

9 33. Word meaning and world knowledge

10 1. Introduction

11 2. Core abstract theories

12 3. Linking word meaning with the theories

13 4. Distinguishing lexical and world knowledge

14 5. References

15

16 *Lexical semantics should be in part about linking the meanings of words with*
17 *underlying theories of the world. But for this to be even remotely possi-*
18 *ble, the theories need to be informed by the insights of cognitive and other*
19 *linguists about the conceptual structure on which language is based. They*
20 *have to be axiomatizations of a kind of abstract topology that, for example,*
21 *includes the domains of composite entities (things made of other things),*
22 *scalar notions, change of state, and causality. Theories of each of these*
23 *domains are sketched briefly, and it is shown how three very common pol-*
24 *ysemous words can be defined or characterized in terms of these theories.*
25 *Finally, there is a discussion of what sort of boundary one can hope to draw*
26 *between lexical knowledge and other world knowledge.*

27 **1. Introduction**

28 We use words to talk about the world. Therefore, to understand what words
29 mean, we should have a prior explication of how we view the world.

30 Suppose we have a formal logical theory of some domain, or some aspect
31 of the world, that is, a set of predicates intended to capture the concepts in
32 that domain and a set of axioms or rules that constrain the possible meanings
33 of those predicates. Then a formal theory of lexical semantics in that domain
34 would be a matter of writing axioms to relate predicates corresponding to
35 the words in the domain to the predicates in the underlying theory of the
36 domain. For example, the word “until” might be anchored in a formal theory
37 of time that provides an axiomatization of intervals and a *before* relation.
38 (See article 29 *Frame semantics* for a similar view, where frames correspond
39 to the domain theory.)

40 For the last forty years researchers in artificial intelligence have made
41 efforts to encode various aspects of world knowledge formally. These efforts
42 have primarily been in commonsense physics in the areas of space, time, and
43 qualitative physics, and, in commonsense psychology, in concepts related to
44 belief and intention. A good review of this work that is old but has not lost
45 its relevance is Davis (1990). Most of this work has focused on narrow areas
46 of commonsense knowledge. But there have been several large-scale efforts
47 to encode knowledge of many domains, most notably, Cyc (Lenat & Guha
48 1990; Cycorp 2008). One might think that this work could form the basis

49 of an effort toward a formal theory of lexical semantics anchored in world
50 knowledge. However, these theories for the most part were not designed with
51 language in mind, and in particular what is missing is precisely some of the
52 linguists' insights described in the previous several articles of this volume.
53 All of this seriously undercuts the utility for lexical semantics of Cyc and
54 similar large ontologies, and indeed of most of the small-scale theories as
55 well.

56 In trying to link words and world, there are a number of bad ways to
57 go about it. For example, we could take our theory of the world to be
58 quantum mechanics and attempt to define, say, verbs of motion in terms of
59 the primitives provided by that theory. A less obviously wrong approach,
60 and one that has sometimes been tried, is to adopt Euclidean 3-space as the
61 underlying model of space and attempt to define, say, spatial prepositions in
62 terms of that. More common is a serious misstep, with respect to language,
63 that many large-scale ontologies take at the start. Cyc begins by enforcing
64 a rigid distinction between tangible and intangible entities, and in other
65 hierarchical ontologies, the top-level split is between physical and abstract
66 entities. Yet this distinction plays very little role in language. We can be
67 in a room, in a social group, in the midst of an activity, in trouble, and in
68 politics. We can move a chair from the desk to the table, move money from
69 one bank account to another, move a discussion from religion to politics, and
70 move an audience to tears. A fundamental distinction between tangibles

71 and intangibles rules out the possibility of understanding the sense of “in”
72 or “move” common to all these uses.

73 Our effort, by contrast, has sought to exploit the insights of linguists
74 such as Gruber (1965), the generative semanticists, Johnson (1987), Lakoff
75 (1987), Jackendoff (see article 31 *Conceptual semantics*), and Talmy (see
76 article 27 *Cognitive semantics: An overview*). Johnson, Lakoff, Talmy and
77 others have used the term “image schemas” to refer to a conceptual frame-
78 work that includes topological relations but excludes, for example, Euclidean
79 notions of magnitude and shape. We have been developing core theories that
80 formalize something like the image schemas, and we have been using these to
81 define or characterize words. Among the theories we have developed are the-
82 ories of composite entities, or things made of other things, the figure-ground
83 relation, scalar notions, change of state, and causality. The idea behind these
84 abstract core theories is that they capture a wide range of phenomena that
85 share certain features. The theory of composite entities, for example, is in-
86 tended to accommodate natural physical objects like volcanos, artifacts like
87 automobiles, complex events and processes like concerts and photosynthesis,
88 and complex informational objects like mathematical proofs. The theory of
89 scales captures commonalities shared by distance, time, numbers, money,
90 and degrees of risk, severity, and happiness. The most common words in
91 English (and other languages) can be defined or characterized in terms of
92 these abstract core theories. Specific kinds of composite entities and scales,

93 for example, are then defined as instances of these abstract concepts, and
94 we thereby gain access to the rich vocabulary the abstract theories provide.

95 We can illustrate the link between word meaning and core theories with
96 the rather complex verb “range”. A core theory of scales provides axioms
97 involving predicates such as *scale*, *<*, *subscale*, *top*, *bottom*, and *at*. Then
98 we are able to define “range” by the following axiom:

$$\begin{aligned} 99 \quad & (\forall x, y, z) \text{range}(x, y, z) \equiv \\ 100 \quad & (\exists s, s_1, u_1, u_2) \text{scale}(s) \wedge \text{subscale}(s_1, s) \wedge \text{bottom}(y, s_1) \\ 101 \quad & \wedge \text{top}(z, s_1) \wedge u_1 \in x \wedge \text{at}(u_1, y) \wedge u_2 \in x \wedge \text{at}(u_2, z) \\ 102 \quad & \wedge (\forall u \in x)(\exists v \in s_1) \text{at}(u, v) \end{aligned}$$

103 That is, x ranges from y to z if and only if there is a scale s with a subscale
104 s_1 whose bottom is y and whose top is z , such that some member u_1 of x is
105 at y , some member u_2 of x is at z , and every member u of x is at some point
106 v in s_1 . Then by choosing different scales and instantiating the *at* relation
107 in different ways, we can get such uses as

108 The buffalo ranged from northern Texas to southern Saskatchewan.

109 The students’ SAT scores range from 1100 to 1550.

110 The hepatitis cases range from moderate to severe.

111 His behavior ranges from sullen to vicious.

112 Many things can be conceptualized as scales, and when this is done, a large
113 vocabulary, including the word “range”, becomes available.

114 It may seem strange for one to embrace logic and the image-schema in-
115 sight in the same framework, because the two are often taken by cognitive
116 linguists to be contradictory. But the use of logic amounts to less than one
117 might at first think. It can be viewed simply as a well-understood way of
118 representing complex information. To use the notation of first-order logic
119 is to adopt a style of representation that provides for predicate-argument
120 relations (so we know the difference between “Dog bites man” and “Man
121 bites dog”), conjunction (so we have the additive effect of two propositions),
122 implication and modus ponens (so we can derive one proposition from oth-
123 ers), and universal instantiation (so we can derive specific instances from
124 general principles). Any adequate representation scheme for knowledge and
125 information must give us at least these features.

126 The use of logic is also often taken to mean that words have strict defi-
127 nitions, and we know strict definitions are usually not possible. This is why
128 I have used the phrase “define or characterize” rather than “define”. In
129 general, we cannot hope to find definitions for words. That is, for very few
130 words p will we find necessary and sufficient conditions, giving us axioms of
131 the form

$$132 \quad (\forall x)p(x) \equiv \dots$$

133 Rather, we will find many necessary conditions and many sufficient condi-
134 tions.

$$135 \quad (\forall x)p(x) \supset \dots$$

136 $(\forall x) \dots \supset p(x)$

137 However, the accumulation of enough such axioms will tightly constrain the
138 possible interpretations of the predicate, and hence the meaning of the word.

