Natural Language Processing

Info 159/259

Lecture 14: Semantic roles (March 11, 2024)

Many slides & instruction ideas borrowed from:
David Bamman, Collin Baker & Dan Jurafsky

Why is syntax important?

- Foundation for semantic analysis (on many levels of representation: semantic roles, compositional semantics, frame semantics)

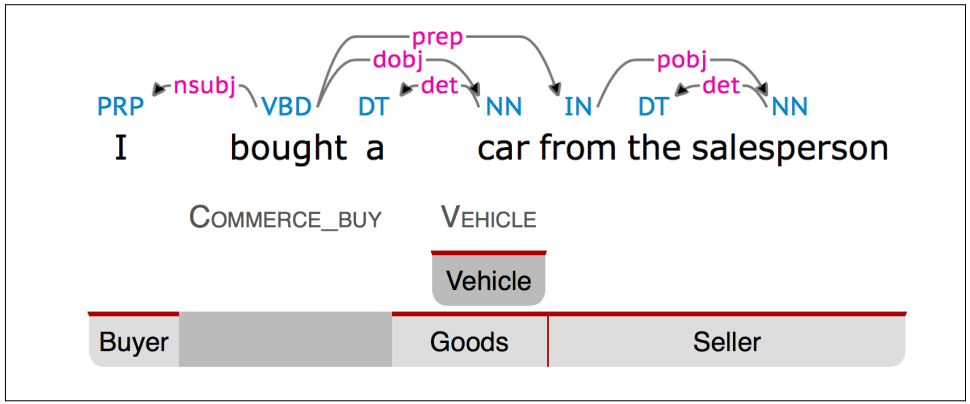

Why is syntax insufficient?

- Syntax encodes the structure of language but doesn't directly address meaning.

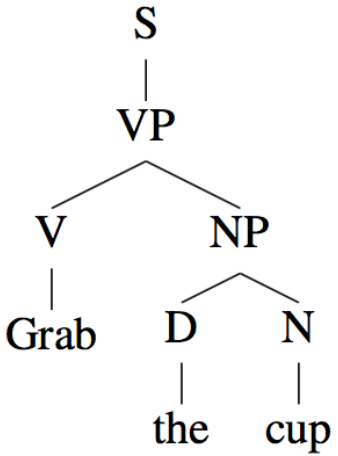

- Syntax alone doesn't ground "grab" in an action to take in the world.

Lexical semantics

- Vector representation that encodes information about the distribution of contexts a word appears in
- Words that appear in similar contexts have similar representations (and similar meanings, by the distributional hypothesis).
- We can represent what individual words “mean” as a function of what other words they’re related to (but that’s still not grounding).

Lexical semantics

grab	1
throw	0.824
pull	0.818
knock	0.799
grabbing	0.789
steal	0.787
pulling	0.764
grabs	0.756
away	0.746
catch	0.74

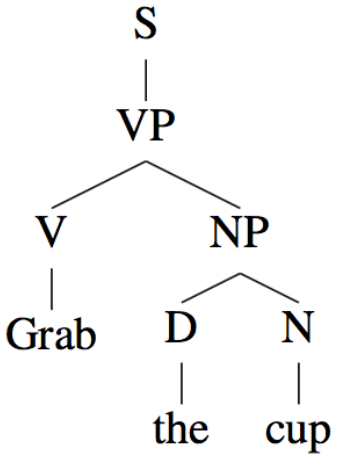

- “Grab” = execute GrabbingFunction()
- “the cup” = object ID 9AF1948A81CD22

Semantics

Lexical semantics is concerned with representing the meaning of words (and their relations)

Logical semantics is concerned with representing the meaning of sentences.

Meaning representation

A meaning representation should be unambiguous; each statement in a meaning representation should have one meaning.

