

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:

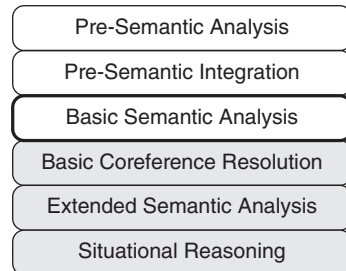
The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.

The title-level DOI for this work is:

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

4

Basic Semantic Analysis



Chapter 2 introduced the basic principles of semantic analysis by LEIAs. As a refresher, the agent builds up the meaning of a sentence by

- syntactically parsing it to identify the verbal head of each clause and its arguments, which are usually noun phrases;
- identifying all available interpretations of the verbal head when used in the given syntactic environment;
- computing all available interpretations of the arguments, which might be multiword constituents (e.g., *a very big, green tree*);
- establishing one or more candidate semantic dependency structures by filling the case role slots (e.g., AGENT, THEME) of available EVENT interpretations with available interpretations of the arguments;
- integrating into the candidate meaning representation(s) the meanings of all other elements of input, such as sentence adverbs, modals, conjunctions, and so on; and
- generating one or more candidate text meaning representations (TMRs) that reflect available analyses; each of these is scored with respect to how well it aligns with the expectations recorded in the lexicon and ontology.

The status of these processes is as follows: They are necessary contributors to Basic Semantic Analysis. They are presented here at a conceptual level, not as an algorithm. In fact, in this book we will not present an algorithm for building a basic semantic analyzer (i.e., implementing stage 3 of analysis) using our knowledge bases and microtheories because (a) many different algorithms can implement this process and (b) such a description would be of interest exclusively to programmers, not the broader readership of this book. To understand why many algorithms can implement the process, consider the following comparison.

Building a basic semantic analyzer is much like putting together a jigsaw puzzle, except that a jigsaw puzzle allows for only one solution, whereas semantic analysis can result in multiple candidate solutions. The puzzle pieces for basic semantic analysis are the syntactic

and semantic descriptions of word senses in the lexicon. Analyzing a sentence involves selecting the best combination of word senses, adjudged using various scoring criteria. Just as a jigsaw puzzle can be approached using many different algorithms—starting with the corners, the outer edges, a set of similarly colored pieces, and so on—so, too, can semantic analysis. In fact, over the years different software engineers in our group have implemented two completely different semantic analyzers based on the same theory, models, and knowledge bases. Although their algorithms organized in different ways how the recorded knowledge and descriptive microtheories were leveraged to analyze inputs, they effectively yielded the same result. For future NLU system developers, this book's main utility is in the descriptions of these knowledge bases and microtheories.¹ As history has shown, knowledge remains useful indefinitely, while system-implementation environments have a relatively short life. With this in mind, we concentrate our presentation around the content of the microtheory of Basic Semantic Analysis and its supporting knowledge bases.

Whereas chapter 2 described the general principles of Basic Semantic Analysis using only very simple examples, this chapter begins a progressively deeper exploration of the kinds of linguistic phenomena that make language at once rich and challenging. The Basic Semantic Analysis described in this chapter attempts to disambiguate the words of input and determine the semantic dependency structure of clauses (essentially, case roles and their fillers) using only the most local context as evidence. But, as this long chapter will reveal, many linguistic phenomena can be treated even at this early stage, when the agent is invoking only a subset of the heuristics and microtheories available to it.

Before proceeding to the phenomena themselves, a few general points deserve mention.

1. Roughly speaking, this stage computes what some linguists call *sentence semantics*: the meaning that can be understood from propositions presented in isolation. For example, the pronoun *he* refers to a male animal, but *which one* must be contextually specified. Similarly, the utterance *I won't!* means *I won't do something*, with that *something* requiring contextual specification. At this stage of processing, the basic analysis (i.e., *male animal* is (ANIMAL (GENDER male))); *do something* is (EVENT)) is recorded in the TMR, along with metadata that includes a call to a procedural semantic routine that can be run later to further specify the meaning.
2. In many cases, different *instances* of a given linguistic phenomenon, such as verb phrase ellipsis or nominal compounding, fall into different functional classes. Some instances can be fully analyzed at this stage, whereas others require heuristics or methods that only become available later.
3. For people, a lot of knowledge about language resides in their construction-rich lexicon; therefore, careful development of the lexicon is key to successful NLU by LEIAs. Any meaning that can be recorded in the lexicon—even using variable-rich constructions—can be processed as part of Basic Semantic Analysis.²

4. One reason for separating Basic Semantic Analysis from its various extensions is that, in many cases, the LEIA does not need a deeper analysis. That is, Basic Semantic Analysis often gives the agent the gist of what the utterance means—which is often enough for it to decide whether digging deeper is needed.

Basic Semantic Analysis covers a large number of linguistic phenomena. The sections below present a representative sampling of them, stopping far short of what needs to be mastered by knowledge engineers working in the environment. The goals of the chapter are

1. to show the benefits of acquiring lexicon and ontology based on how language is actually used, rather than a linguistically pure idealization;
2. to flesh out the notion of *Basic*—as contrasted with *Extended*—Semantic Analysis in the overall functioning of LEIAs; and
3. to underscore just how many semantic phenomena must be treated by natural language processing systems—that is, how much is left to do beyond the tasks most commonly undertaken by mainstream NLP systems.

In wrapping up this introduction, let us make one tactical suggestion, which derives from our experience in giving talks and teaching classes about this material. LEIA modeling is a practical endeavor with one of the key challenges being to balance the desire for human-level agent behavior with the reality that a large amount of rigorous linguistic description remains to be done. It is a fun pastime, but of little practical value, to think up exceptions to every useful generalization about language and the world. We recommend that readers not engage in this because it would distract attention from the considerable amount of information being presented and because fringe cases related both to language (e.g., garden-path sentences³) and to the world (e.g., humans who meow) are just not important enough to demand attention—not until we have gotten much further along in handling the very substantial mass of regular cases. With this in mind, let us move on to the linguistic phenomena handled during Basic Semantic Analysis.

4.1 Modification

The subsections below classify modifiers (primarily adjectives and adverbs, although other syntactic entities can function as modifiers as well) according to how they are described in lexical senses.⁴ Lexical descriptions can include a combination of static descriptors and procedural semantic routines.

4.1.1 Recorded Property Values

The meanings of many modifiers can be described in their lexicon entries as the value of some PROPERTY recorded in the ontology. As background, a PROPERTY is a concept whose

name is an ontological primitive. Each property is constrained by the values for its DOMAIN and RANGE slots. The DOMAIN indicates the types of OBJECTS and/or EVENTS the PROPERTY can apply to, whereas the RANGE indicates the set of its legal fillers.

Properties are divided into three classes based on the fillers for their RANGE slot. The fillers for the RANGE of SCALAR-ATTRIBUTES can be a number, a range of numbers, a point on the abstract scale {0,1}, or a range on that scale. For example,

```
HAPPINESS
  DOMAIN    ANIMAL
  RANGE     {0,1}
```

The fillers for the RANGE of LITERAL-ATTRIBUTES are literals—that is, they are not ontological concepts and are, therefore, not written in small caps. For example,

```
MARITAL-STATUS
  DOMAIN    HUMAN
  RANGE     single, married, divorced, widowed
```

The fillers for the RANGE of RELATIONS are ontological concepts. For example,

```
EXPERIENCER-OF
  DOMAIN    ANIMAL
  RANGE     EVENT
```

Since these ontological descriptions already exist, lexical senses that use properties to describe word meanings need to assert a particular value for the RANGE, but they do not need to repeat the understood constraints on the DOMAIN. For example, the meaning of *happy* is described in the lexicon as .8 on the abstract scale of HAPPINESS.

```
happy-adj1
  def.      experiencing a positive feeling
  ex.       a happy chipmunk
  syn-struct
    adj     $var0
    n       $var1
  sem-struct
    HAPPINESS
      DOMAIN    ^$var1
      RANGE     .8
```

The fact that \$var1 must refer to an ANIMAL is not written in the lexical sense because the fact that HAPPINESS can apply only to ANIMALS is available in the ontology. If \$var1 in a particular input does *not* refer to an ANIMAL, this lexical sense will not be used to analyze the input since the property HAPPINESS cannot apply to it.

Below are specific examples that illustrate how property values are used to describe the meanings of modifiers.

- SCALAR-ATTRIBUTES
 - a happy goat: GOAT (HAPPINESS .8)
 - a severe hailstorm: HAILSTORM (INTENSITY .9)
 - a 19-pound xylophone: XYLOPHONE (WEIGHT 19 (MEASURED-IN POUND))
- LITERAL-ATTRIBUTES
 - a divorcé: HUMAN (GENDER male) (MARITAL-STATUS divorced)
 - an accredited college: COLLEGE (ACCREDITED yes)
- RELATIONS
 - a fearful wolf: WOLF (EXPERIENCER-OF FEAR)
 - a wooden chair: CHAIR (MADE-OF WOOD)
 - to explore endoscopically: INVESTIGATE (INSTRUMENT ENDOSCOPE)
 - to analyze statistically: ANALYZE (INSTRUMENT STATISTICAL-ANALYSIS)
 - a house on fire: PRIVATE-HOME (THEME-OF BURN)

As the lexical sense for happy-adj1 showed, PROPERTY-based descriptions of word senses are recorded in the sem-struct zones of the lexicon. During NLU, they are copied directly into the TMR, meaning that such modifiers present no challenges beyond the potential for ambiguity. For example, *blue* can refer to feeling sad (HAPPINESS .2) or to a color (COLOR blue). In many cases, the meaning of the head noun helps to disambiguate: people and some animals can be sad, whereas it is mostly inanimate objects, as well as some birds and fish, that are colored blue. Of course, nonliteral language can require the opposite interpretations: a *blue person* might describe someone in a blue costume, and a *blue house* might describe a house that invokes feelings of sadness in people who drive by it. However, these examples of extended usage should not sidetrack us from the main point, which is that the meaning of the modified lexical item (such as a noun) tends to provide good disambiguating power for lexical items (such as adjectives) that modify it. During Basic Semantic Analysis, if multiple candidate analyses are possible, they are all generated and scored, and the LEIA waits until later in the process to select among them.⁵

4.1.2 Dynamically Computed Values for Scalar Attributes

All the modifiers discussed so far could be described using static meaning representations recorded in the sem-struct zone of lexical senses. However, the meaning of adverbs that modify scalar attributes—such as *very* and *extremely*—must be computed dynamically, taking into consideration the meaning of the particular adjective they are modifying. For example, *smart* and *dumb* are described using the scalar attribute INTELLIGENCE, with *smart*

having a value of .8 and *dumb* .2. If *very* is added to *smart*, it raises the value from .8 to .9, whereas if *very* is added to *dumb* it decreases the value from .2 to .1. This dynamic modification of property values is carried out using a procedural semantic routine recorded in the lexical sense for *very*. Since this routine is simple, local, and not error-prone, it is carried out as part of Basic Semantic Analysis.⁶

By way of illustration, this section describes two procedural semantic routines involving scalar attributes:

1. *delimit-scale*, for adverbs that modify the abstract value of a scalar attribute: for example, *very experienced*, *not extremely smart*; and
2. *specify-approximation*, for adverbs that indicate the bidirectional stretching of a cited number: for example, *around 10:00*, *about 80 feet wide*.

Delimit-scale. Delimit-scale is the procedural semantic routine that calculates the modified value of a scalar attribute that is expressed as a point on the abstract {0,1} scale. It is called from lexical senses for words like *very*, *extremely*, *quite*, *moderately*, and *somewhat* when they modify a scalar attribute. The function takes three arguments:

- the meaning of the word modified: for example, *intelligent* is (INTELLIGENCE .8), *dim* is (INTELLIGENCE .2);
- which direction along the scale the value should be shifted—toward the mean (e.g., *relatively intelligent*) or toward the extreme (*very intelligent*); and
- the amount the value should be shifted (e.g., .1, .2, .3).

So, *very small* is calculated by taking the value of *small* (SIZE .2) and shifting it to the extreme by .1, returning the value (SIZE .1). Analogously, *moderately small* is calculated by taking the value of *small* (SIZE .2) and shifting it toward the mean by .1, returning the value (SIZE .3).

An interesting situation occurs if one modifies a scalar value such that it is off the scale. For example, *extremely* is defined as shifting the scalar value by .2 toward the extreme, so an *extremely, extremely expensive car* will be calculated as: *expensive* (COST .8) + *extremely* (.2) + *extremely* (.2) = (COST 1.2). The value 1.2 lies outside the {0,1} scale—but this is exactly what we want as an interpretation of *extremely extremely*! To understand why this is so, we must return to the nature of the ontology.

In the LEIA's ontology, properties are associated with different *facets*, including

- *sem*, which introduces the typical selectional restrictions;
- *default*, which introduces a more restricted, highly typical subset of *sem*; and
- *relaxable-to*, which represents an extended interpretation of *sem*.

Whereas the fillers for all these facets are explicitly listed for objects, events, and literal attributes, there is often no need to explicitly state the *relaxable-to* values for scalar attributes: they are simply outside the range of *sem*, with the degree of unexpectedness depending on

how far outside the range they are. So, it is not impossible for a car to cost a million dollars, even though that would be far outside the expectations recorded in the *sem* facet of the property *COST* in the ontological description of *AUTOMOBILE*.

Returning to our example of *extremely, extremely expensive*, a perfectly valid value of 1.2 is being returned, which means that the value lies significantly beyond the expectations recorded in the *sem* facet of *COST*.

Let us reiterate the reason for specifying relative scalar values in the first place. If one text says that the president of France bought a *very expensive car* using government funds, and another text says that the president of Russia bought an *extremely, extremely expensive car* using government funds, a person—and, accordingly, a LEIA—should be able to answer the question, “Who bought a more expensive car, the president of France or the president of Russia?” The answer is the president of Russia—at least if the authors of the texts were reasonably consistent in their evaluations of cost.

Of course, scalar values between 0 and 1 have a concrete meaning only if one knows the objective range of values for the given property applied to the given object or event. That is, an *expensive car* implies a different amount of money than an *expensive jump rope* or an *expensive satellite*. Moreover, an expensive Kia is less expensive than an expensive BMW—something that the LEIA can reason about if the associated information is either recorded in its knowledge bases or is available in resources (e.g., on the internet) that it can consult at runtime.

Specify-approximation. Some lexical items indicate that the cited number is an approximation: for example, *around, about, approximately*. Approximating is a useful human skill, fortunately accompanied by related expressive means in language. When you say you will arrive at *around 3:00*, you are intentionally being vague. If pushed to estimate what you mean, if you are a typical person, you might say *give or take 10 minutes*, but if you are a habitually late person, you might reply *probably no earlier than 3:20*. The estimations described in this section are population-wide generalizations that do not incorporate the modeling of any particular individual—though that is an interesting problem that is worth pursuing in applications such as personal robotic assistants. The fact that calculated numerical values are estimates is reflected in TMR by appending the metaproperty value “approx. +” to a calculated property value. This feature allows the LEIA to reason as follows: If it hears *Louise came at around 3:00* it will resolve the time to *2:50–3:10 (approx. +)*. If it later hears that Louise came at 3:12, it will not flag a discrepancy; it will simply use the more precise time as the actual answer. Note that flagging involves extralinguistic reasoning that must be specially incorporated, if needed, into the reasoning engines of particular application systems.