139 This, by the way, gives us a different perspective on the notion of se-
140 mantic primitives. Our theories should be as elegant as possible, and thus
141 they will have as few “central” predicates as possible. These will give the
142 semblance of a small set of semantic primitives, and in fact are similar to
143 those usually proposed. But in our approach we do not attempt to reduce
144 all concepts to undefinable primitive predicates. Rather, strictly speaking,
145 every predicate is primitive, but its set of possible interpretations is more
146 or less tightly constrained by the axioms it participates in (see article 17
147 *Lexical decomposition*; 19 *Lexical conceptual structure*; 21 *Sense relations*).

148 A further feature required of our logic breaks down the rigidity of formal
149 logic that cognitive linguists sometimes react against. There must be some
150 mechanism for defeasibility; we have to be able to state inferences that are
151 normally true but can be defeated in particular contexts. There are many
152 such logics (e.g., McCarthy 1980; Ginsberg 1987; Shoham 1987). In Hobbs
153 et al. (1993) and Hobbs (2004) it is argued that interpretation of discourse
154 is a matter of coming up with the best proof of the content of an utterance
155 and the fact of its occurrence, using a method of defeasible inference known
156 as abduction. This provides a means of evaluating possibly contradictory
157 “proofs” to determine the best proof, or interpretation. Thus there may be a

158 large number of possible inferences that one may draw in any given context,
 159 but only some of them will be a part of the best interpretation. The mystery
 160 of how words acquire their manifold shades of meaning in different contexts
 161 thereby translates into the problem of how we choose the best interpretation,
 162 or, in a sense, how we select the right set of inferences to draw from the use
 163 of a word in context. This is far from a solved problem, but recasting
 164 meaning and interpretation in this way gives us a formal, computational
 165 way of approaching the problem.

166 Defeasibility in the logic gives us an approach to prototypes (see article
 167 28 *Prototype theory*; Rosch 1975). Categories correspond to predicates and
 168 are characterized by a set of possibly defeasible inferences, expressed as
 169 axioms, among which are their traditional defining features. For example,
 170 bachelors are unmarried and birds fly.

$$171 \quad (\forall x) \text{bachelor}(x) \supset \text{unmarried}(x)$$

$$172 \quad (\forall x) \text{bird}(x) \wedge \text{etc}_1(x) \supset \text{fly}(x)$$

173 where $\text{etc}_1(x)$ indicates the defeasibility of the axiom. Each instance of a
 174 category has a subset of the defeasible inferences that hold in its particular
 175 case. The more prototypical, the more inferences. In the case of the penguin,
 176 which is not a prototypical bird, the defeasible inference about flying is
 177 defeated. In this view, the basic level category is the predicate with the
 178 richest set of associated axioms. For example, there is more gain in useful
 179 knowledge from learning an animal is a dog than from learning a dog is a

180 boxer.

181 Similarly, defeasible inference lends itself to a treatment of novel metaphor.
182 In metaphor, some properties are transferred from a source to a target, and
183 some are not. When we say Pat is a pig, we draw inferences about manner
184 and quantity of eating from “pig”, but not about four-leggedness or species
185 membership. The latter inferences are defeated by the other things we know.
186 Hobbs (1992) develops this idea.

187 Taking abstract core theories as basic may seem to run counter to a
188 central tenet of cognitive linguistics, namely, that our understanding of many
189 abstract domains is founded on spatial metaphor. It is certainly true that
190 the field of spatial relationships, along with social relationships, is one of
191 the domains babies have to figure out first. But I think that to say we
192 figure out space first and then transfer that knowledge to other domains
193 is to seriously underestimate the difficulty of figuring out space. There are
194 many ways one could conceptualize space, e.g., via Euclidean geometry. But
195 in fact it is the topological concepts which predominate in a baby’s spatial
196 understanding. A one-year-old baby fascinated by “in” might put a necklace
197 into a trash can and a Cheerio into a shoe, despite their very different sizes
198 and shapes. In spatial metaphor it is generally the topological properties
199 that get transferred from the source to the target. In taking the abstract
200 core theories as basic, we are isolating precisely the topological properties
201 of space that are most likely to be the basis for understanding metaphorical

202 domains.

203 If one were inclined to make innateness arguments, one position would be
204 that we are born with a instinctive ability to operate in spatial environments.
205 We begin to use this immediately when we are born, and when we encounter
206 abstract domains, we tap into its rich models. The alternative, more in line
207 with our development here, is that we are born with at least a predisposition
208 towards instinctive abstract patterns – composite entities, scales, change,
209 and so on – which we first apply in making sense of our spatial environment,
210 and then apply to other, more abstract domains as we encounter them. This
211 has the advantage over the first position that it is specific about exactly what
212 properties of space might be in our innate repertoire. For example, the scalar
213 notions of “closer” and “farther” are in it; exact measures of distance are
214 not. A nicely paradoxical coda for summing up this position is that we
215 understand space by means of a spatial metaphor. I take Talmy’s critique
216 of the “concreteness as basic” idea as making a similar point (see article 27
217 *Cognitive semantics: An overview*).

218 Many of the preceding articles have proposed frameworks for linking
219 words to an underlying conceptual structure. These can all be viewed as
220 initial forays into the problem of connecting lexical meaning with world
221 knowledge. The content of this work survives translation among the vari-
222 ous frameworks that have been used for examining it, and survives recast-
223 ing it as a problem of explicitly encoding world knowledge, specifically, a

224 theory of image schemas explicating such concepts as composite entities,
225 figure-ground, scales, change of state, causality, aggregation, and granular-
226 ity shifts—an abstract theory that can be instantiated in many different,
227 more specialized domains. The core theories we are developing are not so
228 much theories about *particular* aspects of the world, but rather abstract
229 frameworks that are useful in making sense of a number of different kinds of
230 phenomena. Levin and Rappaport Hovav (see article 19 *Lexical conceptual*
231 *structure*) say, “All theories of event structure, either implicitly or explic-
232 itly, recognize a distinction between the primitive predicates which define
233 the range of event types available and a component which represents what
234 is idiosyncratic in a verb’s meaning.” The abstract theories presented here
235 are an explication of the former of these.

236 This work can be seen as an attempt at a kind of deep lexical semantics.
237 Not only are the words “decomposed” into what were once called primitives,
238 but also the primitives are explicated in axiomatic theories, enabling one to
239 reason deeply about the concepts conveyed by the text.

240 **2. Core abstract theories**

241 **2.1 Composite entities**

242 Composite entities are things made of other things. A composite entity is
243 characterized by a set of components, a set of properties of these components,

244 and a set of relations among the components and between the components
245 and the whole. The concept of composite entity captures the minimal com-
246 plexity something must have in order for it to have structure. It is hard
247 to imagine something that cannot be conceptualized as a composite entity.
248 For this reason, a vocabulary for talking about composite entities will be
249 broadly applicable.

250 The elements of a composite entity can themselves be viewed as com-
251 posite entities, and this gives us a very common example of shifting granu-
252 larities. It allows us to distinguish between the *structure* and the *function*
253 of an entity. The function of an entity as a component of a larger composite
254 entity is its relations to the other elements of the larger composite entity,
255 its environment, while the entity itself is viewed as indecomposable. The
256 structure of the entity is revealed when we decompose it and view it as a
257 composite entity itself. We look at it at a finer granularity.

258 An important question any time we can view an entity both functionally
259 and structurally is how the functions of the entity are implemented in its
260 structure. We need to spell out the structure-function articulations.

261 For example, a librarian might view a book as an indecomposable entity
262 and be interested in its location in the library, its relationship to other books,
263 to the bookshelves, and to the people who check the book out. This is a
264 functional view of the book with respect to the library. We can also view it
265 structurally by inquiring as to its parts, its content, its binding, and so on.

