Linguistics for the Age of AI

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

6

Extended Semantic Analysis

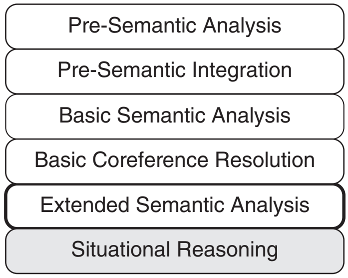

During Extended Semantic Analysis the LEIA looks beyond the local dependency structure (i.e., the main event in a clause and its arguments) in an attempt to resolve outstanding ambiguities, incongruities, and underspecifications that were identified during Basic Semantic Analysis. Like all processing so far, Extended Semantic Analysis uses methods that are applicable to texts in all domains. It does not involve Situational Reasoning, which will be invoked, if needed, later. Extended Semantic Analysis is triggered in the following situations:

1. Multiple TMR candidates received a high score because Basic Semantic Analysis could not resolve some ambiguities (section 6.1).
2. All TMR candidates received a low score because Basic Semantic Analysis encountered incongruities (section 6.2).
3. Data in the basic TMR—namely, calls to procedural semantic routines—indicate that more analysis of a specific kind is needed. Most often, an underspecified concept requires further specification (section 6.3).[1]
4. The TMR is a nonpropositional fragment that must be incorporated into the larger context (section 6.4).

Extended Semantic Analysis, like Basic Coreference Resolution, addresses difficult linguistic phenomena. For some of them, a complete solution will be beyond the state of the art for quite a while. But this is not necessarily detrimental to the agent's overall functioning.[2] Consider a real-life example: You cross paths with a colleague walking across campus, have a quick chat, and she wraps it up by saying, "Sorry, I've got to run to a dean thing." *Dean thing* is a nominal compound that leaves the semantic relation between the nouns unspecified. The *thing* in question could be about a dean, organized by a dean, required by a dean, or for deans only. Do you care which? Probably not. The speaker's point is that she has a good reason to cut the conversation short. Underspecification is a useful design feature of language, and it makes no sense to build agents who will not stop until they have tried long and hard to concretize every vague utterance.

6.1 Addressing Residual Ambiguities

During this stage, the LEIA's main knowledge source for resolving residual ambiguity (i.e., choosing from among multiple high-scoring candidate interpretations) is the ontology. The agent attempts to understand the context by looking for ontological connections between candidate interpretations of words. Consider the following minimal pair of examples:

(6.1) The police arrived at the port before dawn. They arrested the pirates with no bloodshed.

(6.2) The police arrived at the secret computer lab before dawn. They arrested the pirates with no bloodshed.

What comes to mind as the meaning of *pirates* in each case? Most likely, seafaring bandits for the first, and intellectual property thieves for the second. This is because a port suggests maritime activity, whereas a computer lab suggests intellectual activity. The reason why the agent cannot recognize the preferred reading of *pirates* during Basic Semantic Analysis is that the deciding clue—the location of the event—is in a different dependency structure. That is, when the *arrest* sentences are processed in isolation, both readings of *pirate* are equally possible because they both refer to types of HUMAN, and all HUMANs can be arrested. To disambiguate, the agent needs to extend its search space to the preceding sentence.[3]

Five types of ontological knowledge have proven useful for disambiguating such inputs. All of these heuristics involve relations between OBJECTs since it is OBJECT-to-OBJECT relations that were not covered by the dependency-based (largely OBJECT-to-EVENT) disambiguation of Basic Semantic Analysis. The heuristics are applied in the order in which they are presented.

6.1.1 The Objects Are Linked by a Primitive Property

OBJECTs in the LEIA's ontology are described by dozens of properties. Some of these, such as LOCATION and HAS-OBJECT-AS-PART, link OBJECTs to other OBJECTs, asserting their close ontological affinity. The following pair of examples shows how this knowledge is useful for disambiguation.

(6.3) "What a nice big stall!" "Well, that's a very big horse!"

(6.4) The horse was being examined because of a broken tooth.

When a LEIA encounters *horse* in an input, it must determine whether it refers to an animal, a sawhorse, or a piece of gymnastic equipment. When it considers the animal-oriented analysis HORSE, all of the concepts shown in the ontology excerpt below (as well as many others) are understood to be potential participants in the context.

HORSE		
AGENT-OF	sem	TROT, CANTER, GALLOP, BUCK-EVENT, …
COLOR	sem	white, black, gray, bay, chestnut, buckskin,[4] …
LOCATION	sem	BARN, **ANIMAL-STALL**, RIDING-ARENA, …
HAS-OBJECT-AS-PART	sem	HOOF, MANE, TAIL, HEAD, LEG, **TOOTH**, …

So, when analyzing (6.3), the LEIA recognizes that both HORSE and ANIMAL-STALL are in the candidate space; and when analyzing (6.4), it recognizes that both HORSE and TOOTH are in the candidate space. Finding these correlations helps to disambiguate *both* of the words in each context simultaneously—after all, *stall* can also mean a booth for selling goods, and *tooth* can also refer to a tool part.

6.1.2 The Objects Are Case Role Fillers of the Same Event

Another way to detect close correlations between OBJECTS is through a mediating EVENT. That is, the LEIA might be able to find an EVENT for which some interpretations of the OBJECTS in question fill its case role slots. Returning to our seafaring bandit example (6.1), the ontology contains a WATER-TRAVEL-EVENT for which PIRATE-AT-SEA is a typical filler of the AGENT case role, and PORT is a typical filler of both the SOURCE and DESTINATION case roles, as shown below.

WATER-TRAVEL-EVENT		
AGENT	default	SAILOR, **PIRATE-AT-SEA**
	sem	HUMAN
	relaxable-to	ANIMAL
THEME	sem	WATER-VESSEL
LOCATION	sem	BODY-OF-WATER
SOURCE	default	**PORT**
	sem	GEOGRAPHIC-ENTITY
DESTINATION	default	**PORT**
	sem	GEOGRAPHIC-ENTITY

Finding these fillers, the LEIA concludes that WATER-TRAVEL-EVENT is the ontological context of the utterance and selects PIRATE-AT-SEA (not INTELLECTUAL-PROPERTY-THIEF) as the analysis of *pirate*, and PORT (not PORT-WINE) as the analysis of *port*. The success of this search strategy depends on the coverage of the ontology at any given time—that is, it is essential that the ontology *have* an event that associates seafaring bandits and ports using its case roles.

6.1.3 The Objects Are Linked by an Ontologically Decomposable Property

The properties discussed in section 6.1.1, LOCATION and HAS-OBJECT-AS-PART, are primitives in the ontology. However, it is convenient—both for knowledge acquisition and for the agent's reasoning over the knowledge—to record some information using properties that are shorthand for more complex ontological representations. We call these *ontologically* *decomposable* properties because rules for their expansion must be specified in knowledge structures appended to the ontology. Consider the excerpt from the ontological description of INGEST that was introduced in chapter 2:

INGEST		
AGENT	sem	ANIMAL
	relaxable-to	SOCIAL-OBJECT
THEME	sem	FOOD, BEVERAGE, INGESTIBLE-MEDICATION
	relaxable-to	ANIMAL, PLANT
	not	HUMAN

Using the simple slot-filler formalism of the nonscript portion of the ontology, it is not possible to record who eats what—that horses eat grass, hay, oats, and carrots, whereas koalas eat only eucalyptus leaves.[5] That is, the portion of the knowledge structures below indicated in square brackets cannot be easily accommodated using the knowledge representation strategy adopted for the broad-coverage (nonscript) portion of the ontology.

HORSE
AGENT-OF INGEST [THEME GRASS, HAY, OAT-FEED, CARROT]

KOALA
AGENT-OF INGEST [THEME EUCALYPTUS-LEAF]

The reasons why property values cannot, themselves, be further specified by nested property values are both historical and practical. Historically speaking, the ontology was acquired decades ago, in service of particular goals (mostly disambiguating language inputs) and supported by a particular acquisition/viewing interface. Practically speaking, there are reasons to uphold this constraint. Namely, it simplifies not only the human-oriented work of knowledge acquisition, management, and visualization but also an agent's reasoning over the knowledge. Much more could be said about this decision within the bigger picture of knowledge representation and automatic reasoning, but we leave that to another time. The point here is that it is possible both to uphold the decision to allow only simple slot fillers (along with all of its benefits) and to provide the agent with more detailed knowledge. In fact, there are at least two ways to record such knowledge: as ontological scripts (described in sections 2.3.1 and 2.8.2) and using decomposable properties. We consider these in turn.

If an agent needs extensive specialist knowledge about certain animals—for example, to generate a computer simulation of their behavior or reason about it—then full ontological scripts must be recorded. For example, one could acquire an INGESTING-BY-KOALAS script, which would not only assert that the AGENT of this event is KOALA and that its THEME is EUCALYPTUS-LEAF but also describe many more details about this process: how the koala gathers the leaves, how long it chews them, how many it eats per day, and so on. In short, an ontology could contain many descendants of INGEST (INGESTING-BY-KOALAS, INGESTING-BY-HORSES, INGESTING-BY-WHALES) that provide extensive information about what and how different kinds of animals eat. However, unless all of these new events are going to provide much more information than simply what each animal eats, it is inefficient to create concepts for all of them.

A more streamlined solution for recording *who eats what* is to create an ontologically decomposable property like TYPICALLY-EATS that directly links an animal to what it eats. This allows knowledge acquirers to record in the ontology information like the following:

HORSE
  TYPICALLY-EATS   default   GRASS, HAY, OAT-FEED, CARROT, APPLE

KOALA
  TYPICALLY-EATS   default   EUCALYPTUS-LEAF

This is a shorthand for

ANIMAL
  AGENT-OF   sem   INGEST-#1

INGEST-#1
  THEME   default   INGESTIBLE

in which the notation -#1 indicates coreference between *ontological instances* of the concept (i.e., it is a method of indicating coreference among knowledge structures in static knowledge resources). The shorthand TYPICALLY-EATS is connected to its expansion using a rule encoded in the analysis algorithm.

Consider how such object associations can help in disambiguation. Given the input *The cow was eating grass*, the knowledge COW (TYPICALLY-EATS GRASS) allows the agent to simultaneously disambiguate *cow* as the animal COW (not a derogatory reference to a woman) and *grass* as the lawn material GRASS (not marijuana).

