Semantic Graphs

CSE 40657/60657: Natural Language Processing

Representing Meaning

1. “The boy wants the girl to believe him.”
2. “The boy desires the girl to believe him.”
3. “The boy desires to be believed by the girl.”
4. “The boy has a desire to be believed by the girl.”
5. “The boy’s desire is for the girl to believe him.”
6. “The boy is desirous of the girl believing him.”

Representing Meaning

- All of these sentences have the same logical meaning
- Can we represent this meaning with a single formal structure?
- Can we abstract away from morphological and syntactic variability?

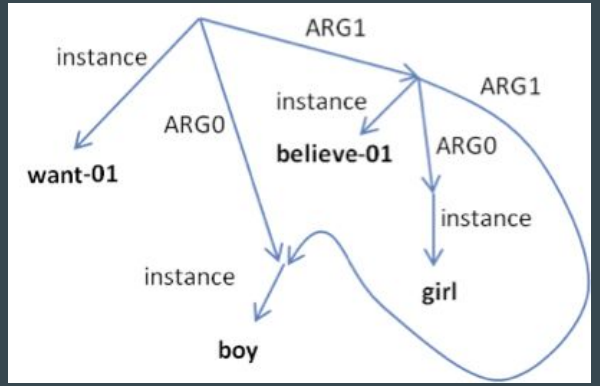

“The boy wants the girl to believe him.”

Diathesis Alternations

- Active vs. passive voice
  - Active: Grammatical subject is the agent
  - Passive: Grammatical object is the agent
- Parse trees do not capture this similarity

1. “John broke the window.”
2. “The window was broken by John.”

Both sentences imply that there as act of breaking and that John is the breaker and the window is the thing broken.

Break(John, window)

Diathesis Alternations

- Verbs often have more arguments beyond subjects and objects
- The verb “break” seems to have multiple ways of realizing its arguments
  - Agent (subject): the thing doing the breaking
  - Theme (object): the thing broken
  - Instrument: the thing used to do the breaking
- These are called “diathesis alternations”

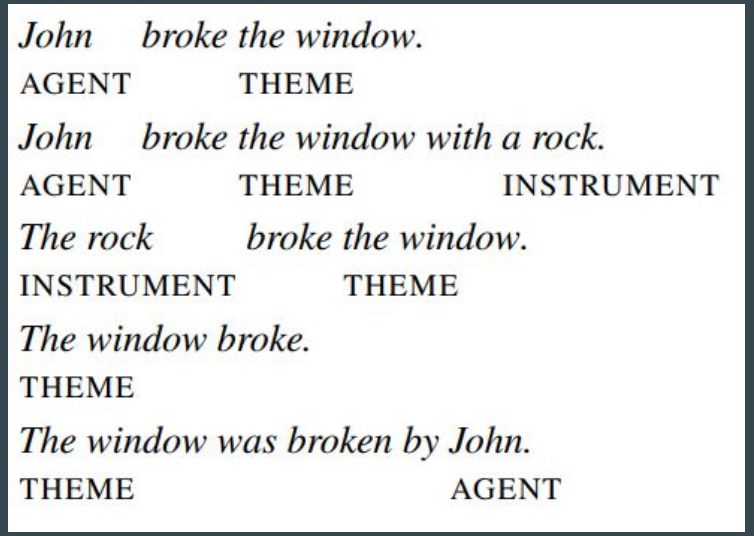

Why Semantics Matters

- Syntax is not enough to link linguistic elements to non-linguistic knowledge of the world
- Coreferences (pronouns)
- Event extraction
- Question answering
  - “Maharani is a vegetarian restaurant.”
  - “Is Maharani a vegetarian restaurant?”
  - “Does Maharani serve vegetarian food?”
  - It would be nice if we had a representation where we could flip a switch to make a question a statement or command

Example: Statements to Questions

“John is walking to the store.”	“Is John walking to the store?”
“John has walked to the store.”	“Has John walked to the store?”
“John will walk to the store.”	“Will John walk to the store?”
“John walked to the store.”	“Did John walk to the store?”
“John went to the store.”	“Did John go to the store?”

Semantic Ambiguity

- The ambiguity in these examples has to do with coordination
- Some meanings are more probable than others, even though the parallel use of “her” in the last two items makes the incorrect reading seem more likely
- A model of semantics could help disambiguate cases like these

Students Cook & Serve Grandparents

On Thursday, September 9, Gorman School hosted the first annual Grandparent's Day.

