 42

1.6.8 Coreference in the Knowledge-Lean Paradigm 44

1.6.9	Dialog Act Detection	46
1.6.10	Grounding	48
1.6.11	More on Empirical NLP	50
1.6.12	Manual Corpus Annotation: Its Contributions, Complexities, and Limitations	52
1.7	Further Exploration	55
2	A Brief Overview of Natural Language Understanding by LEIAs	59
2.1	Theory, Methodology, and Strategy	59
2.2	A Warm-Up Example	64
2.3	Knowledge Bases	68
2.3.1	The Ontology	69
2.3.2	The Lexicon	73
2.3.3	Episodic Memory	77
2.4	Incrementality	77
2.5	The Stages of NLU and Associated Decision-Making	79
2.5.1	Decision-Making after Pre-Semantic Analysis	84
2.5.2	Decision-Making after Pre-Semantic Integration	85
2.5.3	Decision-Making after Basic Semantic Analysis	85
2.5.4	Decision-Making after Basic Coreference Resolution	86
2.5.5	Decision-Making after Extended Semantic Analysis	87
2.5.6	Decision-Making after Situational Reasoning	87
2.6	Microtheories	88
2.7	“Golden” Text Meaning Representations	93
2.8	Deep Dives	94
2.8.1	The LEIA Knowledge Representation Language versus Other Options	94
2.8.2	Issues of Ontology	97
2.8.3	Issues of Lexicon	102
2.8.4	Paraphrase in Natural Language and the Ontological Metalanguage	111
2.9	Further Exploration	115
3	Pre-Semantic Analysis and Integration	117
3.1	Pre-Semantic Analysis	117
3.2	Pre-Semantic Integration	118
3.2.1	Syntactic Mapping: Basic Strategy	118
3.2.2	Recovering from Production Errors	124
3.2.3	Learning New Words and Word Senses	124
3.2.4	Optimizing Imperfect Syn-Maps	126
3.2.5	Reambiguating Certain Syntactic Decisions	128
3.2.6	Handling Known Types of Parsing Errors	129
3.2.7	From Recovery Algorithm to Engineering Strategy	131
3.3	Managing Combinatorial Complexity	134
3.4	Taking Stock	138
3.5	Further Exploration	139

4	Basic Semantic Analysis	141
4.1	Modification	143
4.1.1	Recorded Property Values	143
4.1.2	Dynamically Computed Values for Scalar Attributes	145
4.1.3	Modifiers Explained Using Combinations of Concepts	149
4.1.4	Dynamically Computed Values for Relative Text Components	150
4.1.5	Quantification and Sets	152
4.1.6	Indirect Modification	158
4.1.7	Recap of Modification	159
4.2	Proposition-Level Semantic Enhancements	160
4.2.1	Modality	160
4.2.2	Aspect	162
4.2.3	Non-Modal, Non-Aspectual Matrix Verbs	163
4.2.4	Questions	163
4.2.5	Commands	164
4.2.6	Recap of Proposition-Level Semantic Enhancements	165
4.3	Multicomponent Entities Recorded as Lexical Constructions	165
4.3.1	Semantically Null Components of Constructions	169
4.3.2	Typical Uses of Null-Semming	169
4.3.3	Modification of Null-Semmed Constituents	170
4.3.4	Utterance-Level Constructions	172
4.3.5	Additional Knowledge Representation Requirements	173
4.3.6	Recap of Constructions	175
4.4	Indirect Speech Acts, Lexicalized	175
4.5	Nominal Compounds, Lexicalized	176
4.6	Metaphors, Lexicalized	180
4.6.1	Past Work on Metaphor	180
4.6.2	Conventional Metaphors	182
4.6.3	Copular Metaphors	184
4.6.4	Recap of Metaphors	185
4.7	Metonymies, Lexicalized	185
4.8	Ellipsis	186
4.8.1	Verb Phrase Ellipsis	186
4.8.2	Verb Phrase Ellipsis Constructions	187
4.8.3	Event Ellipsis: Aspectual + NPOBJECT	189
4.8.4	Event Ellipsis: Lexically Idiosyncratic	189
4.8.5	Event Ellipsis: Conditions of Change	190
4.8.6	Gapping	192
4.8.7	Head Noun Ellipsis	192
4.8.8	Recap of Ellipsis	192
4.9	Fragmentary Utterances	193
4.10	Nonselection of Optional Direct Objects	193
4.11	Unknown Words	193
4.11.1	Completely Unknown Words	194
4.11.2	Known Words in a Different Part of Speech	196
4.12	Wrapping Up Basic Semantic Analysis	198
4.13	Further Exploration	199