6.1.4 The Objects Are Clustered Using a Vague Property

The approach just described requires knowledge acquirers to introduce specific, decomposable properties, provide rules for their semantic expansion, and record the associated sets of concepts. The resulting knowledge is very useful but takes time to acquire.[6] A faster and cheaper type of knowledge acquisition is to have a vague RELATED-TO property[7] that can hold a large set of associated concepts. For example:

- Objects related to the outdoor space of someone's house include GRASS, WEED, FENCE, TREE, LAWNMOWER, SWIMMING-POOL, DRIVEWAY, GARDEN-HOSE, BUSH, PICNIC-TABLE, LAWN-CHAIR, JUNGLE-GYM.
- Objects related to a kitchen include UTENSIL, BLENDER, PANTRY, COUNTERTOP, SAUCEPAN, DISHWASHER, KITCHEN-TOWEL.

This shorthand not only is a useful knowledge-engineering strategy but also reflects the concept-association behavior of people. For example, a person asked to name ten things associated with a horse might well include among them *saddle*. The association is recalled without the person actually going through a semantic expansion like "a saddle is the thing a person sits on when riding a horse."

To emphasize, we are talking about inventories of related *concepts*, not ambiguous *words*. So, although various word-based resources—for example, the results of statistical word clustering or the results of human word-association experiments—can be useful to help detect such associations, the knowledge must be (manually) encoded as concepts in order to become part of an agent's ontology and unambiguously inform its reasoning.

6.1.5 The Objects Are Linked by a Short Ontological Path That Is Computed Dynamically

If the above approaches fail to disambiguate an input, the agent can try to establish what the utterance is about by searching for the shortest ontological path between all candidate interpretations of the OBJECTS in the local context. The problem with this last-ditch strategy, however, is that it is difficult to achieve high-quality shortest-path calculations in an ontology. The main reason for this is that shortest-path calculations depend on the effective assignment of traversal costs for different kinds of properties. For example, traversing an IS-A link will have a lower cost than traversing a HAS-OBJECT-AS-PART link because concepts linked by IS-A (e.g., DOG and CANINE) are more similar than concepts linked by HAS-OBJECT-AS-PART (e.g., DOG and EAR). The key to using this strategy is to apply it only if it results in a short, very low-cost path between concepts. If the path is not very low-cost, then this is not a reliable heuristic.[8]

To recap, so far we have seen five ontology-search strategies that an agent can use to resolve residual ambiguity. All of them rely on identifying closely related ontological OBJECTS in the immediately surrounding context. Optimizing the definition of the *immediately surrounding* context is as difficult as automatically determining the window of coreference for coreference resolution.

6.1.6 Reasoning by Analogy Using the TMR Repository

Another source of disambiguating heuristics is the TMR repository, which is a knowledge resource that records the agent's memories of past language-to-meaning mappings. Remembered TMRs can serve as a point of comparison for reasoning by analogy.[9] Reasoning by analogy is a big topic that we will touch on only to the extent needed for the goal at hand: lexical disambiguation.

A difficult problem in lexical disambiguation is the frequency with which a sentence can potentially have both a literal and a metaphorical reading. For example, *strike back* can mean "to hit physically" or "to retaliate nonphysically (e.g., verbally)." If the LEIA has previously encountered the expression *strike back*, the remembered meaning representations can serve as a vote for the associated analysis. However, slick as this approach might sound—and psychologically plausible as well—it is anything but straightforward to implement. There are at least three complications.

Complication 1. In different domains, different disambiguation decisions will be correct. For example, if a LEIA analyzes many texts about boxing and then turns its attention to texts about office interactions, it should not interpret every spat that involves confrontational language as an instance of physical assault simply because it has many TMRs about people punching each other's lights out. It follows that reasoning by analogy requires a nontrivial prerequisite: marking each remembered TMR with a domain in which it is applicable. This makes the applicability of this method rather problematic.

Complication 2. If the remembered TMRs are to be useful targets of reasoning by analogy, then they must not only belong to the same domain as the TMR being disambiguated but also be correct. But generating correct TMRs for every single input is beyond the state of the art. This means that a TMR repository is likely to contain a combination of correct and not-completely-correct TMRs. The most reliable way to ensure the quality of the repository would be to have people check and correct all the TMRs. However, this is realistic only for small repositories. One automatic method of assessing TMR quality is using the agent's own confidence estimates in its interpretations, which are computed and stored as a matter of course. However, for reasons explained in section 9.3, those estimates are not always reliable. Another method of automatically assessing the quality of TMRs relies on heuristics that must be computed outside the NLU module. Namely, if an input requires some action by the agent, and if the agent responds appropriately to it, then there is a good chance that the agent correctly understood it. Of course, automatically determining that the agent's action was appropriate requires task-level reasoning beyond what we detail in this book. Moreover, since not every input gives rise to an observable action by the agent, this is far from an all-purpose solution to evaluating the quality of the TMRs in the TMR repository.

Complication 3. The TMR repository might not contain any analyses relevant to the given input. This raises the questions, "Should different agents share their TMR repositories?" and "How can we best utilize similarity measures to exploit close but not exact matches?" As regards the latter, whereas "I'm going to kill him" does not usually refer to murder (though it can), "I'm going to smack him upside the head" may or may not involve physical violence. So the ontological similarity between HIT and KILL does not necessarily support reasoning by analogy in this instance.

If the domain-independent methods of resolving residual ambiguity described in the preceding subsections do not cover a particular input, then domain-specific methods (described in chapter 7) must be brought to bear. The reason we do not start with the latter is that committing to a specific domain largely erases the word-sense ambiguity problem to begin with. In fact, avoiding word-sense ambiguity by developing narrow-domain applications is a widely practiced strategy for developing agent systems. Although this strategy can work quite well for narrow domains, it will not advance the state of the art in making agents perform at the level of their human counterparts. After all, even when people are engaging in a narrowly defined task, they *will* engage in off-topic conversation—that is just part of being human.

6.1.7 Recap of Methods to Address Residual Ambiguity

- Prefer interpretations of OBJECTs that are linked by a primitive property in the ontology. For example, to analyze *The horse was being examined because of a broken tooth*, use the ontological knowledge HORSE (HAS-OBJECT-AS-PART TOOTH).
- Prefer interpretations of OBJECTs that fill case role slots of the same EVENT in the ontology. For example, to analyze *The police arrived at the port before dawn. They arrested the pirates with no bloodshed*, use the ontological knowledge WATER-TRAVEL-EVENT (AGENT PIRATE-AT-SEA) (DESTINATION PORT).
- Prefer interpretations of OBJECTs that are linked by an ontologically decomposable property. For example, to analyze *The horse wants some grass*, use the ontological knowledge HORSE (TYPICALLY-EATS GRASS), which expands, via a recorded reasoning rule, to HORSE (AGENT-OF INGEST (THEME GRASS)).
- Prefer interpretations of OBJECTs that are linked by the vague ontological property RELATED-TO. For example, to analyze *I need to tack up the horse; where's the bridle?*, use the ontological knowledge HORSE (RELATED-TO BRIDLE).
- Prefer interpretations of OBJECTs that are linked by a short ontological path. For example, to analyze *I need to tack up the horse; where's the bridle?*—assuming that the needed RELATED-TO information was *not* recorded—use the path HORSE (THEME-OF TACK-UP-HORSE (INSTRUMENT BRIDLE)).
- Use reasoning by analogy against the TMR repository. For example, if every past analysis of *strike back* generated the nonphysical interpretation, that is a vote in favor of the nonphysical interpretation for the new input—assuming that none of the complications discussed above confound the process.

6.2 Addressing Incongruities

Incongruity describes the situation when no analysis of an input aligns with the expectations recorded in the LEIA's knowledge bases. The subsections below describe four sources of incongruities—metonymy, preposition swapping, idiomatic creativity, and indirect modification—and the methods LEIAs use to resolve them.

6.2.1 Metonymy

In a metonymy, one entity stands for another. For example, in (6.5), *the spiky hair* refers to a particular person with spiky hair.

(6.5) The spiky hair just smiled at me.

Speakers of each language know which metonymic associations can exist between a named entity and what it stands for. (By contrast, metaphors can establish novel relations between entities.)

Metonymy leads to a sortal incongruity during Basic Semantic Analysis. This means that an event head and its dependents fail to combine in a way that aligns with ontological expectations. In (6.5), the problem is that HAIR is not a valid AGENT of SMILE-EVENT. Speakers of English readily understand the indirect reference because we know that people can be referred to metonymically by their physical features, clothing, or items closely associated with them.

Just as people are aware of typical metonymical relationships, so, too, must be LEIAs. To maintain an inventory of canonical metonymical replacements, our model introduces a dedicated knowledge resource, the LEIA's Metonymic Mapping Repository.[10] A subset of its content is illustrated by (6.6)—(6.10).

(6.6) [Producer for product]
Then your father bought an Audi with a stick shift. (COCA)

(6.7) [Social group for its representative(s)]
And for her heroic efforts, the ASPCA awarded her a gold medal. (COCA)

(6.8) [Container for the substance in it]
… A large pot boiled lid-rattlingly on the stove. (COCA)

(6.9) [Clothing for the person wearing it]
I want to dance with the big belt buckle.

(6.10) [Artist for a work of art]
In addition to the Rembrandts, there are five Vermeers, nearly a dozen Frans Halses, and the list goes on … (COCA)

Carrying out such replacements is a simple and high-confidence method for dealing with the most typical metonymies.

If an input containing a potential metonymy does not match a recorded construction, then the agent attempts to determine how the kind of entity named in the text is related to the kind of entity expected by the ontology. For example, (6.11) says that either SPECTACLES or a set of DRINKING-GLASSes (the two meanings of *glasses* in the lexicon) are the AGENT of BORROW. But the ontology says that only HUMANS can BORROW things.

(6.11) The big glasses borrowed my bike.

So the agent must determine whether either SPECTACLES or a set of DRINKING-GLASSes might be standing in for a HUMAN—and, if so, based on what relation(s). It does so using the same kind of ontological search described in section 6.1.5 (Onyshkevych, 1997). Stated briefly, the agent computes the weighted distance between HUMAN and both of these concepts. The cumulative score for each reading is a function of the length of the path and of the cost of traversing each particular relation link. The shortest path turns out to be between HUMAN and SPECTACLES, so the metonymy can be resolved as *HUMAN RELATION (SPECTACLES (SIZE .8))*.

6.2.2 Preposition Swapping

Prepositions are a common source of performance errors by native and nonnative speakers alike.[11] In each of the examples below, the first preposition choice is canonical and the second is not—but it was attested in examples in the COCA corpus.

- translate language-X into [to] language-Y
- abide by [with] X
- be absolved of [from] X

Considering that English is the current worldwide *lingua franca*, with many speakers having nonnative fluency, it is a high priority for LEIAs to accommodate this type of close-but-not-perfect input. For example, it is not uncommon for subtitles of foreign films to be of high quality overall but to show the occasional odd preposition. The question is, How to detect instances of preposition swapping?