All Grandparents were invited to a school wide pancake breakfast. Upper grade students served as excellent chefs, as well as taking responsibility for serving the food and the clean up after-

Meaning Representation Banks?

The Penn TreeBank is nice because the task is on whole sentences. This is opposed to treating tasks as being separate, like prep phrase attachment, verb-argument dependencies, etc. Those smaller tasks are naturally solved as a byproduct of whole-sentence parsing and are solved better than when approached in isolation.

A meaning representation bank could do the same thing for semantics, for tasks like named entity recognition, coreference resolution, semantic relations, discourse connectives, temporal entities, etc.

Vauquois Triangle

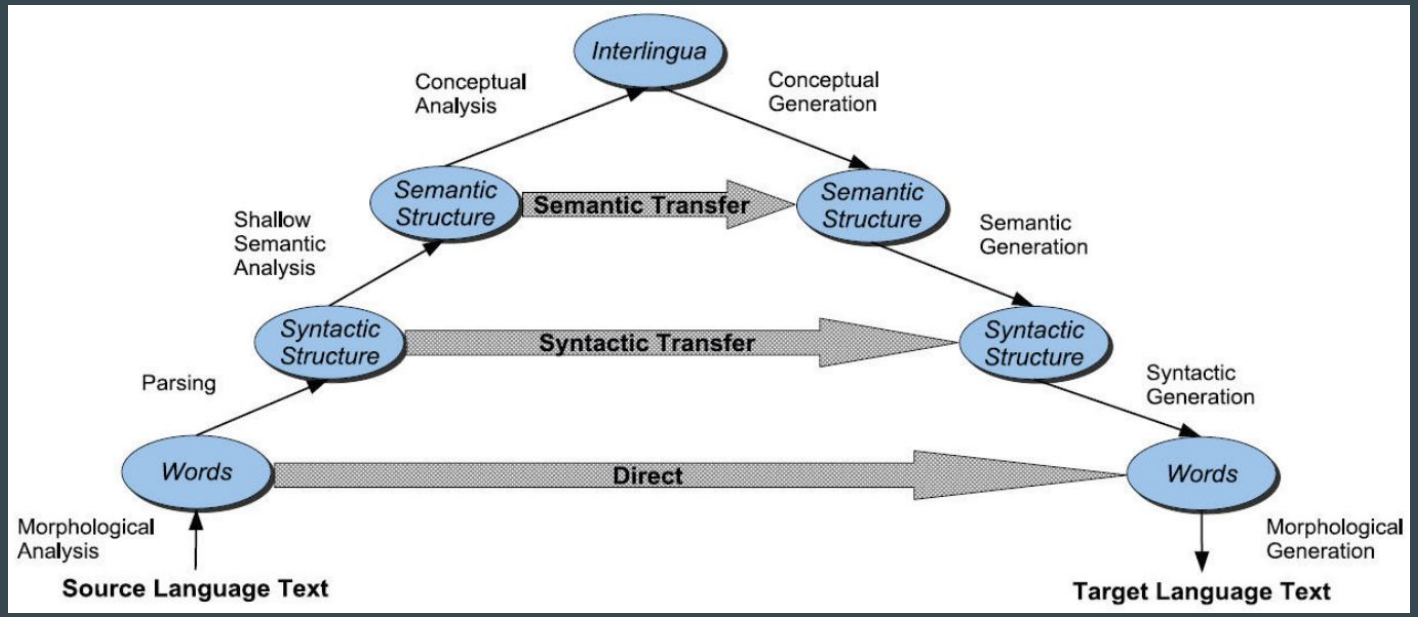

Okay, so we want to model semantics

- We started with sequences of words and n-gram language models to model surface forms
- Then we moved to trees and grammars to model syntax, which is shallower than semantics
- For semantics... what?