Meaning representation

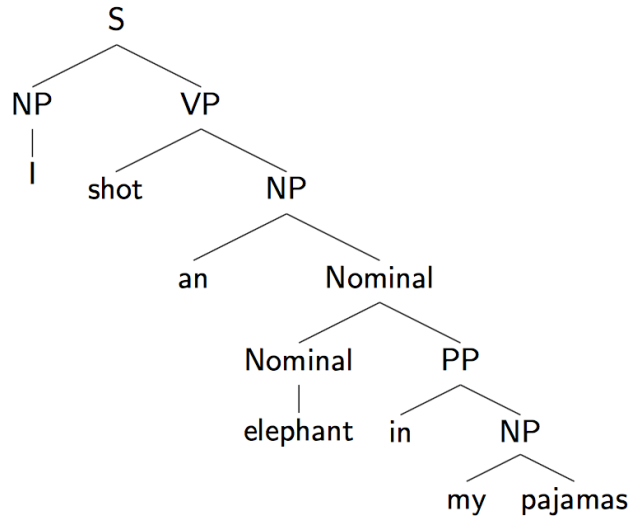

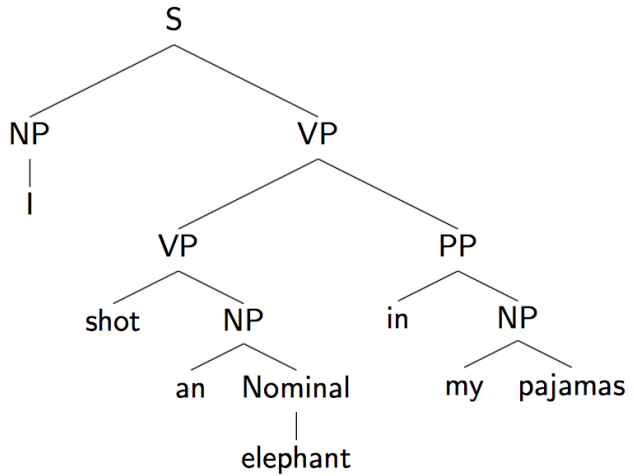

Syntax resolves some ambiguity

The Office (2005)

“Once every hour, someone is involved in an Internet scam”

“That person is Michael Scott”

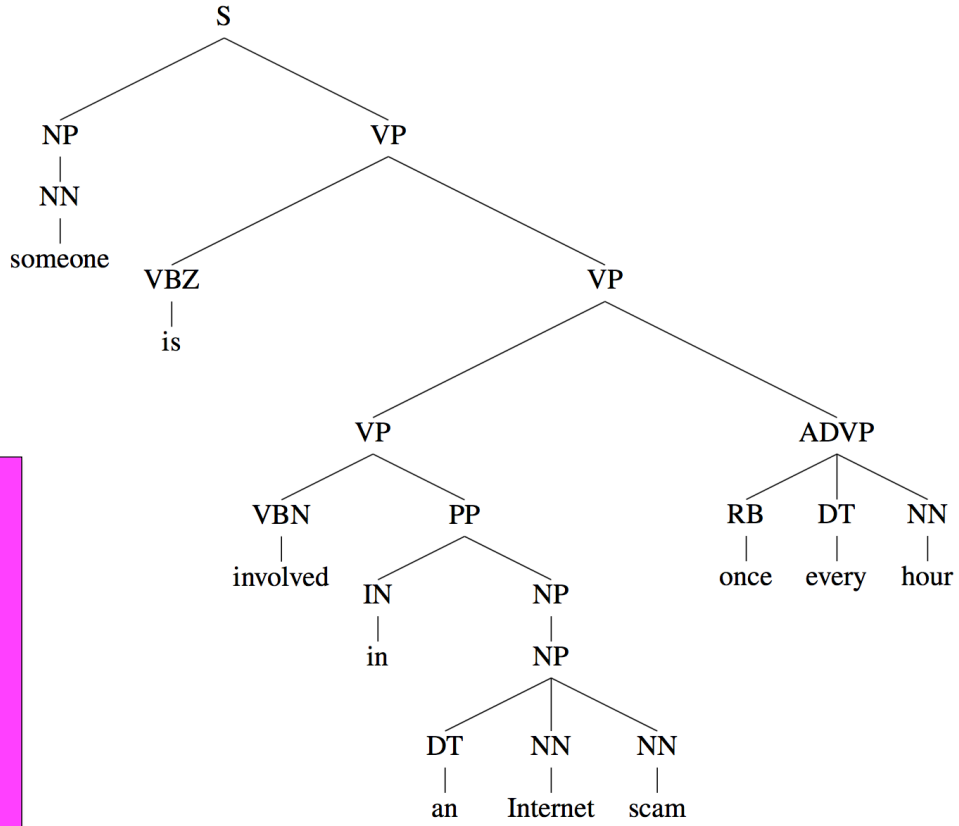

Same structure for “someone” meaning:

- Some person for some scam
- One person in the same scam (Michael Scott)

First-order logic (FOL)

- We want to represent every sentence as an unambiguous proposition in FOL

Sentence	*Luke was fighting with Darth Vader*
FOL	FIGHT(LUKE, VADER)

- Predicate-Arguments

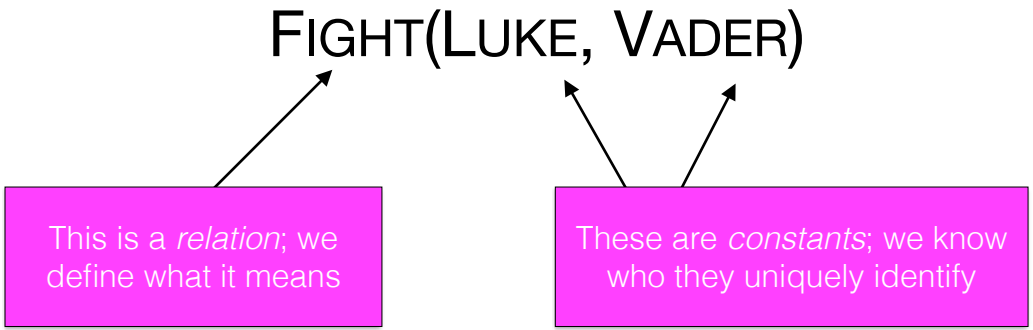
FIGHT(LUKE, VADER)
This is a *relation*; we define what it means
These are *constants*; we know who they uniquely identify

First-order logic (FOL)

- How we map a natural language sentence to FOL is the task of semantic parsing; but we define the FOL relations and entities to be sensitive to what matters in our model.

Sentence	*Luke battled Vader*
Sentence	*Luke fought with Vader*
Sentence	*Skywalker dueled with Darth Vader*
Sentence	*Luke was fighting with Darth Vader*
FOL	FIGHT(LUKE, VADER)

First-order logic (FOL)

- How we map a natural language sentence to FOL is the task of semantic parsing; but we define the FOL relations and entities to be sensitive to what matters in our world model.

Maybe in our star wars model we want to preserve the difference between *fighting* and *dueling*

Sentence	*Luke battled Vader*
Sentence	*Luke fought with Vader*
Sentence	*Luke was fighting with Darth Vader*
FOL	FIGHT(LUKE, VADER)

Sentence	*Skywalker dueled with Darth Vader*
FOL	DUEL(LUKE, VADER)

First-order logic (FOL)

- How we map a natural language sentence to FOL is the task of semantic parsing; but we define the FOL relations and entities to be sensitive to what matters in our world model.

For a robot model, there is only one possible "grabbing" kind of action, so subtle differences don't matter.

Sentence	*Grab the cup*
Sentence	*Snatch the cup!*
Sentence	*Take the cup*
FOL	GRAB(ROBOT, CUP)

Relations

- N-ary relations hold among FOL terms (constants, variables, functions).

Unary (property)	HUMAN(LUKE), ROBOT(C-3PO)
binary relation	FIGHTS(LUKE, VADER)
ternary	GIVES(OBI-WAN, LUKE, LUKE_LIGHTSABER_1)
…	*…*

Event semantics

Pat gives Sal a book

$\exists x.book(x) \wedge \text{GIVE}(Pat,Sal,x)$

Event semantics

Yesterday, Pat gave Sal a book reluctantly

$\exists x.book(x) \wedge \text{GIVE}(\text{Pat, Sal, x, yesterday, reluctantly})$

- One option: extend the arity of the relation (require more arguments)
- But that's not great because we need a separate predicate for every possible combination of arguments (even those that aren't required).

