

Meaning Representations Chapter 17

Lecture #11

November 2012

1

Big Transition

- First we did words (morphology)
- Then we looked at syntax
- Now we're moving on to meaning. Where some would say we should have started to begin with.

- Now we look at meaning representations – representations that link linguistic forms to knowledge of the world.

2

Meaning

- Language is useful and amazing because it allows us to encode/decode...
 - Descriptions of the world
 - What we're thinking
 - What we think about what other people think
- Don't be fooled by how natural and easy it is... In particular, you do not ever
 - Utter word strings that match the world
 - Say what you're thinking
 - Say what you think about what other people think

3

Meaning

- You're simply uttering linear sequences of words such that when other people read/hear and understand them they come to know what you think of the world.

4

Meaning

- So... I can stand up here and bounce waves of compressed air against your eardrums and have the effect of
 - Making you laugh, cry or go to sleep
 - Telling you how to make a soufflé
 - Describing the weather, or a double play, or a glass of wine to you.
- These are not easy tasks. They are amazing tasks. They just look easy.

5

Meaning Representations

- We're going to take the same basic approach to meaning that we took to syntax and morphology

- We're going to create **representations** of linguistic inputs that capture the meanings of those inputs.

- **But unlike parse trees and the like, these representations aren't primarily descriptions of the structure of the inputs...**

6

Meaning Representations

- In most cases, they're simultaneously descriptions of the meanings of utterances and of some potential state of affairs in some world.
- What could this mean...
 - **representations** of linguistic inputs that capture the meanings of those inputs

7

Meaning Representations

- What are some of the linguistic concepts we want to capture?
 - Categories/entities (Roger, Microsoft, vegetarian, UofD, desk)
 - Events (walking to class, eating lunch)
 - Time (9:30am, next week, 2015)
 - Aspect (Kathy knows how to run. Kathy is running. Kathy ran to class yesterday.)
 - Beliefs, Desires, Intentions
- How? What is most important? This means lots of different things to lots of different philosophers.
- We're not going to go there. For us it means
 - Representations that permit or facilitate **semantic processing**.

Semantic Processing

- Ok, so what does that mean?
- What we take as a meaning representation is a representation that serves the core practical purposes of a program that is doing semantic processing.
- Representations that
 - Permit us to reason about their truth (relationship to some world)
 - Is the blue block on the red block?
 - Permit us to answer questions based on their content
 - What is the tallest building in the world.
 - Permit us to perform inference (answer questions and determine the truth of things we don't actually know)
 - If the blue block is on the red block, and the red block is in the room, then the blue block is in the room.

9

Semantic Processing

- Touchstone application is always question answering
 - Can I answer questions involving the meaning of some text or discourse?
 - What kind of representations do I need to mechanize that process?

10

Sample Meaning Representations

I have a car.


- First-Order Predicate Calculus
- Semantic Networks
- Conceptual Dependency
- Frame-based representation

11

Common Meaning Representations

- FOPC:
 $\exists x, y \text{Having}(x) \wedge \text{Haver}(S, x) \wedge \text{HadThing}(y, x) \wedge \text{Car}(y)$

- Semantic Net:



12

- Conceptual Dependency Diagram:
 - Car
 - ↑ Poss-By
 - Speaker
- Frame
 - Having
 - Haver: S
 - HadThing: Car
- All represent 'linguistic meaning' of I have a car **and** state of affairs in some world
- All consist of structures, composed of symbols representing objects and relations among them

13

What requirements must meaning representations fulfill?

- Verifiability: The system should allow us to compare representations to facts in a Knowledge Base (KB)
 - Cat(Huey)
- Ambiguity: The system should allow us to represent meanings unambiguously
 - German teachers has 2 representations
- Vagueness: The system should allow us to represent vagueness
 - He lives somewhere in the south of France.

14

Initial Simplifying Assumptions

- Focus on literal meaning
 - Conventional meanings of words
 - Ignore context

15

Canonical Form

- Inputs that mean the same thing have the same representation.
 - Huey eats kibble.
 - Kibble, Huey will eat.
 - What Huey eats is kibble.
 - It's kibble that Huey eats.
- Alternatives
 - Four different semantic representations
 - Store all possible meaning representations in KB

16

Canonical Form: Pros and Cons

Advantages

- Simplifies reasoning tasks
- Compactness of representations: don't need to write inference rules for all different "paraphrases" of the same meaning

Disadvantages

- Complicates task of semantic analysis

17

Inference

- Draw valid conclusions based on the meaning representation of inputs and its store of background knowledge.

