

CS674 Natural Language Processing

- Last class
 - Word sense disambiguation
 - » Supervised inductive learning approaches
 - » Issues for WSD evaluation
- Today
 - Word sense disambiguation
 - » Finish decision lists approach
 - » Weakly supervised
 - » Unsupervised learning
 - » Dictionary-based approaches
 - » SENSEVAL

Decision list example

- Binary decision: fish *bass* vs. musical *bass*

Rule		Sense
<i>fish</i> within window	⇒	bass¹
<i>striped bass</i>	⇒	bass¹
<i>guitar</i> within window	⇒	bass²
<i>bass player</i>	⇒	bass²
<i>piano</i> within window	⇒	bass²
<i>tenor</i> within window	⇒	bass²
<i>sea bass</i>	⇒	bass¹
<i>play/V bass</i>	⇒	bass²
<i>river</i> within window	⇒	bass¹
<i>violin</i> within window	⇒	bass²
<i>salmon</i> within window	⇒	bass¹
<i>on bass</i>	⇒	bass²
<i>bass are</i>	⇒	bass¹

Learning decision lists


- Consists of *generating* and *ordering* individual tests based on the characteristics of the training data
- **Generation**: every feature-value pair constitutes a test
- **Ordering**: based on accuracy on the training set

$$abs \left(\log \frac{P(\text{Sense}_1 | f_i = v_j)}{P(\text{Sense}_2 | f_i = v_j)} \right)$$

- Associate the appropriate sense with each test

Weakly supervised approaches

- Problem: Supervised methods require a large sense-tagged training set
- Bootstrapping approaches: Rely on a small number of labeled **seed** instances



Repeat:

1. train *classifier* on *L*
2. label *U* using *classifier*
3. add *g* of *classifier*'s best *x* to *L*

Generating initial seeds

- Hand label a small set of examples
 - Reasonable certainty that the seeds will be correct
 - Can choose prototypical examples
 - Reasonably easy to do
- **One sense per collocation** constraint (Yarowsky 1995)
 - Search for sentences containing words or phrases that are strongly associated with the target senses
 - » Select *fish* as a reliable indicator of *bass*₁
 - » Select *play* as a reliable indicator of *bass*₂
 - Or derive the collocations automatically from machine readable dictionary entries
 - Or select seeds automatically using collocational statistics (see Ch 6 of J&M)

One sense per collocation

Klucsevsek **plays** Giuliani or Titano piano accordions with the more flexible, more difficult free **bass** rather than the traditional Stradella **bass** with its preset chords designed mainly for accompaniment.

We need more good teachers – right now, there are only a half a dozen who can **play** the free **bass** with ease.

An electric guitar and **bass player** stand off to one side, not really part of the scene, just as a sort of nod to gringo expectations perhaps.

When the New Jersey Jazz Society, in a fund-raiser for the American Jazz Hall of Fame, honors this historic night next Saturday, Harry Goodman, Mr. Goodman's brother and **bass player** at the original concert, will be in the audience with other family members.

The researchers said the worms spend part of their life cycle in such **fish** as Pacific salmon and striped **bass** and Pacific rockfish or snapper.

Associates describe Mr. Whitacre as a quiet, disciplined and assertive manager whose favorite form of escape is **bass fishing**.

And it all started when **fishermen** decided the striped **bass** in Lake Mead were too skinny.

Though still a far cry from the lake's record 52-pound **bass** of a decade ago, "you could fillet these **fish** again, and that made people very, very happy," Mr. Paulson says.

Saturday morning I arise at 8:30 and click on "America's best-known **fisherman**," giving advice on catching **bass** in cold weather from the seat of a bass boat in Louisiana.

Yarowsky's bootstrapping approach

- Relies on a **one sense per discourse** constraint: The sense of a target word is highly consistent within any given document
 - Evaluation on ~37,000 examples

Word	Senses	Accuracy	Applicability
<i>plant</i>	living/factory	99.8%	72.8%
<i>tank</i>	vehicle/container	99.6%	50.5%
<i>poach</i>	steal/boil	100.0%	44.4%
<i>palm</i>	tree/hand	99.8%	38.5%
<i>axes</i>	grid/tools	100.0%	35.5%
<i>sake</i>	benefit/drink	100.0%	33.7%
<i>bass</i>	fish/music	100.0%	58.8%
<i>space</i>	volume/outer	99.2%	67.7%
<i>motion</i>	legal/physical	99.9%	49.8%
<i>crane</i>	bird/machine	100.0%	49.1%
Average		99.8%	50.1%

Yarowsky's bootstrapping approach

To learn disambiguation rules for a polysemous word:

1. Find all instances of the word in the training corpus and save the contexts around each instance.
2. For each word sense, identify a small set of training examples representative of that sense. Now we have a few labeled examples for each sense. The unlabeled examples are called the *residual*.
3. Build a classifier (decision list) by training a supervised learning algorithm with the labeled examples.
4. Apply the classifier to all the examples. Find members of the residual that are classified with probability $p >$ threshold, and add them to the set of labeled examples.
5. *Optional*: Use the one-sense-per-discourse constraint to filter and/or augment the new examples.
6. Go to Step 3. Repeat until the residual set is stable.

