CS474 Natural Language Processing

- Last class
  - Word sense disambiguation
    - Supervised machine learning methods (finish)
    - Weakly supervised (bootstrapping) methods
    - Issues for WSD evaluation
- Today
  - Word sense disambiguation
    - SENSEVAL
    - Unsupervised methods
  - Intro to the noisy channel model

WSD Evaluation

- Precision
  - # of correct senses predicted / # of words in the test set for which the algorithm made a prediction
- Recall
  - # of correct senses predicted / # of words in the test set

WSD Evaluation

- Metrics
  - Precision
    - Nature of the senses used has a huge effect on the results
    - E.g. results using coarse distinctions cannot easily be compared to results based on finer-grained word senses
  - Partial credit
    - Worse to confuse musical sense of *bass* with a fish sense than with another musical sense
    - Exact-sense match → full credit
    - Select the correct broad sense → partial credit
    - Scheme depends on the organization of senses being used

CS474 Natural Language Processing

- Last class
  - Word sense disambiguation
    - Supervised machine learning methods (finish)
    - Weakly supervised (bootstrapping) methods
    - Issues for WSD evaluation
- Today
  - Word sense disambiguation
    - SENSEVAL
    - Unsupervised methods
  - Intro to the noisy channel model

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 at 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

CS474 Natural Language Processing

- Last class
  - Word sense disambiguation
    - » Supervised machine learning methods (finish)
    - » Weakly supervised (bootstrapping) methods
    - » Issues for WSD evaluation
- Today
  - Word sense disambiguation
    - » SENSEVAL
    - » Unsupervised methods
  - Intro to the noisy channel model

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 an 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.

CS474 Natural Language Processing

- Last class
  - Word sense disambiguation
    - » Supervised machine learning methods (finish)
    - » Weakly supervised (bootstrapping) methods
    - » Issues for WSD evaluation
- Today
  - Word sense disambiguation
    - » SENSEVAL
    - » Unsupervised methods
  - Intro to the noisy channel model

The pronunciation subproblem

[spooky music][music stops]

Head Knight of Ni: Ni!

Knights of Ni: Ni! Ni! Ni! Ni! Ni!

Arthur: Who are you?

Head Knight: We are the Knights Who Say…'Ni'! …

We are the keepers of the sacred words: 'Ni', 'Peng', and 'Neee-wom'!

The pronunciation subproblem

- Given a series of phones, compute the most probable word that generated them.
- Simplifications
  - Given the correct string of phones
    - » Speech recognizer relies on probabilistic estimators for each phone, so it's never entirely sure about the identification of any particular phone
  - Given word boundaries
- "I [ni]…"
  - [ni] → *neat, the, need, new, knee, to,* and *you*
  - Based on the (transcribed) Switchboard corpus
- Contextually-induced pronunciation variation

Probabilistic transduction

- surface representation → lexical representation
- string of symbols representing the pronunciation of a word in context → string of symbols representing the dictionary pronunciation
  - [er] → *her, were, are, their, your*
  - exacerbated by **pronunciation variation**
    - » *the* pronounced as THEE or THUH
    - » some aspects of this variation are systematic
- sequence of letters in a mis-spelled word → sequence of letters in correctly spelled words
  - *acress* → *actress, cress, acres*

Noisy channel model

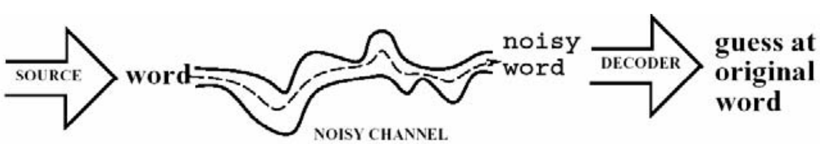

- Channel introduces noise which makes it hard to recognize the true word.
- **Goal:** build a model of the channel so that we can figure out how it modified the true word…so that we can recover it.

Decoding algorithm

- Special case of **Bayesian inference**
  - Bayesian classification
    - » Given observation, determine which of a set of classes it belongs to.
    - » Observation
      - ◆ string of phones
    - » Classify as a
      - ◆ word in the language