

Introduction to Semantics

Announcements

- Part of next Tuesday's class will focus on questions about the midterm. Please send topics that you would like to have re-explained.
- Other midterm news: Sample questions provided by Wed of this week; in-class review in 2nd half of next Tuesday's class
- Homework 1 statistics:
 - Mean: 90.5
 - Median: 94
 - High: 100
 - Low (nonzero): 45

Today

- Reading:
 - Today: C 17, Speech and Language
 - Next time: C 19, C 20.1-20.8 Speech and Language
- Introduction to semantics
 - What is it?
 - Meaning representations
 - Lexical semantics

Which topics from the syllabus would you like to have reviewed?

Language modeling

Supervised machine learning

POS tagging and HMMs

Viterbi algorithm

Context free grammars

dependency parsing

neural nets

word embeddings



Are there other topics (or more specific areas) that you would like reviewed?

I'd like a seafood restaurant in Sag Harbor

Tap to Edit >


OK, check it out:

MAPS

The Dock House
 Restaurant · 4.8 miles
 ★★★★★ (126) on Yelp · \$\$


Page at 63 Main
 Restaurant · 4.7 miles
 ★★★★★ (497) on OpenTable · \$\$\$

Iceberg Seafood
 Seafood · 4.7 miles
 No Reviews



SIRI:
One shot Question and answer

What's required?



Watson and jeopardy

Bang Bang" his "Silver hammer came down upon her head"

The image shows a Jeopardy! game board with three contestants: Ken (\$200), Watson (\$4,000), and Brad (\$600). The question is "Maxwell's silver hammer" with a green bar indicating Watson's 96% correct answer rate. Other options are "FRANK SINATRA" (11%) and "Brown" (7%).

