

INX199 – Assignment 1

Mapping Words to Meanings

Due October 26, 2005 10:10am
75 Marks (7.5% of course mark)

1 Introduction

Learning the meaning of words seems like a trivial task. You put an apple in front of a toddler, point to it, and say “apple!” After a few attempts, if the baby can see and hear properly, she will learn the mapping between the word *apple* and the fruit it refers to. Simple, isn’t it? Not always. For one, children usually hear multi-word utterances that describe an event or observation, not a single word. So they have to figure out which word in the utterance refers to which part of the event to be described. This is complicated by ambiguity—the fact that most words have more than one meaning. Moreover, not all the utterances describe events that are observable by the child, for example, *I saw John in the bookstore yesterday*. And even if they do, there is usually more than one interpretation of the event that the utterance is describing. For example, *John is eating an apple*, *John likes apples*, and *John’s shirt is green* may all describe different aspects of the same scene. The problem of finding the right meaning for an utterance is often referred to as *referential uncertainty*.

In this assignment, we use a simple method, inspired by Siskind (1996),¹ to map words to their meaning. Let’s forget about both ambiguity and referential uncertainty for now. We assume that children hear a sentence, and hypothesize a meaning for the event the sentence is describing, which we call an *utterance meaning*. We represent each utterance meaning as a combination of symbols. For example, the correct (simplified) meaning of the sentence *John went to school* is represented as $GO(\mathbf{John}, TO(\mathbf{school}))$. In this representation, predicates (corresponding to verbs, prepositions, adjectives) are shown in capital letters, and entities (corresponding to nouns) are shown in boldface.

¹Siskind, J. M. (1996). “A computational study of cross-situational techniques for learning word-to-meaning mappings.” *Cognition*, 61:39–91.

2 Word-to-meaning mapping

We simulate the word learning process as follows. The learner receives a number of sentences, each paired with an utterance meaning, and tries to guess the possible meanings for each of the words that appear in a sentence. By hearing more and more sentences, the learner polishes the hypothesized meanings mapped to each word, until she converges on the right meaning for each word in the language.

As an example, consider the following sentence and its meaning:

John went to school.

GO(**John**, TO(**school**))

If this is the first sentence the learner has ever heard, then she hypothesizes that each of the words *John*, *went*, *to*, and *school* may be mapped to zero or more meaning components from the set {GO, **John**, TO, **school**}. She can't know at this point which possible subset of meaning components is actually mapped to each word. However, using some intuitions and enough input data, the learner (hopefully!) over time revises the set of meaning fragments mapped to each word in an appropriate way. In the next section, we will describe the intuitions we rely on in this assignment.

3 Main assumptions

In our method, we make a number of assumptions about how children learn word-to-meaning mappings. Some of these assumptions are not always valid, but make our task easier and more manageable.

3.1 Cross-situational inference

One way a child might determine the meaning of a word is to find something in common across all observed uses of that word. For example, imagine that the learner hears two utterances *John lifted the ball* and *Mary lifted the block*. And imagine that the learner recognizes that CAUSE(**John**, GO(**ball**, UP)) is the meaning of the former and that CAUSE(**Mary**, GO(**block**, UP)) is the meaning of the latter. From the first sentence, the learner could hypothesize the set {CAUSE, **John**, GO, **ball**, UP} as possible meaning fragments for the word *lifted*. From the second sentence, the learner could form the set {CAUSE, **Mary**, GO, **block**, UP} of possible meaning fragments for *lifted*. Since *lifted* occurs in both sentences, only meaning fragments that are associated with *both* sentences can potentially be part of its meaning. So, by intersecting

these two sets, the learner can infer that the meaning of the word *lifted* cannot contain any meaning fragments other than CAUSE, GO, and UP.