As a default strategy, the LEIA approximates the implied range using a 7% rule, which works pretty well in many contexts. For example:

- About 5 gallons ($5 * .07 = .35$) resolves to between 4.65 and 5.35 gallons.
- About 150 lbs. ($150 * .07 = 10.5$) resolves to between 139.5 and 160.5 lbs.

This maximally simple rule would, of course, benefit from a follow-up rounding procedure; however, it would need to be rather sophisticated since something as simple as round-up to the next whole number would clearly not work.

As mentioned earlier, all such calculations are appended with the marker “approx. +” in the TMR to make it clear that the text itself included an approximation.

The 7% rule is too coarse-grained on at least two counts. First, the actual number from which the approximation derives is important in terms of what the approximation actually means. In most cases, one approximates from numbers like 10, 25, 100, or 5,000,000. It’s unusual to say *about 97 feet* or *around 8.24 pounds*. If, however, someone does use such a turn of phrase, its interpretation must be different from what is returned by the 7% rule. We have not yet pursued such pragmatically nonnormative cases.

The second way in which the 7% rule is too coarse-grained is that approximations work in idiosyncratic ways for certain semantic classes. Here we take just a few examples:

- *Heights of people.* Using the 7% rule, *about 6 feet tall* would give a range of over 5 inches on either side, which is far too broad. Instead, fixing the approximation at 1 or 2 inches either way seems closer to what people do.
- *Ages.* Interpreting the approximation of a person’s age depends on how old the person is, with the 7% rule working poorly for children but better for adults. For example, the 7% rule would make a baby who is *about 5 days old* “5 days +/- 8.4 hours,” and it would make a child who is *about 5 years old* “5 years +/- 3.5 months.” Clearly “give or take a day” and “give or take a year” are better approximations. As a person gets older, the 7% rule works better: a person who is about 80 years old would be roughly 75–85, and a person who is about 50 years old would be 46.5–53.5. In this case, as for heights, it seems more direct to simply set the buffer for the approximation of given age ranges rather than try to force the 7% rule.
- *Clock time.* The 7% rule is reasonable for round clock times but much less so for more precise clock times. Rather than employ it, we are using a different approach to calculating approximate clock times:
 - around (hour, ½ hour, *noon*, or *midnight*) = +/- 10 minutes
 - around (:05, :10, :15, :20, etc.) = +/- 5 minutes
 - around (other) = +/- 2 minutes

Whereas special cases like these are easy to treat individually, creating a full inventory of special cases would be a tedious process of questionable utility. However, if it *were* deemed useful, it could be outsourced to anyone with the patience to work out the details. After all, this is not a matter of linguistics; it is one of world knowledge.

4.1.3 Modifiers Explained Using Combinations of Concepts

The meaning of some modifiers is best explained using combinations of concepts. Consider the meaning of *overboard* in *A sailor threw some food overboard*. As explained in the last chapter, *overboard* tells us that before the THROW event, the food was in some surface water vehicle, and after the THROW event, it is in the body of water in which that vehicle is floating. In short, *overboard* tells us about the SOURCE and DESTINATION of a MOTION-EVENT (THROW is an ontological descendant of MOTION-EVENT). Accordingly, the lexical entry for this physical sense of *overboard* specifies

- that the event being modified must be a MOTION-EVENT; if it is not, the usage is metaphorical (e.g., *She went overboard in decorating for the party*);
- that its SOURCE is a SURFACE-WATER-VEHICLE; and
- that its DESTINATION is a BODY-OF-WATER.

This lexical sense allows the agent to generate the following TMR for *A sailor threw some food overboard*. Note that the representation of *some food* is simplified to just *food* to avoid introducing set notation before its time.

THROW-1

AGENT	SAILOR-1
THEME	FOOD-1
SOURCE	SURFACE-WATER-VEHICLE
DESTINATION	BODY-OF-WATER
TIME	<find-anchor-time

If the input indicates the source or destination of the MOTION-EVENT explicitly, as in *A sailor threw some food overboard into the Mediterranean Sea*, then the associated constraint is replaced by the actual value provided.

THROW-2

AGENT	SAILOR-2
THEME	FOOD-2
SOURCE	SURFACE-WATER-VEHICLE
DESTINATION	SEA-5
TIME	<find-anchor-time

SEA-5

HAS-NAME	'Mediterranean Sea'
DESTINATION-OF	THROW-2

Although *overboard into* might sound less typical than *overboard* by itself, a destination is actually included in many examples in the COCA corpus: “overboard into the wind-whipped waves,” “overboard into Ellicott Bay,” “overboard into the sea,” and others.

Overboard is not an unusually complex singleton. Plenty of modifiers are most productively described using multiple concepts. Consider several more examples:

- *Ad infinitum*, when used to modify a SPEECH-ACT (*Grizelda talks ad infinitum!*), adds two features to the SPEECH-ACT: (DURATION 1) and (EVALUATIVE .1).⁷ Both of these are measured on the abstract scale {0,1}. The first means “for an extremely long time,” and the second one reflects the speaker’s negative evaluation of how long the person is talking.
- *Occasional*, when used to modify an event (*an occasional run*), adds two features: iterative aspect (ITERATION multiple) and a typical time interval that is pretty long (TIME-INTERVAL .7).
- *Ambulatory*, when used to modify a person, indicates that the person can walk. So *an ambulatory patient* is analyzed as MEDICAL-PATIENT (AGENT-OF WALK (POTENTIAL 1)). This meaning representation includes the highest value of potential modality on the abstract scale {0,1}, which indicates that the person is able to walk.
- *Argumentative*, when used of a person, indicates that the person argues regularly and that this is viewed negatively by the speaker. So *an argumentative coworker* is analyzed as COLLEAGUE (AGENT-OF ARGUE) (ITERATION multiple) (EVALUATIVE .2). Although arguing is not inherently negative, being argumentative is.
- *Abusive*, when used of a person, indicates that the person is the agent of repeating ABUSE events. So *an abusive spouse* is SPOUSE (AGENT-OF ABUSE (ITERATION multiple)). Here it is not necessary to indicate the negative evaluation because ABUSE is always negative and is ontologically specified as such.

In sum, describing modifiers using combinations of already-available concepts gives the LEIA more reasoning power than inventing endless properties for every modifier used in a language.

4.1.4 Dynamically Computed Values for Relative Text Components

Relative text components are indicated by expressions like *the preceding paragraph*, *this section’s heading*, and *X and Y, respectively*. We will use the latter (*the latter*—a relative text expression!) for illustration. Consider the pair of sentences (4.1a) and (4.1b), which are a bit stilted but are useful because they generate simple and concise TMRs.⁸

- (4.1) a. The doctor and the nurse ordered soup and spaghetti.
 b. The doctor and the nurse ordered soup and spaghetti, respectively.

The first sentence says nothing about who ordered what. It says only that the set comprised of the doctor and the nurse ordered the set comprised of soup and spaghetti. It could be that both people ordered both dishes, that they jointly ordered one portion of each dish to share, or that each one ordered a different dish. (See section 4.1.5 for more on sets.) However, add the adverb *respectively* and everything changes: The doctor ordered the soup and the nurse ordered the spaghetti.

The procedural semantic routine for *respectively* first computes the meaning of the input without the modifier in order to determine whether the TMR shows the necessary types of parallelism. That is, if a set of two elements fills the AGENT case role, then a set of two elements must fill the THEME case role, as in our example.

The doctor and the nurse ordered soup and spaghetti.

ORDER-IN-RESTAURANT-1

AGENT	SET-1
THEME	SET-2

SET-1⁹

ELEMENTS	PHYSICIAN-1, NURSE-1
CARDINALITY	2

SET-2

ELEMENTS	SOUP-1, SPAGHETTI-MEAL-1
CARDINALITY	2

If there is no such parallelism, then further analysis of *respectively* is not undertaken as it will not work. Instead, the input will be left incompletely analyzed, the TMR will receive a low score, and the LEIA will have the opportunity later to attempt to recover from the failure—for example, by asking its human collaborator for clarification.

To incorporate the meaning of *respectively* into this TMR the LEIA must carry out the following five steps.

1. Verify that the EVENT to which this modifier is attached (here: ORDER-IN-RESTAURANT) has more than one property (here: AGENT and THEME), whose fillers are sets with the same cardinality, which must be greater than 1. If these conditions do not hold, the function stops because its reasoning will fail.

2. Create n new instances of the EVENT, such that n equals the cardinality of the sets, and remove the original single instance from the TMR, leaving

ORDER-IN-RESTAURANT-2

ORDER-IN-RESTAURANT-3

3. Add all case roles from the original instance, but without fillers:

ORDER-IN-RESTAURANT-2

AGENT
THEME

ORDER-IN-RESTAURANT-3

AGENT
THEME

4. Pair up the members of the sets across their respective case roles, following the word order of the input. This results in the following TMR.

ORDER-IN-RESTAURANT-2

AGENT	PHYSICIAN-1
THEME	SOUP-1

ORDER-IN-RESTAURANT-3

AGENT	NURSE-1
THEME	SPAGHETTI-MEAL-1

5. Copy any properties of the event that are not involved in the sets—e.g., adverbs of manner—into all new event frames. (Not applicable in our example.)

This is a good time to reiterate one of the benefits of the Ontological Semantics approach to language modeling. Once this meaning procedure has been implemented for one language, it can be ported to the lexicons and analysis systems of other languages, resulting in a substantial savings of time and effort. For example, although in Russian the adverb meaning *respectively*, *sootvetsvenno*, must be preposed rather than postposed—as shown by (4.2)—the basic TMR created without the adverb will be exactly the same as for English, so the procedural semantic routine to compute the adverb’s meaning can be exactly the same as well.

(4.2) Anna i Ivan zakazali, sootvetsvenno, sup i spagetti.
 Anna and Ivan ordered, respectively, soup and spagetti.

4.1.5 Quantification and Sets

Although quantification has attracted a remarkable amount of scholarship within the field of formal semantics, for language understanding by LEIAs, it is no more important than dozens of other linguistic phenomena. Not having become a special priority to date, the associated microtheory—while covering the needs of our current application systems—is understood to require further development.

The notation used to describe sets and quantifiers is shown below. The comment after each semicolon indicates the valid types of fillers for the given property.

SET

MEMBER-TYPE	; an ontological concept
ELEMENTS	; a list of concept instances
NOT-ELEMENTS	; a list of concept instances or a pointer to another set
QUANT	; 0–1, indicating a relative amount or percentage
CARDINALITY	; any integer
RELATIVE-TO-NORM	; 0–1 (used to describe words like ‘too’ in ‘too many’)

HAS-SUBSET	; a pointer to another set
SUBSET-OF	; a pointer to another set
COMPLETE	; yes or no
ORDINALITY	; any integer
SET-TYPE	; conjunctive or disjunctive
<i>any property</i>	; any valid filler for the given property

The following are examples of how sets are represented in TMRs.

Wolves

SET -1

MEMBER-TYPE	WOLF
CARDINALITY	> 1

too many gray wolves

SET-2

MEMBER-TYPE	WOLF
CARDINALITY	> 1
RELATIVE-TO-NORM	> 0.6
COLOR	gray

two gray wolves

SET-3

MEMBER-TYPE	WOLF
CARDINALITY	2
COLOR	gray

both wolves

SET-4

MEMBER-TYPE	WOLF
CARDINALITY	2
COMPLETE	yes

80% of wolves

SET-5

MEMBER-TYPE	WOLF
QUANT	.8

a lot of wolves

SET-6

MEMBER-TYPE	WOLF
QUANT	.7 < > .9

Two of the wolves are hungry

SET-7

MEMBER-TYPE	WOLF
CARDINALITY	> 2
HAS-SUBSET	SET-8

SET-8

MEMBER-TYPE	WOLF
CARDINALITY	2
EXPERIENCER-OF	HUNGER-1
SUBSET-OF	SET-7

all of the wolves except for Whiskers

SET-9

MEMBER-TYPE	WOLF
QUANT	1
NOT-ELEMENTS	WOLF-1

WOLF-1

HAS-PERSONAL-NAME	'Whiskers'
-------------------	------------

Gray wolves are sleeping

SET-10

MEMBER-TYPE	WOLF
CARDINALITY	> 1
COLOR	gray
EXPERIENCER-OF	SLEEP-1

A few of the properties used to describe sets deserve further comment.

- **RELATIVE-TO-NORM** indicates a state relative to what is considered optimal: for example, *too many wolves*. A value of 0.5 indicates the optimal state, with lower values indicating 'too little/few' and higher values 'too much/many.'
- **COMPLETE** indicates whether a set has been exhaustively described. It is overtly specified only if the value is 'yes.' It is used to described strings like *both* and *all three*.
- Any property defined in the ontology can be included in a set description. Our examples show **COLOR** and **EXPERIENCER-OF**.

The fact that any property, including **RELATIONS**, can be included in a set description opens up a much larger discussion regarding (a) what it means to fully interpret a set reference and (b) when and how this is best done by LEIAs.

Compare the inputs *Gray wolves are sleeping* and *Both gray wolves are sleeping*. Whereas *gray wolves* leaves the number of wolves unspecified, *both gray wolves* tells us that there are exactly two. The latter allows for a much more precise meaning representation to be generated: namely, the agent can generate and store in memory exactly two different instances of **WOLF**, each of which is engaging in a different instance of **SLEEP**. This provides more precise information for downstream reasoning.

The full analysis of *Both gray wolves are sleeping* is generated in two stages. The first stage records the meaning using the set notation just described, resulting in the following initial TMR:

SET-1	
MEMBER-TYPE	WOLF
CARDINALITY	2
COMPLETE	yes
COLOR	gray
EXPERIENCER-OF	SLEEP-1

The need for an enhanced analysis is triggered by the fact that the set has a specific cardinality.¹⁰ In all such cases, the agent can generate the listed number of instances of the given MEMBER-TYPE and attribute to each one of them all of the properties applied to the set. (The fact that this type of set expansion might not be desirable for very large sets, especially if they are not described by any additional properties—e.g., *10,000 grains of sand*—is a separate decision that we will not explore here.)

The initial TMR above can be automatically expanded into the more explicit TMR below:

SET-1	
ELEMENTS	WOLF-1, WOLF-2
CARDINALITY	2
COMPLETE	yes
SLEEP-1	
EXPERIENCER	WOLF-1
WOLF-1	
MEMBER-OF	SET-1
COLOR	gray
EXPERIENCER-OF	SLEEP-1
SLEEP-2	
EXPERIENCER	WOLF-2
WOLF-2	
MEMBER-OF	SET-1
COLOR	gray
EXPERIENCER-OF	SLEEP-2

The functional difference between the basic and reasoning-enhanced set notation might not be immediately obvious because when you—a powerful human reasoning machine—read *Both gray wolves are sleeping*, you automatically see in your mind’s eye two different gray wolves sleeping. The agent, by contrast, must carry out the reasoning processes we described to make this explicit. This is all done as part of Basic Semantic Analysis.