The first thing to say is the obvious: Not any preposition can be swapped for any other. So the preposition-swapping algorithm must be tightly constrained. According to our current algorithm, in order for the agent to hypothesize preposition swapping, all of the following must hold:

1. The lexicon must contain a fixed expression (i.e., an idiom or construction) that matches the input lexically and syntactically except for the preposition choice. So we are not talking about free combinations of prepositions and their complements.
2. All of the semantic constraints for that fixed expression must be met. In the example *translate X into [to] Y*, X must be a language or text and Y must be a language. These constraints are specified in the lexical sense for the construction *translate X into Y*. (Note that there is a different sense for *translate* <u>*to*</u> which means "result in," as in *Saving an extra $50 a month translates to $600 a year*.)
3. The preposition pair belongs to a list of preposition pairs that we have determined to be, or hypothesize to be, subject to swapping. These pairs either contain prepositions with similar meanings (*in/into*, *into/to*, *from/out of*, *by/with*) or contain at least one preposition that is extremely semantically underspecified, such as *of*.

A natural question is, If the LEIA successfully processes an input using this preposition-swapping repair method, should the attested preposition be recorded in the lexicon and treated ever after as canonical? The answer is no for three reasons. First, when LEIAs *generate* language, they should not generate less-preferred versions. Second, this recovery procedure works on the fly, so there is no reason to record the less-preferred version. Third, resorting to a recovery procedure models the additional cognitive load of processing unexpected input, which will result in a penalty to the overall confidence score of the TMR. This means that successful analyses that do *not* require recovering from unexpected input will be preferred, as they should be.

In addition, the LEIA's history of language analyses (recorded in the TMR repository) could be consulted when evaluating the likelihood of, and scoring penalty for, a preposition-swapping analysis. For example, if the TMR repository contains multiple examples of a particular preposition swap, the LEIA could reduce the penalty for that swap to a fraction of the norm. After all, maybe diachronic language change over the next couple of decades will result in *translate to* becoming a perfectly natural alternative to *translate into* for language-oriented contexts.

6.2.3 Idiomatic Creativity

The creative use of idioms[12] may or may not trigger extended processing. Let us begin with a case in which extended processing will not be triggered because no incongruity will be detected. Imagine that you are in your backyard entertaining a guest, and two deer sidle up, stomping on your freshly seeded grass. You clap your hands and make some noise, but they ignore you—they are fearless, suburban deer. So you go inside, grab your trumpet (lucky thing, you play the trumpet), and burst into a fanfare, at which time the deer bound out of sight into the woods. Your guest says, *Wow, you've killed two deer with one trumpet!* You laugh, but your companion LEIA won't get the joke—at least not at this stage of analysis. After all, everything lines up semantically. The event KILL requires an ANIMAL as the AGENT (you as a HUMAN fit), it requires a nonhuman ANIMAL as the THEME (the DEER fit), and it allows for a physical object to be the INSTRUMENT (the TRUMPET fits, even though it is not the preferred instrument of killing, which is WEAPON). Even if this utterance were taken out of context, any human would know it must have an indirect meaning: people just don't use trumpets to kill deer. Some day, LEIAs will have to have this depth of world knowledge as well. For now, no incongruity will be flagged for this example.

Not so, however, for the example, *You've killed two scratches with one rug!*, which might be said when a single throw rug works to cover two gouges in a wooden floor. This *will* lead to an incongruity because there is no meaning of *kill* that expects its THEME to be SCRATCH-MAR.

A similar split obtains between the examples *Don't put all your eggs in one boat* versus *Don't put all your eggs in one portfolio of statewide munis*—both of which were attested in the COCA corpus. In the first case, there will be no incongruity since eggs can be put into a boat. (We explain later how the idiomatic usage can be detected in a different way.) But in the second, there *will* be an incongruity because *portfolio of statewide munis* is an ABSTRACT-OBJECT, not a PHYSICAL-OBJECT; therefore, it is not a suitable DESTINATION for the physical event TRANSFER-POSITION. These pairs of examples illustrate how selectional constraints can flag an incongruity and suggest that the input might include idiomatic creativity.

If the input might be a play on an idiom, the agent must first identify the lexical sense that records the canonical form of that idiom. Although some global notion of fuzzy matching could be invoked, this is risky since *close but not quite* typically means that the input simply doesn't match the idiom. For example, *kick the pail* does not mean *die*, even though *pail* is a synonym of *bucket*.

There are two stages to processing creative idiom usages, detecting them and semantically analyzing them, which we consider in turn.

6.2.3.1 Detecting creative idiom use We prepare agents to detect creative idiom use in two ways: (1) by writing lexical senses that anticipate particular kinds of variations on particular idioms and (2) by implementing lexicon-wide rules that cover generic types of idiomatic creativity. We consider these in turn.

Writing lexical senses that anticipate particular kinds of variations on particular idioms. Many individual idioms allow for variations that people know or can easily imagine. The most reliable way to prepare agents to detect and analyze such variations is to record them in the lexicon. Table 6.1 illustrates such anticipatory lexical acquisition using results from an informal corpus study of idiom variation in the COCA corpus. Column 1 presents canonical forms of idioms, which will be recorded as one lexical sense, and column 2 presents variable-inclusive constructions, recorded as another lexical sense.[13] For example, the lexicon contains two senses of the verb *drop* to cover the data in row 1: one for the fixed form "at the drop of a hat" and the other for the variable-inclusive form "at the drop of a [N+]." Column 3 presents corpus-attested variations on the idioms, whose full examples are presented as (6.12)–(6.29).

Table 6.1
Canonical and variable-inclusive forms of idioms recorded as different lexical senses

Canonical form	Variable-inclusive form	Attested variations in cited examples
at the drop of a hat	at the drop of a [N+]	sixteenth note, hot dog, pin, pacifier, beaker, hare, backbeat
put all [one's] eggs in one basket	put all [one's] eggs in one [N+]	portfolio of statewide munis, blender, boat
	put all [one's] eggs in the [N+]	the stock market basket
	put all [one's] [N+] in one basket	
NEG judge a book by its cover	NEG judge a [N+] by its cover	bike, star
	NEG judge a book by its [N+]	title
[get information, be told] straight from the horse's mouth	[get information, be told] straight from the [N+]'s mouth	moose's, looney's
[get, be given] a dose of [one's] own medicine	[get, be given] a dose of [one's] own [N+]	frank talk
kill two birds with one stone	kill [NUM] [N+] with one [N+]	two topics/column

Note: [N+] indicates a head noun plus potential modifiers. [NUM] indicates a number.

(6.12) Each singer could turn the emotional temperature up or down *at the drop of a sixteenth note*. (COCA)

(6.13) Iris said that Judy Garland could cry *at the drop of a hot dog*. (COCA)

(6.14) I could break into *sobs at the drop of a pin*. (COCA)

(6.15) You take 100 pictures *at the drop of a pacifier*. (COCA)

(6.16) As trained scientists are wont to do *at the drop of a beaker*, he postulated a plausible theory. (COCA)

(6.17) I would run from New York and Columbia, like a hound *at the drop of a hare*. (COCA)

(6.18) These 12 tracks boast a startlingly powerful sound, shifting *at the drop of a backbeat* from a whispered seduction to a raging fury. (COCA)

(6.19) You would be foolish to *put all your eggs in one portfolio of statewide munis*. (COCA)

(6.20) Well, if you're running a startup that sells to other startups, you might be *putting all your eggs in one blender*. (COCA)

(6.21) I don't *put all my eggs in one boat*. (COCA)

(6.22) Diversifying, spreading your wealth around, not *putting all your eggs in the stock market basket*, is going to pay off. (COCA)

(6.23) By now we all know better than to *judge a bike by its cover*, but the Time RXR provides a deceptively smooth ride, especially as its angular, aggressive look screams race-stiff. (COCA)

(6.24) Don't *judge a star by its cover*. One of Kepler's seismic discoveries is the Sun-like star Kepler-37, which lies about 220 light-years from Earth in the constellation Lyra. (COCA)

(6.25) Sometimes you're not supposed to *judge a book by its title*, but in these types of books, there's an awful lot to the title. (COCA)

(6.26) Last week, I talked to both of them to get the story of moosebread, or "moose food," as they call it, *straight from the moose's mouth*. (COCA)

(6.27) We're going to hear it *straight from the loony's mouth*. (COCA)

(6.28) Mr. Vajpayee said India was not prepared to do that, and the president *got a dose of his own frank talk* as he sat at that state dinner last night and India's president took issue with the U.S. leader[']s description of the Indian subcontinent. (COCA)

(6.29) Put the two together and you *kill two topics with one column*. (COCA)

Let us work through the first example in the table. The canonical form of the idiom is shown in column 1: *at the drop of a hat*. Column 2 indicates that *at the drop of [any head noun with optional modifiers]* is an acceptable play on the idiom. However, note that the string *at the drop of* is immutable. Of course, like any aspect of knowledge acquisition, the decision about how to best formulate the idiom-extension template is best informed by a combination of intuition and corpus evidence.

In the second example, three different extended templates are all considered possible. They allow for different elements to be variable, but not all at the same time.

Earlier we said that the idiomaticity of *put all one's eggs in one boat* could not be detected on the basis of semantic incongruity because there is no incongruity—one can put eggs in a boat. So, how will the system know to consider an idiomatic interpretation? As long as we list the sense *put all [one's] eggs in one [N+]* in the lexicon, the Basic Semantic Analyzer will generate the idiomatic interpretation alongside the literal one; choosing between them will be undertaken during Situational Reasoning.

Note that a core aspect of acquiring idioms is listing all of their known variations, not just the one that pops to mind first. For example, although *let the cat out of the bag* is the most canonical form of this idiom, it often occurs as *the cat is out of the bag.* Since this variant is so common, we hypothesize that it is part of people's lexical stock and is most appropriately recorded as a separate lexical sense, leaving the more creative flourishes for dynamic analysis.

As concerns the agent's confidence in detecting idioms, the fixed senses offer the most confidence, whereas the variable-inclusive senses are more open to false positives. Our next method of detecting idiom play—using lexicon-wide rules—has broader coverage but is also more open to false positives.

Implementing lexicon-wide rules that detect generic types of idiomatic creativity. So far, we have experimented with three such rules.

Rule 1. Allow for two fixed NPs in the idiom to be swapped, as in (6.30).