Contestant	Score
Ken	\$200
Watson	\$4,000
Brad	\$600

Answer	Percentage
Maxwell's silver hammer	96%
FRANK SINATRA	11%
Brown	7%

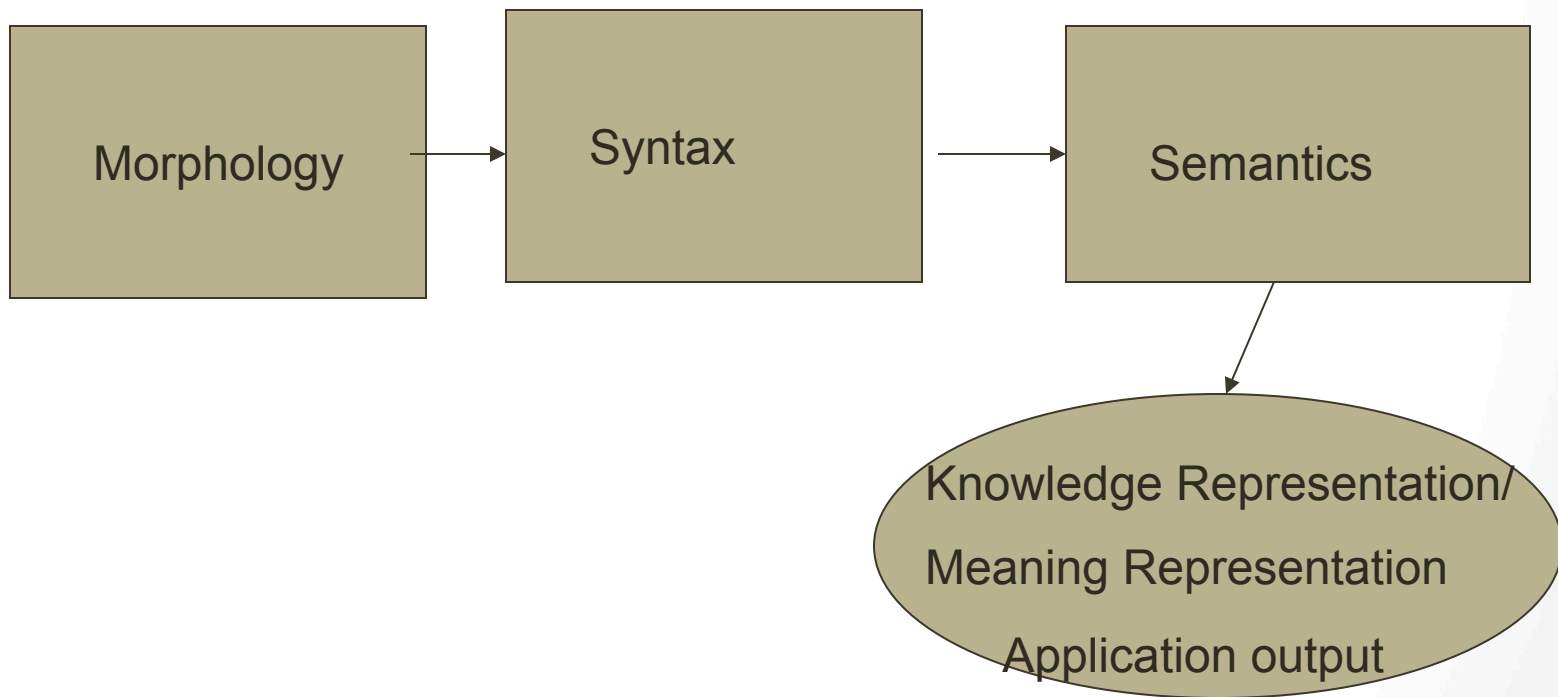
Question-Answering/Dialog

- “Match”

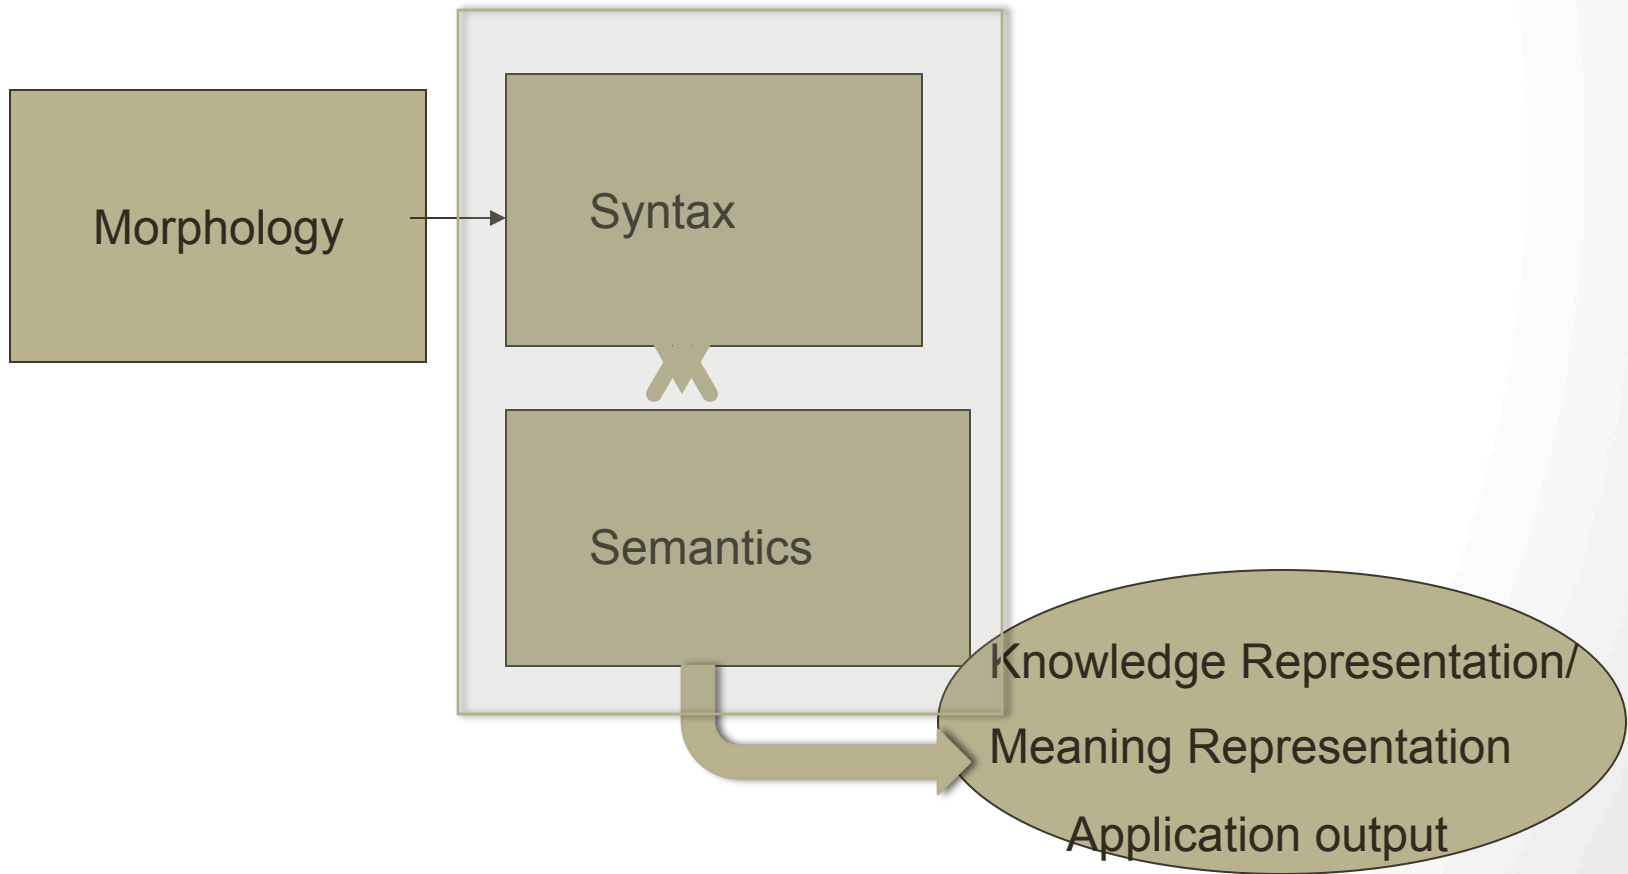
Where does the information for answers come from?

