Semantic Role Labeling (SRL)

INF5830
Fall 2015

Semantic Roles

- Origins in the linguistic notion of *case* [Fillmore 1968]
- Classify arguments of predicates into a set of participant *types*
- Describe the semantic relation between the arguments of the verb and the situation described by the verb
  - *The boy threw the red ball to the girl*
  - The boy – the participant responsible for the action, the "doer"
  - the red ball –the affected entity, "undergoer"
  - the girl – endpoint in a change of location
- A variety of semantic roles have been proposed:
  - AGENT
  - PATIENT
  - INSTRUMENT
  - BENEFICIARY
  - SOURCE
  - etc.

Semantic Roles and syntax

- Semantic roles are often indicated by syntactic position
  - AGENT: subject
  - PATIENT: direct object
  - INSTRUMENT: object of *with*
  - BENEFICIARY: object of *for*
  - SOURCE: object of *from*
- Above generalizations are **preferences** at best
  - *The hammer hit the window*
  - *The ball was passed to Mary from John*
  - *John went to the movie with Mary*
  - *John bought the car for $20K*

Problems for semantic roles

- No real consensus about role inventory
  - granularity
  - atomicity
- Difficult to formulate formal definitions of role types
- $\Rightarrow$ more fine-grained roles, relative to "frames"
  [Fillmore 1968, Fillmore 1977]
- $\Rightarrow$ generalized semantic roles [Dowty 1991]
  - PROTO-AGENT, PROTO-PATIENT

Semantic roles in NLP

- How might semantic role information benefit NLP applications?

Semantic roles in NLP

- How might semantic role information benefit NLP applications?
- Question Answering
  [Narayanan and Harabagiu 2004, Shen and Lapata 2007]
  - **Q:** What year did the U.S. buy Alaska?
  - **A:** . . . before Russia sold Alaska to the United States in 1867.

Semantic roles in NLP

- How might semantic role information benefit NLP applications?
- Question Answering
  [Narayanan and Harabagiu 2004, Shen and Lapata 2007]
  - **Q:** What year did the U.S. buy Alaska?
  - **A:** . . . before Russia sold Alaska to the United States in 1867.
- Information Extraction [Surdeanu et al. 2003]: generalization for template-systems, e.g., Acquisitions-and-Mergers:
  - *Apple bought Cisco*
  - *Apple acquired Cisco*
  - *Cisco was taken over by Apple*

Semantic Role Labeling (SRL)

- Task: determine the semantic relations between a predicate and its associated participants
- pre-specified list of *semantic roles*

1. identify role-bearing constituents
2. assign correct semantic role

- [*The girl on the swing*]$_{AGENT}$[*whispered*]$_{PRED}$ *to* [*the boy beside her*]$_{REC}$

Overview of today's lecture

- Resources
  - FrameNet
  - PropBank
- SRL approaches
  - Pioneering: [Gildea and Jurafsky 2002]
  - Overview: [Màrquez et al. 2008]
  - Dependency analysis: [Johansson and Nugues 2008]
- Project, part B

FrameNet

- Based on Fillmore's frame semantics
- Roles are specific to frames, which are invoked by multiple words
- Database of specific frames developed manually
- Sentences that employ these frames selected from the British National Corpus (BNC) and annotated by linguists for semantic roles
- Initial version: 67 frames, 1462 target words, 49013 sentences, 99232 role fillers

Frame Examples

- **apply heat**: situation involving a cook, food and a heating instrument
  evoked by *bake, blanch, boil, broil, brown, simmer*, etc.
- **change position on a scale**: situation involving the change of an items's position on a scale (the attribute) from a starting point (initial value) to an end point (final value)
  evoked by *decline, decrease, gain, rise*, etc.
- **damaging**: situation involving an agent that affects a patient in such a way that the patient (or some sub-region of the patient) ends up in a non-canonical state
  evoked by *damage, sabotage, scratch, tear, vandalise*, etc.

Frame Annotation Examples

- Verbs:
  - [$_{Cook}$Matilde] **fried** [$_{Food}$the catfish] [$_{HeatingInstrument}$in an iron skillet]
  - [$_{Item}$Colgate's stocks] **rose** [$_{Difference}$\$3.64] to [$_{FinalValue}$\$49.94]
- Nouns:
  - . . . the **reduction** of [$_{Item}$debt levels] to [$_{Value_2}$\$25] from [$_{Value_1}$\$2066]
- Adjectives:
  - [$_{Sleeper}$They] were **asleep** [$_{Duration}$for hours]

PropBank

- Adds a layer of semantic roles to the syntactic trees of the Penn Treebank
- Semantic roles are specific to each individual verb to avoid agreeing on a universal set
- Consistent across uses of a single verb (sense)
- But the same tags are used (Arg0, Arg1, Arg2, . . .)
  - inspired by [Dowty 1991]
  - Arg0 $\approx$ proto-Agent
  - Arg1 $\approx$ proto-Patient
  - . . .
  - variety of ArgM's (Arg#>5): TMP, LOC, DIR, MNR, etc.