(6.30) The man has a bad accent, he tells McClane it was raining "dogs and cats" instead of cats and dogs, and he refers to the elevator as "the lift." (COCA)

NP swapping might be the result of misspeaking, misremembering, or trying for comedic effect.[14]

Rule 2. If there are six or more fixed words in an idiom, allow for any one of them to be replaced, as in (6.31) and (6.32).

(6.31) [A play on the six-word idiom 'wake up and smell the coffee']
If you're an idler, wake up and smell the bushes burn. (COCA)

(6.32) [A play on the six-word idiom 'be a match made in heaven']
By contrast, the group feels its blend of dance sounds and political lyrics is a musical marriage made in heaven. (COCA)

The fixed-word threshold of six attempts to balance the desire to detect as much idiom play as possible against inviting too many false positives.

Rule 3. Allow for modifiers. Example (6.32) illustrates this type of wordplay: *a* musical *marriage made in heaven*.

6.2.3.2 Semantically analyzing creative idiom use Once creative idiom use has been detected, it must be semantically analyzed. The three steps of semantic analysis described below apply to all instances of idiomatic creativity, whether the idiom play is explicitly accounted for by a lexical sense with variables or is detected using lexicon-wide rules. There are slight differences in processing depending on the detection method, but we will mention them only in passing since they are too fine-grained to be of general interest. If the creative idiom use matches a lexical sense that specifically anticipates it, then (a) the procedural semantic routine comprising the three analysis steps below is recorded in the meaning-procedures zone of that sense; (b) that procedural semantic routine can be tweaked, if needed, to accommodate that particular instance of idiomatic creativity; and (c) the confidence in the resulting analysis is higher than when lexicon-wide rules are applied.

We will describe the three steps of analysis using example (6.15): *You take 100 pictures at the drop of a pacifier.*

Step 1. Generate the TMR using the recorded meaning of the basic form of the idiom. Our example plays on the recorded idiom *at the drop of a hat*, which means *very quickly*. So our example means that the person takes one hundred photographs very quickly (SPEED .9).

```
TAKE-PHOTOGRAPH-1
    AGENT           HUMAN-1[15]
    THEME           SET-1
    SPEED           .9

SET-1
    MEMBER-TYPE     PHOTOGRAPH
    CARDINALITY     100
```

Step 2. Explicitly record in the TMR the SPEECH-ACT that is implicitly associated with any utterance. Understanding this step requires a bit of background. The meaning of every utterance is, theoretically speaking, the THEME of a SPEECH-ACT. The AGENT of that SPEECH-ACT is the speaker (or writer) and the BENEFICIARY of that SPEECH-ACT is the hearer (or reader). In general, we do not have the agent generate a SPEECH-ACT frame for every declarative statement because it is cumbersome; however, the SPEECH-ACT does implicitly exist. In the case of plays on idioms, making the SPEECH-ACT explicit is just what is needed to offer a template for recording property values that would otherwise have no other way to attach to the TMR.

```
SPEECH-ACT-1
    THEME           TAKE-PHOTOGRAPH-1

TAKE-PHOTOGRAPH-1
    AGENT           HUMAN-1
    THEME           SET-1
    SPEED           .9

SET-1
    MEMBER-TYPE     PHOTOGRAPH
    CARDINALITY     100
```

Step 3. Add two properties to the SPEECH-ACT: a RELATION whose value is the meaning of the creatively altered constituent, and the feature-value pair ‘WORDPLAY yes’. For our example, this says that the utterance reflects some form of wordplay (which might, but need not, involve humor) involving a baby’s pacifier. Putting all the pieces of analysis together yields the following final TMR.

SPEECH-ACT-1
- THEME TAKE-PHOTOGRAPH-1
- **RELATION PACIFIER-FOR-BABY-1**
- **WORDPLAY yes**

TAKE-PHOTOGRAPH-1
- AGENT HUMAN-1
- THEME SET-1
- SPEED .9

SET-1
- MEMBER-TYPE PHOTOGRAPH
- CARDINALITY 100

Although we have implemented the above algorithm, we haven’t rigorously tested it on a corpus because variations on idioms—although entertaining and not unimportant for agent systems—are just not all that common, at least in available corpora (as observed by Langlotz, 2006, pp. 290–291, as well). However, work on idiomatic variation has broader implications. Idioms are just one type of construction, and constructions of all kinds are open to variation. So our approach to handling idiomatic variability by a combination of listing variable-inclusive senses and implementing lexicon-wide rules applies to the variability of nonidiomatic constructions as well.

6.2.4 Indirect Modification Computed Dynamically

As explained in section 4.1.6, most cases of indirect modification—for example, *responsible decision-making*, *rural poverty*—are best handled by lexical senses that anticipate, and then make explicit, the implied meaning. However, lexicon acquisition takes time and resources. This means that it is entirely possible that the lexicon will contain some sense(s) of a modifier but not every needed sense.

For example, say the lexicon contains only one sense of *responsible*, which expects the modified noun to refer to a HUMAN, as in *responsible adult* or *responsible dog owner*. If the LEIA encounters the input *responsible decision-making*, there will be an incongruity since RESPONSIBILITY-ATTRIBUTE can only apply to HUMANS. This will result in a low-scoring TMR that will serve as a flag for additional processing.

The good news about this state of affairs is that many instances of indirect modification share a similarity: they omit the reference to the agent of the action. So, a *vicious experiment* is an experiment whose agent(s) are vicious; an *honorable process* is a process whose agent(s) are honorable; and a *friendly experience* is an experience whose participant(s) are friendly. (We intentionally do not pursue a depth of semantic analysis that would distinguish between a person *behaving* viciously and a person having the general attribute of *being* vicious. It is early to pursue that grain size of description throughout a broad-coverage system.) In all such examples, the type of entity that was elided can be reconstructed with the help of knowledge recorded in the ontology.

- *Vicious experiment*: Although the attribute HOSTILITY (which is used to represent the meaning of *vicious*) can apply to any ANIMAL, the AGENT of EXPERIMENTATION must be HUMAN, so the elided entity in *vicious experiment* must be HUMAN.
- *Honorable process*: Although the AGENT of carrying out a PROCESS can be any ANIMAL, the attribute MORALITY (which is used to represent the meaning of *honorable*) applies exclusively to HUMANs, so the elided entity in *honorable process* must be HUMAN.
- *Friendly experience*: Since the property FRIENDLINESS can apply to any ANIMAL, and the AGENT of a LIVING-EVENT (used to represent the meaning of *experience*) can also be any ANIMAL, the elided entity in *friendly experience* is understood as ANIMAL.

Let us look in yet more detail at this ellipsis-reconstruction rule using the example *bloodthirsty chase*.

adj.$_{ModifiesAnimal}$ EVENT (e.g., bloodthirsty chase)

```
EVENT
   AGENT      ANIMAL-#1        ; an ontological instance of ANIMAL

ANIMAL-#1
   PROPERTY-IN-QUESTION     value-in-question
```

Rendered in plain English, this rule says, “If an adjective is supposed to modify an ANIMAL but it is being used to modify an EVENT, then introduce an ANIMAL into the meaning representation, make it the AGENT of that EVENT, and apply the modifier’s meaning to it.”

The TMR for *bloodthirsty chase* will convey that there is a CHASE event whose AGENT is an unspecified ANIMAL who wants to KILL (‘wanting to kill’ is the analysis of *bloodthirsty*). Although this representation reliably resolves the initially detected incongruity, it leaves a certain aspect of meaning—who is chasing whom—underspecified. This information may or may not be available in the context, as shown by the juxtaposition between (6.33) and (6.34).

(6.33) Lions regularly engage in bloodthirsty chases.

(6.34) The lion and rabbit were engaged in a bloodthirsty chase.

Strictly language-oriented reasoning can correctly analyze (6.33) but not (6.34). The first works nicely because the sense ‘X engages in Y’ maps the meaning of X to the AGENT slot of the EVENT indicated by Y. In essence, it interprets ‘X engages in EVENT-Y’ as ‘X

EVENT-Ys' (here, *Lions chase*). When our *bloodthirsty* conversion rule is applied, the correct TMR interpretation will be generated. *Whom* lions chase is not indicated in the context. If knowledge about whom lions typically chase is available in the ontology, the agent will look for it only if prompted to do so by some application-specific goal.

The problem with (6.34) is that world knowledge is needed to understand that the lion and the rabbit are not collaborating as coagents chasing somebody else. There is no linguistic clue suggesting that the set should be split up. Specific knowledge engineering in the domain of predation would be needed to enable this level of analysis.

So far, "insert an agent" is the only indirect-modification rule for which we have compelling evidence. But literature offers creative singletons, such as Ross McDonald's gem, *She rummaged in the purse and counted five reluctant tens onto the table*. Although one might assume that incongruous modifiers should always be applied to the nearest available referring expression, we think it premature to jump to that conclusion.

6.2.5 Recap of Treatable Types of Incongruities

- Metonymy: *The spiky hair* (i.e., the person with the spiky hair) *just smiled at me*.
- Preposition swapping: *He was absolved from* (rather than *of*) *responsibility*.
- Idiomatic creativity: *Don't put all your eggs in the stock market basket* (instead of *in one basket*).
- Indirect modification: *The lion engaged in a bloodthirsty chase* (the lion, not the chase, was bloodthirsty).

6.3 Addressing Underspecification

Underspecification is detected when the basic TMR includes a call to a procedural semantic routine that has not yet been run.[16] This section considers three sources of underspecification: nominal compounds that were not covered by lexical senses, missing values in events of change, and underspecified comparisons.

6.3.1 Nominal Compounds Not Covered by Lexical Senses

Section 4.5 described two classes of NN compounds that are fully treated during Basic Semantic Analysis thanks to lexical senses that anticipate them:

- Fixed, frequent compounds that are recorded as head entries: for example, *attorney general*, *drug trial*, *gas pedal*.
- Compounds containing one element that is fixed (and therefore serves to anchor the compound in the lexicon) and one element that is semantically constrained. For example, one sense of the noun *fishing* expects an NN structure composed of any type of FISH followed by the word *fishing*; it can analyze inputs such as *trout fishing* and *salmon fishing*.

If an NN compound does not belong to either of the above classes, then during Basic Semantic Analysis all combinations of meanings of N1 and N2 are linked by the most generic relation, called RELATION. These candidate interpretations are evaluated and scored both within the clause structure (during Basic Semantic Analysis) and with respect to coreference (during Basic Coreference Resolution). Deeper analysis of the candidate interpretations is provided by the four strategies described below.