- Databases
 - <http://newyork.citysearch.com>
- The web
- Textual resources: encyclopedias, journals, etc.
- Dictionaries
- Inference

One View of NL Architecture



One View of NL Architecture



Another View: Understanding in Watson

- **Massive parallelism**
 - Exploit massive parallelism in the consideration of multiple interpretations and hypotheses.
- **Many experts**
 - Facilitate the integration, application, and contextual evaluation of a wide range of loosely coupled probabilistic question and content analytics.
- **Pervasive confidence integration**
 - No component commits to an answer; all components produce features and associated confidences, scoring different question and content interpretations.
- **Integrate shallow and deep knowledge**
 - Balance the use of strict semantics and shallow semantics, leveraging many loosely formed ontologies.

What is Semantics?

- Meaning Representation
- Semantic Parsing or Semantic Interpretation
- Word meaning disambiguation
- Entailment and inference
- Information extraction

Meaning Representation

- To represent questions
- To represent knowledge drawn from text

What Can Serve as a Meaning Representation?

- Anything that allows us to
 - Answer questions (**What is the best French restaurant in the East Village?**)
 - Determine truth (**Is Columbia Univ in NYC?**)
 - Draw inferences (**If One Dine Restaurant is on the top floor of one World Trade Center and one World Trade Center is the tallest building in New York City, then One Dine is the topmost restaurant in New York City.**)

What kinds of meaning do we want to capture?

- Categories/entities
 - Gustav, Shaan, Diana, Asian cuisine, vegetarian
- Events
 - taking a taxi, destruction of buildings in the Bahamas
- Time
 - Oct 30, next week, in 2 months
- Aspect
 - Kathy knows how to run. Kathy is running. Kathy ran to the restaurant in 5 min.
- Beliefs, Desires and Intentions (BDI)

Meaning Representations

- 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
 - FOPC:

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

A Standard Representation: 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

Semantic (Thematic) Roles

- Subcat frames link arguments in surface structure with their semantic roles
 - Agent: **George** hit Bill. Bill was hit by **George**.
 - Patient: George hit **Bill**. **Bill** was hit by George.
- The claim of a theory of semantic roles is that these arguments of predicates can be usefully classified into a small set of semantically contentful classes
 - And that these classes are useful for explaining lots of things

Common semantic roles

- **Agent**: initiator or doer in the event
- **Patient**: affected entity in the event; undergoes the action
 - Sue sent the letter.
- **Theme**: object in the event undergoing a change of state or location, or of which location is predicated
 - The ice melted
- **Experiencer**: feels or perceives the event
 - Bill likes pizza.
- **Stimulus**: the thing that is felt or perceived

Common semantic roles

- **Goal:**
 - Bill ran to Copley Square.
- **Recipient** (may or may not be distinguished from Goal):
 - Bill gave the book to Mary.
- **Benefactive** (may be grouped with Recipient):
 - Bill cooked dinner for Mary.
- **Source:**
 - Bill took a pencil from the pile.
- **Instrument:**
 - Bill ate the burrito with a plastic spork.
- **Location:**
 - Bill sits under the tree on Wednesdays

Common semantic roles

Try for yourself!

1. The submarine sank a troop ship.
2. Doris hid the money in the flowerpot.
3. Emma noticed the stain.
4. We crossed the street.
5. The boys climbed the wall.
6. The chef cooked a great meal.
7. The computer pinpointed the error.
8. A mad bull damaged the fence on Jack's farm.
9. The company wrote me a letter.
10. Jack opened the lock with a paper clip.