Does Huey eat kibble?
 thing(kibble)
 Eat(Huey,x) ^ thing(x)

18

Expressiveness

- Must accommodate wide variety of meanings

19

Predicate-Argument Structure

- Represents concepts and relationships among them
 - Nouns as concepts or arguments (red(ball))
 - Adjectives, adverbs, verbs as predicates (red(ball))
- Subcategorization (or, argument) frames specify number, position, and syntactic category of arguments
 - NP likes NP
 - NP likes Inf-VP
 - NP likes NP Inf-VP

20

Fillmore's Theory about Universal Cases

- Fillmore – there are a small number of semantic roles that an NP in a sentence may play with respect to the verb.
- A major task of semantic analysis is to provide an appropriate mapping between the syntactic constituents of a parsed clause and the semantic roles (cases) associated with the verb.

21

Major Cases Include

- **Agent** – doer of the action, entails intentionality
- **Experiencer** – doer when no intentionality
- **Theme** – thing being acted upon or undergoing change
- **Instrument** – tool used to do the action
- **Beneficiary** – person/thing for whom the event is performed
- **To/At/From Loc/Poss/Time** – location or possession or time representations

22

Some Sentences and their cases

- John opened the door with a key.
- The door was opened by John.
- The door was opened with a key.
- A key opened the door.
- The door opened.
- John gave Mary the book.
- John gave the book to Mary.

Let's identify the cases in these sentences – notice any syntactic regularities in the case assignment.

23

Some Sentences and their cases

- John opened the door with a key.
- The door was opened by John.
- The door was opened with a key.
- A key opened the door.
- The door opened.
- John gave Mary the book.
- John gave the book to Mary.

– Agent – doer of action, attributes intention

24

Some Sentences and their cases

- John opened the door with a key.
- The door was opened by John.
- The door was opened with a key.
- A key opened the door.
- The door opened.
- John gave Mary the book.
- John gave the book to Mary.

- Agent – doer of action, attributes intention
- Theme – thing being acted upon or undergoing change

25

Some Sentences and their cases

- John opened the door with a key.
- The door was opened by John.
- The door was opened with a key.
- A key opened the door.
- The door opened.
- John gave Mary the book.
- John gave the book to Mary.

- Agent – doer of action, attributes intention
- Theme – thing being acted upon or undergoing change
- Instrument – tool used to do the action

26

Some Sentences and their cases

- John opened the door with a key.
- The door was opened by John.
- The door was opened with a key.
- A key opened the door.
- The door opened.
- John gave Mary the book.
- John gave the book to Mary.

- Agent – doer of action, attributes intention
- Theme – thing being acted upon or undergoing change
- Instrument – tool used to do the action
- To-Poss –

27

Some Sentences and their cases

- John opened the door with a key.
- The door was opened by John.
- The door was opened with a key.
- A key opened the door.
- The door opened.
- John gave Mary the book.
- John gave the book to Mary.

Intuition – syntactic choices are largely a reflection of underlying semantic relationships.

28

Semantic Analysis

- A major task of semantic analysis is to provide an appropriate mapping between the syntactic constituents of a parsed clause and the semantic roles associated with the verb.

29

Factors to Complicate

- Ability of syntactic constituents to indicate several different semantic roles
 - E.g., Subject position; agent versus instrument versus theme
 - John broke the window.
 - The rock broke the window.
 - The window broke.
- Large number of choices available for syntactic expression of any particular syntactic role
 - E.g., agent and theme in different configurations
 - John broke the window.
 - It was the window that John broke.
 - The window was broken by John.

30

Factors to Complicate (cont.)

- Prepositional ambiguities – it is the case that a particular preposition does not always introduce the same role
 - E.g., preposition “by” may indicate either agent or instrument
 - The door was opened by John.
 - The door was opened by a key.
- Optionality of a given role in a sentence
 - John opened the door with a key.
 - The door was opened by John.
 - The door was opened with a key.
 - A key opened the door.
 - The door opened.

31

How bad is it?

- It seems that semantic roles are playing “musical chairs” with the syntactic constituents. That is, they seem to “sit down” in any old syntactic constituent and one or more of them seem to be left out at times!
- Actually, it isn't as bad as it may seem!
- There is a great deal of regularity – consider the following set of rules....

32

Some Rules

If \exists Agent it becomes Subject
Else If \exists Instrument it becomes Subject
Else If \exists Theme it becomes Subject

Agent preposition is BY
Instrument preposition is BY if no agent, else WITH

Some Rules:

- Some verbs may have exceptions
- No case can appear twice in the same clause
- Only NP's of same case can be conjoined
- Each syntactic constituent can fill only 1 case

33

What's missing???