266 In spelling out the structure-function articulations, we might say something
267 about how its content, its size, and the material used in its cover determines
268 its proper location in the library.

269 A composite entity can serve as the *ground* against which some external
270 *figure* can be located or can move (see 27 *Cognitive semantics: An overview*).

271 A primitive predicate *at* expresses this relation. In

272 $at(x, y, s)$

273 *s* is a composite entity, *y* is one of its elements, and *x* is an external entity.

274 The relation says that the figure *x* is at a point *y* in the composite entity *s*,
275 which is the ground.

276 The *at* relation plays primarily two roles in the knowledge base. First,
277 it is involved in the “decompositions” of many lexical items. We saw this
278 above in the definition of “range”. There is a very rich vocabulary of terms
279 for talking about the figure-ground relation. This means that whenever a
280 relation in some domain can be viewed as an instance of the figure-ground
281 relation, we acquire at a stroke a rich vocabulary for talking about that
282 domain.

283 This gives rise to the second role the *at* predicate plays in the knowledge
284 base. A great many specific domains have relations that are stipulated to
285 be instances of the *at* relation. There are a large number of axioms of the
286 form

287 $(\forall x, y, s)r(x, y, s) \supset at(x, y, s)$

288 It is in this way that many of the metaphorical usages that pervade
289 natural language discourse are accommodated. Once we characterize some
290 piece of the world as a composite entity, and some relation as an *at* relation,
291 we have acquired the whole locational way of talking about it. Once this is
292 enriched with a theory of time and change, we can import the whole vocab-
293 ulary of motion. For example, in computer science, a data structure can be
294 viewed as a composite entity, and we can stipulate that if a pointer points to
295 a node in a data structure, then the pointer is *at* that node. We have then
296 acquired a spatial metaphor, and we can subsequently talk about, for exam-
297 ple, the pointer *moving around* the data structure. Space, of course, is itself
298 a composite entity and can be talked about using a locational vocabulary.

299 Other examples of *at* relations are

300 A person at an object in a system of objects:

301 John is at his desk.

302 An object at a location in a coordinate system:

303 The post office is at the corner of 34th Street and Eighth
304 Avenue.

305 A person's salary at a particular point on the money scale:

306 John's salary reached \$75,000 this year.

307 An event at a point on the time line:

308 The meeting is at three o'clock.

309 **2.2 Scales**

310 The theory of scales was mentioned in the introduction. It provides the
311 basic vocabulary for talking about partial orderings, including *scale*, $<$,
312 *subscale*, total ordering, *top*, *bottom*, *reverse*, and intervals. The theory
313 also explicates monotone-increasing scale-to-scale functions (“the more X ,
314 the more Y ”), the construction of composite scales, and the characterization
315 of qualitatively high and low regions of a scale.

316 A scale is a composite entity, so we can talk about an entity being *at* a
317 point on the scale. An obvious example of a scale is the scale of nonnegative
318 integers. The cardinality of a set can be defined in the standard way:

319 $card(\phi) = 0$

320 $(\forall x, s) x \notin s \supset card(\{x\} \cup s) = card(s) + 1$

321 We can then define cardinality to be an *at* relation, where N is the scale of
322 nonnegative integers:

323 $(\forall s, n) card(s) = n \supset at(s, n, N)$

324 This gives us access to the rich vocabulary of spatial relationships when
325 talking about cardinality, allowing us to say things like

326 The population of Cairo *reached* 15 million this year.

327 Many scales are composite. A scale s is a composite of scales s_1 and s_2
328 if its elements are the ordered pairs $\langle x, y \rangle$ where x is in s_1 and y is in s_2 .

329 The ordering in s has to be consistent with the orderings in s_1 and s_2 ; if x_1
330 is less than x_2 in s_1 , and y_1 is less than y_2 in s_2 , then $\langle x_1, y_1 \rangle$ is less than
331 $\langle x_2, y_2 \rangle$ in s . The converse is not necessarily true; the composite scale may
332 have more structure than that inherited from its component scales. We need
333 composite scales to deal with complex scalar predicates, such as *damage*.
334 When something is damaged, it no longer fulfills its function in a goal-
335 directed system. It needs to be repaired, and repairs cost. Thus, there are
336 (at least) two ways in which damage can be serious, first in the degradation
337 of its function, second in the cost of its repair. These are independent
338 scales. Damage that causes a car not to run may cost next to nothing to fix,
339 and damage that only causes the car to run a little unevenly may be very
340 expensive.

341 It is very useful to be able to isolate the high and low regions of a scale.
342 We can do this with operators called *Hi* and *Lo*. The *Hi* region of a scale
343 includes its top; the *Lo* region includes its bottom. The points in the *Hi*
344 region are all greater than any of the points in the *Lo* region. Otherwise,
345 there are no general topological constraints on the *Hi* and *Lo* regions. In
346 particular, the bottom of the *Hi* region and the top of the *Lo* region may
347 be indeterminate with respect to the elements of the scale. The *Hi* and *Lo*
348 operators provide us with a coarse-grained structure on scales, useful when
349 greater precision is not necessary or not possible.

350 The absolute form of adjectives frequently isolate *Hi* and *Lo* regions of

351 scales. A totally ordered Height Scale can be defined precisely, but frequently
352 we are only interested in qualitative judgments of height. The word “tall”
353 isolates the *Hi* region of the Height Scale; the word “short” isolates the *Lo*
354 region. A Happiness Scale cannot be defined precisely. We cannot get much
355 more structure for a Happiness Scale than what is given to us by the *Hi*
356 and *Lo* operators. The *Hi* and *Lo* operators can be iterated, to give us the
357 concepts “happy”, “very happy”, and so on.

358 In any given context, the *Hi* and *Lo* operators will identify different
359 regions of the scale. That is, the inferences we can draw from the fact
360 that something is in the *Hi* region of a scale are context-dependent; indeed,
361 inferences are always context-dependent. But two important constraints
362 on the *Hi* and *Lo* regions relate them to distributions and functionality.
363 The *Hi* and *Lo* regions must be related to common distributions of objects
364 on the scale in an as-yet nonexistent qualitative theory of distributions. If
365 something is significantly above average for the relevant set, then it is in the
366 *Hi* region. The regions must also be related to goal-directed behavior; often
367 something is in the *Hi* region of a scale precisely because that property aids
368 or defeats the achievement of some goal in a plan. For example, saying that
369 a talk is long often means that it is longer than the audience’s attention
370 span, and thus the goal of conveying information is defeated. Often when
371 we call someone tall, we mean tall enough or too tall for some purpose.

372 **2.3 Change of state**

373 A predicate of central importance is the predicate *change*. This is a relation
374 between situations, or conditions, or predications, and indicates a change
375 of state. A change from p being true of x to q being true of x , using an
376 ontologically promiscuous notation that reifies states and events (see Hobbs
377 1985; article 35 *Event semantics*), can be represented

378
$$\text{change}(e_1, e_2) \wedge p'(e_1, x) \wedge q'(e_2, x)$$

379 This says that there is a change from the situation e_1 of p being true of x
380 to the situation e_2 of q being true of x . A very common pattern involves a
381 change of location:

382
$$\text{change}(e_1, e_2) \wedge \text{at}'(e_1, x, y, s) \wedge \text{at}'(e_2, x, z, s)$$

383 That is, there is a change from the situation e_1 of x being at y in s to the
384 situation e_2 of x being at z in s .

385 When there is a change, generally there is some entity involved in both
386 the start and end states; there is something that is changing— x in the above
387 formulas.

388 The predicate *change* possesses a limited transitivity. There was a
389 change from Bill Clinton being a law student to Bill Clinton being Pres-
390 ident, because they are two parts of the same ongoing process, even though
391 he was governor in between. There was a change from Bill Clinton being
392 President to George W. Bush being president. But we probably do not want

393 to say there was a change from Bill Clinton being a law student to George
394 W. Bush being President. They are not part of the same process.