Semantic Representations

- First-order logic
- Semantic graphs
  - Abstract Meaning Representation (AMR)
- Frames

Dual perspective: represents meaning of language and state of affairs in world, allowing us to link the two

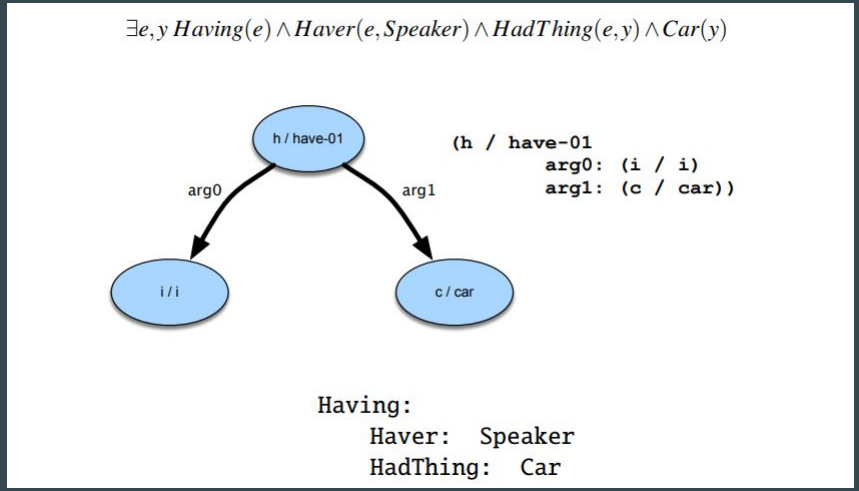

"I have a car."

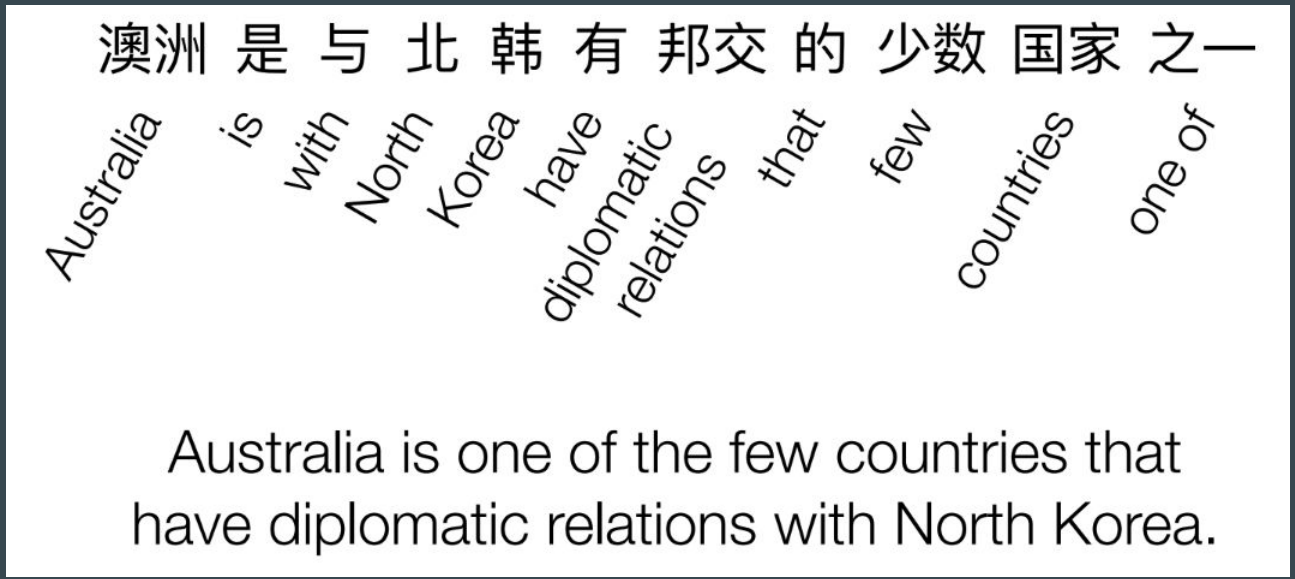
澳洲 是 与 北 韩 有 邦交 的 少数 国家 之一
Australia
is
with
North
Korea
have
diplomatic
relations
that
few
countries
one of
Australia is one of the few countries that
have diplomatic relations with North Korea.