Strategy 1. *Using ontological constructions*. Some combinations of concepts have a prototypical relationship. For example, TEMPORAL-UNIT + EVENT means that the event occurs at the given time, so *Tuesday flight* is analyzed as FLY-EVENT (TIME TUESDAY). Similar analyses apply to *morning meeting*, *weekend getaway*, and so on. Since both components of such constructions are concepts, the constructions cannot be anchored in the lexicon. Instead, they reside in a dedicated knowledge resource, the Ontological Construction Repository, that is consulted at this stage of processing.

Ontological constructions can be further categorized into those showing *unconnected* constraints and those showing *connected* constraints. The latter category indicates that the candidate meanings of one of the nouns must be tested as a property filler of the candidate meanings of the other noun.

Examples of Unconnected Constraints

If N_1 is TEMPORAL-UNIT and N_2 is EVENT, then the interpretation is N_2 (TIME N_1).
Tuesday flight: FLY-EVENT (TIME TUESDAY)

If N_1 is ANIMAL-DISEASE or ANIMAL-SYMPTOM and N_2 is HUMAN (not MEDICAL-ROLE), then the interpretation is N_2 (EXPERIENCER-OF N_1).
polio sufferer: HUMAN (EXPERIENCER-OF POLIO)

If N_1 is SOCIAL-ROLE and N_2 is SOCIAL-ROLE, then the interpretation is HUMAN (HAS-SOCIAL ROLE $N_{1,}$ N_2).
physician neighbor: HUMAN (HAS-SOCIAL-ROLE PHYSICIAN, NEIGHBOR)

If N_1 is FOODSTUFF and N_2 is PREPARED-FOOD, then the interpretation is N_2 (CONTAINS N_1).
papaya salad: SALAD (CONTAINS PAPAYA-FRUIT)

Examples of Connected Constraints

If N_1 is EVENT and N_2 is ANIMAL, and if N_2 is a *default* or *sem* AGENT of N_1, then N_2 (AGENT-OF N_1).
cleaning lady: HUMAN (GENDER female) (AGENT-OF CLEAN-EVENT)

If N_1 is EVENT and N_2 is an ontologically recorded *sem* or *default* INSTRUMENT of N_1, then N_1 (INSTRUMENT N_2).
cooking pot: COOK (INSTRUMENT POT-FOR-FOOD)

If N_1 is OBJECT and N_2 is a filler of the HAS-OBJECT-AS-PART slot of N_1, then N_2 (PART-OF-OBJECT N_1).
oven door: DOOR (PART-OF-OBJECT OVEN)

If N_1 is EVENT and N_2 is EVENT, and if N_2 is a filler of the HAS-EVENT-AS-PART slot of N_1, then N_2 (PART-OF-EVENT N_1).
ballet intermission: INTERMISSION (PART-OF-EVENT BALLET)

If N_2 is EVENT and N_1 is a *default* or *sem* THEME of N_2, then N_2 (THEME N_1).
photo exhibition: EXHIBIT (THEME PHOTOGRAPH)

If N_2 is described in the lexicon as HUMAN (AGENT-OF EVENT-X) (e.g., *teacher* is HUMAN (AGENT-OF TEACH)) and N_1 is a *default* or *sem* THEME of X (e.g., PHYSICS is a *sem* filler for the THEME of TEACH), then the NN analysis is HUMAN (AGENT-OF X (THEME N_1)).
physics teacher: HUMAN (AGENT-OF TEACH (THEME PHYSICS))
home inspector: HUMAN (AGENT-OF INSPECT (THEME PRIVATE-RESIDENCE))
stock holder: HUMAN (AGENT-OF OWN (THEME STOCK-FINANCIAL))

If N_1 is PHYSICAL-OBJECT and N_2 is PHYSICAL-OBJECT and N_1 is a *default* or *sem* filler of the MADE-OF slot of N_2, then N_2 (MADE-OF N_1).
denim skirt: SKIRT (MADE-OF DENIM)

If N_2 is PROPERTY and N_1 is a legal filler of the DOMAIN of N_2, then N_2 (DOMAIN N_1).
ceiling height: HEIGHT (DOMAIN CEILING)

These constructions not only offer high-confidence analyses of the semantic relation inferred by the NN but also can help to disambiguate the component nouns. For example, although *papaya* can mean PAPAYA-FRUIT or PAPAYA-TREE, in *papaya salad* it can be disambiguated to PAPAYA-FRUIT in order to match the associated construction above.

It is important to emphasize that these rules seek only high-confidence ontological relations, defined using the *default* and *sem* facets of the ontology. If a compound is semantically idiosyncratic enough that it would fit only the *relaxable-to* facet of recorded ontological constraints, then it is not handled at this point in analysis. For example, although the LEIA would be able to analyze the clausal input *He teaches hooliganism* using the *relaxable-to* facet of the THEME of TEACH (which permits anything to be taught), it would not analyze the corresponding NN compound, *hooliganism teacher*, using the NN rule that covers *science teacher* or *math teacher* because the NN analysis rule is more constrained. It requires N1 to satisfy the *default* or *sem* THEME of an EVENT—in this case, TEACH. So, when analyzing *hooliganism teacher*, the agent will leave the originally posited generic RELATION between the analyses of N1 (*hooliganism*) and N2 (*teacher*).

One might ask why we record any NN constructions in the lexicon since they could all, in principle, be recorded as more generic ontological constructions. For example, rather than record the construction "FISH fishing" in the lexical sense fishing-n1, we could record the construction "FISH FISH-EVENT" in the Ontological Construction Repository. In the latter scenario, when the system encountered the word *fishing*, it would recognize it as FISH-EVENT, resulting in the same analysis. The reason for the split has primarily to do with (a) convenience for acquirers and (b) the desire to post analyses at the earliest processing stage possible. If a given word, like *fishing*, is often used in compounds, and if it has no synonyms (or only a few synonyms that can readily be listed in its *synonyms* field of the lexical sense), then it is simpler and faster for the acquirer to record the information in the lexicon under *fishing* rather than switch to the Ontological Construction Repository and seek concept-level generalizations. Moreover, when the compound is recorded in the lexicon, it will be analyzed early, during Basic Semantic Analysis.

Strategy 2. *Recognizing NN paraphrases of N + PP constructions*. In some cases, a nominal compound is a paraphrase of an N + PP construction that is already recorded in the lexicon.[17] For example, one lexical sense of the noun *chain* expects an optional PP headed by *of*, and it expects the object of the preposition to mean a STORE or RESTAURANT. This covers inputs like *the chain of McDonald's restaurants*, whose TMR will be as follows:

```
SET-1
    MEMBER-TYPE    RESTAURANT-1
    CARDINALITY    > 1

RESTAURANT-1
    HAS-NAME       'McDonald's'
```

Recording the meanings of typical N + PP constructions in the lexicon is done as a matter of course, since it assists with the difficult challenge of disambiguating prepositions. A pre-runtime lexicon sweep translates these N + PP constructions into corresponding NN constructions. Continuing with our example, this automatic conversion generates the NN construction STORE/RESTAURANT + *chain*, which covers the input *the McDonald's restaurant chain*, generating the same TMR as shown above.

The system computes these NN constructions prior to runtime, rather than storing them permanently in the lexicon, so that the constructions match the inventory of N + PP lexical senses, even if the form, scope, or inventory of the latter changes. If, for example, a new societal trend developed by which churches and schools could be organized into chains—giving rise to turns of phrase like *chain of churches* and *chain of elementary schools*—then knowledge acquirers would need to expand the lexical sense for *chain* + PP by allowing the object of the preposition to mean not only STORE and RESTAURANT but also CHURCH and SCHOOL.

Strategy 3. *Detecting a property-based relationship in the ontology.* In some cases, the meanings of the nouns in an NN compound are directly linked by some ontological property. For example, *hospital procedure* is lexically ambiguous since *procedure* can mean either HUMAN-EVENT (i.e., any event carried out by a person that involves particular subevents in a particular order) or MEDICAL-PROCEDURE. But since the ontology contains the description MEDICAL-PROCEDURE (LOCATION HOSPITAL), this analysis simultaneously disambiguates *procedure* and selects the correct relation between the concepts.

Strategy 4. *Identifying a short (but not direct) property-based path in the ontology.* In other cases, the meanings of the nouns in the NN are ontologically connected, but along a path that involves multiple properties. This is true, for example, of *hospital physician*. The interpretation with the shortest ontological path is PHYSICIAN (AGENT-OF MEDICAL-PROCEDURE (LOCATION HOSPITAL)). But there are actually a lot of ways in which HOSPITAL and PHYSICIAN could be linked using an ontological search. For example, since a hospital is a PLACE and a physician is a HUMAN, and since HUMANS go to PLACES, then the physician could be the AGENT of a MOTION-EVENT whose DESTINATION is HOSPITAL. Similarly, since a hospital is a PHYSICAL-OBJECT, and since a PHYSICIAN is a HUMAN, and since any HUMAN can DRAW practically any PHYSICAL-OBJECT, then the PHYSICIAN could be the AGENT of a DRAW event whose THEME is HOSPITAL. The list of such analyses could go on and on. But the point is this: the use of an essentially elliptical structure like an NN compound requires that the speaker give the listener a fighting chance of figuring out what's going on. Using the compound *hospital physician* to mean a hospital that a physician is sketching is simply not plausible. That lack of plausibility is nicely captured by ontological distance metrics. The ontological path that goes from PHYSICIAN all the way up to HUMAN and from HOSPITAL all the way up to PHYSICAL-OBJECT is much longer than the path of our preferred reading.

Now, one could argue that PHYSICIAN (LOCATION HOSPITAL) is not the most semantically precise analysis possible, which is true. If we wanted a better analysis, we could create a construction that expected LOCATION followed by WORK-ROLE (which is a subclass of SOCIAL-ROLE), which would output meaning representations like the one below:

HUMAN
  HAS-WORK-ROLE $^N_2$
  AGENT-OF WORK-ACTIVITY-1

WORK-ACTIVITY-1
  LOCATION $^N_1$

This construction would precisely analyze inputs like *hospital physician*, *bakery chef*, and *college teacher* as people fulfilling the listed work roles at the listed places. The point is that the agent will only attempt unconstrained ontology-based reasoning if there is no recorded construction to provide a more precise analysis.

The four analysis strategies for NN compounds just described, in addition to the lexically based ones described as part of Basic Semantic Analysis, still do not exhaust the analysis space for NNs. If an NN has not yet been treated, then the generic RELATION posited during Basic Semantic Analysis will remain, and the agent's last chance to generate a more specific analysis will be during Situational Reasoning.

So far we have concentrated on the analysis of two-noun compounds, but the approach can be extended to treating larger compounds. A LEIA's first step in treating any compound containing three or more nouns occurs much earlier than this. During Pre-Semantic Integration, the LEIA reambiguates the syntactic parser's bracketing of the internal structure of compounds containing three or more nouns. Then, during the various stages of semantic analysis, it seeks out islands of highest confidence among pairs of nouns and finally combines those partial analyses.