PropBank

- Annotation process:
  1. rule-based argument tagger on corpus (83% acc on pilot data)
  2. tagger output manually corrected, verb-by-verb basis
  3. differences between annotators resolved
- Annotated over 1M words of Wall Street Journal text with existing gold standard parse trees
- Statistics:
  - 43594 sentences
  - 3324 unique verbs
  - 99265 propositions (verbs+roles)
  - 262281 role assignments

PropBank Examples

- Predicate **accept$_1$ "take willingly"**
  - Arg0: acceptor
  - Arg1: thing accepted
  - Arg2: accepted-from
  - Arg3: attribute
- [$_{Arg0}$He] [$_{ArgM-mod}$would] [$_{ArgM-neg}$n't] **accept** [$_{Arg1}$anything of value] [$_{Arg2}$from those he was writing about].
- Predicate **kick$_1$ "drive or impel with the foot"**
  - Arg0: kicker
  - Arg1: thing kicked
  - Arg2: instrument (defaults to foot)
- [$_{Arg0}$John] tried [$_{Arg0}$*trace*] to **kick** [$_{Arg1}$the football].

PropBank Polysemy

- Polysemous verbs have more than one role assignment
- Predicate **decline$_1$ "go down incrementally"**
  - Arg1: entity going down
  - Arg2: amount gone down by EXT
  - Arg3: start point
  - Arg4: end point
- . . . [$_{Arg1}$its income] **declining** [$_{Arg2-EXT}$42%] [$_{Arg4}$to $2,420].
- Predicate **decline$_2$ "demure, reject"**
  - Arg0: agent
  - Arg1: rejected thing
- [$_{Arg0}$A spokesman] **declined** [$_{Arg1}$*trace* to elaborate].

NomBank

- Argument structure for **nouns**
- Extension of PropBank
  - same Wall Street Journal data
  - same set of semantic roles
- but for nouns?
  - nominalizations of verbs (*destruction*)
  - nominalizations of adjectives (*ability*)
- based on verb senses for verbal nominalizations
- adjectival nominalizations manually coded

NomBank Examples

- Noun **complaint** (based on *complain.01*)
  - Arg0: agent
  - Arg1: topic
  - Arg2: recipient
- Noun example:
  There have been no [$_{Arg0}$customer] **complaints** [$_{Arg1}$about that issue].
- Verb example:
  [$_{Arg0}$They] **complained** [$_{Arg1}$about that issue].

NomBank Examples

- Hyphenated modifiers
  - captures relations within hyphenated words
  - first segment: *H*0, segment after first hyphen: *H*1, segment after *N*th hyphen: *HN*
- *This is a time of self-criticism*
  REL-*H*1 = *self-criticism*, Arg1-*H*0 = *self-criticism*
- *a second daily Chicago-Paris flight*
  REL = *flight*, Arg4-*H*0 = *Chicago-Paris*, Arg3-*H*1 = *Chicago-Paris*, ArgM-TMP = *daily*

Approaches to SRL – overview

- Supervised methods: training data used to train a classifier
  - majority of systems
  - work on FrameNet and PropBank resources
  - shared tasks
- Unsupervised methods: lexical information (large corpora) used to develop classifier
  - few systems

FrameNet SRL

Daniel Gildea and Daniel Jurafsky (2002): "Automatic labeling of semantic roles". *Computational Linguistics* 28(3):245-288.

- Task: Given an input sentence, a target word and a frame, assign all constituents with their semantic roles.
  - locate relevant constituents
  - assign correct semantic roles
- Based on FrameNet examples (BNC)
- Assumed correct frames, the task was to assign roles
- Automatically produced syntactic analyses using Collins (1997) statistical parser
- Results:
  - 80.4% correct role assignment
  - Increased to 82.1% when frame-specific roles were collapsed to 16 more general thematic categories

SRL and parsing

- Syntactic analysis helps identify semantic roles by exploiting generalizations from syntax-semantics linking
  - agent is usually subject
- Needed to identify the true subject
  - *The girl with the dog ate the cookie*
  - "The girl" is the agent, not "the dog"
- Gildea & Jurafsky use constituent parses