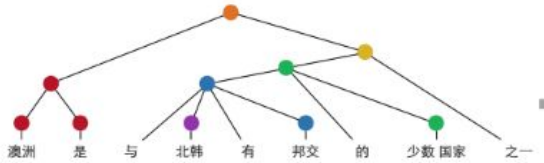

translate

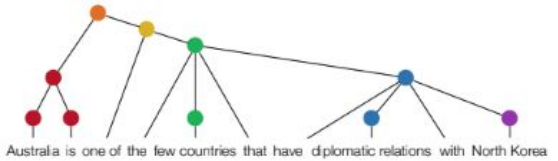

parse

language model

澳洲是与北韩有邦交的少数国家之一

Australia is one of the few countries that have diplomatic relations with North Korea

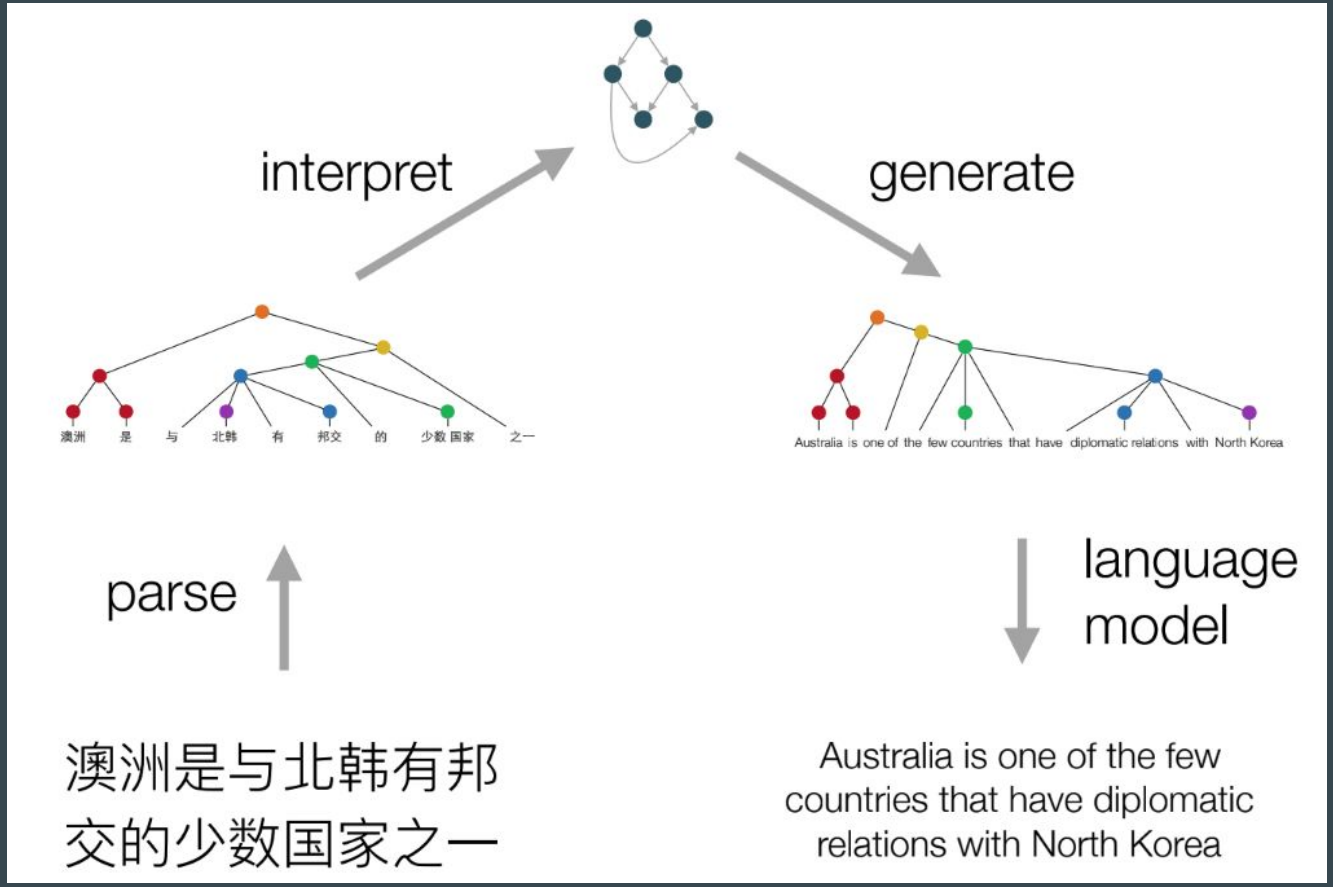
interpret
generate
澳洲 是 与 北韩 有 邦交 的 少数 国家 之一
Australia is one of the few countries that have diplomatic relations with North Korea
parse
language model
澳洲是与北韩有邦交的少数国家之一
Australia is one of the few countries that have diplomatic relations with North Korea

Sentences as Graphs

- Try convert these sentences into graphs
  a. “John went to the store.”
  b. “John gave Mary the book.”
  c. “The boy wants the girl to believe him.”
- What are the vertices, and what are the edges?