What role is "the money" in "Doris hid the money in the flowerpot"?

Agent

Patient

Theme

Stimulus

None of the above

What role is "the stain" in "Emma noticed the stain"?

Agent

Patient

Theme

Stimulus

What role is "me" in "The company wrote me a letter"

Patient

Theme

Recipient

Benefactor

Linking of thematic roles to syntactic positions

- John opened the door
- AGENT THEME
- The door was opened by John
- THEME AGENT
- The door opened
- THEME
- John opened the door with the key
- AGENT THEME INSTRUMENT

Deeper Semantics

- From the WSJ...
 - He melted her reserve with a husky-voiced paean to her eyes.
 - If we label the constituents **He** and **her reserve** as the **Melter** and **Melted**, then those labels lose any meaning they might have had.
 - If we make them **Agent** and **Theme** then we can do more inference.

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*
- *The astronaut married the star.*

Problems

- What exactly is a role?
- What's the right set of roles?
- Are such roles universals?
- Are these roles atomic?
 - I.e. Agents
 - Animate, Volitional, Direct causers, etc
- Can we automatically label syntactic constituents with thematic roles?

First Order Predicate Calculus

- Not ideal as a meaning representation and doesn't do everything we want -- but better than many...
 - Supports the determination of truth
 - Supports compositionality of meaning
 - Supports question-answering (via variables)
 - Supports inference

Abstract Meaning Representation (AMR)

- Aiming to capture a deep level of semantics
- Original goal: as interlingua for machine translation
- Also now see: AMR parsers, generation from AMR, summarization using AMR

Example AMR

- Roles are given theory-neutral names
(d / describe-01
 - :arg0 (m/man)
 - :arg1 (m2 / mission)
 - :arg2 (d /disaster))
- :arg0 the describer, :arg1 the thing described, :arg2 what it is describing
- The man described the mission as a disaster.
As the man described it, the mission was a disaster

Temporal Representations

- How do we represent time and temporal relationships between events?

It seems only yesterday that Martha Stewart was in prison but now she has a popular TV show. There is no justice.

- Where do we get temporal information?
 - Verb tense
 - Temporal expressions
 - Sequence of presentation
- Linear representations: Reichenbach '47