Eisenstein 2017

Event semantics

We can make the event a variable of its own, and then use it as an argument in other relations.

$\exists e,x.\text{GIVE-EVENT}(e)$
$\wedge\ \text{GIVER}(e,\text{Pat})$
$\wedge\ \text{GIFT}(e,x)$
$\wedge\ \text{BOOK}(x)$
$\wedge\ \text{RECIPIENT}(e,\text{Sal})$
$\wedge\ \text{TIME}(e,\text{yesterday})$
$\wedge\ \text{MANNER}(e,\text{reluctanctly})$

Eisenstein 2017

Event semantics

The event is central, and relations are predicated of the event. Each argument of an event holds its own relation.

$\exists e,x.\text{GIVE-EVENT}(e)$
$\wedge\ \text{GIVER}(e,\text{Pat})$
$\wedge\ \text{GIFT}(e,x)$
$\wedge\ \text{BOOK}(x)$
$\wedge\ \text{RECIPIENT}(e,\text{Sal})$
$\wedge\ \text{TIME}(e,\text{yesterday})$
$\wedge\ \text{MANNER}(e,\text{reluctanctly})$

In model-theoretic semantics, each of these has a denotation in a world model

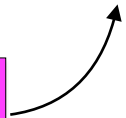

Eisenstein 2017

Event semantics

Sasha broke the window

$$\exists e,y.\text{BREAKING-EVENT}(e) \wedge \text{BREAKER}(e,\text{Sasha}) \wedge \text{BROKEN-THING}(e,y) \wedge \text{WINDOW}(y)$$

Pat opened the door

$$\exists e,y.\text{OPENING-EVENT}(e) \wedge \text{OPENER}(e,\text{Pat}) \wedge \text{OPENED-THING}(e,y) \wedge \text{DOOR}(y)$$

Event semantics

In model-theoretic semantics, each of these has some denotation in the world model.

Example: WINDOW has a identifier in some knowledge base (e.g., Wikidata) uniquely identifying its properties.

$\exists$e,y.BREAKING-EVENT(e)
$\wedge$ BREAKER(e,Sasha)
$\wedge$ BROKEN-THING(e,y)
$\wedge$ WINDOW(y)

$\exists$e,y.OPENING-EVENT(e)
$\wedge$ OPENER(e,Pat)
$\wedge$ OPENED-THING(e,y)
$\wedge$ DOOR(y)

Event semantics

This requires a comprehensive representation of the world

∃e,y.BREAKING-EVENT(e)
∧ BREAKER(e,Sasha)
∧ BROKEN-THING(e,y)
∧ WINDOW(y)

∃e,y.OPENING-EVENT(e)
∧ OPENER(e,Pat)
∧ OPENED-THING(e,y)
∧ DOOR(y)

Shallow semantics

Sasha broke the window

∃e,y.EVENT(e)
∧ CAUSER-OF-ACTION(e,Sasha)
∧ RECIPIENT-OF-ACTION(e,y)
∧ “window”(y)

∃e,y.BREAKING-EVENT(e)
∧ BREAKER(e,Sasha)
∧ BROKEN-THING(e,y)
∧ WINDOW(y)

Pat opened the door

∃e,y.EVENT(e)
∧ CAUSER-OF-ACTION(e,Pat)
∧ RECIPIENT-OF-ACTION(e,y)
∧ “door”(y)

∃e,y.OPENING-EVENT(e)
∧ OPENER(e,Pat)
∧ OPENED-THING(e,y)
∧ DOOR(y)

These roles have a lot in common: direct causal responsibility for the events.