Additional aspects of set-based reasoning can also be handled during Basic Semantic Analysis, as long as they are prepared for by recording appropriate constructions in the lexicon. Below are two examples of lexical senses that record constructions involving sets. The nature of each construction is made clear in the definition and example slots. Each sense includes a call to a specific procedural-semantic routine that will compute

construction-specific aspects of meaning. Whereas the procedural-semantic routine for the first example is rather simple, the procedural-semantic routine for the second is not. (Note that numbered instances of the variable *refsem* are used for certain instances of coreference within a lexical sense.)

out_of-prep2

def. The construction 'NUM out of NUM N'

ex. three out of four dentists

syn-struct

num \$var1

prep \$var0

num \$var2

n \$var3

sem-struct

SET

MEMBER-TYPE ^\$var3

QUANT refsem1

meaning-procedures

refsem1 = (value ^\$var1 / value ^\$var2)

The TMR for *Nine out of ten people* that is generated using the sense out_of-prep2 is

SET-1

MEMBER-TYPE HUMAN

QUANT .9

The second example involves a construction recorded as a lexical sense of *and*. Since the meaning procedure is complicated, we use subscripts to indicate how its elements align with the example sentence listed in and-conj31.

and-conj31

def. The construction 'the ORD and ORD most ADJ N'

ex. the fourth and fifth most important companies

syn-struct

det \$var1 (root 'the')

num \$var2 (type ordinal)

and \$var0

num \$var3 (type ordinal)

adv \$var4 (root 'most')

adj \$var5

N \$var6

sem-struct

SET

MEMBER-TYPE ^\$var6

ELEMENTS refsem1, refsem2

^\$var1 null-sem+

^\$var4 null-sem+

meaning-procedures

refsem1 =

^\$var6-#1_{CORPORATION} (^\$var5_{IMPORTANCE} set-#1.ordinality.^\$var2_{FOURTH})

value refsem2 =

^\$var6-#2_{CORPORATION} (^\$var5_{IMPORTANCE} set-#1.ordinality.^\$var3_{FIFTH})

Let us trace how the example for this lexical sense, *the fourth and fifth most important companies*, will be analyzed. The sem-struc description generates the following TMR chunk:

SET

MEMBER-TYPE	CORPORATION
ELEMENTS	refsem1, refsem2

The refsem calculations, for their part, are as follows. They include a dot notation that is used in many ways throughout the LEIA's knowledge bases.

refsem1

CORPORATION-1

IMPORTANCE	SET-1.ORDINALITY.4
------------	--------------------

refsem2

CORPORATION-2

IMPORTANCE	SET-1.ORDINALITY.5
------------	--------------------

When these chunks are combined, they yield the final TMR:

SET-1

MEMBER-TYPE	CORPORATION
ELEMENTS	CORPORATION-1, CORPORATION-2

CORPORATION-1

IMPORTANCE	SET-1.ORDINALITY.4
------------	--------------------

CORPORATION-2

IMPORTANCE	SET-1.ORDINALITY.5
------------	--------------------

An additional example will illustrate how lexical senses like these accommodate many variations on a theme. The TMR for *the second and third most popular TV shows* that is generated using the lexical sense above is as follows:

SET-1

MEMBER-TYPE	TELEVISION-PROGRAM
ELEMENTS	TELEVISION-PROGRAM-1, TELEVISION-PROGRAM-2

TELEVISION-PROGRAM-1

POPULARITY	SET-1.ORDINALITY.2
------------	--------------------

TELEVISION-PROGRAM-2

POPULARITY	SET-2.ORDINALITY.3
------------	--------------------

Writing lexical senses for constructions like these, and implementing the associated procedural semantic routines, involves straightforward knowledge engineering. It is not difficult; it just takes time. By contrast, some aspects of set-related reasoning cannot be handled by lexical senses alone. For example, sometimes set-related information is distributed across multiple clauses, as in the following examples from the Cornell Natural Language Visual Reasoning corpus (CNLVR; Suhr et al., 2017):

(4.3) There are two towers with the same height but their base is not the same in color. (CNLVR)

(4.4) There is a box with 4 items of which one is a blue item and 3 are yellow items. (CNLVR)

The interpretation of such sentences must be distributed over multiple stages of processing to take care of coreference resolution across clauses as well as set-oriented reasoning beyond what can be recorded in lexical constructions. Preparing agents to analyze such inputs is challenging. For example, in (4.3) the singular noun phrase *their base* actually refers to two different bases, and *not the same in color* requires comparison of the values for COLOR for the bases of the two towers. Similarly, in (4.4), *of which* establishes a reference relation with the *4 items*, and the head nouns associated with *one* (item) and *3* (items) are elided.

Another challenge presented by sets involves interpreting the scope of modifiers in inputs such as “the ADJ N and N.” In some cases, the modifier applies only to the first noun: in *the old men and children*, the children are clearly not old. In other cases, the intended meaning is ambiguous: in *the old men and women*, the women may or may not be old.

4.1.6 Indirect Modification

Indirect modification occurs when a modifier syntactically modifies one constituent but semantically modifies another. Most cases of indirect modification are best handled by dedicated lexical senses that reconstruct the intended meaning. For example:¹¹

- When *married* modifies an event—e.g., *married sex*, *married conflicts*—it introduces into the context a married couple that is the agent or experiencer of the event.
- When *responsible* modifies an event—e.g., *responsible decision-making*, *responsible parenting*—it introduces into the context one or more human agents who are carrying out the given event responsibly.
- When *rural* modifies an abstract noun—e.g., *rural poverty*, *rural income*—it introduces into the context an unspecified set of people who live in a rural area. The semantic relationship between those people and the modified noun is not explicitly indicated and must be dynamically computed (just as with nominal compounds).
- When *sad* modifies anything but an animal, its meaning depends on the meaning of the noun it modifies. For example, when *sad* modifies a temporal expression—e.g., *sad time*, *sad year*—it introduces into the context one or more people who are sad during that time. When it modifies an abstract object—e.g., *sad song*, *sad news*—it means that

the object makes people feel sad. And when it modifies a person's body part or expression—e.g., *sad eyes*, *sad smile*—it means that the associated person is sad.

The meaning representations recorded in the sem-structs of these senses include entities—most often, humans—that are not explicitly mentioned in the local dependency structure, though they might appear somewhere in the preceding context. Where things get interesting is with respect to specifying who, exactly, these entities are in contexts in which that information is provided. This can often be done using textual coreference procedures. For example:

- Interpreting *The committee is tasked with responsible decision-making* requires that the generic HUMAN(s) posited in the meaning representation for this sense of *responsible* be coreferred with the committee.
- Interpreting *Look into her sad eyes and you'll understand everything* requires that the generic HUMAN posited in the meaning representation for this sense of *sad* be coreferred with *her*.
- Interpreting *Married conflicts affect young couples' relationships* requires that the generic set of HUMANS posited in the meaning representation for this sense of *married* be coreferred with the *young couples* mentioned subsequently.

The procedural semantic routines needed to resolve these coreferences are recorded in the meaning-procedures zones of the associated lexical senses. They are called during the next stage of processing, Basic Coreference Resolution (chapter 5).

4.1.7 Recap of Modification

- Recorded property values
 - Scalar attributes—*accomplished person*: HUMAN (EXPERTISE-ATTRIBUTE .8)
 - Literal attributes—*divorced man*: HUMAN (GENDER male) (MARITAL-STATUS divorced)
 - Relations—*fearful fox*: FOX (EXPERIENCER-OF FEAR)
- Dynamically computed values for scalar attributes—*very smart*: *smart* (INTELLIGENCE .8) + *very* (increase value by .1) = (INTELLIGENCE .9)
- Modifiers explained using combinations of concepts—*argumentative person*: HUMAN (AGENT-OF ARGUE (ITERATION multiple) (EVALUATIVE .2))
- Dynamically computed values for relative text components—*The doctor and the nurse ordered soup and spaghetti, respectively* is effectively analyzed as the meaning of [*The doctor ordered soup*] and [*The nurse ordered spaghetti*]
- Quantification and sets—*Many wolves*: SET (MEMBER-TYPE WOLF) (QUANT .7 < > .9)
- Indirect modification—*sad smile*: SMILE (AGENT HUMAN (HAPPINESS-ATTRIBUTE .2))

4.2 Proposition-Level Semantic Enhancements

The meaning of simple propositions can be enhanced in various ways. These include introducing values of modality or aspect; using a nonbasic (i.e., nondeclarative) mood, such as the imperative or interrogative; using the proposition as the complement of a non-modal, non-aspectual matrix verb; and combining any of these. Given the simple proposition *The firefighter swims*, whose bare-bones¹² TMR is

SWIM-1	
AGENT	FIREFIGHTER-1
TIME	<i>find-anchor-time</i>

we can add various semantic enhancements like the following:

- (4.5) [Add modality]
The firefighter *wants to* <*is trying to, can't, won't*>swim.
- (4.6) [Add aspect]
The firefighter *is* swimming <*is starting to swim, finished* swimming>.
- (4.7) [Add a non-modal, non-aspectual matrix verb]
The gardener *sees* the firefighter swimming.
- (4.8) [Add the interrogative mood]
Who is swimming? *Is* the firefighter swimming?
- (4.9) [Add the imperative mood]
Swim!
- (4.10) [Combine several of the above]
Did the gardener *tell* the firefighter *to try to start* swimming?

We will consider each of these proposition-level enhancements in turn.

4.2.1 Modality

Ontological Semantics distinguishes ten types of modality, listed in table 4.1.¹³ Each modal meaning is described in a MODALITY frame with the following properties:

- TYPE: effort, epistemic, and so on;
- VALUE: any value or range on the abstract scale {0,1};
- SCOPE: the ontological concept instance of the head of the proposition; and
- ATTRIBUTED-TO: indicates the individual responsible for reporting the modal meaning. By default, it is the speaker. For example, a speaker who says *Martina loves skydiving* is taking responsibility for Martina's positive evaluation of it. Martina might actually hate it.

In our firefighter example, two modality-enhanced versions are juxtaposed below.

Table 4.1

Types of modality used in Ontological Semantics

Modality type	Informal definition	Sample realizations
Effort	indicates effort expended	try to, not bother to
Epistemic	indicates factivity	did, didn't, will, won't
Epiteutic	indicates success/failure	succeed in, fail to, barely manage to
Evaluative	indicates an assessment	be pleased with, disapprove of
Intentional	indicates intention	intend to, be planning to
Obligative	indicates obligation, requirement	must, have to, not be required to
Permissive	indicates permission	may, may not, can, can't
Potential	indicates ability, potential	can, can't, be unable to
Volitive	indicates want, desire	want do, be dying to, not have any desire to
Belief	indicates belief	believe, disbelieve

The firefighter wants to swim.

MODALITY-1

TYPE volitive
 VALUE 1
 SCOPE SWIM-1
 ATTRIBUTED-TO HUMAN-1 ("speaker")
 TIME *find-anchor-time*

SWIM-1

AGENT FIREFIGHTER-1

The firefighter wants to try to swim.

MODALITY-1

TYPE volitive
 VALUE 1
 SCOPE MODALITY-2
 ATTRIBUTED-TO HUMAN-1 ("speaker")
 TIME *find-anchor-time*

MODALITY-2

TYPE effort
 VALUE 1
 SCOPE SWIM-1
 ATTRIBUTED-TO HUMAN-1 ("speaker")

SWIM-1

AGENT FIREFIGHTER-1

4.2.2 Aspect

Aspect reflects specific time-related characteristics of an event. Formally, it scopes over a proposition, just like modality. It divides into PHASE and ITERATION, which have the following value sets:

- PHASE: *begin* (indicates the start of an event), *continue* (indicates its progression), *end* (indicates its completion), and *begin-continue-end* (represents the action viewed as a whole); and
- ITERATION: *single* (the event occurs once), *multiple* (it occurs multiple times), *any number or range of numbers*.

Aspect interacts with tense in complex ways that differ considerably across languages. Developing an overarching microtheory of tense and aspect—including time-tracking events throughout discourses—is a big job that we have not yet undertaken. When that microtheory comes on agenda, it will be informed by recent work on representing the relative times of events in corpus annotation (e.g., Mani et al., 2005).

Currently, TMRs explicitly show values of aspect in two cases: (a) when the value of PHASE can be unambiguously identified from verbal features—for example, present-progressive verb forms (*Paavo is running*) have “PHASE continue”—and (2) when they are instantiated by the lexical descriptions of specific words and phrases. For example, the verb *begin* is described in the lexicon as adding “PHASE begin” to the meaning of the event it scopes over. Similarly, the adverb *repeatedly* is described as adding “ITERATION multiple” to the meaning of the event it scopes over. Two more versions of our firefighter example illustrate the representation of aspect.

The firefighter is starting to swim.

ASPECT-1	
PHASE	begin
SCOPE	SWIM-1
TIME	<i>find-anchor-time</i>

SWIM-1	
AGENT	FIREFIGHTER-1

The firefighter wants to start swimming.

MODALITY-1	
TYPE	volitive
VALUE	1
SCOPE	ASPECT-1
ATTRIBUTED-TO	HUMAN-1 (“speaker”)
TIME	<i>find-anchor-time</i>

ASPECT-1	
PHASE	begin
SCOPE	SWIM-1
SWIM-1	
AGENT	FIREFIGHTER-1

4.2.3 Non-Modal, Non-Aspectual Matrix Verbs

Many verbs that convey meanings apart from modality and aspect take propositions as their complements. In such cases, the meaning of the matrix verb has a **THEME** slot that is filled by the meaning of the complement. This is illustrated below using the verbs *see* and *hear*.

The gardener saw the firefighter swimming.

INVOLUNTARY-VISUAL-EVENT-1

EXPERIENCER	GARDENER-1
THEME	SWIM-1
TIME	<find-anchor-time

SWIM-1

AGENT	FIREMAN-1
-------	-----------

The gardener heard the boy playing the trumpet.

INVOLUNTARY-AUDITORY-EVENT-1

EXPERIENCER	GARDENER-1
THEME	PLAY-MUSICAL-INSTRUMENT-1
TIME	<find-anchor-time

PLAY-MUSICAL-INSTRUMENT-1

AGENT	HUMAN-1
THEME	TRUMPET-1

HUMAN-1

GENDER	male
AGE	< 16 (MEASURED-IN YEAR)

4.2.4 Questions

There are many types of questions, including yes/no questions, *wh*-questions, choice questions, and tag questions. All questions set up a **REQUEST-INFO** frame whose **THEME** slot is filled by what is being asked. In many cases, representing what is asked requires the use of dot notation. Below are two examples of question frames that use this notation. The first example asks for the **AGENT** of **SWIM**, represented as **SWIM-1.AGENT**. The second example asks whether or not the statement is true—that is, it asks for the value of episodic modality scoping over the proposition (if it is true, the value is 1; if it is false, the value is 0).

Who is swimming?