395 A state cannot change into the same state without going through an
396 intermediate different state.

397 The concept of *change* is linked with time in the obvious way. If state
398 e_1 changes into state e_2 , then e_2 cannot be before e_1 . My view is that the
399 relation between change and time is much deeper, cognitively. The theory of
400 change of state suggests a view of the world as consisting of a large number of
401 more or less independent, occasionally interacting processes, or histories, or
402 sequences of events. x goes through a series of changes, and y goes through
403 a series of changes, and occasionally there is a state that involves a relation
404 between the two. We can then view the time line as an artificial construct,
405 a regular sequence of imagined abstract events—think of them as ticks of a
406 clock in the National Institute of Science and Technology—to which other
407 events can be related by chains of copresence. Thus, I know I went home at
408 six o'clock because I looked at my watch, and I had previously set my watch
409 by going to the NIST Web site. In any case, there is no need to choose
410 between such a view of time and one that takes time as basic. They are
411 inter-definable in a straightforward fashion (Hobbs et al. 1987).

412 For convenience, we define one-argument predicates *changeFrom* and
413 *changeTo*, suppressing one or the other argument of *change*.

414 **2.4 Cause**

415 Our treatment of causality (Hobbs 2005) rests on a distinction between
416 causal complexes and the predicate *cause*. When we flip a switch and the
417 light comes on, we say that flipping the switch caused the light to come on.
418 But many other factors were involved. The wiring and the light bulb had to
419 be intact, the power had to be on in the city, and so forth. We say that all
420 these other states and events constitute the causal complex for the effect.
421 A causal complex for an effect is the set of all the eventualities that must
422 happen or hold in order for the effect to occur. The two principal properties
423 of causal complexes are that when all the eventualities happen, the effect
424 happens, and that every eventuality in the causal complex is required for the
425 effect to happen. These are strictly true, and the notion of causal complex
426 is not a defeasible one.

427 The “cause” of an effect, by contrast, is a distinguished element within
428 the causal complex, one that cannot normally be assumed to hold. It is
429 often the action that is under the agent’s immediate control. It is only
430 defeasibly true that when a cause occurs the effect also occurs. This inference
431 can be defeated because some of the other states and events in the causal
432 complex that normally hold do not hold in this particular case. The notion
433 of *cause* is much more useful in commonsense reasoning because we can
434 rarely if ever enumerate all the eventualities in a causal complex. Most of
435 our commonsense causal knowledge is expressed in terms of the predicate

436 *cause*.

437 The concept *cause* has the expected properties, such as defeasible transi-
438 tivity and consistency with temporal ordering. But we should not expect to
439 have a highly developed theory of causality *per se*. Rather we should expect
440 to see causal information distributed throughout our knowledge base. For
441 example, there is no axiom of the form

$$442 \quad (\forall e_1, e_2) \text{cause}(e_1, e_2) \equiv \dots$$

443 defining *cause*. But there will be many axioms of the forms

$$444 \quad p'(e_1, x) \supset q'(e_2, x) \wedge \text{cause}(e_1, e_2)$$

$$445 \quad r'(e_3, x) \supset p'(e_1, x) \wedge \text{cause}(e_1, e_3)$$

446 expressing causal connections among specific states and events; e.g., *p*-like
447 events cause *q*-like events or *r*-like events are caused by *p*-like events. We
448 don't know precisely what causality is, but we know lots and lots of examples
449 of things that cause other things.

450 Some would urge that causes and effects can only be events, but it seems
451 to me that we want to allow states as well, since in

452 The slipperiness of the ice caused John to fall.

453 the cause is a state. Moreover, intentional agents are sometimes taken to be
454 the unanalyzed causes of events. In

455 John lifted his arm.

456 John is the cause of the change of position of his arm, and we probably don't
457 want to have to coerce this argument into some imagined event taking place
458 inside John. Physical forces may also act as causes, as in

459 Gravity causes the moon to circle the earth.

460 The world is laced with threads of causal connection. In general, two
461 entities x and y are causally connected with respect to some behavior p
462 of x , if whenever p happens to x , there is some corresponding behavior q
463 that happens to y . Attachment of physical objects is one variety of causal
464 connection. In this case, p and q are both *move*. If x and y are attached,
465 moving x causes y to move. Containment is similar.

466 A particularly common variety of causal connection between two entities
467 is one mediated by the motion of a third entity from one to the other. This
468 might be called, somewhat facetiously, a “vector boson” connection. In par-
469 ticle physics, a vector boson is an elementary particle that transfers energy
470 from one point to another. Photons, which really are vector bosons, medi-
471 ate the causal connection between the sun and our eyes. Other examples of
472 such causal connections are rain drops connecting a state of the clouds with
473 the wetness of our skin and clothes, a virus transmitting disease from one
474 person to another, and utterances passing information between people.

475 Containment, barriers, openings, and penetration are all with respect to
476 paths of causal connection. Force is causality with a scalar structure (see
477 article 27 *Cognitive semantics: An overview*).

478 The event structure underlying many verbs exhibits causal chains. In-
479 struments, for example, are usually vector bosons. In the sentence,

480 John pounded the nail with a hammer for Bill.

481 the underlying causal structure is that the agent John causes a change in
482 location of the instrument, the hammer, which causes a change in location
483 of the object or theme, the nail, which causes or should cause a change in
484 the mental or emotional state of the beneficiary, Bill.

485 $Agent \text{--}cause\text{--} \rightarrow change(at(Instr, x, s), at(Instr, Object, s))$
486 $\text{--}cause\text{--} \rightarrow change(at(Object, y_1, s), at(Object, y_2, s))$
487 $\text{--}cause\text{--} \rightarrow change(p_1(Beneficiary), p_2(Beneficiary))$

488 Much of case grammar and work on thematic roles can be seen as a matter
489 of identifying where the arguments of verbs fit into this kind of causal chain
490 when we view the verbs as instantiating this abstract frame (see Jackendoff
491 1972; article 18 *Thematic roles*; article 19 *Lexical conceptual structure*).

492 In addition, in this theory we define such concepts as *enable*, *prevent*,
493 *help*, and *obstruct*. There are also treatments of attempts, success, failure,
494 ability, and difficulty.

495 With this vocabulary, we are in a position to characterize more precisely
496 the intuitive notions of state, event, action, and process. A state is a static
497 property that does not involve a change (at the relevant granularity), such
498 as an *at* relationship, $at(x, y, s)$. To be up, for example, is a state. An event
499 is a change of state, a common variety of which is a change of location:

500 $change(e_1, e_2) \wedge at'(e_1, x, y, s) \wedge at'(e_2, x, z, s)$

501 For example, the verb “rise” denotes a change of location of something to a
502 higher point. An action is the causing of an event by an intentional agent:

503 $cause(a, e) \wedge change'(e, e_1, e_2) \wedge at'(e_1, x, y, s) \wedge at'(e_2, x, z, s)$

504 The verb “raise” denotes an action by someone of effecting a change of
505 location of something to a higher point. A process is a sequence of events
506 or actions. For example, to fluctuate is to undergo a sequence of risings and
507 fallings, and to pump is to engage in a sequence of raisings and lowerings.
508 We can coarsen the granularity on processes so that the individual changes
509 of state become invisible, and the result is a state. This is a transformation
510 of perspective that is effected by the progressive aspect in English. Thus,
511 fluctuating can be viewed as a state.