SRL as constituent classification

- Treat task as a classification of parse tree nodes
  - For each predicate (verb), label each node in the parse tree as either not a role or one of the semantic roles
- Any machine learning algorithm may be employed
- The real work is in the feature engineering!
- This was the largest contribution of [Gildea and Jurafsky 2002]

Features for SRL

- Three general types of features in SRL [Màrquez et al. 2008]:
  1. features that characterize the candidate argument and its context
  2. features that characterize the verb predicate and its context
  3. features that capture the relation (syntactic or semantic) between the candidate and the predicate

Features for SRL

- **Phrase type:** The syntactic label of the candidate role filler, e.g., NP
  - Different roles tend to be realized by different syntactic categories
- **Parse tree path:** The path in the parse tree between predicate and candidate role filler
  - captures the syntactic relation of a constituent to the rest of the sentence
  - **V** $\uparrow$ **VP** $\uparrow$ **S** $\downarrow$ **NP**
  - **V** $\uparrow$ **VP** $\uparrow$ **S** $\downarrow$ **NP** $\downarrow$ **PP** $\downarrow$ **NP**

Features for SRL

- **Position:** Records whether the candidate role filler precedes or follows the predicate
  - *The girl ate the cookie*
- **Voice:** Records whether the predicate is in active or passive voice
  - *The cookie was eaten by the girl*
- **Head word:** records the head word of the candidate role filler
  - G&J use head finding rules
  - dependency analysis?
- **Governing category:** applied to NPs only, two possible values: S (subjects) or VP (objects)

Probability estimation

- G&J used simple Bayesian method with smoothing to classify parse nodes
- 80% training set, 10% test set, 10% tuning set
- Probability of a semantic role r given the features h (head), pt (phrase type), gov, $position$, $voice$, t (predicate):

$$P(r|h, pt, gov, position, voice, t) = \frac{\#(r, h, pt, gov, position, voice, t)}{\#(h, pt, gov, position, voice, t)}$$

- Sparse data
  - condition on subsets of the features

Other techniques

- Collapsing roles into 18 abstract thematic roles
- Additional features for subcategorization frame
- Abstraction over lexical heads: clustering, WordNet, bootstrapping from (automatically) annotated corpus data

CoNLL shared tasks

- CoNLL04, CoNLL05
- Task:
  - identifying arguments of verbs in a sentence
  - labeling the arguments with their semantic roles
- Gold standard data set: PropBank
- Data:
  - training data: train systems
  - development data: tune systems
  - test data: calculate precision, recall, f-score
    (correct argument requires correct span and role)
    - Precision: (# roles correctly assigned) / (# roles assigned)
    - Recall: (# roles correctly assigned) / (total # of roles)
    - F-score: harmonic mean of precision and recall

CoNLL shared tasks

- CoNLL05: a wide variety of learning approaches
  - Maximum entropy (8 teams)
  - Support Vector Machines (7 teams)
  - SNoW (1 team) (ensemble of enhanced Perceptrons)
  - Decision trees (1 team)
  - AdaBoost (2 teams) (ensemble of decision trees)
  - Nearest neighbour (2 teams)
  - Combination of approaches (2 teams)
- Best results:

WSJ test			Brown test		
P	R	F	P	R	F
82.28	76.78	79.44	73.38	62.93	67.75

Issues in SRL

- How to integrate syntactic parsing, WSD, and role assignment so they all aid eachother
- How to use SRL in down-stream applications
  - Q&A
  - Machine Translation
  - Text Mining

CoNLL08, CoNLL09 shared tasks

- Addresses the integration of syntactic and semantic information
- Syntactic and semantic parsing of English (2008) and several other languages (2009)
- Dependency representations
  - constituent-to-dependency conversion
  - PropBank and NomBank
  - common representation for syntactic and semantic information

CoNLL08 shared task

- Semantic dependencies: semantic role assigned to syntactic head of constituent
- Heads have already been recognized (syntax)
  - "the head of a semantic argument is assigned to the token inside the argument boundaries whose head is a token outside the argument boundaries"
  - Example: [$_{Pred}$sold] [$_{Arg1}$1214 cars] [$_{ArgM-LOC}$in the U.S.]