3.2 Principle of compositionality

Cross-situational inference helps us find the *potential* fragments of a word's meaning, by ruling out some fragments that don't occur across different uses of the word. But it cannot provide evidence that a particular fragment *must* be part of a word's meaning. For example, applying cross-situational inference to the above situation, the learner could infer that **ball** could not be part of the meaning of the word *lifted*, since it is absent from the situation surrounding the second use of that word. However, using cross-situational inference alone, the learner could not determine that CAUSE *must* be part of the meaning of *lifted*. That is, CAUSE is *possibly* a meaning component of *lifted* since it occurs as a meaning component of all utterances that include the word *lifted*. But it is not *necessarily* a meaning component of *lifted*, since it could be part of the meaning of another word in those utterances.

Suppose that human language has the property that all components of the meaning of an utterance must be derived from the meanings of words in that utterance. This property is called compositionality, because meanings of larger phrases are *composed from* meanings of their component words and phrases. Now consider again the utterance *John lifted the ball* paired with the meaning CAUSE(**John**, GO(**ball**, UP)) in the learner's mind. Imagine that, by applying cross-situational inference to other utterances, the learner could rule out CAUSE as a component of the meanings of the words *John*, *the*, and *ball*. Under the compositionality constraint, the learner could infer that CAUSE must be a part of the meaning of *lifted*, since it is part of the meaning of *John lifted the ball* yet not part of the meaning of any other word in that utterance.

3.3 Principle of exclusivity

Further suppose that human language has the property that the words in an utterance must contribute non-overlapping portions of the utterance meaning. The learner could use this property, via a principle of exclusivity, to perform an additional form of inference in the following fashion. Imagine that a learner heard the utterance *John walked* and hypothesized the meaning WALK(**John**) for that utterance. Further imagine that the learner already determined, via a combination of cross-situational inference and compositionality, that *John* must mean **John**. Applying only cross-situational inference and compositionality to this data, the learner could not determine whether the meaning fragments of *walked* are {WALK, **John**} or just {WALK}. However, under the assumption that *John* and *walk* contribute non-overlapping portions of the

utterance meaning WALK(**John**), knowing that *John* means **John** rules out {WALK, **John**} as a potential meaning of *walked*, and so the meaning of walk must be simply {WALK}.

4 Word learning method

We will formalize the intuitions described in the previous section into a set of rules for inferring word meaning. But before we get to the rules, we define a few terms which we will use in our word learning method, which we'll call "the learner." Our method maintains two sets of meaning fragments for each word: a **Possible** set, which contains all meaning fragments that can *potentially* be part of the meaning of the word, and a **Necessary** set, which contains all the meaning fragments that *must* be part of the meaning of a word. If the two sets are identical, then we say that the meaning of the word has been successfully learned.

4.1 The first step

When the learner hears a word for the first time, it puts all the meaning fragments in the utterance meaning into the **Possible** set of the new word. For example, if the first time the learner hears the word *notebook* is in the sentence *Tim wrote in the black notebook*, which comes with the utterance meaning WRITE(**Tim**, IN(BLACK(**notebook**))), then the **Possible** set for *notebook* will be {WRITE, **Tim**, IN, BLACK, **notebook**}. However, the **Necessary** set for the new word is empty, since the learner cannot determine which of these meanings must be part of the meaning of the word. Using the inference rules given in the next section, the learner is able to change and refine each of the **Possible** and **Necessary** sets for the words it hears in the input so that they eventually converge to the correct, identical set that represents the meaning of each word.

4.2 Constraining rules

At each point in learning, the **Possible** and **Necessary** sets for each word are either initialized (as described just above), or possibly modified using the following rules (in the order given).

Cross-situational inference:

Rule 1. *For each word in the sentence, remove from its **Possible** set any meaning symbols that do not appear in the utterance meaning.*

For example, suppose that the learner is part-way through the learning process and already possesses the following words and their corresponding sets:

	Necessary	Possible
<i>apple</i>	{}	{EAT, Mary, apple}
<i>ball</i>	{ball}	{ball, arm}
<i>eat</i>	{EAT}	{EAT, Mary, apple}
<i>John</i>	{John}	{John, ball}
<i>Mary</i>	{}	{EAT, Mary, apple}
<i>the</i>	{}	{WANT, arm}
<i>took</i>	{CAUSE}	{CAUSE, WANT, GO, TO, arm}

Now suppose that the method receives the input:

John took the ball.