Shallow semantics

Sasha broke the window

Agent: Sasha
Theme: window

$\exists e,y.$BREAKING-EVENT(e)
$\wedge$ BREAKER$(e,$Sasha$)$
$\wedge$ BROKEN-THING(e,y)
$\wedge$ WINDOW(y)

Pat opened the door

Agent: Pat
Theme: door

$\exists e,y.$OPENING-EVENT(e)
$\wedge$ OPENER$(e,$Pat$)$
$\wedge$ OPENED-THING(e,y)
$\wedge$ DOOR(y)

Thematic roles

Thematic roles capture the semantic commonality among arguments for different relations (predicates)

- John broke the window
- The window was broken by John

Different syntactic roles, but the same thematic role.

Shallow semantics

Sasha broke the window

∃e,y.EVENT(e)
∧ AGENT(e,Sasha)
∧ THEME(e,y)
∧ “window”(y)

∃e,y.BREAKING-EVENT(e)
∧ BREAKER(e,Sasha)
∧ BROKEN-THING(e,y)
∧ WINDOW(y)

Pat opened the door

∃e,y.EVENT(e)
∧ AGENT(e,Pat)
∧ THEME(e,y)
∧ “door”(y)

∃e,y.OPENING-EVENT(e)
∧ OPENER(e,Pat)
∧ OPENED-THING(e,y)
∧ DOOR(y)

Thematic roles

Agent	The volitional causer of an event
Experiencer	The experiencer of an event
Force	The non-volitional causer of the event
Theme	The participant most directly affected by an event
Result	The end product of an event
Content	The proposition or content of a propositional event
Instrument	An instrument used in an event
Beneficiary	The beneficiary of an event
Source	The origin of the object of a transfer event
Goal	The destination of an object of a transfer event

Thematic roles

Agent	*The waiter* spilled the soup.
Experiencer	*John* has a headache.
Force	*The wind* blows debris from the mall into our yards.
Theme	Only after Benjamin Franklin broke *the ice*...
Result	The city built *a regulation-size baseball diamond*...
Content	Mona asked *"You met Mary Ann at a supermarket?"*
Instrument	He poached catfish, stunning them with *a shocking device*...
Beneficiary	Whenever Ann makes hotel reservations for *her boss*...
Source	I flew in from *Boston*.
Goal	I drove to *Portland*.

Thematic roles

- John broke the window
- The window was broken by John
- John broke the window with a rock
- The rock broke the window
- The window broke

Agent	The volitional causer of an event
Experiencer	The experiencer of an event
Force	The non-volitional causer of the event
Theme	The participant most directly affected by an event
Result	The end product of an event
Content	The proposition or content of a propositional event
Instrument	An instrument used in an event
Beneficiary	The beneficiary of an event
Source	The origin of the object of a transfer event
Goal	The destination of an object of a transfer event

Thematic roles

The thematic roles for verbs generally are predictable by the syntactic position of the argument (specific to each verb class). Some allow for consistent alternations:

Doris gave the book to Cary

Doris gave Cary the book

Thematic roles

- Thematic roles are very useful but difficult to formally define AGENT, THEME, etc.
- At the same time, they may be too coarse for some applications.

Thematic roles

Intermediary instruments can be subjects

- The cook opened the jar with the new gadget
- The new gadget opened the jar

- Shelly ate the sliced banana with a fork
- *The fork ate the sliced banana

Enabling instruments cannot

Levin and Rappaport Hovav 2005; SLP3

Propbank

- Sentences from the Penn Treebank annotated with abstract forms of arguments (proto-roles), along with lexical entries for each sense of a verb.

https://propbank.github.io

Propbank

(22.11) **agree.01**

Arg0: Agreer
Arg1: Proposition
Arg2: Other entity agreeing

Ex1: [$_{\text{Arg0}}$ The group] *agreed* [$_{\text{Arg1}}$ it wouldn't make an offer].
Ex2: [$_{\text{ArgM-TMP}}$ Usually] [$_{\text{Arg0}}$ John] *agrees* [$_{\text{Arg2}}$ with Mary] [$_{\text{Arg1}}$ on everything].