Representing Time

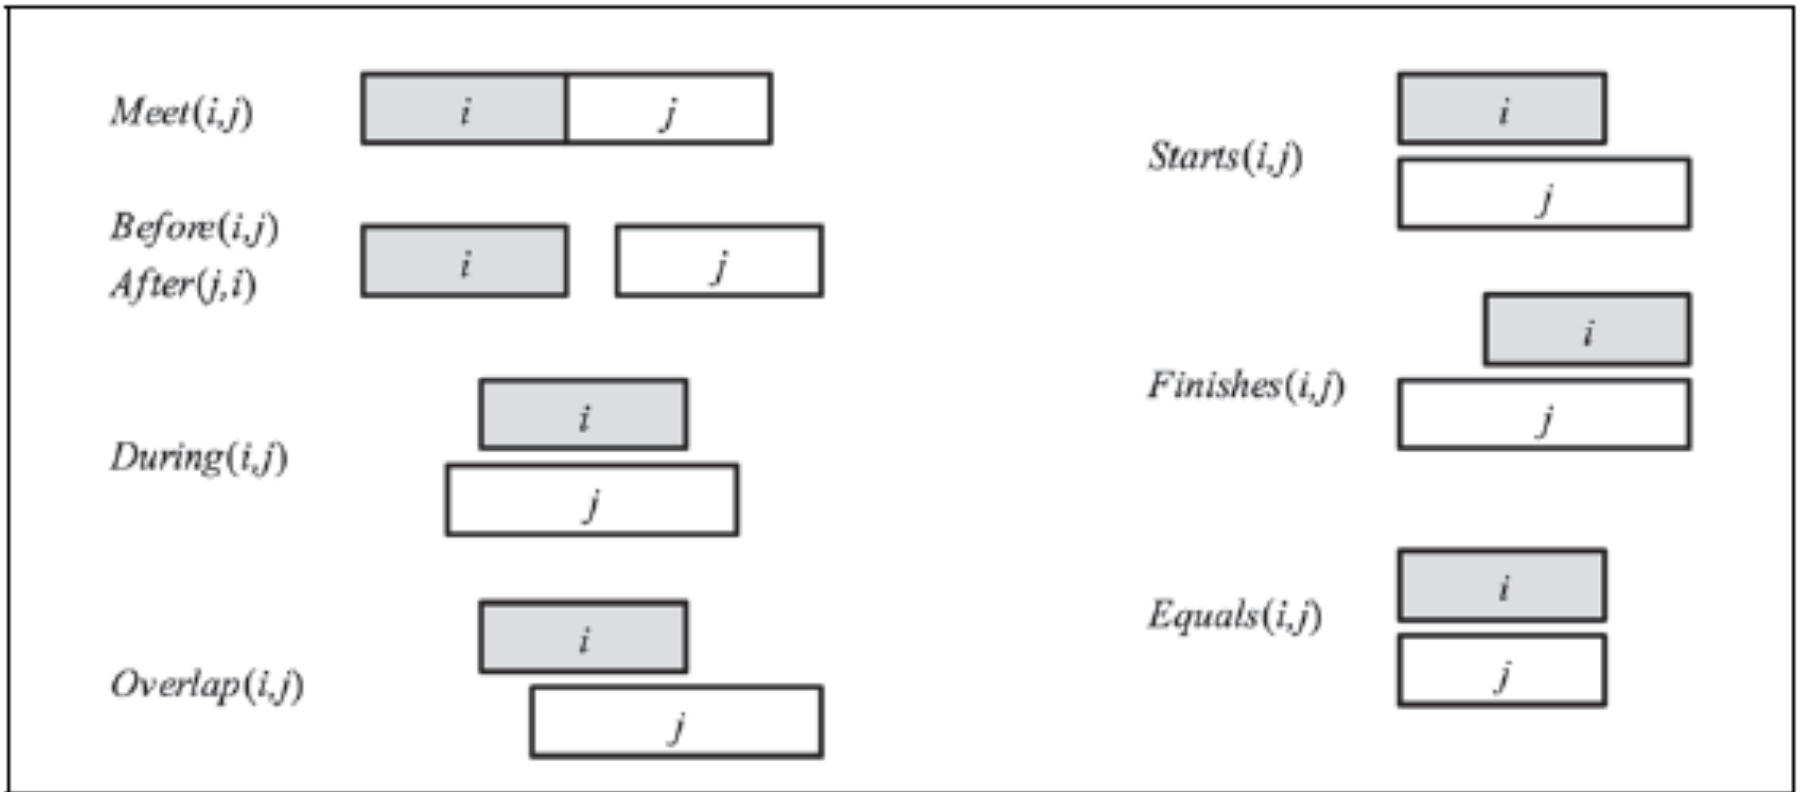


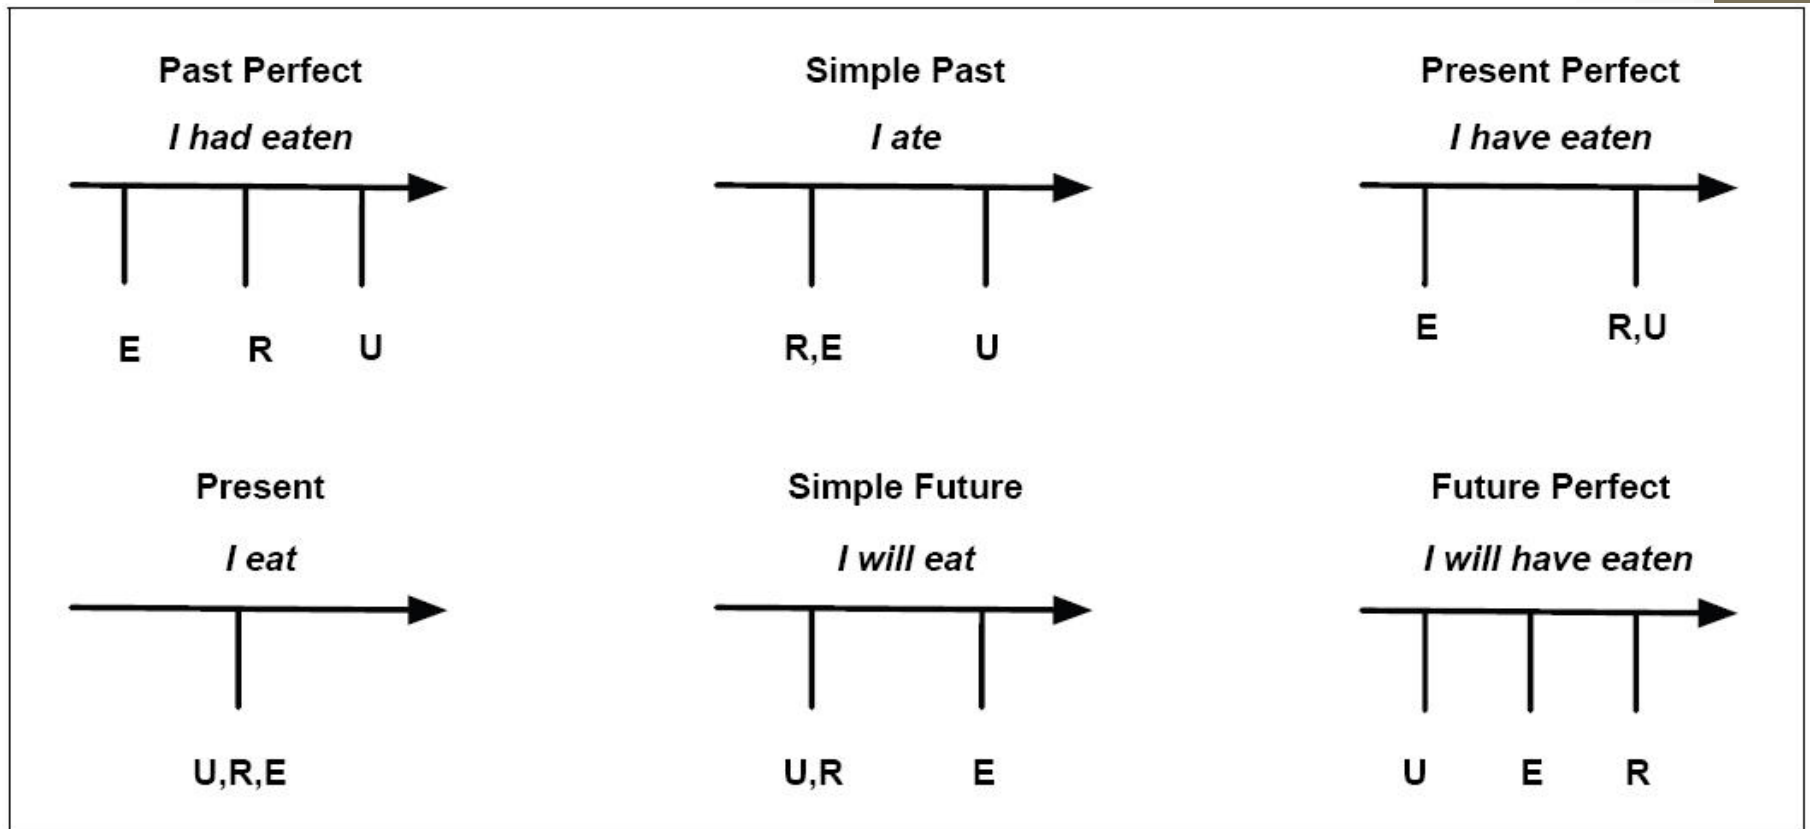
Figure 12.2 Predicates on time intervals.

Representing Time

Utterance time (U): when the utterance occurs

Reference time (R): the temporal point-of-view of the utterance

Event time (E): when events described in the utterance occur



Verbs and Event Types: Aspect

- **Statives**: states or properties of objects at a particular point in time
I am hungry.
- **Activities**: events with no clear endpoint
I am eating.
- **Accomplishments**: events with durations and endpoints that result in some change of state
I ate dinner.
- **Achievements**: events that change state but have no particular duration – they occur in an instant
I got the bill.

Beliefs, Desires and Intentions

- Very hard to represent internal speaker states like believing, knowing, wanting, assuming, imagining
 - Not well modeled by a simple DB lookup approach so..
 - 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

Believes(George, dance(George))

Knows(Bill,Believes(George,dance(George)))

- **Mutual belief:** I believe you believe I believe....
 - Practical importance: modeling belief in dialogue
 - Clark's **grounding**

Take Away

- Semantics means many things
 - Meaning representations
 - Semantic parsing or semantic role labeling
 - Determining word meaning
 - Use of semantic information for an end application
 - Inference
 - Temporal relations: tense, aspect
 - Information extraction
 - Belief, desire and intent
- Work in semantics can fall into any of these areas

Lexical Semantics

- Grounding for representation of word meaning
 - Used for word disambiguation
 - Semantic interpretation
- Pre-cursor to representations used in distributional semantics, vector representations and neural nets
- Topics
 - Homonymy, Polysemy, Synonymy
 - Online resources: WordNet

Preliminaries

- What's a word?
 - Definitions we've used over the class: Types, tokens, stems, roots, inflected forms, etc...
 - **Lexeme**: An entry in a lexicon consisting of a pairing of a form with a single meaning representation
 - **Lexicon**: A collection of lexemes