REQUEST-INFO-1

AGENT	HUMAN-1 ("speaker")
THEME	SWIM-1.AGENT
BENEFICIARY	HUMAN-2 ("hearer")

SWIM-1

TIME	find-anchor-time
------	------------------

Is the fireman swimming? or The fireman is swimming, isn't he?

REQUEST-INFO-1

AGENT HUMAN-1 ("speaker")
THEME MODALITY-1.VALUE
BENEFICIARY HUMAN-2 ("hearer")

MODALITY-1

TYPE epistemic
SCOPE SWIM-1

SWIM-1

AGENT FIREMAN-1
TIME *find-anchor-time*

Most interrogative inputs are recognized as such during syntactic analysis, at which point the feature “interrogative+” is incorporated as metadata into the nascent TMR. Then, at this stage of Basic Semantic Analysis, that feature value is converted into an instance of the concept REQUEST-INFO. Special cases include tag questions and indirect questions. Tag questions are recorded in the lexicon as constructions whose syntactic components are illustrated in table 4.2. Indirect questions are discussed in section 4.4.

4.2.5 Commands

Commands, also called imperatives, are propositions in the imperative mood. They are recognized by the syntactic parser during Pre-Semantic Analysis. During Basic Semantic Analysis, the feature “imperative+” triggers the instantiation of a REQUEST-ACTION concept whose THEME is the action in question and whose main case role—usually AGENT—is the hearer, as illustrated below.¹⁴

Swim!

REQUEST-ACTION-1

AGENT HUMAN-1 ("speaker")
THEME SWIM-1
BENEFICIARY HUMAN-2 ("hearer")

SWIM-1

AGENT HUMAN-2 ("hearer")

Table 4.2
Examples of syntactic components of tag-question constructions

Clause	,	Aux	(neg)	pronoun	?
The Murphys won't go	,	will		they	?
Leslie can ski	,	can	't	she	?
Horatio should eat the nuts	,	should	n't	he	?

Eat two bananas, now!

REQUEST-ACTION-1

AGENT	HUMAN-1 ("speaker")
THEME	INGEST-1
BENEFICIARY	HUMAN-2 ("hearer")

INGEST-1

AGENT	HUMAN-2 ("hearer")
THEME	SET-1
URGENCY	1

SET-1

MEMBER-TYPE	BANANA
CARDINALITY	2

4.2.6 Recap of Proposition-Level Semantic Enhancements

- Modality: The firefighter *wants to* <*is trying to, can't, won't*>swim.
- Aspect: The firefighter *is starting* to swim <*finished* swimming>.
- Non-modal, non-aspectual matrix verbs: The gardener *sees* the firefighter swimming.
- Questions: *Is* the firefighter swimming? *Who* is swimming?
- Commands: *Swim!*
- Combinations: The firefighter *wants to start* swimming.

4.3 Multicomponent Entities Recorded as Lexical Constructions

Knowledge of a language is, in large part, knowledge of how words fit together to express particular meanings. Even memorizing 10,000 individual words of English would not prepare a non-English speaker to express in a normal, idiomatic way ideas like, "Seconds, anyone?" "Are we there yet?" "Time's up." or "I will if you will." In fact, the lion's share of the work in learning a new language is memorizing tens of thousands of instances of how things are said.¹⁵ Whereas some of these are fixed expressions, many are constructions—that is, templates containing a combination of particular words and slots for variables. If LEIAs are to ever achieve human-level language capabilities, they need a very large, construction-packed lexicon that mirrors the lexical knowledge possessed by native speakers.

Unfortunately, the word *construction* comes with baggage, so we need to take a short detour to address terminology. Recently, *construction grammar* has become a popular sub-field of theoretical linguistics (Hoffman & Trousdale, 2013; see section 1.4.3.1). In keeping with its theoretical status, it attempts to account for certain aspects of human language processing—in particular, the form-to-meaning mapping of linguistic structures. We couldn't agree more that the latter is key to the study of language. And, as we have already

demonstrated, the LEIA's lexicon records exactly such form-to-meaning mappings. However, where we part ways with construction grammarians is in the specific interpretation of the word *construction*.¹⁶ For them, even non-argument-taking words are constructions in the sense that the syntactic form maps to the semantic interpretation. For us, by contrast, constructions must contain multiple constituents.

The constituents in a construction can be any combination of specific words and variable categories. Examples include:

- Idiomatic expressions: for example, *SUBJ kick the bucket*, *SUBJ spill the beans*
- *Verb + particle* collocations: *SUBJ give in* <*give out*, *give up*>
- Common expressions whose meanings are presumably remembered by people, not recomputed each time: *Have a nice day*. *Are we there yet?* *Salt, please*.
- Proverbs and sayings: *Nothing ventured, nothing gained*. *It takes one to know one*.
- Semantically non-compositional nominal compounds: *training data*, *trial balloon*
- Semantically constrained nominal compounds: FISH + *fishing* (i.e., any type of FISH followed by the word *fishing*) means a FISHING-EVENT whose THEME is that type of fish, for example, *trout fishing*
- Elliptical constructions: for example, NP V ADV_{COMPAR} *than* NP is recorded as a sense of *than*, which covers inputs like *Betty jumps higher than Lucy* _____. The lexical sense includes a procedural semantic routine that copies the meaning of V into the elided slot, effectively rendering our example, *Betty jumps higher than Lucy jumps*.
- All argument-taking words: adjectives, adverbs, verbs, prepositions, and so on.

The last class might come as a terminological surprise: Why would regular argument-taking words be considered constructions? Because the way they are recorded in the LEIA's lexicon, they fulfill our definition of construction: that is, they are multipart entities defined by a combination of required words and variable slots. Consider again the first verbal sense of *eat*, which was initially presented in section 2.2.

eat-v1

def.	to ingest
ex.	He was eating (cheese).
syn-struc	
subject	\$var1
v	\$var0
directobject	\$var2 (opt +)
sem-struc	
INGEST	
AGENT	^\$var1
THEME	^\$var2 (sem FOOD)

The verb *eat* is a required word in this construction, though it can appear in any inflectional form. Syntactically, it requires a subject and permits an optional direct object. Semantically, the subject must be a valid AGENT of INGEST, which the ontology indicates must be some type of ANIMAL. The direct object, if selected, must be some type of FOOD. Stated differently, this sense of *eat* is not treated as an isolated word mapped to the concept INGEST. Instead, it is described with its expected syntactic dependents, its expected set of case roles, and the semantic constraints on the fillers of those case roles. For this reason, it is a construction.

We further divide constructions into *lexical constructions* and *non-lexical constructions*. Lexical constructions must contain at least one specific word (which can, however, be used in different inflectional forms), which anchors the construction in the lexicon. Lexical constructions can contain any number of other required words and/or syntactic constituents.¹⁷ All of the examples above are of lexical constructions.

Non-lexical constructions, by contrast, contain only category types. For example, the syntactic construction called *object fronting* places the NP serving as the direct object in the sentence-initial position. This allows for *This I like!* to be used as an emphatic alternative to *I like this!* Similarly, the ontological construction FRUIT + TREE → TREE (HAS-OBJECT-AS-PART FRUIT) allows the nominal compound *apple tree* to be analyzed as TREE (HAS-OBJECT-AS-PART APPLE) (see section 6.3.1). Non-lexical constructions must be recorded in separate knowledge bases since they have no required word to serve as an anchor in the lexicon.

Lexical constructions, as we define them, are a supercategory of what descriptive linguists and practitioners of NLP call *multiword expressions*.¹⁸ Multiword expressions require multiple specific words, possibly with some variable slot(s) as well. For example, the idiomatic verb phrase *kick the bucket* is considered a multiword expression whose subject slot is a variable.

When LEIAs process constructions, their goal is the same as for any input: to compute the input's full contextual meaning. This has little to do with the most popular threads of work on multiword expressions over the past thirty years by descriptive linguists and practitioners of NLP. Descriptive linguists have primarily pursued classification, including analyzing the degree to which multiword expressions are fixed versus variable. Practitioners of NLP, for their part, have pursued the automatic detection of multiword expressions and their translation in machine translation systems. It is noteworthy that neither detecting nor translating multiword expressions directly addresses their meaning since even a correct translation achieved using statistical methods does not imply that the expression has been understood in a way that would support reasoning by intelligent agents.¹⁹

Let us return to lexical constructions as recorded in the LEIA's lexicon. The top-level distinction is between constructions that can be treated as lexemes with white spaces and those that cannot. The first category, which includes entities like *vice president*, *stock market*, and *nothing ventured, nothing gained*, is trivial. The components must occur in the listed order, they do not permit modifiers or other elements to intervene between them,

and only the last word, if any, is subject to inflection. We record such entities in the lexicon as multipart head words with an underscore indicating each white space. This approach is simple and works nearly perfectly—only *nearly* because in rare cases an expletive, speaker correction, or interruption might occur between the elements. (This can also, by the way, happen in the middle of regular words: *decon (ouch!) struction*.) As with all unexpected input, such deviations must be handled by recovery procedures, which amount to a sequence of attempts to relax certain constraints, such as the expectation that only a blank space can intervene between components of a multiword head entry.

Other constructions, as we have said, can have any combination of particular words and variable constituents. Consider the idiom *X pays homage to Y*. It is recorded in the lexicon as a sense of the verb *pay*.

pay-v16

def. The construction 'pay homage to', which maps to PRAISE

ex. They paid homage to the king.

syn-struc

subject \$var1
v \$var0
directobject \$var2 (root 'homage')

pp

prep \$var3 (root 'to')
obj \$var4

sem-struc

PRAISE

AGENT ^\$var1

THEME ^\$var4

^\$var2 null-sem+

^\$var3 null-sem+

Syntactically, this sense permits an NP of any legal form to fill the subject and object-of-preposition slots, but the direct object must be the noun *homage* and the preposition must be *to*. These constraints are indicated by appending the feature “root ‘homage’” to the direct object slot and the feature “root ‘to’” to the preposition slot. Semantically, *pay homage to* is interpreted using the ontological concept PRAISE. The subject of *pay homage to* fills the AGENT slot of PRAISE and the object of the preposition fills the THEME slot. The constraints on the AGENT and THEME are drawn from the ontological description of PRAISE.

Much more could be said about the expressive power of the metalanguage used to record lexical constructions, but those details are more appropriate for knowledge engineers than for the general reader. Suffice it to reiterate that (a) constructions can be composed of any sequence of lexical, syntactic, or ontological categories, which can be constrained by any of the morphological, syntactic, or semantic features used in the system, and (b) *lexical* constructions require that at least one particular word be fixed so that it can anchor the construction in the lexicon.

4.3.1 Semantically Null Components of Constructions

Some components of some constructions do not carry independent meaning, such as *the bucket* in the idiom *kick the bucket*. Such components are marked with the feature “null-sem+” in the sem-struc zone of the associated construction so that they are not analyzed compositionally.

Although null-semming might seem to be a perfect solution for getting rid of non-compositional elements, it has one complication. Occasionally, a null-semmed element is modified, and its meaning must somehow be attached to the meaning of the construction overall. For example, in (4.11) the modifier *goddamned* is not modifying *the bucket*; it is expressing the speaker’s frustration at the person’s having died before paying back the money.

- (4.11) My neighbor kicked the goddamned bucket before he paid me back the money he owed me!

There *is* a solution to this modification problem, which we will turn to once we have described typical uses of null-semming in the LEIA’s lexicon.

4.3.2 Typical Uses of Null-Semming

The most obvious use of null-semming is in canonical idioms like *kick the bucket*, but this mechanism has broader uses as well. For example, it can remove wordy reformulations of simpler meanings (as in examples (4.12)–(4.15) below), turns of phrase that primarily serve a discourse (rather than a semantic) function (4.16), and aspects of meaning that are so fine-grained or difficult to explain that they are, for the time being, not being chased down (4.17). In the examples, ^ indicates “the meaning of the given string” (we borrow this convention from the sem-struc zones of lexical senses).

- (4.12) ~~The fact that~~ the guy’s religious will have nothing to do with it. (COCA)
^((the guy’s religious) will have nothing to do with it)
- (4.13) But ~~the thing is, is that~~ it was relaxed. (COCA)
^(but it was relaxed)
- (4.14) ~~It’s just that~~ he’ll lose in a heartbeat. (COCA)
^(he’ll lose in a heartbeat)
- (4.15) And here we are in ~~the month of~~ April. (COCA)
^(And here we are in April)
- (4.16) ~~It turned out that~~ he had a small baggie of marijuana. (COCA)
^(he had a small baggie of marijuana)
- (4.17) She ~~couldn’t help but~~ laugh. (COCA)
^(She laughed)

Let us linger for a moment on the last example: What *is* the meaning of *She couldn't help but laugh*? Perhaps something like: “Irrespective of whether or not she wanted to laugh, she did so because, given some unspecified properties of the situation, it would have been too difficult for her not to laugh.” However, even if this were deemed a reasonable analysis, it is still only an English paraphrase, which is a far cry from a formal representation that could support useful automatic reasoning. The formal representation would be quite complex and it is not clear what goal it would serve for LEIAs in the foreseeable future. So spending time trying to specify it would be a poor use of limited resources.

4.3.3 Modification of Null-Semmed Constituents

As mentioned above, although null-semming can be a useful, elegant solution to handling some relatively superfluous text elements, it has its downside: in some cases, the null-semmed elements can be the target of modification. From the COCA corpus and other sources, we found quite a variety of examples:

- *Pay homage to* means PRAISE, but the nature of the homage can be specified as *silent*, *deliberate*, or *the strictest*.
- *Put a spell on* means BEWITCH, but that spell can be specified as *protective*, *love*, or *magic*.
- *Raise eyebrows* means SURPRISE, but “raised some serious eyebrows” appears multiple times on the internet.

If we null-sem the elements *homage*, *spell*, and *eyebrows* in these constructions, their modifiers will be left hanging. The solution is to anticipate such eventualities.

We will consider two cases by way of example: the previously mentioned *kick the god-damned bucket*, which illustrates proposition-level modification, and *put a ADJ spell on*, which illustrates modification of a meaningful element in a semi-idiomatic turn of phrase.

Kick the bucket is fully idiomatic since there is neither kicking nor a bucket involved. The only modifiers that are typically permitted in this expression are *proverbial* and expletives like *bloody* or *goddamned*. *Proverbial* can be included as an optional modifier in the *kick the bucket* sense with no semantic modifications needed. As for the expletives, they are best handled in a separate sense of the construction that requires an expletive adjective. The associated semantic description indicates that a very low value of evaluative modality (namely, .1 on the {0,1} scale), attributed to the speaker, scopes over the DIE event. In other words, the speaker who is reporting the event of dying is very unhappy about it.

One might ask, Isn't there a generalization that expletives can syntactically modify non-compositional elements of all idioms? And, if so, shouldn't the LEIA have a general rule to this effect, rather than having to rely on additional, expletive-inclusive senses of individual idioms? Perhaps, but there is a practical problem with implementing such a rule:

Where would it live in the system? Developing language-understanding capabilities for LEIAs is a practical endeavor. The system needs to work, the knowledge bases need to be inspectable, and knowledge engineers need to be able to trace what is happening and why. The program that uses lexical senses to build TMRs expects certain kinds of information in those senses and should not rely on fix-up rules with a high potential of getting lost in the code—even if one believes that such rules might exist in the minds of humans. Of course, specific decisions about practicality and engineering can change in different project configurations. One might invent an architecture that elegantly houses all kinds of fix-up rules. The approach just described for dealing with expletives in constructions reflects our current preference, which is based on real-world experience with organizing knowledge bases and engines that employ them.

