CS 294-5: Statistical Natural Language Processing

Semantics I

Dan Klein

Feedback

- Your comments:
  - Like lectures, prefer to have slides
  - Assignments educational, but too much Java hacking
  - Sections useful, want more of them
  - Readings not so useful?
- My comments:
  - I'm really impressed with the quality of the work!
  - I've enjoyed this class immensely

Some Honors (HW1)

Course Updates

- One more missed class: Nov 1 (sorry!)
- In exchange, a bunch of sections:
  - Oct 27: Agenda-based parsing
  - Nov 10: The EM algorithm
  - Nov 17: Machine translation
  - TBD: Java tricks? Too late?
  - TBD: CRFs and M3Ns
- Fernando Pereira visit and talk on Oct 27 (next Wednesday!)

Semantics

I'd like to buy a flight from Chicago to Denver for under $200

- Once we've got a syntactic parse, then what?

Information Extraction

- Information extraction is basically role-filling
  - The slots are particular to the application
    - Air reservation: departure_city, arrival_city, departure_time
    - Financial: acquired_company, hired_employee
- Classic information extraction systems (e.g. MUC entries), maximally distilled:
  - Use verbs to identify which frame is present
  - Fill the slots using syntactic and semantic cues
  - Frames can extend across sentences (integration)

I'd like to buy a flight from Chicago to Denver for under $200 ⇒

PURCHASE_REQUEST
DEPT_CITY: Chicago
ARRV_CITY: Denver
DEPT_TIME: ???
DATE: ???
PRICE_LIMIT: $200

Semantic Roles

- Semantic roles:
  - Verbs (and some nouns) express events
  - Arguments fill roles in those events
  - Semantic role theory models how roles pattern, how they relate to the syntax
- Granularity of roles
  - Proto-agent, proto-patient (think subject and object)
  - Fillmore's case theory had 9 (agent, patient, location, experiencer, etc)
  - Can subdivide them forever!
  - Extreme view: each verb has its own set of roles
    - buyer, bought_thing, seller, sold_thing
    - PropBank works like this
  - Middle view: roles are particular to a "semantic frame" like transaction
    - Frames can be evoked by various verbs, but not too many
    - FrameNet (here at Berkeley!) works like this

So where's the model?

- Not much work on frame filling
  - … aside from years of IE systems, of course
  - First broad coverage PropBank / FrameNet system was Gildea and Jurafsky 02
  - How does it work?
    - Go node by node, predicting the roles
    - P(role|verb) is the baseline
    - How to do better? (You tell me!)

Is this Semantics?

- It's certainly a step closer!
  - You could imagine extending such a model to make inferences between sentences
  - Can extract relational data
  - You can do IE with such a system (sort of)
- It's *part* of lexical semantics
- What's missing?
  - Quantifiers, negation, coordination, reference ambiguity, modality, tense and aspect…
  - … most of what you learn about in an intro semantics course!

Modeling Compositional Semantics

- We have no statistical model of compositional semantics
- In applications which extract structured data, the last step is always rule-driven
- For the rest of today and next class, we're going to sketch a logical approach to compositional semantics
  - … at least you'll know what we're trying to replace
  - … this is an extension of the lambda-translation approach from the second class (except this time deeper and more interactive)

Phenomena to Model

- Proper names
- Simple verbs
- Quantifiers
  - Subject quantifiers
  - Object quantifiers
  - Reverse scope
  - Generalized quantifiers
- Adjectives and adverbs
- wh- movement (easy and hard!)
- Conjunction and plurals
- Tenses
- Propositional attitudes