(22.12) **fall.01**

Arg1: Logical subject, patient, thing falling
Arg2: Extent, amount fallen
Arg3: start point
Arg4: end point, end state of arg1
Ex1: [$_{\text{Arg1}}$ Sales] *fell* [$_{\text{Arg4}}$ to $25 million] [$_{\text{Arg3}}$ from $27 million].
Ex2: [$_{\text{Arg1}}$ The average junk bond] *fell* [$_{\text{Arg2}}$ by 4.2%].

Propbank

- Verb-specific argument structures lets us map the commonalities among the different surface forms
    - [$_{Arg0}$ Big Fruit Co.] increased [$_{Arg1}$ the price of bananas].
    - [$_{Arg1}$ The price of bananas] was increased again [$_{Arg0}$ by Big Fruit Co.]
    - [$_{Arg1}$ The price of bananas] increased [$_{Arg2}$ 5%].

- [$_{Arg1}$ 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].

FrameNet

- Propbank maps argument structure for individual verb senses
- FrameNet maps argument structure for frames, which are evoked by a lexical unit (typically a verb)

https://framenet.icsi.berkeley.edu/fndrupal/framenet_data

Frames

AI

- Schank and Abelson 1975, 1977
- Minksky 1974

Linguistics

- Fillmore 1975, 1982, Tannen 1979

Cognitive Psychology

- Rumelhart 1975, 1980

Sociology

- Goffman 1975

Media Studies

- Entman 1993

Frames

John went into a restaurant. He ordered a hamburger and coke. He asked the waitress for the check and left.

(Schank & Abelson 75)

Frames

- “A frame is a data-structure for representing a stereotyped situation” (Minsky 1975)

- By the term ‘frame’ I have in mind any system of concepts related in such a way that to understand any one of them you have to understand the whole structure in which it fits; when one of the things in such a structured is introduced … all of the others are automatically made available.” (Fillmore 1982)

Commercial Transaction

commercial_transaction

- John bought the car at the dealership
- The car was bought by John
- John's purchase of the car
- The sale of the car cleared their inventory.

Commercial Transaction

Buyer

- John bought the car at the dealership
- The car was bought by John
- John's purchase of the car
- The sale of the car cleared their inventory.

Commercial Transaction

Thing bought

- John bought the car at the dealership
- The car was bought by John
- John's purchase of the car
- The sale of the car cleared their inventory.

Semantic Frame

APPLY_HEAT

- Lexical units:

bake.v, barbecue.v, blanch.v, boil.v, braise.v, broil.v, brown.v, char.v, coddle.v, cook.v, deep fry.v, fry.v, grill.v, microwave.v, parboil.v, plank.v, poach.v, roast.v, saute.v, scald.v, sear.v, simmer.v, singe.v, steam.v, steep.v, stew.v, toast.v

- Core Frame Elements:

Cook	The Cook applies heat to the Food.
Food	Food is the entity to which heat is applied by the Cook.
Heating instrument	The entity that directly supplies heat to the Foo
Container	The Container holds the Food to which heat is applied.
Temperature setting	The Temperature_setting of the Heating_instrument for the Food.

Semantic Frame

DESTROY

- Lexical units:

annihilate.v, annihilation.n, blast.v, blow up.v, demolish.v, demolition.n, destroy.v, destruction.n, destructive.a, devastate.v, devastation.n, dismantle.v, dismantlement.n, lay waste.v, level.v, obliterate.v, obliteration.n, raze.v, ruin.v, take out.v, unmake.v, vaporize.v

- Core Frame Elements:

Cause	The event or entity which is responsible for the destruction of the Patient.
Destroyer	The conscious entity, generally a person, that performs the intentional action that results in the Patient's destruction.
Patient	The entity which is destroyed by the Destroyer.