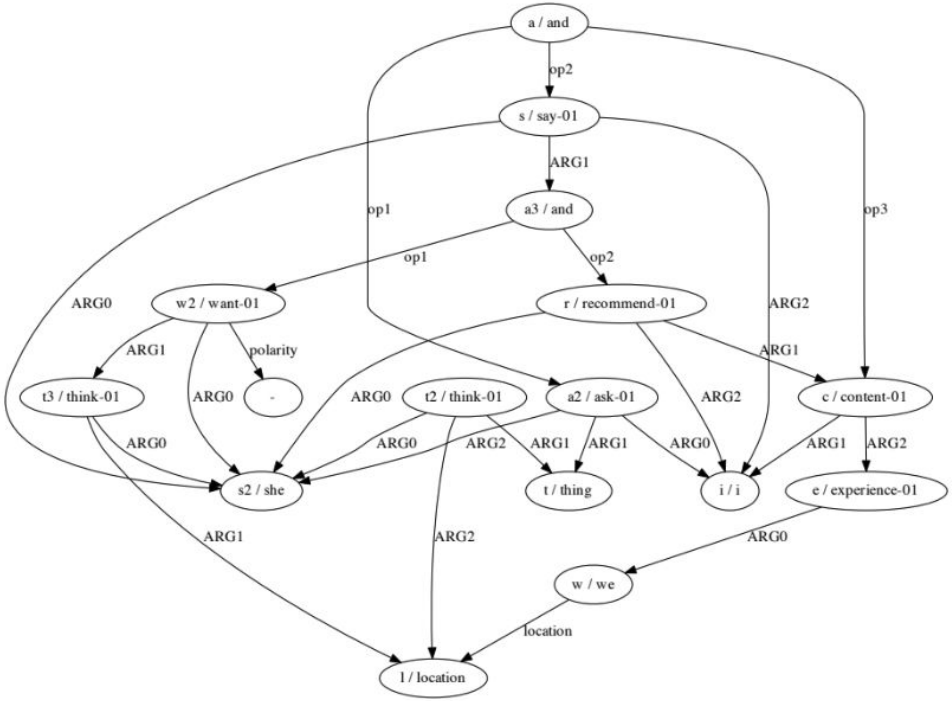

I asked her what she thought about where we'd be and she said she doesn't want to think about that, and that I should be happy about the experiences we've had (which I am).

Neo-Davidsonian Event Representations

- The verb “eat” seems to be able to take a changing number of arguments
- But edges only connect two vertices at a time (they have a fixed arity of 2)
- Solution: treat the verb as a variable
- Accordingly, AMRs treat the verb as a vertex

“I ate a turkey sandwich.”
“I ate a turkey sandwich at my desk.”
“I ate at my desk.”
“I ate lunch.”
“I ate a turkey sandwich for lunch.”
“I ate a turkey sandwich for lunch at my desk.”
“Eating a turkey sandwich is nutritious.”

Eat(Speaker, TurkeySandwich, Lunch, Desk, ...?)

$\exists$ e Eating(e) $\wedge$ Eater(e, Speaker) $\wedge$ Eaten(e, TurkeySandwich) $\wedge$ Meal(e, Lunch) $\wedge$ Location(e, Desk)

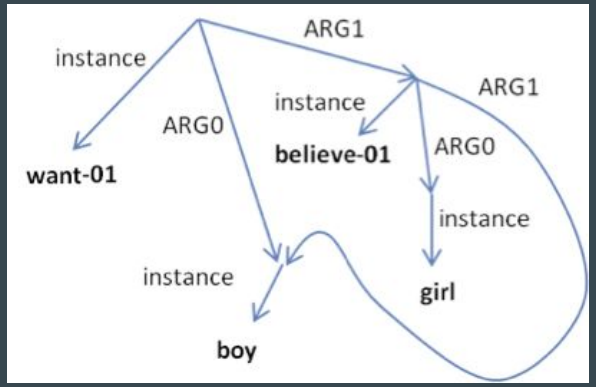

“The boy wants the girl to believe him.”