Consider the compound *ceiling height estimation*. The candidate bracketings are [[ceiling height] estimation] and [ceiling [height estimation]]. The analysis of the first bracketing will receive a very high score using two rules introduced above.

[ceiling height]

Rule: If N_2 is PROPERTY and N_1 is a legal filler of the DOMAIN of N_2, then N_2 (DOMAIN N_1).

Here: HEIGHT is a PROPERTY and CEILING is a legal filler of it, so HEIGHT (DOMAIN CEILING).

[[ceiling height] estimation]

Rule: If N_2 is EVENT and N_1 is a *default* or *sem* THEME of N_2, then N_2 (THEME N_1).

Here: ESTIMATE is an EVENT and HEIGHT is a *sem* THEME of ESTIMATE, so ESTIMATE (THEME HEIGHT).

Putting these two analyses together, the TMR for *ceiling height estimation* is

```
ESTIMATE-1
    THEME       HEIGHT-1

HEIGHT-1
    DOMAIN      CEILING-1
```

By contrast, the analysis of the second bracketing analysis, [ceiling [height estimation]], will receive a much lower score because there is no high-confidence rule to combine CEILING with ESTIMATE (CEILING is not a *sem* or *default* filler of the THEME case role of the event ESTIMATE).

Although it would be unwise to underestimate the potential complexity of processing large compounds, it is reasonable to assume that multinoun compounds require an extension to the algorithm presented here rather than a fundamentally different approach.

An important question is, *How well do our NN analysis strategies work*? For that answer, see section 9.2.1.

6.3.2 Missing Values in Events of Change

Events of change are events that describe a change in the value of some property: for example, *speed up*, *lose weight*, *increase*. Descriptions of events of change often convey two property values from which a third can be calculated. It is likely that people actually do this calculation, at least if the information is important to them, so LEIAs should as well. Consider some examples from the *Wall Street Journal* (1987–1989; hereafter WSJ), all of which present two values from which a third can be calculated:

(6.35) In 1985, 3.9 million women were enrolled in four-year schools. Their number increased by 49,000 in 1986. $_{\text{(WSJ)}}$

(6.36) Interco shot up 4 to 71¾ after a Delaware judge barred its poison pill defense against the Rales group's hostile \$74 offer. $_{\text{(WSJ)}}$

(6.37) An index of longterm Treasury bonds compiled by Shearson Lehman Brothers Inc. rose by 0.79 point to 1262.85. $_{\text{(WSJ)}}$

At this stage of processing, the LEIA can carry out these calculations and save them to memory along with the stated information. The functions for calculating are listed as meaning procedures in the lexical senses for the words indicating the events of change, such as *increase*, *shoot up*, and *rise* from our examples. (See McShane, Nirenburg, & Beale, 2008, for a more in-depth treatment of events of change).

6.3.3 Ungrounded and Underspecified Comparisons

Here we present the microtheory of ungrounded and underspecified comparisons as a whole, even though different classes of comparisons are analyzed to different degrees across stages 3–6 of NLU (3: Basic Semantic Analysis, 4: Coreference Resolution, 5: Extended Semantic Analysis, and 6: Situational Reasoning).

The first thing to say is that this microtheory is at a less advanced stage of development than some of our other ones. Although corpus analysis has informed it, we have not yet rigorously vetted it against a corpus. Still, this microtheory reflects a nontrivial modeling effort and nicely illustrates the distribution of labor across the modules of semantic and pragmatic analysis. Specifically, it underscores (a) that different types of heuristics become available at different stages of processing and (b) that an agent can decide how deeply to pursue the intended meaning of an input. For example, if *My car is better than that one* is used as a boast, then it doesn't matter which particular properties the speaker has in mind. However, if the speaker is advising the interlocutor about the latter's upcoming car purchase, then the properties in question absolutely *do* matter. Does the car handle well in snow? Have heated seats? An above-average extended warranty? The point is that, to function like people, agents need to judge how deeply to analyze inputs on the basis of their current interests, tasks, and goals.

Our current microtheory of comparatives classifies them according to two parameters: how/where the compared entities are presented, and how precise the comparison is.[18] We first define the value sets for these properties without examples and then illustrate their combinations with examples.

Values for *how/where the compared entities are presented*

1. They are both included in a comparative construction that is recorded in the lexicon.
2. The comparison involves a single entity, in what we call an *inward-looking* comparison (in contrast to an *outward-looking* comparison, in which two different entities are compared). So far, the most frequently encountered inward-looking comparisons involve either the change of an entity's property value over time or a counterfactual.
3. The compared-with entity is either located elsewhere in the linguistic context (i.e., not in a construction) or is not available in the linguistic context at all. These eventualities are combined because the agent engages in the same search process either way. Notably, searching for the point of comparison invokes some of the same features as coreference resolution: semantic affinity (comparability), text distance (the point of comparison should not be too far away), and the understanding that the point of comparison might not be in the linguistic context at all.

Values for *how precise the comparison is*

1. Specific: A specific property is referred to, such as INTELLIGENCE or HEIGHT.
2. Vague: The comparison is expressed as a value of evaluative modality (e.g., *better*, *worse*) or as a simile (*Your smile is like a moonbeam*).
3. Vague with an explanation: Either type of vague comparison mentioned above can be followed by an explanation of what is meant. Semantically, the explanation can either identify the particular property value(s) in question (which is quite useful), or it can supplement the vague comparison with an equally vague explanation (e.g., the comparison can be followed by a metaphor: *Your smile is like a moonbeam: it lights up my heart*). In practical terms, the explanation can be easy to detect because it participates in a construction with the comparison, or it can be difficult to detect because, in principle, any text that follows a vague comparison may or may not explain it.

The permutations of these feature values result in the nine classes of comparatives shown in table 6.2. The table includes an indication of which modules can be invoked to analyze associated examples. We say "*can be* invoked" because the agent can, at any time in NLU, decide to forgo deeper analysis of an input. In operational terms, this means that it can choose not to launch a procedural semantic routine that is recorded in the nascent TMR.[19]

Table 6.2
Classes of comparative examples and when they are treated during NLU

Class	Target	Precision	Basic Semantic Analysis	Basic Coreference Resolution	Extended Semantic Analysis	Situational Reasoning
1	The compared entities are in a comparative construction.	Specific	✓			
2		Vague	✓			✓
3		Vague + explanation	✓	✓	✓	✓
4	The comparison involves a single entity (it is inward-looking).	Specific	✓			
5		Vague	✓			✓
6		Vague + explanation	✓	✓	✓	✓
7	The point of comparison is located elsewhere in the text or is not available in the language context.	Specific	✓	✓	✓	✓
8		Vague	✓	✓	✓	✓
9		Vague + explanation	✓	✓	✓	✓

We will now work through each of the nine classes of comparatives, providing further details and examples.

Class 1. *The compared entities are in a comparative construction, and the comparison is precise*. Examples of this type are fully analyzed during Basic Semantic Analysis thanks to constructions recorded in the lexicon. In some cases, these constructions include calls to procedural semantic routines to compose the meanings of the many variable elements on the fly (something that is quite common for constructions overall). Two of the many comparative constructions recorded in the lexicon are shown below. Both are recorded as senses of their only invariable word: *than*. In (6.38), the property referred to is INTELLIGENCE, whereas in (6.39), it is AESTHETIC-ATTRIBUTE.

(6.38) [Subj Verb Comparative *than* NP Auxiliary/Modal/Copula]
Animals are smarter than we are. (COCA)

(6.39) [Subj Verb Comparative *than* NP]
"You think she's prettier than Mama?" (COCA)

The fact that the basic, construction-based proposition in (6.39) is scoped over by both modality ('you think') and an interrogative does not require multiple constructions. Instead, these proposition-level enhancements are handled by general rules.

Class 2. *The compared entities are in a comparative construction, and the comparison is vague.*

(6.40) [Subj Verb Comparative *than* NP]
I don't sleep because my real life is better than my dreams. $_{\text{(COCA)}}$

(6.41) [Subj Verb Modifier *but* Subj Verb Comparative]
A southerly breeze is adequate but a west wind is better. $_{\text{(COCA)}}$

(6.42) [Subj *be like* NP]
Tell Jack Hanna his life is like a zoo and he'll say, "Thanks!" $_{\text{(COCA)}}$

The difference in the TMRs resulting from class 1 and class 2 is that the TMRs for class 2 examples include a call to a procedural semantic routine that can, if run, attempt to concretize the vagueness. In some cases, the procedural semantic routine is recorded in the lexical sense for the construction itself. For example, "NP *be like* NP" is always vague in that it does not specify which properties and values are implied when comparing two nominals. In other cases, the procedural semantic routine is attached to the lexical description of a vague comparative word used in the construction, such as *better* or *worse*. No matter the source of these procedural semantic routines, they basically say, "This meaning is vague. It may or may not be important/useful to concretize it. This determination cannot be made until Situational Reasoning, when the agent knows its task and goals. Therefore, carry the call to this procedural semantic routine in the TMR until that stage. At that stage, determine whether a more precise interpretation is needed. If it is, use all available heuristic evidence to try to compute it. If that fails, ask the interlocutor for help."

In most cases, vague expressions are meant to be vague and their interpretation can be left as such. In addition, in many cases even the speaker/writer might be hard-pressed to come up with the specific connotation. For example, comparing the zookeeper Jack Hanna's life to a zoo was likely as much a witticism as anything else. However, the very inclusion of the call to the procedural semantic routine in the TMR carries information: it asserts that the agent is aware that the utterance and its interpretation are vague.

Class 3. *The compared entities are in a comparative construction, and the comparison is vague with an explanation*. As described earlier, it can be tricky to determine whether the text that follows a vague comparison actually concretizes it. Currently, the only way an agent can determine this is if the explanation participates in a construction with the comparison, as in the following examples.

(6.43) [Subj_1 Verb *like* NP {, : ;—} $\text{Clause}_{\text{Subj1/Not-Comparative}}$]
A career is like a flower; it blooms and grows.

(6.44) [Subj Verb *like* NP {,—} Modifier(s)]
… When he's on the basketball court, he moves like a rabbit, all quick grace and long haunches. $_{\text{(COCA)}}$

According to the construction used for (6.43),

- The "Subj Verb *like* NP" clause must be followed by another clause, but that latter clause cannot also be of the form "Subj Verb *like* NP." (This excludes, e.g., *A career is like a flower; it is like a rose*).
- The clauses must be joined by a non-sentential punctuation mark. This requirement will result in some losses (an explanation could be presented as a new sentence), but we hypothesize that those losses are justified by the reduction in false positives.
- The subject of the second clause must be coreferential with the subject of the first. (This excludes, e.g., *A career is like a flower; I am happy about that*.)