5 Basic Coreference Resolution 201

- 5.1 A Nontechnical Introduction to Reference Resolution 201
 - 5.1.1 Definitions 202
 - 5.1.2 An Example-Based Introduction 203
 - 5.1.3 A Dozen Challenges 205
 - 5.1.4 Special Considerations about Ellipsis 210
 - 5.1.5 Wrapping Up the Introduction 211
- 5.2 Personal Pronouns 212
 - 5.2.1 Resolving Personal Pronouns Using an Externally Developed Engine 213
 - 5.2.2 Resolving Personal Pronouns Using Lexico-Syntactic Constructions 213
 - 5.2.3 Semantically Vetting Hypothesized Pronominal Coreferences 215
 - 5.2.4 Recap of Resolving Personal Pronouns during Basic Coreference Resolution 217
- 5.3 Pronominal Broad Referring Expressions 217
 - 5.3.1 Resolving Pronominal Broad RefExes Using Constructions 218
 - 5.3.2 Resolving Pronominal Broad RefExes in Syntactically Simple Contexts 219
 - 5.3.3 Resolving Pronominal Broad RefExes Indicating Things That Must Stop 221
 - 5.3.4 Resolving Pronominal Broad RefExes Using the Meaning of Predicate Nominals 223
 - 5.3.5 Resolving Pronominal Broad RefExes Using Selectional Constraints 224
 - 5.3.6 Recap of Resolving Pronominal Broad RefExes 226
- 5.4 Definite Descriptions 227
 - 5.4.1 Definite Description Processing So Far: A Refresher 227
 - 5.4.2 Definite Description Processing at This Stage 228
 - 5.4.2.1 Rejecting coreference links with property value conflicts 228
 - 5.4.2.2 Running reference-resolution meaning procedures listed in lexical senses 228
 - 5.4.2.3 Establishing that a sponsor is not needed 229
 - 5.4.2.4 Identifying bridging references 229
 - 5.4.2.5 Creating sets as sponsors for plural definite descriptions 231
 - 5.4.2.6 Identifying sponsors that are hypernyms or hyponyms of definite descriptions 232
 - 5.4.3 Definite Description Processing Awaiting Situational Reasoning 233
 - 5.4.4 Recap of Definite Description Processing at This Stage 233
- 5.5 Anaphoric Event Coreference 234
 - 5.5.1 What is the Verbal/EVENT Head of the Sponsor? 235
 - 5.5.2 Is There Instance or Type Coreference between the Events? 237
 - 5.5.3 Is There Instance or Type Coreference between Objects in the VPs? 238
 - 5.5.4 Should Adjuncts in the Sponsor Clause Be Included in, or Excluded from, the Resolution? 239
 - 5.5.5 Should Modal and Other Scopers Be Included in, or Excluded from, the Resolution? 239
 - 5.5.6 Recap of Anaphoric Event Coreference 240
- 5.6 Other Elided and Underspecified Events 241
- 5.7 Coreferential Events Expressed by Verbs 242
- 5.8 Further Exploration 244

6	Extended Semantic Analysis	247
6.1	Addressing Residual Ambiguities	248
6.1.1	The Objects Are Linked by a Primitive Property	248
6.1.2	The Objects Are Case Role Fillers of the Same Event	249
6.1.3	The Objects Are Linked by an Ontologically Decomposable Property	249
6.1.4	The Objects Are Clustered Using a Vague Property	251
6.1.5	The Objects Are Linked by a Short Ontological Path That Is Computed Dynamically	252
6.1.6	Reasoning by Analogy Using the TMR Repository	252
6.1.7	Recap of Methods to Address Residual Ambiguity	254
6.2	Addressing Incongruities	254
6.2.1	Metonymy	254
6.2.2	Preposition Swapping	256
6.2.3	Idiomatic Creativity	257
6.2.3.1	Detecting creative idiom use	258
6.2.3.2	Semantically analyzing creative idiom use	261
6.2.4	Indirect Modification Computed Dynamically	262
6.2.5	Recap of Treatable Types of Incongruities	264
6.3	Addressing Underspecification	264
6.3.1	Nominal Compounds Not Covered by Lexical Senses	264
6.3.2	Missing Values in Events of Change	270
6.3.3	Ungrounded and Underspecified Comparisons	270
6.3.4	Recap of Treatable Types of Underspecification	279
6.4	Incorporating Fragments into the Discourse Meaning	279
6.5	Further Exploration	284
7	Situational Reasoning	285
7.1	The OntoAgent Cognitive Architecture	286
7.2	Fractured Syntax	287
7.3	Residual Lexical Ambiguity: Domain-Based Preferences	290
7.4	Residual Speech Act Ambiguity	290
7.5	Underspecified Known Expressions	291
7.6	Underspecified Unknown Word Analysis	291
7.7	Situational Reference	292
7.7.1	Vetting Previously Identified Linguistic Sponsors for RefExes	293
7.7.2	Identifying Sponsors for Remaining RefExes	296
7.7.3	Anchoring the TMRs Associated with All RefExes in Memory	297
7.8	Residual Hidden Meanings	297
7.9	Learning by Reading	299
8	Agent Applications: The Rationale for Deep, Integrated NLU	301
8.1	The Maryland Virtual Patient System	301
8.1.1	Modeling Physiology	304
8.1.2	An Example: The Disease Model for GERD	306
8.1.3	Modeling Cognition	311
8.1.3.1	Learning new words and concepts through language interaction	312
8.1.3.2	Making decisions about action	313

8.1.4	An Example System Run	317
8.1.5	Visualizing Disease Models	320
8.1.5.1	Authoring instances of virtual patients	320
8.1.5.2	The knowledge about tests and interventions	321
8.1.5.3	Traces of system functioning	324
8.1.6	To What Extent Can MVP-Style Models Be Learned from Texts?	324
8.1.7	To What Extent Can Cognitive Models Be A Automatically Elicited from People?	328
8.2	A Clinician's Assistant for Flagging Cognitive Biases	331
8.2.1	Memory Support for Bias Avoidance	333
8.2.2	Detecting and Flagging Clinician Biases	335
8.2.3	Detecting and Flagging Patient Biases	339
8.3	LEIAs in Robotics	343
8.4	The Take-Home Message about Agent Applications	347
9	Measuring Progress	349
9.1	Evaluation Options—and Why the Standard Ones Don't Fit	350
9.2	Five Component-Level Evaluation Experiments	354
9.2.1	Nominal Compounding	355
9.2.2	Multiword Expressions	358
9.2.3	Lexical Disambiguation and the Establishment of the Semantic Dependency Structure	362
9.2.4	Difficult Referring Expressions	365
9.2.5	Verb Phrase Ellipsis	366
9.3	Holistic Evaluations	369
9.4	Final Thoughts	381
	Epilogue	383
	Notes	385
	References	397
	Index	415