CoNLL08 shared task

- Data format (extended CoNLL-format)
  - sentences separated by blank line
  - one token per line
  - at least 11 fields, separated by whitespace

Number	Name	Description
1	ID	token counter
2	FORM	(unsplit) word form
3	LEMMA	lemma of form
4	GPOS	gold PoS-tag
5	PPOS	predicted PoS-tag
6	SFORM	tokens split at hyphens
7	SLEMMA	lemma of split forms
8	PPOSS	predicted PoS of split forms
9	HEAD	syntactic head
10	DEPREL	syntactic dependency relation
11	PRED	semantic predicate
12...	ARG	columns with argument labels

CoNLL08 shared task

- Data format (extended CoNLL-format)
- variable towards the end with columns for argument labels for each semantic predicate following textual order

ID	FORM	...	HEAD	DEPREL	PRED	ARG	ARG
2	sold	...	0	ROOT	sold.01	_	
3	1214	...	4	NMOD	_	_	_
4	cars	...	2	OBJ	_	A1	_
5	in	...	2	ADV	_	AM-LOC	_
6	the	...	7	DET	_	_	_
7	U.S.	...	5	PMOD	_	_	_
8	and	...	2	CONJ	_	_	_
9	they	...	5	PMOD	_	_	A0
10	made	...	5	PMOD	make.01	_	_

CoNLL08 shared task

- Data format (extended CoNLL-format)
- Extra rows for tokens split on hyphens

ID	FORM	...	SLEMMA	...	HEAD	DEPREL	PRED	ARG
3	second	...	second	...	8	NMOD	_	_
4	daily	...	daily	...	8	NMOD	_	AM-TMP
5	Chicago-Paris	...	chicago	...	7	NAME	_	A4
6	_	...	-	...	7	HYPH	_	_
7	_	...	paris	...	8	NMOD	_	A3
8	flight	...	flight	...	2	OBJ	flight.01	_

CoNLL08 shared task: example system

[Johansson and Nugues 2008]:

- syntactic and semantic subcomponents
- Semantic model: pipeline of classifiers
  - predicate identification
  - predicate disambiguation
  - argument identification
  - argument classification
- nouns and verbs treated separately

CoNLL08 shared task: example system

[Johansson and Nugues 2008]:

- Features: dependency formulations of phrase-structure features ++
  1. features that characterize the candidate argument and its context: ArgPos, ArgWord, LeftWord, LeftPos, RightWord, RightPos, Function, etc.
  2. features that characterize the verb predicate and its context: PredLemmaSense, PredPos, PredWord
  3. features that capture the relation (syntactic or semantic) between the candidate and the predicate RelPath, PosPath, e.g., *I want him to sleep*: `IM`↑`OPRD`↑`OBJ`↓

Project B

- CoNLL08 data set
  - **train** – open and **closed**
  - **devel** – open and **closed**
  - **test** – open and **closed**
- Data licensing
- Scikit learn: machine learning in Python
- Focus on the task of **argument classification**, i.e. assume gold standard argument identification
- Main components:
  - feature extraction
  - classification
  - evaluation

Project B

- Data processing:
  - extract semantic arguments
  - extract features for these arguments
  - output correct format
- Baseline system: classifier that uses the following features (taken from the Johansson & Nugues article). You may restrict yourself to verbal predicates:

PredLemmaSense	The lemma and sense number of the predicate, e.g., *give.01*
ArgPos	The (predicted) PoS-tag of the argument
PredPos	The (predicted) PoS-tag of the predicate
Function	The grammatical function of the argument

Project B

- Feature engineering
  - take inspiration from the literature
  - add at least 4 new features
  - evaluate
- Choose between one of the following two

  Machine learning algorithm

  Nominal predicates
- Final testing on held-out data

- David Dowty. 1991.
  Thematic proto-roles and argument selection. *Langauge*, 67(3):547–619.
- Charles Fillmore. 1968.
  The case for case. In E. Bach and R. Harms, editors, *Universals in Linguistic Theory*. Holt, Rinehard and Winston, New York.
- Charles Fillmore. 1977.
  The case for case reopened. In *Syntax and Semantics*, volume 8.
- Daniel Gildea and Daniel Jurafsky. 2002.
  Automatic labeling of semantic roles. *Computational Linguistics*, 28:245–288.
- Richard Johansson and Pierre Nugues. 2008.
  Dependency-based syntactic-semantic analysis with propbank and nombank. In *Proceedings of the Twelfth Conference on Computational Natural Language Learning*, CoNLL '08, pages 183–187.
- Lluís Màrquez, Xavier Carreras, Kenneth C. Litkowski, and Suzanne Stevenson. 2008.
  Semantic role labeling: an introduction to the special issue. *Computational Linguistics*, 34:145–159.
- S Narayanan and S Harabagiu. 2004.
  Question answering based on semantic structures. In *In Proceedings of COLING 2004*.

- Dan Shen and Mirella Lapata. 2007.
  Using semantic role to improve question answering. In *In Proceedings of EMNLP 2007*.
- M Surdeanu, S Harabagiu, J Williams, and P Aarseth. 2003.
  Using predicate-argument structures for information extraction. In *In Proceedings of ACL*.