CAUSE(**John**, GO(**ball**, TO(**John**)))

By applying Rule 1, the meaning fragment **arm** is removed from the **Possible** set for *ball*, since **arm** is not part of the meaning of the sentence *ball* has just appeared in. Similarly, WANT and **arm** will be removed from the **Possible** sets for both *took* and *the*. The resulting word list will be as follows:

	Necessary	Possible
<i>apple</i>	{}	{EAT, Mary, apple}
<i>ball</i>	{ball}	{ball}
<i>eat</i>	{EAT}	{EAT, Mary, apple}
<i>John</i>	{John}	{John, ball}
<i>Mary</i>	{}	{EAT, Mary, apple}
<i>the</i>	{}	{}
<i>took</i>	{CAUSE}	{CAUSE, GO, TO}

The compositionality principle:

Rule 2. *If a meaning symbol in the utterance meaning is part of the **Possible** set of only one of the words in the sentence, then add it to the **Necessary** set of that word.*

Consider the previous example, in which the utterance meaning contains the meaning symbols GO and TO. These symbols are in the **Possible** set for *took*, but not in its **Necessary** set, nor in the **Possible** set for the words *John*, *the*, and *ball*. The method can then infer that GO and TO must be part of the meaning of the word *took*. After applying Rule 2 to the previous table, the word list and their corresponding sets will be as follows:

	Necessary	Possible
<i>apple</i>	{}	{EAT, Mary, apple }
<i>ball</i>	{ ball }	{ ball }
<i>eat</i>	{EAT}	{EAT, Mary, apple }
<i>John</i>	{ John }	{ John , ball }
<i>Mary</i>	{}	{EAT, Mary, apple }
<i>the</i>	{}	{}
<i>took</i>	{CAUSE, GO, TO}	{CAUSE, GO, TO}

The principle of exclusivity:

Rule 3. *If a meaning symbol that appears in the utterance meaning is part of the Necessary set of one of the words in the sentence, remove it from the Possible set of all other words of the sentence.*

At this point, using the principle of exclusivity, the method can make the following inference: since the meaning symbol **ball** appears only once in the utterance meaning, and the word *ball* necessarily contributes this meaning, the word *John* cannot also contain **ball** as part of its meaning. This inference allows the method to remove the conceptual symbol **ball** from the **Possible** set for *John*. Applying Rule 3 to the previous table, we obtain the following word list:

	Necessary	Possible
<i>apple</i>	{}	{EAT, Mary, apple }
<i>ball</i>	{ ball }	{ ball }
<i>eat</i>	{EAT}	{EAT, Mary, apple }
<i>John</i>	{ John }	{ John }
<i>Mary</i>	{}	{EAT, Mary, apple }
<i>the</i>	{}	{}
<i>took</i>	{CAUSE, GO, TO}	{CAUSE, GO, TO}

At this point, for some of the words in the list—*ball*, *John*, *the*, *took*—the **Possible** and **Necessary** sets are identical, which means that the learner has successfully learned the actual meaning-symbol set for these words. (Note that *the* in our meaning representation is not associated with any meaning components—its **Possible** and **Necessary** sets are empty.)

However, the meanings of some words are still not learned—*apple*, *eat*, *Mary* still have non-identical meaning sets, and further input would be needed to appropriately map them to their meanings.

5 Task I: Applying the Learning Method (60 marks)

We provide you with a sequence of eight sentence-meaning pairs. You should apply the learning method described above on this input data. (See below for what to turn in.)

For each input pair (in order), do the following steps, in this order:

- **Initialization:** Form the **Possible** and **Necessary** sets for all the words that have never appeared before in an utterance.
- **Revision of meaning sets:**
 - Apply Rule 1 on all words in the sentence to which it is applicable.
 - Apply Rule 2 on all words in the sentence to which it is applicable.
 - Apply Rule 3 on all words in the sentence to which it is applicable.
- **Repeat:** When all applicable changes have been made in response to this input, move to the next input item and repeat the above steps.