If \exists Agent it becomes Subject
Else If \exists Instrument it becomes Subject
Else If \exists Theme it becomes Subject

How do I know whether or not an agent exists? How about an instrument?

Selectional Restrictions – restrict the types of certain roles to be a certain semantic entity

- Agents must be animate
- Instruments are not animate
- Theme? – type may be dependent on the verb itself.

34

Selectional Restrictions

Selectional Restrictions: constraints on the **types** of arguments verbs take

George assassinated the senator.
***The spider assassinated the fly.**
assassinate: intentional (political?) killing

NOTE: dependence on the particular verb being used!

35

So? What about Case in General?

- You may or may not see particular cases used in semantic analysis.
- In the book, they have NOT used the specific cases.
- But, note, the “roles” they use are derived from the general cases identified in Fillmore's work – they make them verb-specific.
- Semantic analysis is going to take advantage of the syntactic regularities and selectional restrictions to identify the role being played by each constituent in a sentence!

36

Representational Schemes

- Let's go back to the question – what kind of semantic representation should we derive for a given sentence?
- We're going to make use of First Order Predicate Calculus (FOPC) as our representational framework
 - Not because we think it's perfect
 - All the alternatives turn out to be either too limiting or
 - They turn out to be notational variants
 - Essentially the important parts are the same no matter which variant you choose!

37

FOPC

- Allows for...
 - The analysis of truth conditions
 - Allows us to answer yes/no questions
 - Supports the use of variables
 - Allows us to answer questions through the use of variable binding
 - Supports inference
 - Allows us to answer questions that go beyond what we know explicitly

38

FOPC

- This choice isn't completely arbitrary or driven by the needs of practical applications
- FOPC reflects the semantics of natural languages because it was designed that way by human beings
- In particular...

39

Meaning Structure of Language

- The semantics of human languages...
 - Display a basic predicate-argument structure
 - Make use of variables (e.g., indefinites)
 - Make use of quantifiers (e.g., every, some)
 - Use a partially compositional semantics (sort of)

40

Predicate-Argument Structure

- Events, actions and relationships can be captured with representations that consist of predicates and arguments.
- Languages display a division of labor where some words and constituents function as predicates and some as arguments.
- E.g., predicates represent the verb, and the arguments (in the right order) represent the cases of the verb.

41

Predicate-Argument Structure

- Predicates
 - Primarily Verbs, VPs, PPs, adjectives, Sentences
 - Sometimes Nouns and NPs
- Arguments
 - Primarily Nouns, Nominals, NPs
 - But also everything else; as we'll see it depends on the context

42

Example

- John gave a book to Mary
- Giving(John, Mary, Book)
- More precisely
 - Gave conveys a three-argument predicate
 - The first argument is the giver (agent)
 - The second is the recipient (to-poss), which is conveyed by the NP in the PP
 - The third argument is the thing given (theme), conveyed by the direct object

43

Not exactly

- The statement
 - The first arg is the subject can't be right.
- Subjects can't be givers.
- We mean that the meaning underlying the subject phrase plays the role of giver.

44

More Examples

What about situation of missing/additional cases?

- John gave Mary a book for Susan.
Giving(John, Mary, Book, Susan)
- John gave Mary a book for Susan on Wednesday.
Giving(John, Mary, Book, Susan, Wednesday)
- John gave Mary a book for Susan on Wednesday in class.
Giving(John, Mary, Book, Susan, Wednesday, InClass)

Problem: Remember each of these predicates would be different because of the different number of arguments! Except for the suggestive names of predicates and arguments, there is nothing that indicates the obvious logical relations among them.

45

Meaning Representation Problems

- Assumes that the predicate representing the meaning of a verb has the same number of arguments as are present in the verb's syntactic categorization frame. This makes it hard to
- Determine the correct number of roles for any given event
- Represent facts about the roles associated with the event
- Insure that all and only the correct inferences can be derived from the representation of an event

46

Better

- Turns out this representation isn't quite as useful as it could be.
 - Giving(John, Mary, Book)
- Better would be one where the "roles" or "cases" are separated out. E.g., consider:

$$\exists x, y \text{ Giving}(x) \wedge \text{Giver}(\text{John}, x) \wedge \text{Given}(y, x) \\ \wedge \text{Giver}(\text{Mary}, x) \wedge \text{Isa}(y, \text{Book})$$

- Note: essentially Giver=Agent, Given=Theme, Giver=To-Poss

47

Predicates

- The notion of a predicate just got more complicated...
 - In this example, think of the verb/VP providing a template like the following
- $$\exists w, x, y, z \text{ Giving}(x) \wedge \text{Giver}(w, x) \wedge \text{Given}(y, x) \wedge \text{Giver}(z, x)$$
- The semantics of the NPs and the PPs in the sentence plug into the slots provided in the template (we'll worry about how in a bit!)