Representing Verbs

- The exact meaning of verb arguments tends to be verb-specific
- We need to slice up verbs into separate senses, but where do we stop?
- PropBank
  - Sentences annotated with semantic roles (includes Penn TreeBank)
  - Arguments are arbitrarily labeled Arg1, Arg2, Arg3, ...
  - Arg0 tends to refer to subjects, Arg1 to objects, Arg2 to instruments, etc.
  - There are general-purpose ArgMs for things with stable meaning like time, location, reason, etc.

1. “I chatted with friends.”
2. “I broke the window with a rock.”

https://verbs.colorado.edu/verb-index/vn3.3/search.php

Verb Frames

- Frames generalize semantic roles to nouns that represent actions
- FrameNet
  - A frame is a background knowledge structure that defines a set of frame-specific semantic roles called frame elements and includes a set of predicates that use these roles
  - Multiple words (verbs or nouns) can map to the same frame and evoke some aspect of the frame

[$_{Arg1}$ The price of bananas] increased [$_{Arg2}$ 5%].
[$_{Arg1}$ The price of bananas] rose [$_{Arg2}$ 5%].
There has been a [$_{Arg2}$ 5%] rise [$_{Arg1}$ in the price of bananas].

https://verbs.colorado.edu/verb-index/vn3.3/search.php

Abstract Meaning Representation (AMR)

- Graph model
  - Rooted
  - Directed
  - Edge-labeled
  - Leaf-labeled
- AMR concepts are either English words (“boy”), PropBank framesets (“want-01”), or special keywords
- Often written as text in Penman notation
  - Variables allow reentrancy

LOGIC format:

∃ w, b, g:
instance(w, want-01) ∧ instance(g, go-01) ∧
instance(b, boy) ∧ arg0(w, b) ∧
arg1(w, g) ∧ arg0(g, b)

AMR format (based on PENMAN):

```
(w / want-01
  :arg0 (b / boy)
  :arg1 (g / go-01
           :arg0 b))
```

GRAPH format:

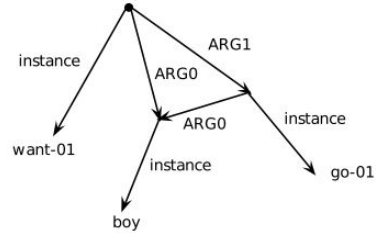

Figure 1: Equivalent formats for representating the meaning of “The boy wants to go”.

What does this AMR mean?

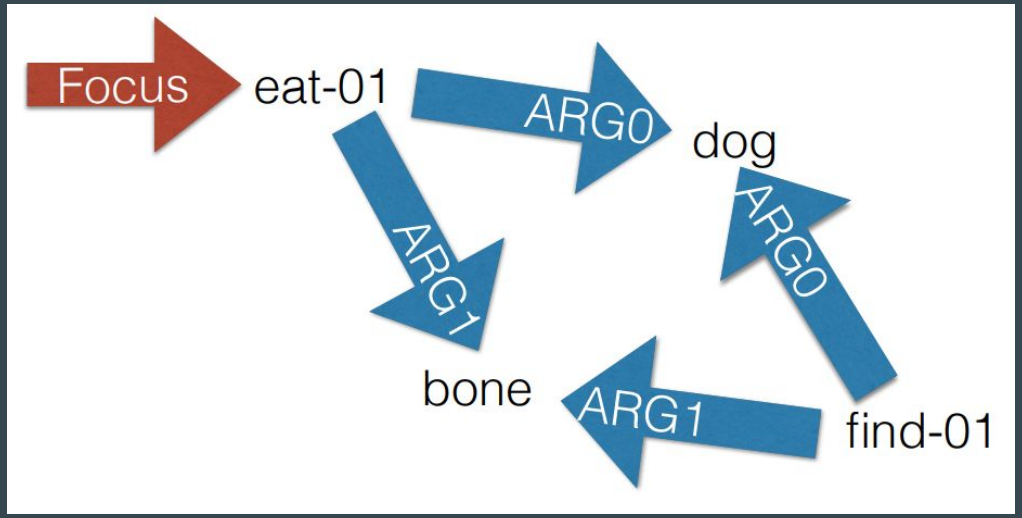

Translate this sentence to AMR

"Rachael Ray finds inspiration in cooking, her family, and her dog."

```
(i / inspire-01
   :ARG0 (a / and
            :op1 (c / cook-01
                    :ARG0 p)
            :op2 (f / family
                   :ARG1-of (h / have-org-role-91
                               :ARG0 p
                               :ARG2 (m / member)))
            :op3 (d / dog :poss p))
   :ARG1 (p / person :name (n / name :op1 "Rachael" :op2 "Ray")))
```