Relationships between word meanings

- Homonymy
- Polysemy
- Synonymy
- Antonymy
- Hypernymy
- Hyponymy
- Meronymy

Homonymy

- **Homonymy:**
 - Lexemes that share a form
 - Phonological, orthographic or both
 - But have unrelated, distinct meanings
 - Clear examples:
 - Bat (wooden stick-like thing) vs
 - Bat (flying scary mammal thing)
 - Or bank (financial institution) versus bank (riverside)
 - Can be homophones, homographs, or both:
 - Homophones:
 - Write and right
 - Piece and peace

Homonymy causes problems for NLP applications

- Text-to-Speech
 - Same orthographic form but different phonological form
 - bass vs bass
- Information retrieval
 - Different meanings same orthographic form
 - QUERY: bat care
- Machine Translation
- Speech recognition
 - Why?

Polysemy

- The **bank** is constructed from red brick
I withdrew the money from the **bank**
- Are those the same sense?
- Or consider the following WSJ example
 - While some banks furnish sperm only to married women, others are less restrictive
 - Which sense of bank is this?
 - Is it distinct from (homonymous with) the river bank sense?
 - How about the savings bank sense?

Polysemy

- A single lexeme with multiple related meanings (bank the building, bank the financial institution)
- Most non-rare words have multiple meanings
 - The number of meanings is related to its frequency
 - Verbs tend more to polysemy
 - Distinguishing polysemy from homonymy isn't always easy (or necessary)

Metaphor and Metonymy

- Specific types of polysemy
- Metaphor:
 - Germany will pull Slovenia out of its economic slump.
 - I spent 2 hours on that homework.
- Metonymy
 - The White House announced yesterday.
 - This chapter talks about part-of-speech tagging
 - Bank (building) and bank (financial institution)

How do we know when a word has more than one sense?

- ATIS examples
 - Which flights serve breakfast?
 - Does America West serve Philadelphia?
- The “zeugma” test:
 - ?Does United serve breakfast and San Jose?

Synonyms

- Word that have the same meaning in some or all contexts.
 - filbert / hazelnut
 - couch / sofa
 - big / large
 - automobile / car
 - vomit / throw up
 - Water / H₂O
- Two lexemes are synonyms if they can be successfully substituted for each other in all situations
 - If so they have the same **propositional meaning**

Synonyms

- But there are few (or no) examples of perfect synonymy.
 - Why should that be?
 - Even if many aspects of meaning are identical
 - Still may not preserve the acceptability based on notions of politeness, slang, register, genre, etc.
- Example:
 - **Water** and **H₂O**

Some more terminology

- Lemmas and wordforms
 - A **lexeme** is an abstract pairing of meaning and form
 - A **lemma** or **citation form** is the grammatical form that is used to represent a **lexeme**.
 - *Carpet* is the lemma for *carpets*
 - *Dormir* is the lemma for *duermes*.
 - Specific surface forms *carpets*, *sung*, *duermes* are called **wordforms**
- The lemma *bank* has two **senses**:
 - Instead, a **bank** can hold the investments in a custodial account in the client's name
 - But as agriculture burgeons on the east **bank**, the river will shrink even more.
- A **sense** is a discrete representation of one aspect of the meaning of a word

Synonymy is a relation between senses rather than words

- Consider the words *big* and *large*
- Are they synonyms?
 - How **big** is that plane?
 - Would I be flying on a **large** or small plane?
- How about here:
 - Miss Nelson, for instance, became a kind of **big** sister to Benjamin.
 - ?Miss Nelson, for instance, became a kind of **large** sister to Benjamin.
- Why?
 - *big* has a sense that means being older, or grown up
 - *large* lacks this sense

Antonyms

- Senses that are opposites with respect to one feature of their meaning
- Otherwise, they are very similar!
 - dark / light
 - short / long
 - hot / cold
 - up / down
 - in / out
- More formally: antonyms can
 - define a binary opposition or at opposite ends of a scale (*long/short, fast/slow*)
 - Be **reversives**: *rise/fall, up/down*

Hyponymy



- One sense is a **hyponym** of another if the first sense is more specific, denoting a subclass of the other
 - *car* is a hyponym of *vehicle*
 - *dog* is a hyponym of *animal*
 - *mango* is a hyponym of *fruit*
- Conversely
 - *vehicle* is a hypernym/superordinate of *car*
 - *animal* is a hypernym of *dog*
 - *fruit* is a hypernym of *mango*

superordinate	vehicle	fruit	furniture	mammal
hyponym	car	mango	chair	dog

Hypernymy more formally

- Extensional:
 - The class denoted by the superordinate
 - extensionally includes the class denoted by the hyponym
- Entailment:
 - A sense A is a hyponym of sense B if being an A entails being a B
- Hyponymy is usually transitive
 - (A hypo B and B hypo C entails A hypo C)

- Why would hypernyms/hyponyms be important to constructing a meaning representation?



