Detecting novel metaphor using selectional preference information

- **Hessel Haagsma and Johannes Bjerva**
- **University of Groningen, The Netherlands**

Outline

1. Types of metaphor
2. Selectional preference violation
3. Approach & implementation
4. Evaluation & results
5. Analysis & discussion

A definition of metaphor

A lexical unit is metaphorical if it has a more basic contemporary meaning in other contexts than in the current context

A definition of metaphor

A lexical unit is metaphorical if it has a more basic contemporary meaning in other contexts than in the current context

- Wide range of metaphor:
  1. 'Do the Greeks **have** a word for it?'
  2. 'only little scientific evidence **supports** the link'

Degrees of metaphoricity

1. **None**
   Literal meaning, most basic, in lexicon
2. **Conventional**
   Metaphorical meaning, non-basic, in lexicon
3. **Novel**
   Metaphorical meaning, non-basic, not in lexicon

university of groningen

Examples

1. **No metaphor**
   ‘The scientists **eat** their sandwiches.’
   eat#1 (take in solid food)

Senses from WordNet 3.1

Examples

1. **No metaphor**
   ‘The scientists **eat** their sandwiches.’
   eat#1 (take in solid food)
2. **Conventional metaphor**
   ‘Firefox is **eating** my memory.’
   eat#5 (use up (resources or materials))

Senses from WordNet 3.1

Examples

1. **No metaphor**
   'The scientists **eat** their sandwiches.'
   eat#1 (take in solid food)
2. **Conventional metaphor**
   'Firefox is **eating** my memory.'
   eat#5 (use up (resources or materials))
3. **Novel metaphor**
   'You wanted to **eat** up my sadness.'
   eat#? (take away/cure/remove)

Senses from WordNet 3.1

Metaphor processing and WSD

- Problem: which is the meaning of this ambiguous word/phrase in this specific context?
- WSD and metaphor processing overlap on conventional metaphors
- Novel metaphor outside of scope WSD
- Improved handling of metaphor can benefit WSD

Outline

1. Types of metaphor
2. **Selectional preference violation**
3. Approach & implementation
4. Evaluation & results
5. Analysis & discussion

Selectional preference violation

- Selectional preferences capture intuitive knowledge about what fits in a certain domain
- Metaphor combines a source and target domain
- Violation of selectional preferences as an indicator of two distinct domains, metaphor

Examples

1. **No metaphor**
   ‘The scientists **eat** their sandwiches.’
   eat#1 (take in solid food)
2. **Conventional metaphor**
   ‘Firefox is **eating** my RAM.’
   eat#5 (use up (resources or materials))
3. **Novel metaphor**
   ‘You wanted to **eat** up my sadness.’
   eat#? (take away/cure/remove)

Senses from WordNet 3.1

Novel metaphor

- Automatically acquired selectional preferences capture frequency, not basicness
- Conventional metaphor sometimes more frequent than literal
  e.g. ‘uncover a treasure’ vs. ‘uncover a secret’
- Assumption: novel metaphors are always infrequent

Outline

1. Types of metaphor
2. Selectional preference violation
3. **Approach & implementation**
4. Evaluation & results
5. Analysis & discussion

Approach

- Gather verb-subject and verb-object pairs from a large, parsed English corpus
- Extract selectional preference metrics
- Generalize over co-occurrence counts
- Use as features in a logistic regression classifier to detect metaphors in the VUAMC

Selectional preference information

- Word-level verb metaphor detection
- Parse Wikipedia dump (1.6B words), extract and count verb-noun pairs
- Calculate conditional probability (CP), log probability (LP), selectional association (SA) and selectional preference strength (SPS)
- CP, LP, SA represent likelihood of verb-noun pair
- SPS represents selectivity of verb

Generalization

- Generalization helps going from word-word pairs to domain-domain pairs
- Three approaches
  1. Pre-trained Brown clusters, from Derczynski et al. (2015), 80-5120 clusters
  2. K-means clustered GloVe embeddings (300D/840B), 400k vocabulary, 80-5120 clusters
  3. Neural net predictor of LP, based on embeddings, single hidden layer, 600 units, ADAM, Dropout