The construction used for (6.44), for its part, requires that the agent be able to identify any nonclausal syntactic entitles that semantically serve as modifiers. As we see, this can be hard, since *all quick grace* and *long haunches* serve as modifiers although they do not have the most typical form of a modifier (adjective, adverb, or prepositional phrase). For all constructions in this class (i.e., vague with an explanation), the explanation is semantically attached to the comparison in the TMR using the property EXPLAINS-COMPARISON. This is specified, of course, in the sem-struc of the comparative construction recorded in the lexicon.

As corpora show, people explain their comparisons quite frequently, which makes recording these kinds of constructions worth the effort. The constructions posited above will clearly overreach, and more knowledge engineering is needed to identify the sweet spot between coverage and precision.

It is noteworthy that, even given an optimal inventory of constructions, the resulting analyses can be unenlightening because of the actual language input. For example, (6.43) uses metaphorical language to describe the vague comparison—not a whole lot of help for automatic reasoning. However, it is still useful for agents to recognize that the text attempted to explain the comparison.

Class 4. *The comparison involves a single entity (it is inward-looking), and the comparison is specific*. The key to processing inward-looking comparisons is being able to automatically detect that the comparison is, in fact, inward-looking—that is, that no external point of comparison need be sought. So far, we have identified three semantic clues for inward-looking comparisons: the use of noncausative CHANGE-EVENTS (6.45), causative CHANGE-EVENTS (6.46), and counterfactuals (6.47).

(6.45) [Subj *gets/grows* Comparative]
a. I noticed, all summer long, I was getting healthier. (COCA)
b. That patch was moving. And it was getting larger. (COCA)
c. Jack felt his grin get bigger. (COCA)
d. The sobs grow louder. (COCA)
e. The centipede grew bolder. (COCA)

(6.46) [Subj *gets/makes* Direct-Object Comparative]

a. You should make the sanctions tougher. (COCA)

b. Also, government subsidies to get industry greener are short term when industry prefers long term commitment. (COCA)

(6.47) [Subj *could (not) / could (not) have* Verb Comparative]

a. "The replacement process could have been easier too," Swift says. (COCA)

b. Surely his heart couldn't beat any faster. (COCA)

As a reminder, CHANGE-EVENTs are events that compare the value of a particular property in the event's PRECONDITION and EFFECT slots. They are realized in language using a very large inventory of words and phrases: *increase*, *decrease*, *lose confidence*, *speed up*, *grow taller*, and so on. The constructions noted in (6.45) and (6.46) involve CHANGE-EVENTs. Their semantic descriptions (which include a procedural semantic routine) allow the agent to generate TMRs like the following—using the example *I was getting healthier*.

```
CHANGE-EVENT-1
    PRECONDITION    HEALTH-ATTRIBUTE-1
    EFFECT          HEALTH-ATTRIBUTE-2
    TIME            <find-anchor-time
    PHASE           continue              ; a shorthand for the full ASPECT frame

HEALTH-ATTRIBUTE-1
    DOMAIN          HUMAN-1
    RANGE           <HEALTH-ATTRIBUTE-2.RANGE

HEALTH-ATTRIBUTE-2
    DOMAIN          HUMAN-1
    RANGE           >HEALTH-ATTRIBUTE-1.RANGE
```

This TMR says that the value of the person's HEALTH-ATTRIBUTE is lower in the PRECONDITION of the CHANGE-EVENT than in its EFFECT; and that is, in fact, what *get healthier* means.

As regards counterfactuals, like those in (6.47), they, too, can be treated by lexicalized constructions. However, since counterfactuals have not to date been a priority of our R&D, we will say nothing further about what the associated meaning representation should look like.

All such lexicalized constructions can be fully analyzed as part of Basic Semantic Analysis.

Class 5. *The comparison involves a single entity (it is inward-looking), and the comparison is vague*. We have already explained how inward-looking comparisons are treated, and we have already explained how vague comparison words "carry along" calls to procedural semantic routines in their TMRs, in case the agent decides to try to concretize the basic interpretation. Those two functionalities need only be combined to treat this class of comparatives, illustrated by the following examples.

(6.48) [An inward-looking, noncausative CHANGE-EVENT with a vague comparison]
After the experiment the waking dreams got worse. (COCA)

(6.49) [A counterfactual with a vague comparison]
It wasn't the best start for the day but it could have been worse. (COCA)

Class 6. *The comparison involves a single entity (it is inward-looking), and the comparison is vague with an explanation*. Like the class above, this one uses already explained functions. We have not come across any examples of this class, but they are easily invented, as the following modification to (6.48) shows.

(6.50) After the experiment the waking dreams got worse: they changed into nightmares.

As discussed earlier, explanations can be detected using various types of constructions. For (6.50), the construction is largely similar to one posited earlier in that it requires the first clause to be followed by non-sentential punctuation, and the subjects of the two clauses must be coreferential.

Vague counterfactuals, by contrast, require a different type of explanatory construction, since counterfactuals are often explained by more counterfactuals, as our invented expansions of (6.49) below show.

(6.51) It wasn't the best start for the day but it could have been worse: my car could have broken down <the train could have been late; my boss could have been in one of his nasty moods>.

Although it would require human-level knowledge and reasoning to understand why one's car breaking down <the train being late; one's boss being in a nasty mood> would make for a bad day, the agent does not need this to hypothesize that the continuation explains the vague counterfactual. What it needs is a construction that expects a vague inward-looking counterfactual to be followed by a non-sentential punctuation mark and then a precise counterfactual. Will this rule always identify only explanations? Probably not. But it serves as a foothold for further work on this microtheory.

Classes 7–9. *The point of comparison is located elsewhere in the text or is not available in the language context at all, and the comparison is either specific (class 7), vague (class 8), or vague with an explanation (class 9).* We group these classes together because this part of the microtheory is, at the time of writing, underdeveloped. Part of the work belongs to the stage of analysis we are focusing on here (Extended Semantic Analysis), part must wait for Situational Reasoning, and much depends on difficult aspects of coreference resolution having worked correctly during the last stage of processing. In short, a lot is required to make the associated examples work.

What differentiates this class from the others is that the point of comparison might be anywhere in the linguistic context or not available at all. The salient features that differentiate examples are as follows:

- How the compared entity is realized: as a full nominal (*this book is better*), a pronoun (*it is better*), or an elliptical expression (e.g., *the second __ is better*).
- How the point of comparison is realized: as a full nominal with the same head as what it is compared with, as a full nominal with a different head from what it is compared with, as a pronoun, as an elliptical expression, or not at all (i.e., it is absent from the linguistic context).
- If applicable, the distance between the linguistically overt compared entities: that is, the point of comparison can be the most proximate preceding nominal, the next one back, and so on.

Below are some examples illustrating different combinations of the abovementioned parameter values.

(6.52) Your force field is good but my teleporting is better. (COCA)

(6.53) Whatever your secret was, you have to agree, mine is better. (COCA)

(6.54) I often tell my clients that the state of mind they want when negotiating or navigating conflict is curiosity, not certainty. If you can manage to be curious when things get tough, that curiosity will be your best friend. Curiosity is better. It's the mode that opens us to discovery. (COCA)

(6.55) I like the sweet potato idea! Way tastier than store bought white potato chips. (COCA)

(6.56) Let's see if we can find what he was reaching for. Here. My reach is better. (COCA)

(6.57) I met a guy last night who brought 80 pounds of screenplays out here in his suitcase. But he didn't bring his skis. I think my gambit is better. (COCA)

(6.58) And you think this is easier?!

(6.59) 50% tastier!

In the last two examples, there is no linguistic point of comparison. We can imagine the first being said as two people struggle to carry a sofa up a skinny and winding stairway, having just argued about various strategies. The last example is typical of advertising: the comparison is so vague that there is nothing legally binding about it.

Having delved deep into this model of processing comparatives, let us now take a step back to the big picture in order to more fully motivate why we present this cross-modular microtheory as part of this module of Extended Semantic Analysis. During this stage of processing, if a LEIA considers it worthwhile to attempt to concretize underspecified comparisons, it can apply additional resolution functions. Those functions still rely exclusively on the agent's broad-coverage knowledge bases. We have identified three such types of reasoning that can be applied at this stage. Working out the full microtheory, however, remains on agenda.

1. *The LEIA can semantically reason about whether the assertion following a vague comparison explains it*. So far, we have prepared the agent to detect explanations for

vague comparisons using lexico-syntactic constructions. However, (a) those constructions might overreach, identifying a text segment as an explanation when it is actually not, and (b) they do not cover all eventualities. Ontologically grounded reasoning could weigh in on this determination. For example, the second sentence in (6.60) explains the vague, inward-looking comparison, but there is no text-level clue to point that out. One needs to know that road salt corrodes car finishes—information that is entirely reasonable to expect in an ontology with moderate coverage of car-related information.

(6.60) Come pring, my car looked a lot worse. Road salt is a bear.

2. *The LEIA can attempt to concretize vague comparisons based on ontological generalizations*. Vague comparisons often rely on people's knowledge of the salient aspects of different entities.

(6.61) Her eyes were like a sunrise. (beautiful; bright)

(6.62) She ran like a deer. (gracefully)

(6.63) He's like a regular giraffe! (very tall)

Vague comparisons are like ellipsis: when using one, the speaker has to give the hearer a fair chance of interpreting it correctly. This means relying on the expectations of a shared ontology, including the canonically distinguishing features of entities.

We can prepare agents to reason about saliency by manually indicating the most salient property values for each concept (which might, by the way, differ in some cases across cultures) and/or by having agents dynamically learn this information from text corpora. For example, a sentence like *She skipped barefoot across the stepping stones as graceful as a deer running …* $_{\text{(COCA)}}$ suggests that a salient property of deer is their gracefulness.

This stage of Extended Semantic Analysis is the appropriate place for carrying out salience-based reasoning about vague comparisons because (a) this extra reasoning will not always be necessary (and, therefore, should not be a part of Basic Semantic Analysis) and (b) to the extent that the ontology indicates the salient properties of entities, this reasoning can be carried out for texts in any domain (i.e., it does not rely on the situational awareness that becomes available only later in the NLU process).

3. *The LEIA can attempt to identify the points of comparison for classes 7–9*. As explained above, this involves (a) leveraging previously established coreference relations, (b) reasoning about which entities in a context are semantically comparable, (c) factoring in the text distance between mentioned entities (since there might be multiple entities in the preceding context that must be considered as candidate targets of the comparison), and (d) leaving open the possibility that the point of comparison is not in the text at all.