Why would hypernyms/hyponyms be important for meaning representation?

II. WordNet

- A hierarchically organized lexical database
- On-line thesaurus + aspects of a dictionary
 - Versions for other languages are under development

Category	Unique Forms
Noun	117,097
Verb	11,488
Adjective	22,141
Adverb	4,601

WordNet

- Where it is:
 - <https://wordnet.princeton.edu/>

Format of Wordnet Entries

The noun “bass” has 8 senses in WordNet.

1. bass¹ - (the lowest part of the musical range)
2. bass², bass part¹ - (the lowest part in polyphonic music)
3. bass³, basso¹ - (an adult male singer with the lowest voice)
4. sea bass¹, bass⁴ - (the lean flesh of a saltwater fish of the family Serranidae)
5. freshwater bass¹, bass⁵ - (any of various North American freshwater fish with lean flesh (especially of the genus *Micropterus*))
6. bass⁶, bass voice¹, basso² - (the lowest adult male singing voice)
7. bass⁷ - (the member with the lowest range of a family of musical instruments)
8. bass⁸ - (nontechnical name for any of numerous edible marine and freshwater spiny-finned fishes)

The adjective “bass” has 1 sense in WordNet.

1. bass¹, deep⁶ - (having or denoting a low vocal or instrumental range)
“a deep voice”; *“a bass voice is lower than a baritone voice”*;
“a bass clarinet”

WordNet Noun Relations

Relation	Also called	Definition	Example
Hypernym	Superordinate	From concepts to superordinates	<i>breakfast</i> ¹ → <i>meal</i> ¹
Hyponym	Subordinate	From concepts to subtypes	<i>meal</i> ¹ → <i>lunch</i> ¹
Member Meronym	Has-Member	From groups to their members	<i>faculty</i> ² → <i>professor</i> ¹
Has-Instance		From concepts to instances of the concept	<i>composer</i> ¹ → <i>Bach</i> ¹
Instance		From instances to their concepts	<i>Austen</i> ¹ → <i>author</i> ¹
Member Holonym	Member-Of	From members to their groups	<i>copilot</i> ¹ → <i>crew</i> ¹
Part Meronym	Has-Part	From wholes to parts	<i>table</i> ² → <i>leg</i> ³
Part Holonym	Part-Of	From parts to wholes	<i>course</i> ⁷ → <i>meal</i> ¹
Antonym		Opposites	<i>leader</i> ¹ → <i>follower</i> ¹

WordNet Verb Relations

Relation	Definition	Example
Hypernym	From events to superordinate events	<i>fly</i> ⁹ → <i>travel</i> ⁹
Troponym	From a verb (event) to a specific manner elaboration of that verb	<i>walk</i> ¹ → <i>stroll</i> ¹
Entails	From verbs (events) to the verbs (events) they entail	<i>snore</i> ¹ → <i>sleep</i> ¹
Antonym	Opposites	<i>increase</i> ¹ ⇔ <i>decrease</i> ¹

WordNet Hierarchies

Sense 3

bass, basso --

(an adult male singer with the lowest voice)

=> singer, vocalist, vocalizer, vocaliser

=> musician, instrumentalist, player

=> performer, performing artist

=> entertainer

=> person, individual, someone...

=> organism, being

=> living thing, animate thing,

=> whole, unit

=> object, physical object

=> physical entity

=> entity

=> causal agent, cause, causal agency

=> physical entity

=> entity

Sense 7

bass --

(the member with the lowest range of a family of musical instruments)

=> musical instrument, instrument

=> device

=> instrumentality, instrumentation

=> artifact, artefact

=> whole, unit

=> object, physical object

=> physical entity

How is “sense” defined in WordNet?

- The set of near-synonyms for a WordNet sense is called a **synset (synonym set)**; it's their version of a sense or a concept
- Example: **chump** as a noun to mean
 - ‘a person who is gullible and easy to take advantage of’

{chump¹, fool², gull¹, mark⁹, patsy¹, fall guy¹, sucker¹, soft touch¹, mug²}

- Each of these senses share this same gloss
- Thus for WordNet, the meaning of this sense of **chump** is this list.

Wordnet example

Questions?