Training data

Verb	Subj.	Obj.	CP-s	LP-s	SPS-s	SA-s	...	Label
maintain	couple	link	0.005	-7.51	0.93	6.20		1
need	we	pilot	0.05	-2.98	0.73	0.17		0

Outline

1. Types of metaphor
2. Selectional preference violation
3. Approach & implementation
4. **Evaluation & results**
5. Analysis & discussion

Evaluation data

- VU Amsterdam Metaphor Corpus (VUAMC), parsed
- Extract all verbs
  - Verb-subject-object: 5,539
  - Verb-subject: 13,466
  - Verb-object: 3,913
- Downside: broad definition of metaphor, highly conventionalized metaphors dominate
- Manual inspection of metaphor type

Classifier

- Logistic regression with L2 regularization
- 10-fold cross-validation
- Separate classifier per dataset
- Back-off to majority class (non-metaphor)

Re-weighting

- Re-weighting of examples to counter class imbalance
  - Subject-verb: 13.0%
  - Verb-object: 34.7%
  - Subject-verb-object: 36.4%
- Assign more weight to minority class examples

university of groningen

Results (1)

Without re-weighting of training data

Data	BL	CP	LP	Pred-LP	SPS	SA	All
Subject	**23,0**	0,0	0,0	0,0	0,0	0,0	1,3
Object	**50,8**	0,0	3,2	1,4	0,0	0,0	2,4
Both	**53,4**	0,0	18,1	0,7	0,0	2,3	32,1

Results (1)

Without re-weighting of training data

Data	BL	CP	LP	Pred-LP	SPS	SA	All
Subject	**23,0**	0,0	0,0	0,0	0,0	0,0	1,3
Object	**50,8**	0,0	3,2	1,4	0,0	0,0	2,4
Both	**53,4**	0,0	18,1	0,7	0,0	2,3	32,1

With re-weighting of training data

Data	BL	CP	LP	Pred-LP	SPS	SA	All
Subject	23,0	24,5	24,5	23,2	20,9	26,4	**33,6**
Object	50,8	**53,4**	45,6	49,2	49,0	51,2	47,6
Both	53,4	54,2	44,3	50,0	50,5	63,8	**57,8**

Results (2)

With Brown clustering

Data	BL	80	160	320	640	1280	2560	5120
Subject	23,0	26,3	**28,8**	27,9	25,9	26,3	26,6	25,3
Object	**50,8**	48,7	47,7	45,3	46,9	44,7	44,6	46,2
Both	53,4	52,7	52,8	53,7	54,3	53,5	54,3	**54,5**

With k-means clustering

Data	BL	80	160	320	640	1280	2560	5120
Subject	23,0	24,2	23,5	**30,7**	28,6	24,4	23,6	22,9
Object	**50,8**	40,4	44,8	45,8	44,2	48,9	48,8	49,8
Both	**53,4**	49,8	48,2	50,4	49,2	47,6	50,4	49,5

Outline

1. Types of metaphor
2. Selectional preference violation
3. Approach & implementation
4. Evaluation & results
5. **Analysis & discussion**

Generalization

- In the current set-up, generalization does not work
  - Brown ≈ k-means ≈ prediction
  - No clear effect of cluster size
- Information loss outweighs generalization gain
- Clusters do not form coherent domains

Error analysis

› Large number of (unresolved) pronouns
› True positives contain many light verbs (*take*, *have*, *make*, *put)*.
› Logistic regression exploits corpus distribution
› One example of novel metaphor:
  • [...] Adam might have **escaped** the file memories for years, [...]

Conclusion

- *Is selectional preference information useful for detecting novel metaphors?*
- Better evaluation data is needed
  - Annotate novel/OOV senses in VUAMC
  - Annotate metaphor on a scale, not binary
  - Use selectional preference violation to discover novel metaphors