512 Detailed expositions of all the core theories can be found at

513 <http://www.isi.edu/hobbs/csk.html>

514 **3. Linking word meaning with the theories**

515 Once we have in place the core theories that capture world knowledge at a
516 sufficiently abstract level, we can begin to construct the axioms that link
517 word meaning to the theories. We illustrate here how that would go, using
518 the words “have”, “remove”, and “remain”. Words have senses, and for
519 each sense the linkage will be different. Here we examine the word senses

520 in WordNet (Miller 1995) and FrameNet (Baker, Fillmore & Cronin 2003),
521 since they are the most heavily used lexical resources in computational lin-
522 guistics. The word sense numbers correspond to their order in the Web
523 interfaces to the two resources:

524 <http://wordnet.princeton.edu/>

525 <http://framenet.icsi.berkeley.edu>

526 **3.1 “Have”**

527 In WordNet the verb “have” has 19 senses. But they can be grouped into
528 three broad “supersenses”. In its first supersense, X has Y means that X is
529 in some relation to Y. The WordNet senses this covers are as follows:

- 530 1. a broad sense, including have a son, having a condition hold
531 and having a college degree
- 532 2. having a feature or property, i.e., the property holding of the
533 entity
- 534 3. a sentient being having a feeling or internal property
- 535 4. a person owning a possession
- 536 7. have a person related in some way: have an assistant
- 537 9. have left: have three more chapters to write
- 538 12. have a disease: have influenza
- 539 17. have a score in a game: have three touchdowns

540 The supersense can be characterized by the axiom

541 $have-s1(x, y) \supset relatedTo(x, y)$

542 In these axioms, supersenses are indexed with s , WordNet senses with w ,
543 and FrameNet senses with f . Unindexed predicates are from core theories.

544 The individual senses are then specializations of the supersense where
545 more domain-specific predicates are explicated in more specialized domains.

546 For example, sense 4 relates to the supersense as follows:

547 $have-w4(x, y) \equiv possess(x, y)$

548 $have-w4(x, y) \supset have-s1(x, y)$

549 where the predicate *possess* would be explicated in a commonsense theory
550 of economics, relating it to the privileged use of the object. Similarly, *have-*
551 *w3*(x, y) links with the supersense but has the restrictions that x is sentient
552 and that the “relatedTo” property is the predicate-argument relation be-
553 tween the feeling and its subject.

554 The second supersense of “have” is “come to be in a relation to”. This
555 is our *changeTo* predicate. Thus, the definition of this supersense is

556 $have-s2(x, y) \equiv changeTo(e) \wedge have-s1'(e, x, y)$

557 The WordNet senses this covers are as follows:

558 10. be confronted with: we have a fine mess

559 11. experience: the stocks had a fast run-up

560 14. receive something offered: have this present

561 15. come into possession of: he had a gift from her

562 16. undergo, e.g., an injury: he had his arm broken in the fight

563 18. have a baby

564 In these senses the new relation is initiated but the subject does not nec-
565 essarily play a causal or agentive role. The particular change involved is
566 specialized in the WordNet senses to a confronting, a receiving, a giving
567 birth, and so on.

568 The third supersense of “have” is “cause to come to be in a relation to”.

569 The axiom defining this is

$$570 \quad \textit{have-s3}(x, y) \equiv \textit{cause}(x, e) \wedge \textit{have-s2'}(e, x, y)$$

571 The WordNet senses this covers are

572 5. cause to move or be in a certain position or condition: have

573 your car ready

574 6. consume: have a cup of coffee

575 8. organize: have a party

576 13. cause to do: she had him see a doctor

577 19. have sex with

578 In all these cases the subject initiates the change of state that occurs.

579 FrameNet has five simple transitive senses for “have”. Their associated

580 frames are

581 1. Have associated

582 2. Possession

583 3. Ingestion

584 4. Inclusion

585 5. Birth

586 The first sense corresponds to the first WordNet supersense:

$$587 \quad \textit{have-f1}(x, y) \equiv \textit{have-s1}(x, y)$$

588 The second sense is WordNet sense 4.

$$589 \quad \textit{have-f2}(x, y) \equiv \textit{have-w4}(x, y)$$

590 The third sense is WordNet sense 6. The fourth sense is a *partOf* relation.

591 It is a specialization of WordNet sense 2.

$$592 \quad \textit{have-f4}(x, y) \equiv \textit{partOf}(x, y)$$

$$593 \quad \textit{have-f4}(x, y) \supset \textit{have-w2}(x, y)$$

594 The fifth sense is WordNet sense 18.

595 3.2 “Remove”

596 If x removes y from z , then x causes a change from the state in which y is

597 at z .

$$598 \quad \textit{remove}(x, y, z) \supset \textit{cause}(x, e_1) \wedge \textit{changeFrom}'(e_1, e_2) \wedge \textit{at}'(e_2, y, z, s)$$

599 This is the “supersense” covering all of the WordNet and FrameNet senses

600 of “remove”.

601 WordNet lists 8 senses of “remove”. In WordNet sense 1, *at* is instan-
602 tiated as physical location. In sense 2, *at* is instantiated as position in an
603 organization, as in “The board removed the VP of operations.” In sense
604 3, *y* is somehow dysfunctional, as in removing trash. In sense 4, *at* is in-
605 stantiated as the membership relation in a set; *y* is removed from set *z*. In
606 sense 5, the change is functional or strategic, as in a general removing his
607 troops from a vulnerable position. In sense 6, *x* and *y* are identical, as in
608 “He removed himself from the contest.” In sense 7, *at* is instantiated as
609 “alive”, as in “The Mafia don removed his enemy.” In sense 8, *y* is abstract
610 and dysfunctional, as in removing an obstacle.

611 FrameNet has two senses of the word. The first is the general meaning,
612 our supersense. In the second sense, *x* is a person, *y* is clothes, and *z* is a
613 body.

614 Note that the supersense gives the topological structure of the meaning
615 of the verb. The various senses are then generated from that by instantiating
616 the *at* relation to something more specific, or by adding domain constraints
617 to the arguments *x*, *y* and *z*.

618 **3.3 “Remain”**

619 There are four WordNet senses of the verb “remain”:

- 620 1. Not change out of an existing state: He remained calm.
- 621 2. Not change out of being at a location: He remained at his

622 post.

623 3. Entities in a set remaining after others are removed: Three
624 problems remain.

625 4. A condition remains in a location: Some smoke remained
626 after the fire was put out.

627 The first sense is the most general and subsumes the other three. We can
628 characterize it by the axiom

$$629 \quad \textit{remain-w1}(x, e) \supset \textit{arg}(x, e) \wedge \neg \textit{changeFrom}(e)$$

630 That is, if x remains in condition e , then e is a property of x (or x is
631 an argument of e), and there is no change from state e holding. By the
632 properties of *changeFrom* it follows that x is in state e , as is presupposed.

633 In the second sense, the property e of x is being in a location.

$$634 \quad \textit{remain-w2}(x, e) \equiv \textit{remain-w1}(x, e) \wedge \textit{at}'(e, x, y)$$

635 The fourth sense is a specialization of the second sense in which the entity
636 x that remains is a state or condition.

$$637 \quad \textit{remain-w4}(x, e) \equiv \textit{remain-w2}(x, e) \wedge \textit{state}(x)$$

638 The third sense is the most interesting to characterize. As in the fourth
639 WordNet sense of “remove”, there is a process that removes elements from
640 a set, and what remains is the set difference between the original and the
641 set of elements that are removed. In this axiom x remains after process e .

642 $remain-w3(x, e)$

643 $\equiv remove-w4'(e, y, s_2, s_1) \wedge setdiff(s_3, s_1, s_2) \wedge member(x, s_3)$

644 That is, x remains after e if and only if e is a removal event by some agent
645 y of a subset s_2 from s_1 , s_3 is the set difference between s_1 and s_2 , and x is
646 a member of s_3 .

647 There are four FrameNet senses of “remain”. The first is the same as
648 WordNet sense 1. The second is the same as WordNet sense 3. The third
649 and fourth are two specializations of WordNet sense 3, one in which the
650 removal process is destructive and one in which it is not.

651 There are two nominalizations of the verb “remain”—“remainder” and
652 “remains”. All of their senses are related to WordNet sense 3. The first
653 WordNet noun sense is the most general.