Semantic representations

Two different perspectives on a commercial transaction

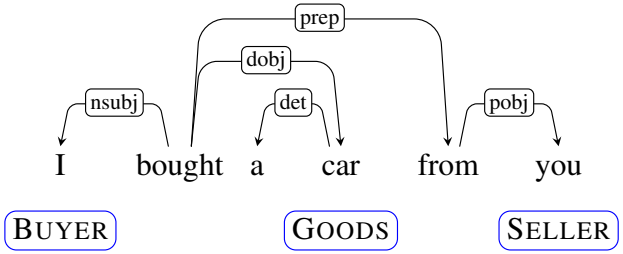

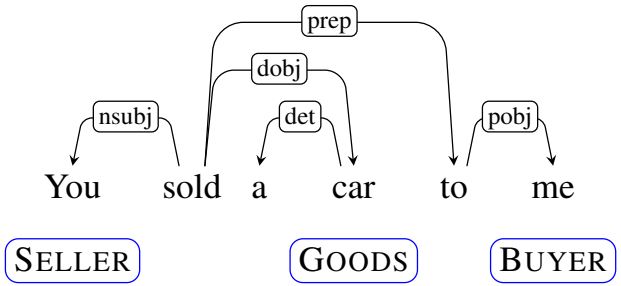

Multilingual frames

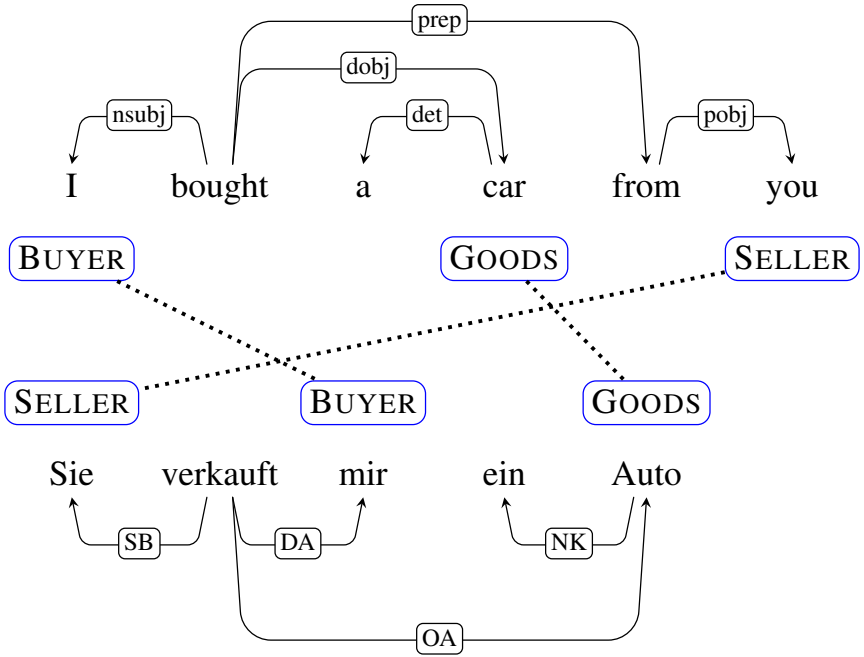

Multilingual frames

- French
- Chinese
- Brazilian Portuguese
- German
- Spanish
- Japanese
- Swedish
- Korean

https://framenet.icsi.berkeley.edu/fndrupal/framenets_in_other_languages

Semantic role labeling

- Input: a sentence
- Output:
    - A list of predicates, each containing:
        - a label (e.g., Framenet frame)
        - a span
        - a set of arguments, each containing:
            - a label (thematic role, FrameNet role)
            - a span

Smith 2017

Semantic role labeling

FrameNet

[You] can't [blame] [the program] [for being unable to identify it]
COGNIZER TARGET EVALUEE REASON

PropBank

[The San Francisco Examiner] issued [a special edition] [yesterday]
ARG0 TARGET ARG1 ARGM-TMP

How would we do this?

Semantic role labeling

```
function SEMANTICROLELABEL(words) returns labeled tree

    parse ← PARSE(words)
    for each predicate in parse do
        for each node in parse do
            featurevector ← EXTRACTFEATURES(node, predicate, parse)
            CLASSIFYNODE(node, featurevector, parse)
```