48

Advantages

- Can have variable number of arguments associated with an event: events have many roles and fillers can be glued on as appear in the input.
- Specifies categories (e.g., book) so that we can make assertions about categories themselves as well as their instances. E.g., $Isa(MobyDick, Novel)$, $AKO(Novel, Book)$.
- Reifies events so that they can be quantified and related to other events and objects via sets of defined relations.
- Can see logical connections between closely related examples without the need for meaning postulates.

49

Additional Material

- The following are some aspects covered in the book that will likely not be covered in lecture!

50

FOPC Syntax

- Terms: constants, functions, variables
 - Constants: objects in the world, e.g. Huey
 - Functions: concepts, e.g. $sisterof(Huey)$
 - Variables: x , e.g. $sisterof(x)$
- Predicates: symbols that refer to relations that hold among objects in some domain or properties that hold of some object in a domain
 - $likes(Huey, kibble)$
 - $cat(Huey)$

51

- Logical connectives permit compositionality of meaning
 - $kibble(x) \rightarrow likes(Huey, x)$
 - $cat(Vera) \wedge weird(Vera)$
 - $sleeping(Huey) \vee eating(Huey)$
- Sentences in FOPC can be assigned truth values, T or F, based on whether the propositions they represent are T or F in the world
 - Atomic formulae are T or F based on their presence or absence in a DB (Closed World Assumption?)
 - Composed meanings are inferred from DB and meaning of logical connectives

52

- $cat(Huey)$
- $sibling(Huey, Vera)$
- $sibling(x, y) \wedge cat(x) \rightarrow cat(y)$
- $cat(Vera)??$
- Limitations:
 - Do 'and' and 'or' in natural language really mean ' \wedge ' and ' \vee '?
 - Mary got married and had a baby.
 - Your money or your life!
 - She was happy but ignorant.
 - Does ' \rightarrow ' mean 'if'?
 - I'll go if you promise to wear a tutu.

53

- Quantifiers: \exists, \forall
 - Existential quantification: There is a unicorn in my garden. Some unicorn is in my garden.
 - Universal quantification: The unicorn is a mythical beast. Unicorns are mythical beasts.
- Inference:
 - Modus ponens:
 - $rich(Harry)$
 - $\forall x \ rich(x) \rightarrow happy(x)$
 - $happy(Harry)$
- Production systems:
 - Forward and backward chaining

54

Temporal Representations

- How do we represent time and temporal relationships between events?
Last year Martha Stewart was happy but soon she will be sad.
- Where do we get temporal information?
 - Verb tense
 - Temporal expressions
 - Sequence of presentation
- Linear representations: Reichenbach '47

55

- Utterance time: when the utterance occurs
 - Reference time: the temporal point-of-view of the utterance
 - Event time: when events described in the utterance occur
- George had intended to eat a sandwich.
E – R – U →
- George is eating a sandwich.
-- E,R,U →
- George had better eat a sandwich soon.
--R,U – E →

56

Verbs and Event Types: Aspect

- Statives: states or properties of objects at a particular point in time
Mary needs sleep.
*Mary is needing sleep. *Need sleep. *Mary needs sleep in a week.
- Activities: events with no clear endpoint
Harry drives a Porsche. *Harry drives a Porsche in a week.

57

- Accomplishments: events with durations and endpoints that result in some change of state
Marlon filled out the form. Marlon stopped filling out the form (Marlon did not fill out the form) vs. Harry stopped driving a Porsche (Harry still drove a Porsche ...for a while)
- Achievements: events that change state but have no particular duration
Larry reached the top. *Larry stopped reaching the top.
*Larry reached the top for a few minutes.

58

Beliefs, Desires and Intentions

- How do we represent internal speaker states like believing, knowing, wanting, assuming, imagining..?
 - Not well modeled by a simple DB lookup approach
 - Truth in the world vs. truth in some possible world
George imagined that he could dance.
George believed that he could dance.
- Augment FOPC with special modal operators that take logical formulae as arguments, e.g. believe, know

59