654 $remainder-w1(x, e) \equiv remain-w3(x, e)$

655 That is, x is the remainder after process e if and only if x remains after
656 e . The other three senses result from specialization of the removal process
657 to arithmetic division, arithmetic subtraction, and the purposeful cutting of
658 a piece of cloth. The noun “remains” refers to what remains ($w3$) after a
659 process of consumption or degradation.

660 **3.4 The nature of word senses**

661 The most common words in a language are typically the most polysemous.
662 They often have a central meaning indicating their general topological struc-

663 ture. Each new sense introduces inferences that cannot be reliably deter-
664 mined just from a core meaning plus contextual factors. They tend to build
665 up along what Brugman (1981), Lakoff (1987) and others have called a ra-
666 dial category structure (see 28 *Prototype theory*). Sense 2 may be a slight
667 modification of sense 1, and senses 3 and 4 different slight modifications
668 of sense 2. It is easy to describe the links that take us from one sense to
669 an adjacent one in the framework presented here. Each sense corresponds
670 to a predicate which is characterized by one or more axioms involving that
671 predicate. A move to an adjacent sense happens when incremental changes
672 are made to the axioms. As we have seen in the examples of this section,
673 the changes are generally additions to the antecedents or consequents of
674 the axioms. The principal kinds of additions are embedding in *change* and
675 *cause*, as we saw in the supersenses of “have”; the instantiation of general
676 predicates like *relatedTo* and *at* to more specific predicates in particular
677 domains, as we saw in all three cases; and the addition of domain-specific
678 constraints on arguments, as in restricting *y* to be clothes in *remove-f2*.

679 A good account of the lexical semantics of a word should not just cata-
680 log various word senses. It should detail the radial category structure of the
681 word senses, and for each link, it should say what incremental addition or
682 modification resulted in the new sense. Note that radial categories provide
683 us with a logical structure for the lexicon, and also no doubt a historical
684 one, but not a developmental one. Children often learn word senses inde-

685 pendently and only later if ever realize the relation among the senses. See
686 article 28 *Prototype theory* for further discussion of issues with respect to
687 radial categories.

688 4. Distinguishing lexical and world knowledge

689 It is perhaps natural to ask whether a principled boundary can be drawn
690 between linguistic knowledge and knowledge of the world. To make this
691 issue more concrete, consider the following seven statements:

- (1) If a string w_1 is a noun phrase and a string w_2 is a verb phrase,
then the concatenation w_1w_2 is a clause.
- (2) The transitive verb “moves” corresponds to the predication
 $move_2(x, y)$, providing a string describing x occurs as its subject
and a string describing y occurs as its direct object.
- (3) If an entity x moves (in sense $move_2$) an entity y , then x causes a
change of state or location of y .
- 692 (4) If an entity y changes to a new state or location, it is no longer in
its old state or location.
- (5) If a physical object x moves a physical object y through a fluid
medium, then x must apply force to y against the resistance of the
medium.
- (6) The function of a barge is to move freight across water.
- (7) A barge moved the wreckage of Flight 1549 to New Jersey.

693 Syntax consists in part of rules like (1), or generalizations of them. One
694 could view the lexicon as consisting of axioms expressing information like
695 (2), specifying for each word sense and argument realization pattern what
696 predication is conveyed, perhaps together with some generalizations of such
697 statements. (Lexical knowledge of other languages would be encoded as sim-
698 ilar axioms, sometimes linking to the same underlying predicates, sometimes
699 different.) Axioms expressing information like (3) link the lexical predicates
700 with underlying domain theories, in this case, theories of the abstract do-
701 mains of causality and change of state. Axioms expressing facts like (4) are
702 internal to domain theories, in this case, the theory of the abstract domain
703 of change of state. Axioms expressing general facts like (5) are part of a
704 commonsense or scientific theory of physics, which can be viewed as a spe-
705 cialization and elaboration of the abstract theories. Axioms expressing facts
706 like (6) encode telic information about artifacts. Statement (7) is a specific,
707 accidental fact about the world.

708 Many have felt that the viability of lexical semantics as a research enter-
709 prise requires a principled distinction between lexical knowledge and world
710 knowledge, presumably somewhere below axioms like (2) and above facts
711 like (7). Many of those who have believed that no such distinction is possi-
712 ble have concluded that lexical semantics is impossible, or at least can only
713 be very limited in its scope.

714 For example, in his discussion of meaning, Bloomfield (1933, 139–140)
715 rules out the possibility of giving definitions of most words.

716 In order to give a scientifically accurate definition of meaning
717 of every form of a language, we should have to have a scientif-
718 ically accurate knowledge of everything in the speakers' world.
719 While this may be possible for certain scientifically well-understood
720 terms like “salt”, we have no precise way of defining words like
721 “love” or “hate” which concern situations that have not been
722 accurately classified – and these latter are in the great majority.

723 He concludes that

724 The statement of meanings is therefore a weak point in language-
725 study, and will remain so until human knowledge advances very
726 far beyond its present state.

727 Lexical semantics is impossible because we would need a theory of the
728 world. Bloomfield goes on to talk about such phenomena as synonymy
729 and antonymy, and leaves issues of meaning at that.

730 More recently, Fodor (1980) similarly argued that lexical semantics would
731 need a complete and correct scientific theory of the world to proceed, and
732 is consequently impossible in the foreseeable future.

733 A counterargument is that we don't need a scientifically correct theory
734 of the world, because people don't have that as they use language to convey

735 meaning. We rather need to capture people's commonsense theories of the
736 world. In fact, there are a number of interesting engineering efforts to en-
737 code commonsense and scientific knowledge needed in specific applications
738 or more broadly. Large ontologies of various domains, such as biomedicine
739 and geography, are being developed for the Semantic Web and other com-
740 putational uses. Cyc (Lenat & Guha 1990) has been a large-scale effort to
741 encode commonsense knowledge manually since the middle 1980s; it now
742 contains millions of rules. The Open Mind Common Sense project (Singh
743 2002) aims at accumulating huge amounts of knowledge rapidly by marshal-
744 ing millions of "netizens" to make contributions; for example, a participant
745 might be asked to complete the sentence "Water can ..." and reply with
746 "Water can put out fires." Many of these projects, including Cyc, involve
747 a parallel effort in natural language processing to relate their knowledge of
748 the world to the way we talk about the world. Might we do lexical semantics
749 by explicating the meanings of words in terms of such theories?

750 Fodor (1983) can be read as responding to this possibility. He argues
751 that peripheral processes like speech recognition and syntactic processing
752 are encapsulated in the sense that they require only limited types of infor-
753 mation. Central processes like fixation of belief, by contrast, can require
754 any knowledge from any domain. He gives the example of the power of ana-
755 logical reasoning in fixation of belief. The body of knowledge that can be
756 appealed to in analogies can not be circumscribed; analogies might involve

757 mappings from anything to anything else. Scientific study of modular pro-
758 cesses is feasible, but scientific study of global processes is not. No scientific
759 account of commonsense reasoning is currently available or likely to be in
760 the foreseeable future; by implication reasoning about commonsense world
761 knowledge is not currently amenable to scientific inquiry, nor is a lexical
762 semantics that depends on it. Syntax *is* amenable to scientific study, but
763 only, according to Fodor, because it is informationally encapsulated.

764 Thus, the debate on this issue often centers on the modularity of syntax.
765 Do people do syntactic analysis of utterances in isolation from world knowl-
766 edge? Certainly at time scales at which awareness functions, there is no
767 distinction in the processing of linguistic and world knowledge. We rarely if
768 ever catch ourselves understanding the syntax of a sentence we hear with-
769 out understanding much about its semantics. For example, in Chomsky’s
770 famous grammatical sentence, “Colorless green ideas sleep furiously,” there
771 is no stage in comprehension at which we are aware that “colorless” and
772 “green” are adjectives, but haven’t yet realized they are contradictory.