Gildea and Jurafsky 2002; SLP3

Semantic role labeling

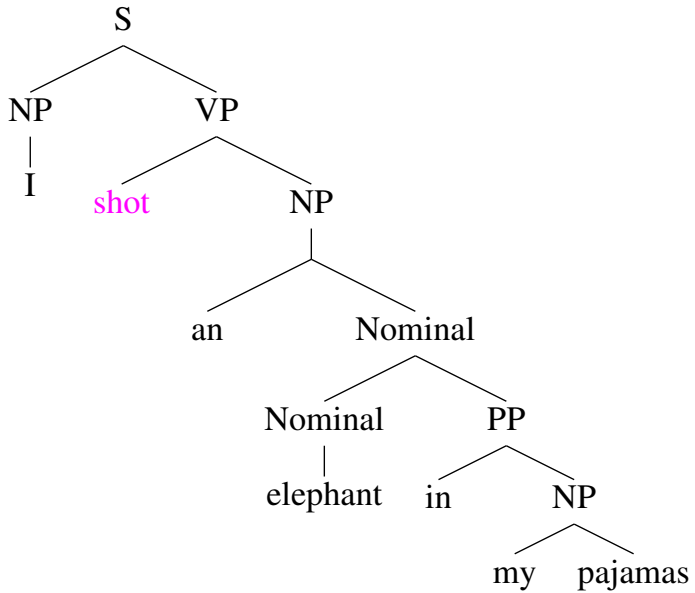

feature
predicate: shot
phrase type = NP
headword of phrase = elephant
path = NP↑S↓VP
voice of verb = active
voice of verb = passive
phrase before verb?
first/last words of phrase

Semantic role labeling

Collobert et al. (2011), Natural Language Processing (Almost) from Scratch

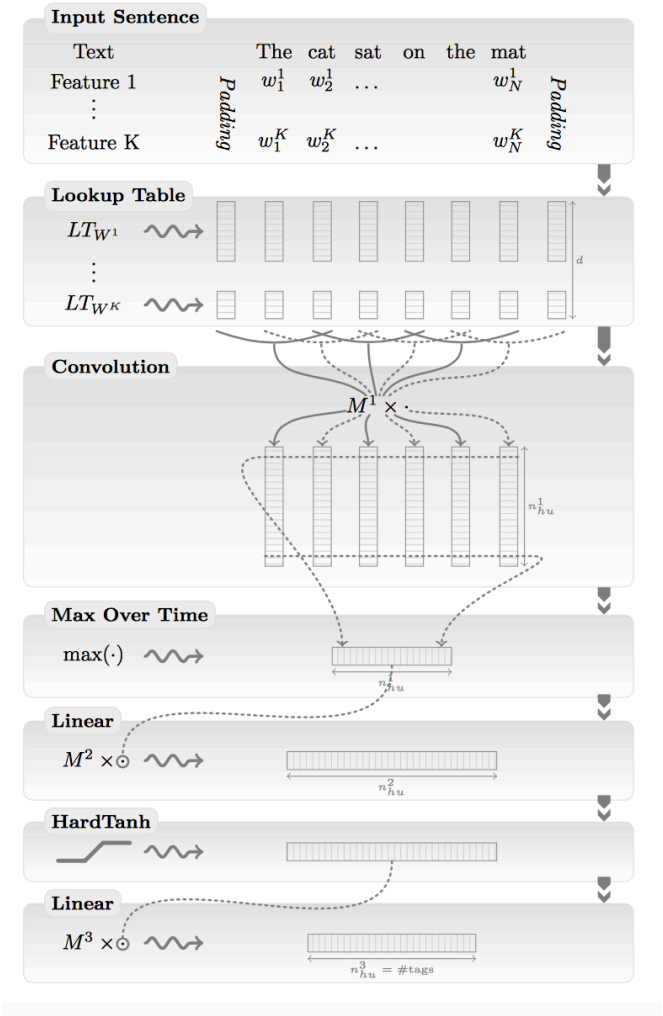

Semantic role labeling

- Sentence-level constraints:
    - Arguments can't overlap
    - For a given predicate, typically only one argument of each type (e.g., ARG0, BUYER)
- Approximate joint decoding (Das et al. 2010)

Smith 2017

Data

- PropBank
  https://propbank.github.io
- FrameNet
  https://framenet.icsi.berkeley.edu/fndrupal/framenet_data