Now we turn to *cast/put a spell on/over* as an example of a semi-idiomatic turn of phrase that can include modification. This construction means BEWITCH (remember, this is an ontological concept, not an English word). The idiomatic aspect of the construction is the verb choice—it must be *put* or *cast*. *Spell*, by contrast, is used in one of its canonical meanings. If we only ever expected the unmodified input *X cast a spell on Y*, then null-semming *a spell* would be a fast and sufficient solution since the ontology already tells us that the INSTRUMENT of BEWITCH is SPELL. However, since *spell* can, in fact, be modified, we choose to leave it as an explicit INSTRUMENT of BEWITCH in the sem-struct of the construction sense, as shown below.

cast-v2

def.	The construction 'cast/put a spell on someone'
ex.	The witch cast a spell on the frog.
comments	Not null-semming \$var2 means that it can be modified in the normal way: 'The witch cast an evil/magical spell on the frog.'
syn-struct	
subject	\$var1
root	\$var0
directobject	\$var2 (root 'spell')
pp	
prep	\$var3 (root 'on')
obj	\$var4
sem-struct	
BEWITCH	
AGENT	^\$var1
THEME	^\$var4
INSTRUMENT	^\$var2
^\$var3	null-sem+
synonyms	put

If *spell* is not modified, then this representation is simply redundant: both the lexicon and the ontology tell us that the INSTRUMENT is SPELL—not a problem. But if *spell* is modified, then the modification can be applied to ^\$var2 in the normal way.

Preparing for modification-oriented eventualities in constructions involves extra work, and that work may or may not be a priority at a given time in system development. This raises the question of what happens if modification is not accounted for in a construction sense but occurs in an input. In this case, overall processing will not fail, but the meaning of the modifier will be lost since the head that it modifies will be ignored.

One additional detail is worth noting. Analysis of the COCA corpus revealed that the constructions discussed above are overwhelmingly used in their base forms—not an unexpected outcome for essentially fixed expressions:

- The collocation *paid* followed by *homage* (within a distance of four words) yielded 138 hits, of which 123 were exactly *paid homage*, 11 inserted a modifier, and 4 included a pronominal indirect object (i.e., *paid him/her homage*).
- The collocation *kicked the* followed by *bucket* (within a distance of four words) yielded 34 hits, all of them exactly *kicked the bucket*.
- The collocation *put a* followed by *spell* (within a distance of three words) yielded 60 hits, 57 of which were exactly *put a spell*.

In conclusion, it is useful to be aware of the potential for modification of construction components and to prepare for them to the degree that resources permit. The fact that they are usually used in their base forms is informative when determining how to best allocate knowledge acquisition resources.

4.3.4 Utterance-Level Constructions

Some full utterances are recorded as senses in the lexicon. There are both theoretically and practically motivated reasons for this, which we consider in turn.

Theoretical motivations. The modeling of LEIAs is inspired by our hypotheses about human cognition. It seems likely that after we have encountered utterances some number of times, we simply remember their analyses: *When are we gonna be there? I'm hungry! Salt, please.* Such utterances do not have to be idiomatic. Any utterance that is frequent for a given person (possibly, in a given physical or social context) likely has a remembered analysis. For example, a running coach might begin every long-distance run with the elliptical reminder “Water?” meaning “Have you brought water to drink during the run?” The mapping between the word “Water?” and the full proposition is likely to be remembered, rather than recomputed, by the runners after the first practice or two. As concerns modeling LEIAs, it makes no difference whether these remembered analyses are recorded in the LEIA’s lexicon or in a separate repository of remembered text meaning representations (see section 6.1.6 for a discussion of the latter). Either way, the LEIA has the associated text-to-meaning mappings in its knowledge substrate.

Practical motivations. On the practical front, recording the meaning of full utterances can provide immediate language support for agent reasoning and action before all of the component linguistic phenomena can be handled using general-purpose methods. For

example, say a robotic agent is capable of visually locating, picking up, and passing to its human collaborator a hammer. A typical way to ask for this to happen is by saying, “Hammer!” So, the elliptical “Hammer!” must be understood as “Give me a hammer.”

If “Hammer!” were to be analyzed from first principles, the agent would need to use goal- and plan-based reasoning supported by domain knowledge (see chapter 7). Specifically, it would need to recognize that the object HAMMER is not a proposition, so there must be an implied EVENT. The agent must determine which hammer-oriented events it is capable of carrying out and, of those, which is the most relevant at the moment of speech. Depending on the domain, this might be complicated and error-prone. However, if “Hammer!” is typically used to mean “Give me a hammer,” then it merits having a lexical sense that records exactly this meaning when this word is used as an independent utterance.

```
hammer-n2
def.      'Hammer!' or 'Hammer.' used to mean 'Give me a hammer.'
ex.       'Hammer!' 'Hammer.'
comments  The punctuation mark indicates that this is an utterance, not a word
syn-struc
  root    $var0
  punct   $var1 (root *period* *exclamation-point*)
sem-struc
  REQUEST-ACTION
    AGENT      HUMAN-#1 ("speaker")
    THEME      TRANSFER-POSSESSION-#1
    BENEFICIARY HUMAN-#2 ("hearer")
  TRANSFER-POSSESSION-#1
    AGENT      HUMAN-#2
    THEME      HAMMER
    BENEFICIARY HUMAN-#1
  ^$var1    null-sem+
```

Different domains have analogous elliptical structures, as in this example about surgery.

(4.18) “Everyone ready? Let’s go. Knife.” The nurse pops the handle of the scalpel into my palm. (COCA)

Recording analyses for the elliptical uses of *Hammer*, *Knife*, and others is not only a convenient strategy, it also models the fact that, in certain application domains, members of human-robotic teams will likely encounter these elliptical utterances so often that they will simply remember the intended meanings, without having to dynamically resolve the ellipsis every time.

4.3.5 Additional Knowledge Representation Requirements

As microtheories go, the microtheory of constructions is quite advanced, meaning that it covers a lot of eventualities. However, language never ceases to surprise, and a recent

evaluation study (McShane, Nirenburg, & Beale, 2015) pointed out the need for more precise knowledge acquisition with respect to constructions along three lines.

1. Exclusion criteria can be needed to avoid false positive construction matches. For example, *X can tell that Y* means UNDERSTAND, as shown by (4.19). However, this idiom does not permit an indirect object. If there *is* an indirect object, then the idiom conveys emphasis, as shown by (4.20).

(4.19) I can tell that you like this film. (COCA)

(4.20) Well, I can tell you that the markets are on edge right now. (COCA)

Two different constructions, recorded in two different lexical senses, are needed to cover these different usages.

2. Constructions are more ambiguous than might be expected, which makes treating polysemy as much a priority for constructions as it is for simple lexemes. Pairs of literal and metaphorical meanings are particularly common, as illustrated by *take a bath* in (4.21) and (4.22) and *take a look at* in (4.23) and (4.24).

(4.21) I scrubbed the tub and took a bath. (COCA)

(4.22) Yes, they took a bath in the stock market, “but not as badly as some people,” Dot-tie says. (COCA)

(4.23) I took a look at his shoes: winterized Air Jordan Six-Rings, gleaming black. (COCA)

(4.24) So as part of that investigation we took a look at her finances ... (COCA)

3. In some cases, constructions overlap. For example, (4.25) uses the construction *let it go*, which means to stop unproductively attending to something. By contrast, (4.26) uses the construction *let it <something>go to hell (in a handbasket)*, which means to let something deteriorate by not attending to it.

(4.25) How can you be so judgmental? Life’s life. Let it go. (COCA)

(4.26) Doggone it, this is our—this is our community. And we’re not going to let it go to hell in a handbasket. (COCA)

The best strategy for handling this last eventuality is for the LEIA to select the largest textspan that matches a recorded construction. That is, if one construction completely subsumes another, then the longer one should be selected if all of its elements are attested in the input.

The reason for pointing out these complications is to emphasize that although constructions do represent relatively stable islands of confidence for language understanding, they are not without their challenges. This is important because, in reading the knowledge-lean NLP literature on constructions (or, more specifically, the subset that they call *multiword expressions*), one can get the impression that the main challenge is automatically detecting them. In reality, that is just the beginning.

4.3.6 Recap of Constructions

- Idioms: Moira has kicked the bucket. (COCA)
- Idioms with indirect modification: And you people tell me the Communists are running rampant in the outlying provinces and that if Mikaso kicks the damned bucket we could lose all ties to the Philippines ... (COCA)
- Wordy formulations: But ~~the thing is~~, is that it was relaxed. (COCA)
- Frequent locutions: Signature, please. (COCA)
- Application-specific frequent locutions: Knife. (COCA)
- Additional knowledge representation requirements involving (a) exclusion criteria, (b) ambiguity, and (c) overlapping constructions.

4.4 Indirect Speech Acts, Lexicalized

Often, questions and commands are expressed using indirect speech acts. An indirect speech act differs from a direct speech act in that its form (the locutionary act) does not align with its function (the illocutionary act). Starting from the basics, an alignment between form and function occurs when

- An assertion is used to make a statement: The cake is fantastic. (COCA)
- A question is used to ask a question: How is the game? (COCA)
- A command is used to issue a command: Get me a coffee. (COCA)

Various types of misalignments can occur for good pragmatic reasons, such as being polite or conveying emphasis, as shown by the following set of examples.

- (4.27) [An assertion used to request information]
 “I need to know what you’re talking about.” (COCA)
- (4.28) [An assertion used to request action]
 “... Quinn, it would be great if you’d bring your own card.” (COCA)
- (4.29) [An assertion used as a command]
 “But now you had better be off home.” (COCA)
- (4.30) [A question used to request action]
 “I have a lost kid here. Can you help us find his parents?” (COCA)
- (4.31) [A question used as an assertion. *Are you kidding me* expresses nonagreement and/or frustration]
 These people are the moral exemplars of the 21st century, are you kidding me?
 (COCA)

- (4.32) [A command used as an assertion. *Dream on!* indicates that the hope, plan, or wish just expressed is unrealistic, will not happen]
 Damn her, she knew exactly what he was trying to do, and would she help him?
Dream on! (COCA)
- (4.33) [A command used as a threat]
 “Come on, big boy,” he yelled. “Make my day!” (COCA)

A convenient aspect of indirect speech acts is that many of them are formulaic and can be recorded in the lexicon as constructions. For example, *I need to know X* and *You need to tell me X* are both typical ways of asking for information about X. By recording such constructions in the lexicon, we give LEIAs the same knowledge of their formulaic nature as a person has.

Of course, every formula that can serve as an indirect speech act also has a direct meaning. For example, you can say, *I need to know Mary’s address* to a person who you think can give you that information or to a person who couldn’t possibly know it by way of explaining why you are shuffling through your address book.

LEIAs treat speech-act ambiguity in the same way as any other type of ambiguity: by recognizing all available interpretations during Basic Semantic Analysis and then waiting until Situational Reasoning to select the contextually appropriate one.

Of course, not all indirect speech acts use well-known conventions. You can say, *The mail just came* with the implication that you want the interlocutor to go fetch it, despite the lack of any linguistic flag to indicate the veiled request. Indirect speech acts of this type are more difficult to detect, and configuring LEIAs to seek them out must be approached judiciously since we would not be well served by paranoid agents who assume that every utterance is a call to action.

4.5 Nominal Compounds, Lexicalized

A nominal compound (hereafter, NN) is a sequence of two or more nouns in which one modifies the other(s): e.g., *glass bank*. Although NNs are often subsumed under general discussions of constructions, they pose sufficiently idiosyncratic issues to merit a separate computational microtheory.²⁰

Most computational work on NNs has focused exclusively on establishing the implied relationship between the nouns without disambiguating the nouns themselves. However, the latter is actually more challenging. For example, *glass bank* can mean a coin storage unit made of glass, a slope made of glass, a storage unit for glass, a financial institution with a prominent architectural feature made of glass, and more. And even though some NNs might seem unambiguous at first reading—e.g., *pilot program* (feasibility study) and *home life* (private life, how one lives at home)—they actually have other available readings: *pilot program* could mean a program for the benefit of airplane pilots, and *home life*

could refer to the length of time that a dwelling is suitable to be lived in (by analogy with *battery life*).

A LEIA's analysis of NNs involves both the contextual disambiguation of the nouns and the establishment of the semantic relationship between them. We start by discussing the treatment of two-noun compounds for which the relevant senses of both nouns are available in the lexicon. Compounds with more than two nouns and/or unknown words are discussed at the end of the section.

As a prelude to describing LEIAs' approach to NN analysis, let us briefly consider best-case NN analysis results reported by others. Some examples are shown in table 4.3. Columns 3 and 4 juxtapose optimal results of LEIA processing with optimal results from three other paradigms: Tratz and Hovy (2010) (T); Rosario and Hearst (2001) (R); and Levi (1979) (L).

The points below explain why LEIA analysis is semantically more comprehensive—albeit much more expensive to operationalize—than the relation-selection approach undertaken by the others.

1. The LEIA analyses include disambiguation of the component nouns along with identification of the relation between them, whereas relation-selection approaches address only the relation itself.
2. The LEIA analyses are written in an unambiguous, ontologically grounded metalanguage (a reminder: strings in small caps are concepts, not English words), whereas the relation-selection approaches use ambiguous English words and phrases.

Table 4.3
Comparison of best-case analyses of NNs across paradigms

	Example	Full NN analysis by OntoAgent	Relation selection from an inventory
1	cooking pot	POT (INSTRUMENT-OF COOK)	perform/engage_in (T)
2	eye surgery	PERFORM-SURGERY (THEME EYE)	modify/process/change (T)
3	cat food	FOOD (THEME-OF INGEST (AGENT CAT))	consumer + consumed (T)
4	shrimp boat	BOAT (INSTRUMENT-OF CATCH-FISH (THEME SHRIMP))	obtain/access/seek (T)
5	plastic bag	BAG (MADE-OF PLASTIC)	substance/material/ingredient + whole (T)
6	court order	ORDER (AGENT LEGAL-COURT)	communicator of communication (T)
7	gene mutation	MUTATE (THEME GENE)	defect (R)
8	papilloma growth	CHANGE-EVENT (THEME PAPILLOMA) (PRECONDITION SIZE (< SIZE.EFFECT))	change (R)
9	headache onset	HEADACHE (PHASE begin)	beginning of activity (R)
10	pet spray	LIQUID-SPRAY (THEME-OF APPLY (BENEFICIARY PET))	for (L)

3. In 2, 7, 8, and 9 of table 4.3, the meaning of the “relation” in relation-selection approaches is actually not a relation at all but, rather, the meaning of the second noun or its hypernym: for example, *growth is-a change*. By contrast, since the LEIA’s treatment involves full analysis of all aspects of the compound, the meaning of each of the nouns is more naturally incorporated into the analysis.
4. The relation-selection approach can merge relations into supersets that are not independently motivated, such as (T)’s *obtain/access/seek*.²¹ For LEIAs, by contrast, every relation available in the independently developed ontology is available for use in compounding analysis—there is no predetermined list of compounding relations. This harks back to an early observation that practically any relationship could be expressed by a nominal compound (Finin, 1980).
5. Relation-selection approaches can include unbalanced sets of relations: for example, *consumer + consumed* (T) has been promoted to the status of a separate relation, but many other analogous correlations have not.
6. Relation-selection approaches occlude the semantic identity between paraphrases. By contrast, LEIAs generate the same meaning representation whether the input is *headache onset*, *the onset/beginning/start of a headache*, or *someone’s headache began/started*.