773 Moreover, psychological studies seem to indicate that syntactic process-
774 ing and the use of world knowledge are intricately intertwined. Much of
775 this work has focused on the use of world knowledge to resolve references
776 and disambiguate ambiguous prepositional phrase attachments. Tanenhaus
777 & Brown-Schmidt (2008) review some of this research that makes use of
778 methods of monitoring eye movements to track comprehension. For exam-

779 ple, they present evidence that subjects access the current physical context
780 while they are processing syntactically ambiguous instructions and integrate
781 it with the language immediately. In terms of our examples, they are us-
782 ing facts like (1) and facts like (7) together. The authors contend that
783 their results “are incompatible with the claim that the language processing
784 includes subsystems (modules) that are informationally encapsulated, and
785 thus isolated from high-level expectations.”

786 Often the line between linguistic and world knowledge is drawn to in-
787 clude selectional constraints within language. Hagoort et al. (2004) used
788 electroencephalogram and functional magnetic resonance imaging data to
789 investigate whether there was any difference between the temporal course of
790 processing true sentences like “Dutch trains are yellow and very crowded”,
791 factually false but sensible sentences like “Dutch trains are white and very
792 crowded”, and sentences that violate selectional constraints like “Dutch
793 trains are sour and very crowded.” The false sentences and the selection-
794 ally anomalous sentences showed a virtually identical peak of activity in the
795 left inferior prefrontal cortex. The authors observed that there is “strong
796 empirical evidence that lexical semantic knowledge and general world knowl-
797 edge are both integrated in the same time frame during sentence interpre-
798 tation, starting at ~300ms after word onset.” However, there is a difference
799 in frequency profile between the two conditions, consisting of a measur-
800 able increase in activity in the 30-70 Hz range (gamma frequency) for the

801 false sentences, and an increase in the 4-7Hz range (theta frequency) in the
802 anomalous condition. The authors conclude that “semantic interpretation
803 is not separate from its integration with nonlinguistic elements of meaning,”
804 but that nevertheless “the brain keeps a record of what makes a sentence
805 hard to interpret, whether this is word meaning or world knowledge.”

806 Thus, if the brain makes a distinction between linguistic and world
807 knowledge, it does not appear to be reflected in the temporal course of
808 processing language.

809 The most common argument in linguistics and related fields for drawing
810 a strict boundary between lexicon and world is a kind of despair that a
811 scientific study of world knowledge is possible. Others have felt it is possible
812 to identify lexically relevant domains of world knowledge that are accessible
813 to scientific study.

814 Linguists investigating “lexical conceptual structure” (e.g., see article 19
815 *Lexical conceptual structure*) are attempting to discover generalizations in
816 how the way an entity occurs in the underlying description of a situation
817 or event in terms of abstract topological predicates influences the way it is
818 realized in the argument structure in syntax. For example, do verbs that
819 undergo dative alternation all have a similar underlying abstract structure?
820 Does the causative always involve embedding an event as the effect in a
821 causal relation, where the cause is the agent or an action performed by
822 the agent? The hypothesis of this work is that facts like (2), which are

823 linguistic, depend crucially on facts like (3), which have a more world-like
824 flavor. However, this does not mean that we have identified a principled
825 boundary between linguistic and world knowledge. One could just as well
826 view this as a strategic decision about how to carve out a tractable research
827 problem.

828 Pustejovsky (1995) pushes the line between language and world far-
829 ther into the world. He advocates representing what he calls the “qualia
830 structure” of words, which includes facts about the constituent parts of an
831 entity (Constitutive), its place in a larger domain (Formal), its purpose and
832 function (Telic), and the factors involved in its origin (Agentive). One can
833 then, for example, use the Telic information to resolve a metonymy like
834 “She began a cigarette” into its normal reading of “She began smoking a
835 cigarette,” rather than any one of the many other things one could do with
836 a cigarette—eating it, rolling it, tearing it apart, and so on. His framework
837 is an attempt to relate facts like (2) about what arguments can appear with
838 what predicates with facts like (6) about the functions and other properties
839 of things. Several places in his book, Pustejovsky suggests that it is im-
840 portant to see his qualia structures as part of lexical semantics, and hence
841 linguistics, as opposed to general commonsense knowledge that is not lin-
842 guistic. But he never makes a compelling argument to this effect. All of his
843 qualia structures and coercion mechanisms are straightforward to express
844 in a logical framework, so there are no formal reasons for the distinction.

845 I think it is best to see this particular carving out of knowledge and inter-
846 pretation processes, as with the study of lexical conceptual structures, as a
847 strategic decision to identify a fruitful and tractable research problem.

848 Pustejovsky's work is an attempt to specify the knowledge that is re-
849 quired for interpreting at least the majority of nonstandard uses of words.
850 Kilgarriff (2001) tests this hypothesis by examining the uses of nine partic-
851 ular words in a 20-million word corpus. 41 of 2276 instances were judged to
852 be nonstandard since they did not correspond to any of the entries for the
853 word in a standard dictionary. Of these, only two nonstandard uses were
854 derivable from Pustejovsky's qualia structures. The others required deeper
855 commonsense knowledge or previous acquaintance with collocations. Kil-
856 garriff's conclusion is that "Any theory that relies on a distinction between
857 general and lexical knowledge will founder." (Kilgarriff 2001: 325)

858 Some researchers in natural language processing have argued that lexical
859 knowledge should be distinguished from other knowledge because it results in
860 more efficient computation or more efficient comprehension and production.
861 One example concerns hyperonymy relations, such as that *car(x)* implies
862 *vehicle(x)*. It is true that some kinds of inferences lend themselves more
863 to efficient computation than others, and inferences involving only monadic
864 predicates are one example. But where this is true, it is a result not of their
865 content but of structural properties of the inferences, and these cut across
866 the lexical-world distinction. Any efficiency realized in inferring *vehicle(x)*

867 can be realized in inferring *expensive(x)* as well.

868 All of statements (1)-(7) are facts about the world, because sentences
869 and their structure and words and their roles in sentences are things in the
870 world, as much as barges, planes, and New Jersey. There is certainly knowl-
871 edge we have that is knowledge about words, including how to pronounce
872 and spell words, predicate-argument realization patterns, alternation rules,
873 subcategorization patterns, grammatical gender, and so on. But words are
874 part of the world, and one might ask why this sort of knowledge should have
875 any special cognitive status. Is it any different in principle from the kind of
876 knowledge one has about friendship, cars, or the properties of materials? In
877 all these cases, we have entities, properties of entities, and relations among
878 them. Lexical knowledge is just ordinary knowledge where the entities in
879 question are words. There are no representational reasons for treating lin-
880 guistic knowledge as special, providing we are willing to treat the entities in
881 our subject matter as first-class individuals in our logic (cf. Hobbs 1985).
882 There are no procedural reasons for treating linguistic knowledge as special,
883 since parsing, argument realization, lexical decomposition, the coercion of
884 metonymies, and so on can all be implemented straightforwardly as infer-
885 ence. The argument that parsing and lexical decomposition, for example,
886 can be done efficiently on present-day computers, whereas commonsense
887 reasoning cannot, does not seem to apply to the human brain; psycholin-
888 guistic studies show that the influence of world knowledge kicks in as early

889 as syntactic and lexical knowledge, and yields the necessary results just as
890 quickly.

891 We are led to the conclusion that any drawing of lines is for the strategic
892 purpose of identifying a coherent, tractable and fruitful area of research.
893 Statements (1)-(6) are examples from six such areas. Once we have iden-
894 tified and explicated such areas, the next question is what connections or
895 articulations there are among them; Pustejovsky's research and work on
896 lexical conceptual structures are good examples of people addressing this
897 question.