Input data:

1. *Mary ate the apple.*
EAT(**Mary**, **apple**)
2. *Mary took the ball to school.*
CAUSE(**Mary**, GO(**ball**, TO(**school**)))
3. *John took the apple.*
CAUSE(**John**, GO(**apple**, TO(**John**)))
4. *Mary took the book.*
CAUSE(**Mary**, GO(**book**, TO(**Mary**)))
5. *John ate the orange.*
EAT(**John**, **orange**)
6. *John took the apple to school.*
CAUSE(**John**, GO(**apple**, TO(**school**)))
7. *John went to school.*
GO(**John**, TO(**school**))
8. *Mary took the orange.*
CAUSE(**Mary**, GO(**orange**, TO(**Mary**)))

Task I: What you should turn in:

(Label all answers with the appropriate numbers.)

1. (40 marks) Provide tables showing what is learned after each input pair:
 - You should turn in eight numbered tables (one per input pair) such as those we showed above in Section 4.2.
 - In each table, show the list of all words processed thus far, **in alphabetical order**, together with their **Possible** and **Necessary** meaning sets **after** the current input item has been processed.
 - In each table, show which rule(s) (if any) have been applied to each word.
2. (5 marks) List the words that have been successfully learned after processing all the input pairs.
3. (10 marks) After all the input pairs have been processed, some word (possibly more than one) has not been learned, even though all other words it co-occurs with in at least one utterance have been learned. Why does this happen? Briefly explain using an example word that is not learned.
4. (5 marks) Assume the learner is now given the input pair:

9. *Mary bit the apple.*
BITE(Mary, apple)

Briefly explain why the word *bit* is so much easier to learn than the other words have been.

6 Task II: “Noisy” Data (15 marks)

The data a child receives from her environment is not always reliable. She may mishear a word, or relate an utterance to an inappropriate meaning. We refer to this as noisy data.

For example, imagine that, due to some misinterpretation, the input item (6) in the sequence above is **replaced with** the following input item:

- 6'. *John took the apple to school.*
GO(John, TO(school))

Explore how this type of noisy data affects the process of learning by answering the following questions.

Task II: What you should turn in:

(Label all answers with the appropriate numbers.)

1. **(3 marks)** Show the output of the method using the modified input data. Provide revised tables 6', 7', and 8' using the input sequence with the new item (6') in place of item (6).
2. **(2 marks)** At the end of the input, have any words been “successfully” but incorrectly learned?
3. **(5 marks)** At what point(s) does the learner make a mistake, and why?
4. **(5 marks)** Do you think the method can recover from these types of mistakes by receiving more input? Briefly explain why or why not.

7 Task III: Referential uncertainty [Extra Credit (15 marks)]

We talked about the referential uncertainty problem at the beginning of this document: it is not always easy to hypothesize the correct meaning for an utterance, since it can describe different aspects of a single event. In our method, we assumed that each input item pairs an utterance with a single hypothesized meaning. Now imagine that a learner may hypothesize more than one meaning for an utterance, such as the following:

John took the apple.

CAUSE(**John**, GO(**apple**, TO(**John**)))

WEAR(**John**, GREEN(**shirt**))

MOVE(**John**, PART_OF(**John**, **arm**))

Task III: What you should turn in: [Optional]

(Label all answers with the appropriate numbers.)

1. **(5 marks)** Is the current learning method able to handle this situation? Briefly explain why or why not. (Your answer may be of the form “Yes, as long as the following is true...”, as long as you explain why your assumption is reasonable.)
2. **(10 marks)** Add a rule or rules to the current method that will allow it to handle this situation without generally requiring a large amount of extra input in order to converge. Your rule should not depend on knowledge beyond the current table keeping track of the **Possible** and **Necessary** sets for words seen thus far.