In sum, for LEIAs there is no isolated NN task that exists outside the overall semantic analysis of a text. LEIAs need to compute the full meaning of compounds along with the full meaning of everything else in the input, with the same set of challenges encountered at every turn. For example:

- *Processing the elided relations in NN compounds is similar to processing lexically underspecified ones.* In NNs (e.g., *physician meeting*), the relation holding between the nouns is elided and must be inferred. However, in paraphrases that contain a preposition (e.g., *meeting of physicians*), the preposition can be so polysemous that it provides little guidance for interpretation anyway. Both of these formulations require the same reasoning by LEIAs to determine the intended meaning—here, that there is a MEETING whose AGENT is a set of PHYSICIANS.
- *Unknown words are always possible.* Encountering out-of-lexicon words is a constant challenge for agents, and it can be addressed using the same types of learning processes in all cases (see chapter 8).
- *Many word combinations are lexically idiosyncratic.* Although past work has considered compounds like *drug trial* and *coffee cup* to be analyzable as the sum of their parts, they are actually semantically idiosyncratic: they represent specific elements of a person’s world model whose full meaning cannot be arrived at by compositional analysis. Ter Stal and van der Vet (1993) are correct that much more lexical listing is called for in treating compounds than the community at large acknowledges; here is a short excerpt from their discussion:

In natural science, (room temperature) means precisely 25 degrees Centigrade. A process able to infer this meaning would have to make deductions involving a concept for room, its more specific interpretation of room in a laboratory, and the subsequent standardisation that has led to the precise meaning given above. All these concepts play no role whatsoever in the rest of the system. That is a high price to pay for the capacity to infer the meaning of (room temperature) from the meanings of *room* and *temperature*. Thus, (room temperature) is lexicalized.

- *Analyses should not introduce semantic ellipsis.* The relation-selection method often introduces semantic ellipsis. For example, (T) analyze *tea cup* as *cup* with the *purpose* of *tea*. But objects cannot have purposes; only events can. So this analysis introduces ellipsis of the event *drink*. Similarly, *shrimp boat* is cited in (T) as an example of *obtain/access/seek*, but the boat is not doing any of this; it is the *instrument* of the *fisherman* who is doing this. If intelligent agents are to be equipped to reason about the world like people do, then they need to be furnished with nonelliptical analyses or else configured to dynamically recover the meaning of those ellipses.

In sum, viewing NN compounds within the context of broad-scale semantic analysis is a different task from what has been pursued to date in descriptive and NLP approaches. Let us turn now to how we are preparing LEIAs to undertake that task.

When a LEIA encounters an NN, it calls a confidence-ordered sequence of analysis functions. The first method that successfully analyzes the NN is accepted as long as it is contextually appropriate—a point we return to at the end of the section.²² There are two lexically supported strategies for analyzing NN compounds, and one default strategy for cases that are not covered by either of these.

The first lexicon-oriented strategy is to record the NN as a head entry: *abnormal psychology*, *gas pedal*, *finish line*, *coffee cup*. As just mentioned, this is actually necessary since a lot of compounds have incompletely predictable meanings. A *coffee cup* is not a cup that can only contain coffee or that does, at the moment, contain coffee; instead, it is a particular type of object that can contain water, coins, or a plant.

The second lexicon-oriented strategy is to use one of the words in a two-word NN as the key for a lexical construction whose other component is a variable. For example:

- One sense of the noun *fishing* expects an NN structure in which *fishing* is N2 and some type of FISH is N1. It covers NNs such as *trout fishing*, *bass fishing*, and *salmon fishing*, analyzing them as

FISHING-EVENT	
THEME	TROUT/BASS-FISH/SALMON

- In compounds of the structure *N detective*, if N is a kind of CRIMINAL-ACTIVITY, then the overall meaning is that the DETECTIVE is the AGENT-OF an INVESTIGATE event whose THEME is N. So, *homicide detective* is analyzed as

DETECTIVE		
	AGENT-OF	INVESTIGATE-#1
INVESTIGATE-#1		
	THEME	MURDER

If, by contrast, the input were *university detective*, then this construction would *not* match, and the NN would be passed on for other types of processing.

If a given NN is not covered by a construction recorded in the lexicon, or if it *is* covered but the associated analysis does not work within the overall semantic dependency structure of the clause, then the agent opts for an underspecified analysis. Specifically, it links all available meanings of N1 with all available meanings of N2 using the most generic relation, RELATION. These candidates are subjected to deeper analysis during Extended Semantic Analysis (section 6.3.1).

4.6 Metaphors, Lexicalized

Metaphors are a frequent occurrence in natural language, and they are essential to people's ability to understand abstract ideas.²³ As Lakoff and Johnson (1980, p. 3) explain, "Our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature." Bowdle and Gentner (2005, p. 193) concur: "A growing body of linguistic evidence further suggests that metaphors are important for communicating about, and perhaps even reasoning with, abstract concepts such as time and emotion. ... Indeed, studies of scientific writing support the notion that far from being mere rhetorical flourishes, metaphors are often used to invent, organize, and illuminate theoretical constructs." As regards the frequency of metaphors, the latter report: "In an analysis of television programs, Graesser et al. (1989) found that speakers used approximately one unique metaphor for every 25 words" (p. 193).

Since LEIAs treat some metaphors (conventional metaphors) during Basic Semantic Analysis and postpone others (novel metaphors) until Extended Semantic Analysis, the discussion of metaphors is divided between this chapter and chapter 6. But before we turn to LEIA-specific issues, let us begin with a short overview of past work on metaphor.

4.6.1 Past Work on Metaphor

Metaphor has been addressed from a broad variety of premises and in different contexts: in rhetoric since Aristotle, literary criticism (e.g., Skulsky, 1986), semiotics (e.g., Eco, 1979), a variety of schools in linguistics (e.g., Lakoff & Johnson, 1980; Steen, 2017), psychology

(e.g., Bowdle & Gentner, 2005), psycholinguistics (e.g., Glucksberg, 2003), philosophy (e.g., Bailer-Jones, 2009; Lepore & Stone, 2010), and neuroscience (e.g., Goldstein et al., 2012).²⁴ The distinction between conventional and novel metaphors has been firmly established in linguistics (e.g., Nunberg, 1987) and psychology (e.g., Gibbs, 1984). Bowdle and Gentner (2005) view the novel-to-conventional metaphor continuum in an etymological perspective and argue that metaphors conventionalize and diachronically lose their metaphoricity. Most metaphors discussed within the popular conceptual metaphor theory (e.g., Lakoff, 1993) are actually conventional and, therefore, presumably exist in a native speaker's lexicon. Even if the early AI approaches to metaphor do not state it overtly, their underlying motivation was to use metaphor processing as a means of bypassing the need for lexical and conceptual knowledge acquisition.

Theorists go beyond the novel/conventional distinction. Steen (2011) introduces a distinction between deliberate and nondeliberate metaphors: "A metaphor is deliberate when addressees must pay attention to the source domain as an independent conceptual domain (or space or category) that they are instructed to use to think about the target of the metaphor" (p. 84). But as he concedes, "the processes leading up to the product of metaphor comprehension ... are largely immaterial to the question of whether their product counts as a deliberate metaphor or not" (p. 85). This corroborates our position: To successfully process input containing conventional metaphors, the hearer does not need to realize that a metaphor is present. Conventional metaphor qua metaphor may be of interest to scholars or as the subject of an entertaining etymological parlor game. But to understand *ballpark* in *ballpark figure*, it is not necessary to know that it is a baseball metaphor.

We hypothesize that people usually process novel and deliberate metaphors in the same manner in which they process unknown lexical units that are not metaphorical—by learning their meanings over time from their use in text and dialog and recording those meanings in their lexicons for later use. In other words, the novel (nonmetaphorical) senses of *pocket* and *bank* in *He pocketed the ball by banking it off two rails* will be learned with the help of knowledge of the domain (billiards) and general knowledge of what can typically be done with a billiard ball. By the same token, the meaning of *albatross* in "*But O'Malley's heaviest albatross is the state of his state*"²⁵ will also be understood based on the hearer's knowledge of the overall context, with no need for the hearer to have read, or even know about the existence of, Coleridge's *The Rime of the Ancient Mariner*.

Of course, building an agent that models an etymologist is a potentially interesting research direction, but it is much more important in agent systems to cover conventional metaphors. We believe that the best way to do this is to view the task as a routine part of the lifelong enhancement of an agent's knowledge resources. An agent will fail to register the aesthetic contribution of an extended metaphor like the following, but this is equally true about many people—after all, not everybody knows about baseball: (A team leader cajoling a team member) *Eric, the bases are loaded; tomorrow's demo is crucial. Please*

stop grandstanding and playing hardball, step up to the plate, join the effort, and lead off with a ballpark figure.

We discussed metaphorical language in a paper entitled *Slashing Metaphor with Occam's Razor* (Nirenburg & McShane, 2016b). Initially, some readers might not have fully understood the meaning of the title. However, most everyone will have guessed that we intended to say something negative about the study of metaphor. Some readers will also have understood that we would justify this attitude on the grounds that the study of metaphor is unnecessary from some point of view. Having read on, readers who still remember the title would realize what it intended to convey: that separating metaphor detection and interpretation from the treatment of other types of figurative language and other semantic anomalies violates the dictum *entities must not be multiplied beyond necessity*.

Now, readers (such as LEIAs) with no training in philosophy may have recognized Occam as a named entity without realizing that *Occam's razor* refers to the above dictum. Such readers would fully understand this paper's title only after having read the previous sentence. The above observations further motivate our contention that delayed interpretation of input is a viable and potentially effort-saving strategy for agents. Some readers will also appreciate the double entendre in the paper's title: the metaphorical use of an action (*slashing*) associated with the physical tool (*razor*) that once served as the source of the metaphor to describe the mental tool (*Occam's razor*). While recognizing this may be a nice bonus, it is not essential for understanding the main argument of the paper. This observation illustrates and motivates our contention that agents can often function well without understanding all of an input. Anybody who has ever communicated in a foreign language can vouch for this. Sometimes incomplete understanding can lead to misunderstandings or embarrassment but, more often than not, it works well enough to achieve success in communication. Of course, the \$64,000 question is how to teach LEIAs to determine what, if any, parts of an input they can disregard with impunity. This is a direction of ongoing work for our team.

4.6.2 Conventional Metaphors

Conventional metaphors are represented as regular word senses in the LEIA's lexicon. Lakoff and Johnson (1980) propose an inventory of conventional-metaphor templates, which are associated with a large number of linguistic realizations that can be recorded in the lexicon as senses of words and phrases. Below are examples of the Ontological Semantic treatments of some of Lakoff and Johnson's templates. Note that one of the challenges in meaning representation is that some of these meanings are vague—which is a knowledge representation challenge not specific to metaphorical language. Our goal here, as in all knowledge engineering, is to provide the agent with an analysis that will support its reasoning about action.

Template: Argument is war

someone attacks someone

CRITICIZE-1	
AGENT	HUMAN-1
THEME	HUMAN-2
INTENSITY	> .8

an idea is on target

IDEA-1	
ACCURACY	1

someone shoots down a proposal

ACCEPT-1	
AGENT	HUMAN-1
THEME	PROPOSAL-1
SCOPE-OF	MODALITY-1
MODALITY-1	
TYPE	epistemic
VALUE	0 ; indicates negation
SCOPE	ACCEPT-1

Template: Time is money

an activity cost someone an hour (this has a negative implication, in contrast to *took* someone an hour)

EVENT-1	
AGENT	HUMAN-1
DURATION	1 (MEASURED-IN HOUR)
SCOPE-OF	MODALITY-1
MODALITY-1	
TYPE	evaluative
VALUE	< .5
SCOPE	EVENT-1

Template: Happy is up, sad is down

someone is in high spirits

HUMAN-1	
HAPPINESS	.8

someone falls into a depression

CHANGE-EVENT-1	
EXPERIENCER	HUMAN-1
PRECONDITION	HAPPINESS-1
EFFECT	HAPPINESS-2

HAPPINESS-1		
DOMAIN		HUMAN-1
RANGE		> .5
HAPPINESS-2		
DOMAIN		HUMAN-1
RANGE		< .2

Template: The mind is a machine

someone grinds out a solution to a problem (analyzed as ‘someone thinks very hard and succeeds in solving a problem’)

SOLVE-1		
AGENT		HUMAN-1
THEME		PROBLEM-1
PRECONDITION		THINK-1
SCOPE-OF		MODALITY-1
THINK-1		
AGENT		HUMAN-1
SCOPE-OF		MODALITY-2
MODALITY-1		
TYPE		epiteuctic
VALUE		1
SCOPE		SOLVE-1
MODALITY-2		
TYPE		effort
VALUE		1
SCOPE		THINK-1

someone runs out of steam

ANIMAL-1		
EXPERIENCER-OF		FATIGUE-1
FATIGUE-1		
INTENSITY		> .8

In sum, conventional metaphors pose the exact same inventory of meaning representation challenges and opportunities as any other lexemes or phrases that happen to not have what linguists would consider a historical source in metaphor.

4.6.3 Copular Metaphors

Many creative (not conventionalized) metaphors are of the form NP1 *is* NP2. Lakoff and Johnson (1980, p. 139) say that creative metaphors “are capable of giving us a new understanding of our experience.” Their example: *Love is a collaborative work of art* (p. 139).

We describe this class using the syntactic term *copular metaphors* because their defining feature is, in fact, syntactic: they use the copular verb *be*.

We prepare LEIAs to analyze copular metaphors using a lexical sense of the verb *be* that

1. expects the syntactic structure NP1 *is* NP2;
2. puts no semantic constraints on the meanings of the NPs;
3. links the meanings of the NPs using the almost vacuous concept JUXTAPOSITION; and
4. includes a meaning procedure that, if called during Extended Semantic Analysis, will attempt to identify the most salient properties of the second NP and apply them to the first NP (see section 6.3.3).

Note that this sense of *be* is just one of many senses of this verb that links two noun phrases. It is the least-specified one semantically, meaning that others will be preferred if their semantic constraints are satisfied. For example, *John is a doctor* will be handled by a sense of *be* that requires NP2 to indicate a SOCIAL-ROLE. That sense will be used to analyze this input as follows:

HUMAN-1

HAS-PERSONAL-NAME	'John'
HAS-SOCIAL-ROLE	PHYSICIAN

The sense used for novel metaphors, which has no semantic constraints, will be used as a fallback when more narrowly constrained senses like this one are not applicable.