Representing Questions

```
(f / find-01
   :ARG0 (g / girl)
   :ARG1 (a / amr-unknown))
```
“What did the girl find?”

```
(r / run-01
   :ARG0 (g / girl)
   :manner (f / fast
              :degree (a / amr-unknown)))
```
“How fast did the girl run?”

```
(s / see-01
   :ARG0 (g / girl)
   :ARG1 (a / amr-unknown
            :ARG1-of (p / purple-02)))
```
“What purple thing did the girl see?”

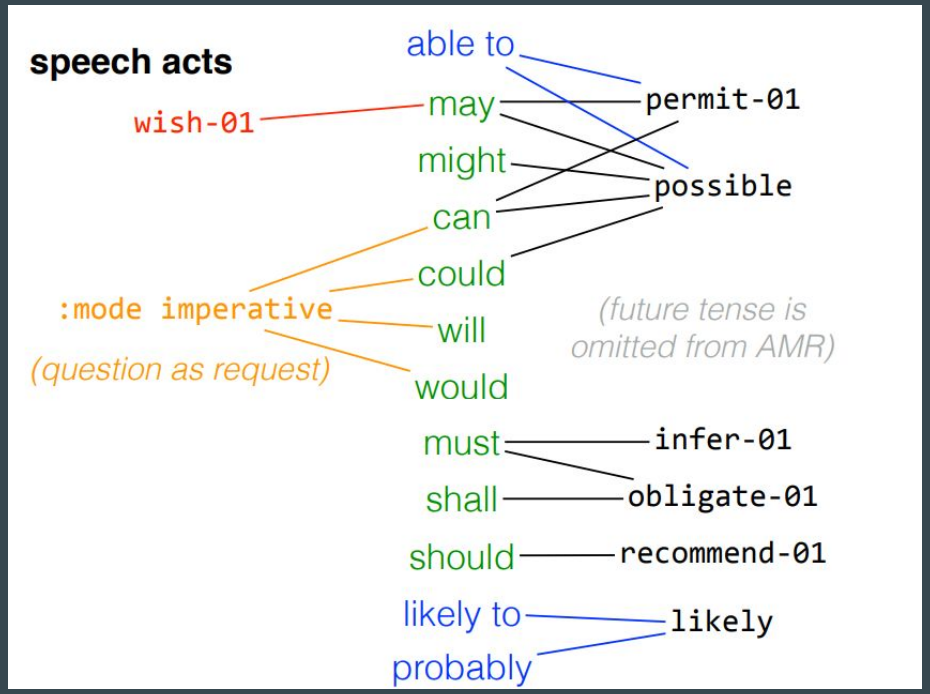
speech acts
wish-01
:mode imperative
(question as request)
able to
may
might
can
could
will
would
must
shall
should
likely to
probably
permit-01
possible
infer-01
obligate-01
recommend-01
likely
(future tense is omitted from AMR)

AMR Properties

- Rooted labeled graphs that are easy for humans and programs to use
- Abstract away syntax; different sentences with the same meaning should have the same AMR
- Not invertible; cannot recover the original surface form
- Make use of framesets from PropBank
- Info about how the AMR was derived is intentionally not preserved, e.g. alignments, ordering of rules applied
- Heavily biased toward English; NOT an interlingua