6.3.4 Recap of Treatable Types of Underspecification

- Many nominal compounds that are not covered by the lexicon are covered by ontological constructions recorded in the Ontological Construction Repository: TEMPORAL-UNIT + EVENT, as in *night flight*.
- Missing values in events of change can be calculated and recorded: *Their earnings grew from $10,000 to $15,000* [change: + $5,000].
- Some types of underspecified comparisons can be made explicit: *John got stronger* (than he was before).

6.4 Incorporating Fragments into the Discourse Meaning

Since LEIAs understand inputs incrementally, they are routinely processing midsentence fragments. Those are not the kinds of fragments we are talking about here.[20] We are talking about fragments that remain nonpropositional when the end-of-sentence indicator is reached. A LEIA's basic approach to analyzing fragments is as follows:[21]

1. Generate whatever semantic interpretation is possible from the fragment itself.
2. Detect the as-yet unfilled needs in that semantic interpretation.
3. Attempt to fill those needs using all available heuristics.
4. Once those needs are filled, verify that the original semantic interpretation is valid. Otherwise, amend it.

This process is best illustrated using an example:

(6.64) "My knee was operated on. Twice." "When?" "In 2014."

The TMR for the first utterance, *My knee was operated on*, is

SURGERY-1		
	THEME	KNEE-1
	TIME	<*find-anchor-time*
HUMAN-1		
	HAS-OBJECT-AS-PART	KNEE-1
	COREF	*identify-speaker*
KNEE-1		
	THEME-OF	SURGERY-1
	PART-OF-OBJECT	HUMAN-1

When the LEIA encounters the input *Twice*, which occurs as an independent sentence, it will find only one lexical sense, which describes this adverb as a verbal modifier that adds the feature ASPECT (ITERATION 2) to the EVENT it modifies. Since no EVENT is available in the local dependency structure, an EVENT instance is posited in the meaning representation without any associated text string.

ASPECT-1	
ITERATION	2
SCOPE	EVENT-1
EVENT-1	
SCOPE-OF	ASPECT-1
textstring	*none*

The feature-value pair *textstring none* triggers the search for a coreferential EVENT (in the same way as *find-anchor-time* triggers the search for the time of speech). The algorithm is currently quite simple: it identifies the main EVENT (i.e., the event to which any subordinate or relative clauses would attach) in the previous clause. The reason why this simple algorithm works pretty well is that understanding sentence fragments would impose too great a cognitive load on the listener if the intended link to the rest of the context were not easily recoverable.

In our context, the search for the most recent event instance will identify SURGERY-1 as a candidate, leading to the following TMR for *My knee was operated on. Twice.*

SURGERY-1	
THEME	KNEE-1
TIME	<*find-anchor-time*
COREF-OF	EVENT-1
HUMAN-1	
HAS-OBJECT-AS-PART	KNEE-1
COREF	*identify-speaker*
KNEE-1	
THEME-OF	SURGERY-1
PART-OF-OBJECT	HUMAN-1
ASPECT-1	
ITERATION	2
SCOPE	EVENT-1
EVENT-1	
SCOPE-OF	ASPECT-1
textstring	*none*
COREF	SURGERY-1

Although the LEIA does not need to pretty-print these results to effectively reason with them, it is easier for people to understand the TMR if we remove the coreference slots and replace EVENT-1 with SURGERY-1. This yields the following TMR:

SURGERY-1	
THEME	KNEE-1
TIME	<*find-anchor-time*
SCOPE-OF	ASPECT-1
HUMAN-1	
HAS-OBJECT-AS-PART	KNEE-1
COREF	*identify-speaker*
KNEE-1	
THEME-OF	SURGERY-1
PART-OF-OBJECT	HUMAN-1
ASPECT-1	
ITERATION	2
SCOPE	SURGERY-1

The next utterance is *When?*, another fragment. For each question word, the lexicon contains a sense that expects the word to be used as a fragmentary utterance. This reflects expectation-driven knowledge engineering—that is, preparing the system for what it actually will encounter, not only what grammar books say is the most typical.

For certain question words (e.g., *When? Where? How?*) the semantic representation (i.e., the sem-struc zone of the lexical sense) posits an EVENT that is flagged with coreference needs like we just saw for *twice*. For other question words (e.g., *Who? How many?*) the semantic representation posits an OBJECT that is similarly flagged for coreference resolution. Returning to our example, for the independent utterance *When?* the procedure seeks out the most recent main EVENT, just like our last meaning procedure did. Formally, the initial, sentence-level meaning representation for *When?* is

REQUEST-INFO-1	
THEME	EVENT-1.TIME
EVENT-1	
textstring	*none*
THEME-OF	REQUEST-INFO-1

When the coreference is resolved and the structure is pretty-printed, it looks like this.

REQUEST-INFO-1	
THEME	SURGERY-1.TIME

The final fragment in our example is *In 2014*. This is a bit more challenging because the preposition *in* is highly polysemous. One rule of thumb used by LEIAs when resolving polysemous words is to select the interpretation that matches the narrowest selectional constraints. In this case, the LEIA selects *in-prep10* (the tenth prepositional sense of *in*) because that sense asserts that the object of the preposition must refer to a MONTH, YEAR, DECADE, or CENTURY. The preprocessor has already provided the knowledge that 2014 is a date, which the LEIA translates (during CoreNLP-to-LEIA tag conversion) into the appropriate ontological subtree that holds all date-related concepts. Since this constraint is met, the LEIA can confidently disambiguate *in* as the property TIME applied to some EVENT. As before, the TMR for *In 2014* refers to an as-yet unresolved EVENT.

YEAR-1
    textpointer   2014
    ABSOLUTE-TIME   (YEAR 2014)
    TIME-OF   EVENT-1

EVENT-1
    textpointer   *none*
    TIME   YEAR-1

When the event is contextually grounded—that is, when it is linked to the SURGERY in question—the meaning representation looks as follows:

YEAR-1
    textpointer   2014
    ABSOLUTE-TIME   (YEAR 2014)
    TIME-OF   SURGERY-1

Putting all these pieces together, we can see what the agent learns from the dialog, *"My knee was operated on. Twice." "When?" "In 2014."*

SURGERY-1
    THEME   KNEE-1
    TIME   YEAR-1
    SCOPE-OF   ASPECT-1

HUMAN-1
    HAS-OBJECT-AS-PART   KNEE-1
    COREF   *identify-speaker*

KNEE-1
    THEME-OF   SURGERY-1
    PART-OF-OBJECT   HUMAN-1

ASPECT-1
    ITERATION   2
    SCOPE   SURGERY-1

YEAR-1
    ABSOLUTE-TIME   (YEAR 2014)

Of course, the entire dialog history (the series of TMRs) is also available to the agent, but the most important thing is what it stores to memory, which is the information shown above.

This example was useful in showing (a) how lexical senses can posit concepts that are not directly attested in the text and (b) how coreference resolution can be carried out with the help of associated meaning procedures. The results of the coreference ground the meaning of the fragment in the context.

There are many variations on this theme, which are treated as a matter of course using constructions in the lexicon that have associated procedural semantic routines.[22] For example, the lexical sense that covers questions of the form "Who Verb-Phrase?" instantiates a REQUEST-INFO frame whose THEME is the AGENT, EXPERIENCER, or BENEFICIARY of the given EVENT. The agent dynamically determines which case role is correct based on the meaning of the EVENT. So the TMRs for the following questions are

Who had surgery?

REQUEST-INFO-1
    THEME   SURGERY-1.EXPERIENCER

SURGERY-1
    TIME   *<find-anchor-time*

Who wrote that book?

REQUEST-INFO-2
    THEME   WRITE-1.AGENT

WRITE-1
    THEME   BOOK-1
    TIME   *<find-anchor-time*

Who got a prize?

REQUEST-INFO-3
    THEME   AWARD-EVENT-1.BENEFICIARY

AWARD-EVENT-1
    THEME   AWARD-1
    TIME   *<find-anchor-time*

The format of the slot filler—EVENT.CASE-ROLE—shows what is expected in the upcoming context. In the first example, the LEIA is waiting for an utterance whose meaning is compatible with an EXPERIENCER—that is, it must be an ANIMAL (which includes HUMANS). So if the next input is an ANIMAL, it will interpret it as the filler of that CASE-ROLE.

Note that the use of fragments is not limited to dialogs, since a speaker can answer his own question and the resulting analysis will be identical. Note also that the agent, during incremental processing, might initially get the wrong analysis, as would be the case if the dialog were *"Who got a prize?" "Antonio said that Mary did."* Initially, the agent might think that Antonio did. This is fine, and is exactly what a person would do if the speaker of the second utterance made a long pause, or coughed or laughed, after the first word.

When fragments are used outside of prototypical language strategies like these, their interpretation must be postponed until Situational Reasoning, when the agent can leverage its understanding of the domain script and the related plans and goals to guide the interpretation.

6.5 Further Exploration

1. Explore idiomatic creativity using the online search engine of the COCA corpus (https://www.english-corpora.org/coca/). Your searches need to be seeded by actual idioms, but the search strings can allow for various types of nonstandard usages. For example, the search string *kill _mc _nn with _mc _nn*, which covers the construction [*kill* + cardinal-number + any-noun + *with* + *cardinal-number* + any-noun], returns hits including the canonical *kill two birds with one stone* as well as *kill three flies with one stone, kill two birds with one workout*, and others.

2. Try to find examples of preposition swapping—and other performance errors—by watching foreign films and TV series with subtitles. We found the subtitles to the Finnish TV series *Easy Living* particularly interesting in this respect since they were of high quality overall with the occasional slip in preposition choice or use of an idiomatic expression.

3. Explore how numerical values in events of change are expressed using the search engine of the COCA corpus. Sample search strings include *increased by _mc*, which searches for [*increased by* + cardinal-number], *by _mc to _mc*, which searches for [*by* + cardinal-number + *to* + cardinal number], and countless more that use different verbs (e.g., *increase*, *decrease*, *go up*, *rise*) and different presentations of numbers (e.g., *by* cardinal-number *to* cardinal-number; *from* cardinal-number *% to* cardinal-number *%*). Consider the following questions about cognitive modeling:

- Do you think that you actually remember all the numbers from such contexts?
- If not, which ones do you remember and in which contexts?
- What should intelligent agents remember and not remember?
- Should they be more perfect than people in this respect (calculating and remembering everything), or should they be more humanlike?

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 6 and recall or invent examples of each phenomenon.