4.6.4 Recap of Metaphors

- Conventionalized metaphors are recorded as lexical senses.
- Nonconventionalized metaphors in copular constructions (*Life is a garden*) are recognized as potential metaphors, analyzed using the property JUXTAPOSITION, and flagged with a procedural semantic routine that will be run during Extended Semantic Analysis.
- Metaphorical meanings not treated by these methods will often give rise to a low-scoring TMR, indicating that something is wrong. That “something” will be explored during Extended Semantic Analysis.

4.7 Metonymies, Lexicalized

Like metaphors, many word senses that can historically be analyzed as metonymies are, in synchronic terms, regular lexical senses that are recorded in the LEIA's lexicon, such as *get a pink slip* (get fired) and *red tape* (excessive bureaucratic requirements). There are

also ontological classes of metonymies: for example, a piece of clothing can be used to refer to the person wearing it (*Give the red shirt a glass of milk*). The latter are not treated at this stage. Instead, semantic analysis results in an incongruity, reflected by a low score for the TMR (there is no sense of *give* that expects a SHIRT as the BENEFICIARY). This low score is a flag to track down the source of the incongruity during Extended Semantic Analysis (section 6.2).

4.8 Ellipsis

Of all the topics treated in this chapter on Basic Semantic Analysis, ellipsis is likely to be the most surprising. After all, ellipsis is not only an aspect of reference resolution—which is a prime example of a *pragmatic* phenomenon—it is a particularly difficult aspect of it. However, if one dispenses with preconceived notions about linguistic modularity and, instead, considers the twin challenges of detecting and resolving ellipsis, it turns out that much of the work can be neatly subsumed under Basic Semantic Analysis. In fact, in some cases *all* of the work can be, as the subsections below will explain.

4.8.1 Verb Phrase Ellipsis

Modal and aspectual verbs can be used either with or without an overt complement, as shown by the juxtaposed (a) and (b) versions of (4.34) and (4.35).

- (4.34) a. “You have to get up.”
 b. “I know you don’t feel like getting up, but you have to ____.” (COCA)
- (4.35) a. “I just started playing.”
 b. “They’ve been playing at least five years, maybe three years, and I just started ____.”
 (COCA)

Structures whose verbal complements are not overt are said to contain verb phrase (VP) ellipsis.

We prepare LEIAs to detect and resolve VP ellipsis by recording lexical senses of modal and aspectual verbs in pairs. Whereas one member of the pair expects an overt VP, the other expects VP ellipsis. The elliptical senses posit a placeholder *EVENT* along with a meaning procedure indicating that it requires downstream coreference resolution (section 5.5). In other words, an input like “I just started” will be analyzed as “I just started *EVENT*”, and the nature of the *EVENT* will be flagged as needing to be tracked down later.

By explicitly preparing for VP ellipsis during lexical acquisition, we solve two problems at once. First, we account for the fact that we, as people, do understand—even outside of context—that sentences with VP ellipsis imply some event; so, too, should LEIAs. Second, we do not need a special process for detecting VP ellipsis. When an elliptical input is processed, it simply uses the lexical sense that expects the ellipsis.

Consider the example *John washed his car yesterday but Jane didn't ____*. Its basic TMR—shown in two parts below for readability's sake—indicates that, prior to forthcoming coreference procedures, all the LEIA knows about Jane from this utterance is that she didn't *do something* (EVENT-1).

John washed his car yesterday but

WASH-1		
AGENT	HUMAN-1	
THEME	AUTOMOBILE-1	
TIME	(<i>combine-time</i> -1 day)	; "yesterday"
HUMAN-1		
HAS-PERSONAL-NAME	'John'	
GENDER	male	
AGENT-OF	WASH-1	
AUTOMOBILE-1		
RELATION	HUMAN-1	; "his"
THEME-OF	WASH-1	
CONTRAST-1		
DOMAIN	WASH-1	; "but"
RANGE	EVENT-1	

Jane didn't

EVENT-1		
CASE-ROLE	HUMAN-2	
SCOPE-OF	MODALITY-1	
COREF	<i>seek-specification</i>	
HUMAN-2		
HAS-PERSONAL-NAME	'Jane'	
GENDER	female	
AGENT-OF	EVENT-1	
MODALITY-1		
TYPE	epistemic	
VALUE	0	; negation
SCOPE	EVENT-1	

To reiterate, the point of this example is to show that during Basic Semantic Analysis the agent detects the VP ellipsis in “Jane didn’t ____” and provisionally resolves it as an unspecified EVENT whose precise meaning will be sought during the next stage of processing, Basic Coreference Resolution.

4.8.2 Verb Phrase Ellipsis Constructions

Identifying the sponsor for an elided VP can be quite difficult, which is why, in the general case, it is postponed until the dedicated Basic Coreference Resolution module (see

section 5.5). However, there exist elliptical constructions in which identifying the elided meaning is quite straightforward. We record these in the lexicon, and the associated lexical senses allow for the ellipsis to be fully resolved right now, during Basic Semantic Analysis. The example in (4.36) shows a construction that indicates that the agent applies maximum effort to carrying out the action.

- (4.36) [Subj_i V as ADV as Pronoun_i can/could]
Agatha wrote back as fast as she could _____. (COCA)

In this example, the modal verb *could* is used without its VP complement, making it an elliptical structure. But the ellipsis is resolved by copying the meaning of the same verb that this expression modifies: that is, Agatha wrote back as fast as she could *write back*. This means that no discourse-level reasoning is needed to resolve the ellipsis; the answer lies in the construction itself.

Deciding how many ellipsis-oriented constructions to record, and how to balance literal and variable elements in them, represents a microcosm of knowledge engineering overall. For example, the input *Boris gives his children as many gifts as he wants to* can be covered by either of the constructions shown in table 4.4. However, the second, more generic one also covers examples like *Boris takes as few as he can*.

In general, the more narrowly defined the construction, the more likely it will give rise to a unique and correct semantic analysis. But recording constructions takes time, and more narrowly specified constructions offer less coverage. So knowledge acquirers must find the sweet spot between the generic and the specific.

Below are some VP ellipsis constructions, presented using an informal notation, along with examples that were automatically extracted from the COCA corpus. The elliptical gaps are indicated by underscores, and the antecedents for the elided categories are underlined.

- (4.37) [V <all, everything, everything and anything, anything and everything, whatever> Pronoun (ADV) AUX]
a. My biggest focus right now is just learning all I can _____. (COCA)
b. The government gobbled up whatever it could _____. (COCA)
- (4.38) [VP as ADV as Pronoun AUX]
Yeah, he wanted me to come pick him up as quickly as I could _____. (COCA)

Table 4.4.
VP ellipsis constructions

Boris	gives	his children	as many	gifts	as	he	wants to
Subj	V	(NP)	as many	N	as	Pronoun	wants to
Subj	V	(NP)	QUANT	(N)	as	Pronoun	MODAL
Boris	takes		as few		as	he	can

- (4.39) [VP as ADV/ADJ as Pronoun AUX]
 a. He said, ‘It’s going to hit me, so I’m going to enjoy life as best as I can ____.’ (COCA)
 b. On this picture, I feel like I got my way as much as I could _____. (COCA)

To conclude this section, certain cases of VP ellipsis can be fully treated—both detected and resolved—during Basic Semantic Analysis thanks to constructions recorded in the lexicon.

4.8.3 Event Ellipsis: Aspectual + NPOBJECT

Another type of event ellipsis occurs when aspectual verbs take an NP complement that refers to an OBJECT rather than an EVENT. Such clauses must involve ellipsis because aspectual meanings can only ever apply to events. Consider the following pair of examples.

- (4.40) He boldly went up to her and started a conversation ... (COCA)
 (4.41) He wrote and directed plays, started ____ a book about the Yiddish language with his grandfather ... (COCA)

Example (4.40) refers to starting a conversation, which is an EVENT, so there is no ellipsis. By contrast, (4.41) refers to starting a book, which is an OBJECT; what is meant is that he started *writing* a book.

We prepare LEIAs to treat *aspectual + NPOBJECT* cases the same way as VP ellipsis: by creating a lexical sense of each aspectual verb that expects its complement to refer to an OBJECT. The semantic description of such constructions includes an underspecified EVENT whose THEME is that OBJECT. So, his starting a book in (4.41) will be analyzed as his being the AGENT of an EVENT (scoped over by “PHASE begin”) whose THEME is BOOK. This is all the text says—and that is the very definition of Basic Semantic Analysis. The rest requires nonlinguistic reasoning that the agent will pursue, if it chooses to, during Extended Semantic Analysis and/or Situational Reasoning (chapters 6 and 7).

4.8.4 Event Ellipsis: Lexically Idiosyncratic

Certain words, when used in certain constructions, always imply a particular kind of EVENT. For example, when someone invites someone else to some place, there is always an implied MOTION-EVENT.

- (4.42) Marino brought the young artist into his sphere, secured several commissions for him, and eventually invited him to Rome. (COCA)

The TMR for the clause *he invited him to Rome* is as follows.

INVITE-1	
AGENT	HUMAN-1
THEME	MOTION-EVENT-1
BENEFICIARY	HUMAN-2
TIME	<find-anchor-time

MOTION-EVENT-1		
AGENT		HUMAN-2
DESTINATION		CITY-1
HUMAN-1		
GENDER		male
HUMAN-2		
GENDER		male
CITY-1		
HAS-NAME		'Rome'

The lexical sense of *invite* that covers this construction explicitly lists the MOTION-EVENT, which explains how it turns up in the TMR. Similarly, when one *forgets* some PHYSICAL-OBJECT that can be carried (i.e., it is an ontologically licensed THEME of a CARRY event), the elided event is, by default, TAKE. There is a lexical sense of *forget* that anticipates, and resolves, this ellipsis.

As discussed earlier with respect to VP ellipsis constructions, the challenge in writing such lexical senses is determining the sweet spot for coverage versus precision. Let us continue with the example of forgetting an OBJECT. Above, we suggested that the object in question must be able to be carried in order for the elided verb to be understood as CARRY. This works for forgetting one's keys, notebook, lunch, and many other objects. It does not cover forgetting one's car or one's file cabinet. But what does it mean to forget one's car or one's file cabinet? It is impossible to say without context since there is no high-confidence default interpretation. Forgetting one's car might mean forgetting to move it to the other side of the street according to alternate-side-of-the-street parking rules, or it might mean forgetting to drive it to school, rather than ride one's bike, in order to help transport something after class. Forgetting one's file cabinet might mean forgetting to look there for a lost object or forgetting to have it transported when moving from one office to another. Since we know that nonholdable objects can occur as the direct object of *forget*, we need to write another lexical sense that anticipates them. This sense, like the aspectual senses discussed earlier, will initially underspecify the event—listing it as simply EVENT—and call a procedural semantic routine that will later attempt to reason more precisely about what it might mean.

4.8.5 Event Ellipsis: Conditions of Change

Events and states cause other events and states.

- (4.43) [An event causes an event; 'rain(s)' is analyzed as RAIN-EVENT]
Heavy rains caused the river and its tributaries to flood ... (COCA)

- (4.44) [A state causes an event]
"There's room for two," Sophie called out. Her excitement made Mr. Hannon laugh. (COCA)

- (4.45) [An event causes a state]
... The disappearance of the valuables made people nervous ... (COCA)
- (4.46) [A state causes a state]
The deaths and the publicity about the state’s raging rivers have taken a toll on commercial rafters’ business. The conditions have made people nervous ... (COCA)

Events and states cannot be caused by objects. However, language permits us to express situations as if objects could cause events and states.

- (4.47) [An object is said to cause an event]
Investigators want to closely inspect the engines to figure out how exactly the birds caused the plane to fail so badly and so fast. (COCA)
- (4.48) [An object is said to cause a state]
And I knew ice cream was something that made people happy. (COCA)

In such cases, the named object participates in an event or state that is the actual cause of another event or state. In our examples, the accident happened because of something the birds did, and ice cream plays some role in an event (eating it) that makes people happy.

We prepare LEIAs to detect elided events in causal clauses using special lexical senses of words and phrases that express causation—e.g., *cause (sth.)*, *make (sth.) happen*, *bring (sth.) about*. The given senses expect the named cause to be, ontologically speaking, an OBJECT, and they explicitly posit an EVENT for which that OBJECT serves as a case role. For example, the sense of *make* that expects the construction *NP_{OBJECT} makes NP V* has a semantic representation that will generate the following TMR for *The onion made her cry*.

CRY-1	
EXPERIENCER	HUMAN-1
CAUSED-BY	EVENT-1
TIME	<find-anchor-time
EVENT-1	
INSTRUMENT	ONION-1
COREFER	seek-specification
HUMAN-1	
GENDER	female

This analysis posits an EVENT without specifying its nature.²⁶ The fact that the event is under-specified is reflected in the TMR by the call to the meaning procedure *seek-specification*. If, later on, the LEIA has reason to believe that the nature of this event is important, it can attempt to track it down—though it will be successful only if it has a sufficient amount of domain and situational knowledge to support the analysis (see chapter 7).

4.8.6 Gapping

Gapping is a type of verbal ellipsis that occurs in structurally parallel coordinate and comparative structures. The following examples illustrate two common gapping constructions, presented informally.

- (4.49) [Subj₁ V DO₁(,) and Subj₂, DO₂]
Of course, thoughts influence actions, and actions, thoughts. (COCA)
- (4.50) [Subj₁ V IO₁ DO₁(,) and Subj₂, DO₂]
The plumber charged us \$200, and the electrician, \$650.

Although gapping is not used all that commonly in English, it makes sense to cover the most frequent eventualities using constructions like those illustrated above, which are anchored in the lexicon on the keyword *and*. These senses are supplied with meaning procedures that can be run right away, during Basic Semantic Analysis. They semantically analyze the overt verb and then reconstruct the elided one as a different instance of the same ontological type. This means that gapping can be detected and fully resolved at this stage of processing for any input that matches a recorded gapping construction.

4.8.7 Head Noun Ellipsis

In English, ellipsis of the head noun in noun phrases is permitted when the head noun follows a number (4.51), follows a quantifier (4.52 and 4.53), or participates in constructions like *someone's own* (4.54).

- (4.51) He had a number of offspring but only two __ were considered worthy contenders for the ducal crown. (COCA)
- (4.52) “Tea, Mr. Smith?” Tracy asked. “Yes, I’d love some __.” (COCA)
- (4.53) Who are the people who worked for him? Unless he didn’t have any __? (COCA)
- (4.54) The voice was not my own __. (COCA)

The lexicon includes a special sense of each applicable word and phrase that anticipates head noun ellipsis. These senses include a call to a procedural semantic routine that guides the LEIA in resolving the ellipsis by attempting to identify the most recent mention of an entity that matches the selectional constraints of the clause’s main verb.