Where AMR stands

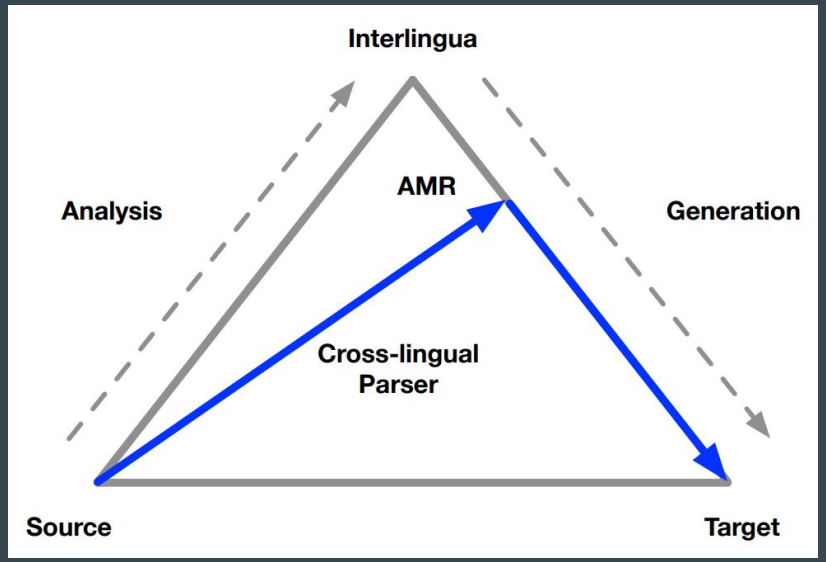

AMR Parsing

- Active research area
  - Because it's 2018, it's all neural
- Approaches
  - Learn alignments, then identify concepts (vertices), then identify relations among concepts (edges)
    - Graph is initially dense with weights for all edges
    - Edges are eliminated based on score and on graph constraints (preserving, simple, spanning, connected, deterministic)
  - Neural network that jointly learns alignments, concepts, and relations (Lyu and Titov)
  - Use neural sequence-to-sequence models to learn to translate sentences to linearized versions of AMRs (Konstas et al. 2017, Viet et al. 2017)
  - Use a neural network that acts like a stack (Stack-LSTM) to learn sequences of operations to transform strings into AMRs (Ballesteros and Al-Onaizan 2017)

Graph Grammars

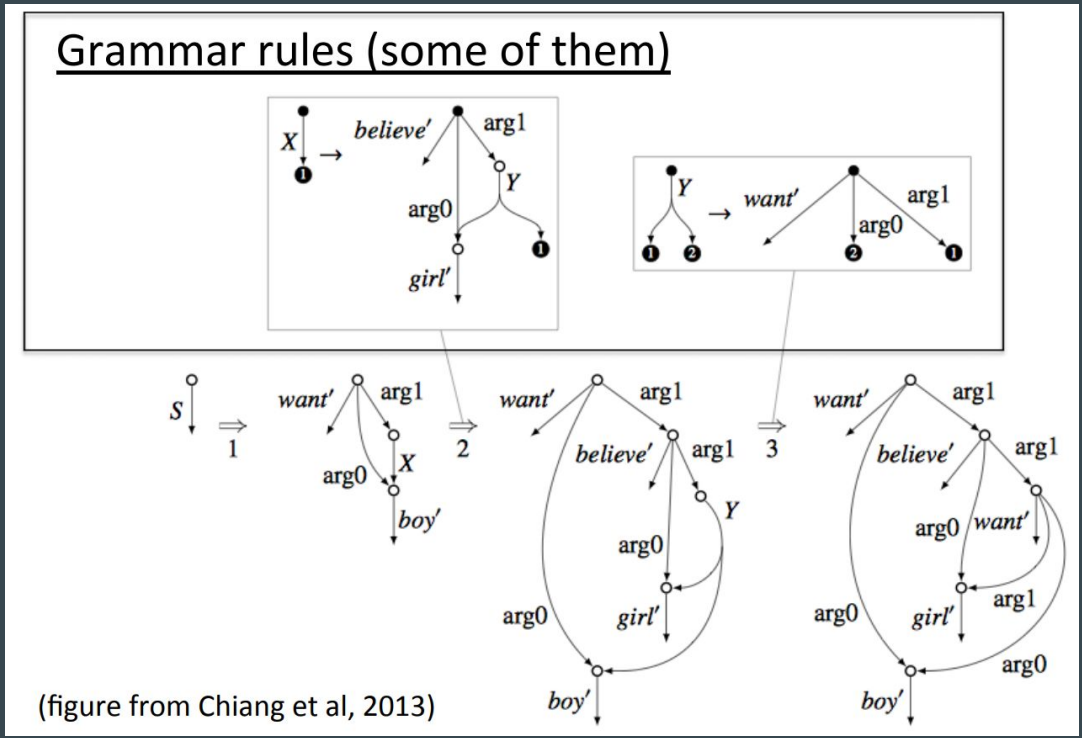

(figure from Chiang et al, 2013)

DAG Automata