Topics for today

- Word sense disambiguation
 - » Weakly supervised
 - » **Unsupervised learning**
 - » **Dictionary-based approaches**
 - » SENSEVAL

Unsupervised WSD

- Rely on **agglomerative clustering** to cluster feature-vector representations (without class/word-sense labels) according to a similarity metric
- Represent each cluster as the average of its constituent feature-vectors
- Label the cluster by hand with known word senses
- Unseen feature-encoded instances are classified by assigning the word sense of the most similar cluster
- Schuetze (1992, 1998) uses a (complex) clustering method for WSD
 - For coarse binary decisions, unsupervised techniques can achieve results approaching those of supervised and bootstrapping methods
 - In most cases approaching the 90% range
 - Tested on a small sample of words

Issues for evaluating clustering

- The **correct senses** of the instances used in the training data **may not be known**.
- The **clusters** are almost certainly **heterogeneous** w.r.t. the sense of the training instances contained within them.
- The **number of clusters** is almost always **different from the number of senses** of the target word being disambiguated.

Dictionary-based approaches

- Rely on machine readable dictionaries
- Initial implementation of this kind of approach is due to Michael Lesk (1986)
 - Given a word W to be disambiguated
 - » Retrieve all of the sense definitions, S , for W from the MRD
 - » Compare each s in S to the dictionary definitions of all the remaining words in the context
 - » Select the sense s with the most overlap with (the definitions of) these context words

Example

- Word: *cone*
- Context: *pine cone*
- Sense definitions
 - pine* 1 kind of evergreen tree with needle-shaped leaves
 - 2 waste away through sorrow or illness
 - cone* 1 solid body which narrows to a point
 - 2 something of this shape whether solid or hollow
 - 3 fruit of certain evergreen trees
- Accuracy of 50-70% on short samples of text from *Pride and Prejudice* and an AP newswire article.

Topics for today

- Word sense disambiguation
 - » Weakly supervised
 - » Unsupervised learning
 - » Dictionary-based approaches
 - » **SENSEVAL**

SENSEVAL-2

- Three tasks
 - Lexical sample
 - All-words
 - Translation
- 12 languages
- Lexicon
 - SENSEVAL-1: from HECTOR corpus
 - SENSEVAL-2: from WordNet 1.7
- 93 systems from 34 teams

Lexical sample task

- Select a sample of words from the lexicon
- Systems must then tag several instances of the sample words in short extracts of text
- SENSEVAL-1: 35 words, 41 tasks
 - 700001 John Dos Passos wrote a poem that talked of `the <tag>bitter</> beat look, the scorn on the lip."
 - 700002 The beans almost double in size during roasting. Black beans are over roasted and will have a <tag>bitter</> flavour and insufficiently roasted beans are pale and give a colourless, tasteless drink.

Lexical sample task: SENSEVAL-1

Nouns		Verbs		Adjectives		Indeterminates	
-n	N	-v	N	-a	N	-p	N
accident	267	amaze	70	brilliant	229	band	302
behaviour	279	bet	177	deaf	122	bitter	373
bet	274	bother	209	floating	47	hurdle	323
disability	160	bury	201	generous	227	sanction	431
excess	186	calculate	217	giant	97	shake	356
float	75	consume	186	modest	270		
giant	118	derive	216	slight	218		
...		
TOTAL	2756	TOTAL	2501	TOTAL	1406	TOTAL	1785

All-words task

- Systems must tag almost all of the content words in a sample of running text
 - sense-tag all predicates, nouns that are heads of noun-phrase arguments to those predicates, and adjectives modifying those nouns
 - ~5,000 running words of text
 - ~2,000 sense-tagged words

Translation task

- SENSEVAL-2 task
- Only for Japanese
- word sense is defined according to translation distinction
 - if the head word is translated differently in the given expressional context, then it is treated as constituting a different sense
- word sense disambiguation involves selecting the appropriate English word/phrase/sentence equivalent for a Japanese word

SENSEVAL-2 results

Language	Task	No. of submissions	No. of teams	IAA	Baseline	Best system
Czech	AW	1	1	-	-	.94
Basque	LS	3	2	.75	.65	.76
Estonian	AW	2	2	.72	.85	.67
Italian	LS	2	2	-	-	.39
Korean	LS	2	2	-	.71	.74
Spanish	LS	12	5	.64	.48	.65
Swedish	LS	8	5	.95	-	.70
Japanese	LS	7	3	.86	.72	.78
Japanese	TL	9	8	.81	.37	.79
English	AW	21	12	.75	.57	.69
English	LS	26	15	.86	.51/.16	.64/.40

SENSEVAL plans

- Where next?
 - Supervised ML approaches worked best
 - » Looking the role of feature selection algorithms
 - Need a well-motivated sense inventory
 - » Inter-annotator agreement went down when moving to WordNet senses
 - Need to tie WSD to real applications
 - » The translation task was a good initial attempt