4.8.8 Recap of Ellipsis

- Verb phrase (VP) ellipsis: Our fondness for sweetness was designed for an ancestral environment in which fruit existed but candy didn’t __. (COCA)
- VP ellipsis constructions: She hides her true identity as long as she can __. (COCA)
- Event ellipsis—Aspectual + NP_{OBJECT}: “Started __ the book yet?” (COCA)

- Event ellipsis—Lexically idiosyncratic: She made friends with a French girl, who invited her ___ to Paris. (COCA)
- Event ellipsis—Conditions of change: The acid in vinegar caused the iron in the steel to combine rapidly with oxygen from the air. (COCA)
- Gapping: Lou likes Coke, and Sherry ___, Pepsi.
- Head noun ellipsis: He was good, getting up to eight skips. At best Annabel got three ___. (COCA)

4.9 Fragmentary Utterances

We define *fragmentary utterances* as nonpropositional, freestanding utterances—in contrast to the midsentence fragments that are analyzed as a matter of course during incremental NLU. Examples of fragmentary utterances are “Large latte” and “Not yet.” During Basic Semantic Analysis, all of the available analyses of such utterances are posited as candidates. Selecting the intended meaning and incorporating it into the larger context is the shared responsibility of Extended Semantic Analysis (section 6.4) and, if needed, Situational Reasoning (chapter 7).

4.10 Nonselection of Optional Direct Objects

Some verbs, such as *read* and *paint*, are optionally transitive. This means that they *can* select a direct object but do not require one. Nonselection of a direct object is not ellipsis but it occasionally presents an interesting problem: the unexpressed object can be needed to interpret a subsequent referring expression, as in (4.55).

(4.55) They won’t be doing any hiring this year apart from replacing those ___ who leave.

The elided noun in the noun phrase [_{NP} those ___ who leave] is *employees*. This concept was implicitly introduced into the context as the filler for the THEME of HIRE. Engineering details aside, the point is this: During Basic Semantic Analysis, if the direct object of an optionally transitive verb is unselected (i.e., not explicit in the input), a special slot for it is created in the TMR. That slot is filled by the generic constraint found in the ontology (here: HIRE (THEME EMPLOYEE)) and is available for later coreference as needed. This is an excellent example of why it is risky to split language processing tasks finely across systems and developers. If one does so, then phenomena like this will more than likely fall between the cracks as nobody’s responsibility.

4.11 Unknown Words

As mentioned earlier, the LEIA’s lexicon contains about 30,000 word senses, making it substantial but far from comprehensive. This means that LEIAs must be able to process

both unknown words and unknown senses of known words. We already described the first stage of treating unknown words: during Pre-Semantic Integration, LEIAs posit syntactically specific, but semantically underspecified, lexical senses for unknown words. Now, during Basic Semantic Analysis, they try to narrow down the meaning of the newly learned word senses using ontological search. The specifics of the process vary depending on (a) the part of speech of the newly learned word sense and (b) whether what is being learned is a completely new string or a known string in a different part of speech. In the latter case, the meaning being learned might, though need not, be related to the known meaning.

4.11.1 Completely Unknown Words

New-word learning focuses on open-class words—currently nouns, adjectives, and verbs, which we consider in turn.

Unknown nouns. Syntactically, simple nouns take no arguments. Semantically, they can refer to an OBJECT, EVENT, or PROPERTY. During Pre-Semantic Integration, the LEIA generates three candidate senses for each unknown noun, one for each of these semantic mappings. Each candidate sense is then evaluated at this stage with two goals: (a) to choose the best of these mappings for the context and (b) if possible, to narrow down that interpretation to a more specific concept in the given branch of the ontology. Consider example (4.56), which contains the unknown noun *hobo*.

(4.56) A hobo came to the door. _(COCA)

The sense of *come* that best fits the context has the syntactic structure “Subject V PP” and the semantic analysis “COME (AGENT ^Subject) (DESTINATION ^PP-obj)”. Since the meaning of the subject must fill the AGENT slot, the best interpretation of *hobo* is OBJECT (rather than EVENT or PROPERTY). However, based on the fact that this OBJECT has to fill the AGENT slot of COME, the agent can narrow it down to ANIMAL, since that is the *sem* filler of this property recorded in the ontology. This results in the following TMR for *A hobo came to the door*.

```
COME-1
AGENT          ANIMAL-1
DESTINATION    DOOR-1
TIME           <find-anchor-time

ANIMAL-1
from-sense     new-noun-object-n1
```

The metadata in this TMR (i.e., the *from-sense* slot) carries a trace that unknown-word processing was carried out, should the LEIA decide to pursue a more fine-grained analysis of this word through learning by reading (Nirenburg et al., 2007) or by interacting with a human collaborator. Now consider example (4.57), in which the only unknown word is *tripe*.

(4.57) Jane was eating tripe with a knife.

The verb *eat* has several senses in our lexicon, all but one of which cover idiomatic constructions that are rejected on lexico-syntactic grounds (e.g., *eat away at*). So the LEIA can immediately narrow down the choice space to the main sense of *eat* (eat-v1), which is optionally transitive. It maps to an INGEST event whose case roles are AGENT and THEME. The THEME is specified as FOOD, which is more constrained than the ontologically listed disjunctive set [FOOD, BEVERAGE, INGESTIBLE-MEDICATION] (i.e., one does not eat a beverage or medication). This leads to the following TMR for *Jane was eating tripe with a knife*.

ASPECT-1		
PHASE		continue
SCOPE		INGEST-1
INGEST-1		
AGENT		HUMAN-1
INSTRUMENT		KNIFE-1
THEME		FOOD-1
SCOPE-OF		ASPECT-1
TIME		<find-anchor-time
HUMAN-1		
AGENT-OF		INGEST-1
HAS-PERSONAL-NAME		'Jane'
GENDER		female
FOOD-1		
THEME-OF		INGEST-1
from-sense		new-noun-object-n1

Unknown adjectives. Syntactically, adjectives modify nouns. Semantically, they map to a PROPERTY, and the meaning of the noun they modify fills the DOMAIN slot of that property. The RANGE, however, depends on the meaning of the adjective itself. For example, the subject noun phrase in (4.58) includes the unknown adjective *confrontational*.

(4.58) A confrontational security guard was yelling.

The LEIA's analysis of *confrontational security guard* is

SECURITY-GUARD-1		
DOMAIN-OF	PROPERTY-1	

PROPERTY-1 is an instance of the most underspecified ontological property (the root of the PROPERTY subtree). If asked to, the LEIA can create the entire set of properties for which SECURITY-GUARD is a semantically acceptable filler for DOMAIN. This would narrow the

interpretation from ‘any property’ to ‘one of a listed set of properties’. However, in many cases—like this one—this set will still be too large to be of much more utility than the generic PROPERTY interpretation.

Unknown verbs. Verbs can take various numbers and types of arguments and complements, which can realize various semantic relations. We will use a transitive verb as our example. Transitive verbs most often use the subject to express the AGENT and the direct object to express the THEME. If the subject cannot semantically fill the AGENT role, then the next case role in line is INSTRUMENT. We can see how these case role preferences play out in example (4.59), where the unknown verb is *nicked*.

(4.59) The truck nicked the tree.

The agent already mapped *nicked* to EVENT during Pre-Semantic Integration. Now it tries to determine the case roles of its arguments using the abovementioned preferences. *The truck* cannot be the AGENT since it is inanimate, but it can be the INSTRUMENT. *The tree*, for its part, can be a THEME. This results in the following TMR:

EVENT-1	
THEME	TREE-1
INSTRUMENT	TRUCK-1
COREF	<i>seek-specification</i>

The question is, can the agent usefully constrain the interpretation of EVENT on the basis of what the ontology says about this combination of case roles and fillers? Not too much, so the EVENT remains underspecified at this stage and can be analyzed more deeply, if the agent chooses to, during Situational Reasoning.

4.11.2 Known Words in a Different Part of Speech

It is not unusual for the lexicon to contain a needed word but not in the part of speech needed for the input. For example, it might have the noun *heat* but not the verb *to heat*. The first thing to say about such situations is that there are many eventualities:

- The lexicon might contain exactly one sense, which luckily is semantically related to the new-part-of-speech sense.
- The lexicon might contain multiple senses, exactly one of which is related to the new-part-of-speech sense.
- The lexicon might contain any number of senses, none of which is related to the new-part-of-speech sense.

This complexity is just one manifestation of the open-world problem. No matter how many words and phrases an agent knows, an input might contain a new one.

To keep useful bounds on this discussion, we will work through one example,

(4.60) A large radiator was heating the room.

with a number of simplifying assumptions:

1. The lexicon contains exactly one nominal sense of *heat*, which refers to temperature.
2. That sense is the one needed to learn the meaning of the verb *to heat*.
3. All the other words in the sentence have just one sense in the lexicon.

If (1) or (3) did not hold (i.e., if there was lexical ambiguity), then the process would iterate over all possibilities and generate multiple candidate analyses. This is not a problem; it is just inconvenient to present. If (2) did not hold, the agent would get the answer wrong and would suspect it only if all TMR candidates got low confidence scores, suggesting some problem in combining semantic heads with their arguments. What is interesting about this example is that the noun *heat* maps not to an OBJECT or EVENT but to an ontological PROPERTY—namely, TEMPERATURE.

Consider the TMR for *A large radiator was heating the room*, which the agent generates using the analysis process described below.

CHANGE-EVENT-1

THEME	ROOM-1
INSTRUMENT	HEATER-1
PRECONDITION	TEMPERATURE-1
EFFECT	TEMPERATURE-2
TIME	< <i>find-anchor-time</i>
SCOPE-OF	ASPECT-1

TEMPERATURE-1

DOMAIN	ROOM-1
RANGE	< TEMPERATURE-2.RANGE

TEMPERATURE-2

DOMAIN	ROOM-1
RANGE	> TEMPERATURE-1.RANGE

HEATER-1

SIZE	.7
INSTRUMENT-OF	CHANGE-EVENT-1

ASPECT-1

PHASE	continue
SCOPE	CHANGE-EVENT-1

1. The agent looks up the word *heat* and finds only a nominal sense, described as TEMPERATURE (RANGE (> .7)).
2. From its inventory of methods to treat different eventualities of new-word-sense learning, it selects the one aimed at learning new verb senses from noun senses that map to a PROPERTY value.

3. The chosen learning method directs the agent to hypothesize that the verb refers to a CHANGE-EVENT involving this property. CHANGE-EVENT is an ontological concept used to describe events whose meaning is best captured by comparing the value of some property in the event's PRECONDITION and EFFECT slots. For example, *speed up* is described as a CHANGE-EVENT whose value of SPEED is lower in the PRECONDITION than in the EFFECT. Analogous treatments are provided in the lexicon for *grow taller* (HEIGHT), *quiet down* (LOUDNESS), *go on sale* (COST), and countless other words and phrases (see McShane, Nirenburg, & Beale, 2008, for details).
4. The agent creates a CHANGE-EVENT TMR frame, along with its PRECONDITION and EFFECT slots.
5. It hypothesizes the direction of change (i.e., the comparison between the range of TEMPERATURE in the PRECONDITION and EFFECT) on the basis of the RANGE of the property in the nominal sense: if the nominal sense has a high value (like the .7 listed in the nominal sense of 'heat'), then it assumes that the direction of change is increase; if the nominal sense has a low value, then it assumes that the direction of change is decrease.
6. It interprets the THEME of the CHANGE-EVENT (here, ROOM-1) as the DOMAIN of the TEMPERATURE frames.
7. It links the meaning of *heater* to the CHANGE-EVENT using the case role INSTRUMENT, since the default case role for subjects (AGENT) cannot apply to inanimates like *heater*.
8. It deals with routine semantic analysis needs, such as the analysis of tense and aspect.

To reiterate, this algorithm works for the case of unknown verbs for which an available nominal sense maps to a SCALAR-ATTRIBUTE. This is only one of many eventualities, all of which need to be fully fleshed out algorithmically and then tested against corpus evidence. Our expectation is that, along with some impressive automatic results, we will encounter many false positives—that is, cases in which the meaning of a word will not be predictable from a morphologically related word form. That, in turn, will motivate further enhancements to the learning algorithms. This is, however, an envelope that we must push hard because automating lexical knowledge acquisition is a high priority for knowledge-based agent systems.

4.12 Wrapping Up Basic Semantic Analysis

Even after all this processing has been carried out, a lot of loose ends remain, such as residual lexical and referential ambiguity, low-scoring TMRs resulting from incongruities, procedural semantic routines that have been posited in the TMR but have not yet been run, and fragmentary utterances that have not yet been incorporated into the discourse structure. All of these can be pursued by a LEIA if it chooses to do so in later stages of

processing—but it might not. By this point, the LEIA has an idea—or several competing ideas—of what the input is about, and the given topic may or may not be relevant to its operation. For example, if the LEIA is a mechanic’s assistant, but the topic is football, and there are multiple humans involved in the conversation, there is no reason for the agent to ponder which meaning of football is intended (soccer or American football), and it should certainly not pester the speaker with clarification questions about it. In short, the conclusion of Basic Semantic Analysis is an important decision point for the LEIA with respect to its language- and task-related reasoning.

4.13 Further Exploration

1. Explore inventories and classifications of metaphors available online, using search terms such as “Lakoff and Johnson metaphor,” “conceptual metaphor,” “structural metaphor,” and “English metaphor list.” When analyzing individual metaphors, think about questions such as the following:

- Does it feel metaphorical to you, or does it seem more like a fixed expression at a distance from its metaphorical roots?
- Do you know the etymology?
- Do you think that knowing the etymology helps you to understand the intended meaning?
- Do you fully understand its intended meaning?
- Is that meaning precise or vague?
- Does the metaphor sound normal/everyday or creative/flowery?
- Can you quickly think of a nonmetaphorical paraphrase, or is the metaphor the default way of expressing the given idea?

2. Use the online version of the COCA corpus (<https://www.english-corpora.org/coca/>) to explore the distribution of speech acts (direct vs. indirect) for constructions that can canonically indicate an indirect speech act. Can you identify any heuristics to predict whether the direct or indirect meaning is intended? For example,

- When does *I need to know* indicate a direct speech act (simply that the speaker needs to know something) versus an indirect speech act (i.e., “Tell me”)?
- When does *Can you ...* indicate a direct speech act (asking about the hearer’s ability to do something) versus an indirect speech act (“Please do it”)?
- When does *It would be great if* indicate a direct speech act (the expression of a desire) versus an indirect speech act (“Please make this happen”)?

3. Use the online version of the COCA corpus (<https://www.english-corpora.org/coca/>) to explore nominal compounds. There's a challenge, however: the interface does not allow you to search for patterns that are as unconstrained as "any noun followed by any noun." Invent search strategies that give you some insights into how nominal compounding works within the constraints of the search engine. For example, think about classes of nouns for which a particular word can serve as an example to seed exploration. For example, *professor_nn* can be used to investigate "social-role + any-noun" (e.g., *professor rank*). How do the hits compare with *carpenter_nn*? *nurse_nn*? *nn_nn chef* (e.g., *genius pastry chef*)?

4. Looking just at the table of contents at the beginning of the book, try to reconstruct what was discussed in each section of chapter 4 and recall or invent examples of each phenomenon.