898 However, all of this does not mean that linguistic insights can be ig-
899 nored. The world can be conceptualized in many ways. Some of them lend
900 themselves to a deep treatment of lexical semantics, and some of them im-
901 pede it. Put the other way around, looking closely at language leads us to a
902 particular conceptualization of the world that has proved broadly useful in
903 everyday life. It provides us with topological relations rather than with the
904 precision of Euclidean 3-space. It focuses on changes of state rather than on
905 correspondences with an *a priori* time line. A defeasible notion of causality
906 is central in it. It provides means for aggregation and shifting granularities.
907 It encompasses those properties of space that are typically transferred to
908 new target domains when what looks like a spatial metaphor is invoked.

909 More specific domains can then be seen as instantiations of these abstract
910 theories. Indeed, Euclidean 3-space itself is such a specialization. Language

911 provides us with a rich vocabulary for talking about the abstract domains.
912 The core meanings of many of the most common words in language can be
913 defined or characterized in these core theories. When the core theory is in-
914 stantiated in a specific domain, the vocabulary associated with the abstract
915 domain is also instantiated, giving us a rich vocabulary for talking about
916 and thinking about the specific domain. Conversely, when we encounter
917 general words in the contexts of specific domains, understanding how the
918 specific domains instantiate the abstract domains allows us to determine the
919 specific meanings of the general words in their current context.

920 We understand language so well because we know so much. Therefore,
921 we will not have a good account of how language works until we have a
922 good account of what we know about the world and how we use that knowl-
923 edge. In this article I have sketched a formalization of one very abstract
924 way of conceptualizing the world, one that arises from an investigation of
925 lexical semantics and is closely related to the lexical decompositions and
926 image schemas that have been argued for by other lexical semanticists. It
927 enables us to capture formally the core meanings of many of the most com-
928 mon words in English and other languages, and it links smoothly with more
929 precise theories of specific domains.

930

931 *I have profited from discussions of this work with Gully Burns, Peter Clark,*
932 *Tim Clausner, Christiane Fellbaum, and Rutu Mulkar-Mehta. This work*

933 was performed in part under the IARPA (DTO) AQUAINT program, con-
934 tract N61339-06-C-0160.

935 5. References

- 936 Baker, Colin F., Charles J. Fillmore & Beau Cronin 2003. The structure of
937 the Framenet database. *International Journal of Lexicography* 16(3),
938 281–296.
- 939 Bloomfield, Leonard 1933. *Language*. New York: Holt, Rinehart & Winston.
- 940 Brugman, Claudia 1981. *The Story of “Over”*. M.A. Thesis. University of
941 California, Berkeley.
- 942 Cycorp 2008. <http://www.cyc.com/>.
- 943 Davis, Ernest 1990. *Representations of Commonsense Knowledge*. San Ma-
944 teo, CA: Morgan Kaufmann.
- 945 Fodor, Jerry A. 1980. Methodological solipsism considered as a research
946 strategy in cognitive science. *Behavioral and Brain Sciences* 3(1), 63–
947 109.
- 948 Fodor, Jerry A. 1983. *The Modularity of Mind. An Essay on Faculty Psy-*
949 *chology*. Cambridge, MA: The MIT Press.
- 950 Ginsberg, Matthew L. (ed.) 1987. *Readings in Nonmonotonic Reasoning*.
951 Los Altos, CA: Morgan Kaufmann.
- 952 Gruber, Jeffrey C. 1965/1976. *Studies in Lexical Relations*. Ph.D. Disserta-
953 tion. MIT, Cambridge. Reprinted in J. S Gruber. *Lexical Structures*

- 954 *in Syntax and Semantics*. Amsterdam: North-Holland, 1976, 1–210.
- 955 Hagoort, Peter, Lea Hald, Marcel Bastiaansen & Karl Magnus Petersson
956 2004. Integration of word meaning and world knowledge in language
957 comprehension. *Science* 304(5669), 438–441.
- 958 Hobbs, Jerry R. 1985. Ontological promiscuity. In: *Proceedings of the*
959 *23rd Annual Meeting of the Association for Computational Linguistics*
960 (=ACL). ACL, University of Chicago, 61–69.
- 961 Hobbs, Jerry R. 1992. Metaphor and abduction. In: A. Ortony, J. Slack &
962 O. Stock (eds.). *Communication from an Artificial Intelligence Per-*
963 *spective. Theoretical and Applied Issues*. Berlin, Heidelberg: Springer,
964 35–58.
- 965 Hobbs, Jerry R. 2004. Abduction in natural language understanding. In:
966 L. Horn & G. Ward (eds.). *Handbook of Pragmatics*. Malden, MA:
967 Blackwell, 724–741.
- 968 Hobbs, Jerry R. 2005. Toward a useful notion of causality for lexical seman-
969 tics. *Journal of Semantics* 22(2), 181–209.
- 970 Hobbs, Jerry R., William Croft, Todd Davies, Douglas Edwards & Ken-
971 neth Laws 1987. Commonsense metaphysics and lexical semantics.
972 *Computational Linguistics* 13(3–4), 241–250.
- 973 Hobbs, Jerry R., Mark Stickel, Douglas Appelt & Paul Martin 1993. Inter-
974 pretation as abduction. *Artificial Intelligence* 63(1–2), 69–142.
- 975 Jackendoff, Ray S. 1972. *Semantic Interpretation in Generative Grammar*.

- 976 Cambridge, MA: The MIT Press.
- 977 Johnson, Mark 1987. *The body in the Mind. The Bodily Basis of Meaning,*
978 *Imagination, and Reason.* Chicago: University of Chicago Press.
- 979 Kilgariff, Adam 2001. Generative lexicon meets corpus data. The case
980 of nonstandard word uses. In: P. Bouillion & F. Busa (eds.). *The*
981 *Language of Word Meaning.* Cambridge: Cambridge University Press,
982 312–328.
- 983 Lakoff, George 1987. *Women, Fire, and Dangerous Things. What Cate-*
984 *gories Reveal About the Mind.* Chicago: University of Chicago Press.
- 985 Lenat, Douglas B. & Ramanathan V. Guha 1990. *Building Large Knowledge-*
986 *based Systems. Representation and Inference in the Cyc Project.* Read-
987 ing, MA: Addison-Wesley.
- 988 McCarthy, John 1980. Circumscription. A form of non-monotonic reason-
989 ing. *Artificial Intelligence* 13(1–2), 27–39.
- 990 Miller, George 1995. WordNet. A lexical database for English. *Communi-*
991 *cations of the ACM* 38(11), 39–41.
- 992 Pustejovsky, James 1995. *The Generative Lexicon.* Cambridge, MA: The
993 MIT Press.
- 994 Rosch, Eleanor 1975. Cognitive representations of semantic categories. *Jour-*
995 *nal of Experimental Psychology. General* 104(3), 192–233.
- 996 Shoham, Yoav 1987. Nonmonotonic logics. Meaning and utility. In: J.
997 MacDermott (ed.). *Proceedings of the International Joint Conference*

998 *on Artificial Intelligence (= IJCAI)* 10. Los Altos, CA: Morgan Kauf-
999 mann, 388–393.

1000 Singh, Push 2002. The public acquisition of commonsense knowledge. In:
1001 *Proceedings of AAAI Spring Symposium on Acquiring (and Using)*
1002 *Linguistic (and World) Knowledge for Information Access*. Palo Alto,
1003 CA: AAAI.

1004 Tanenhaus Michael K. & Sarah Brown-Schmidt 2008. Language processing
1005 in the natural world. In: B. C. M. Moore, L. K. Tyler & W. D.
1006 Marslen-Wilson (eds.). *The Perception of Speech. From Sound to*
1007 *Meaning. Philosophical Transactions of the Royal Society B. Biological*
1008 *Sciences* 363, 1105–1122.

1009 *Keywords:*

1010 Lexical semantics, world knowledge, lexical decomposition.

1011

1012

1013 *Jerry R. Hobbs, Marina del Rey (USA)*