

Computing Meaning: What's Semantics Got to Do with It?

Emily M. Bender
University of Washington

Center for Language and Speech Processing
Johns Hopkins University
September 10, 2013



Computing Meaning: What's Semantics Got to Do with It?

& Syntax

Emily M. Bender
University of Washington

Center for Language and Speech Processing
Johns Hopkins University
September 10, 2013



This talk in a nutshell



- Conventional sentence meaning is only one clue to speaker meaning.
- Understanding the relationship between the two is critical to creating NLU applications.
- Morphosyntax is critical to extracting sentence meaning.
- Creating cross-linguistically portable systems requires understanding cross-linguistic variation in morphosyntax.

Meanings of 'mean'



- A user types in: *pumpkin carving ideas*
 - Search engine (engineer) wants to know what the user **means** by that
 - The user wants the search engine to know what she **means** by that
- Wikipedia says “A jack-o'-lantern is a carved pumpkin, turnip or beet, associated chiefly with the holiday of Samhain and Halloween, and was named after the phenomenon of strange light flickering over peat bogs, called will-o'-the-wisp or jack-o'-lantern.” and an IE system (engineer) wants to know what that **means**.
- ハロウィンを祝う家庭では、カボチャを刻んで怖い顔や滑稽な顔を作る。 **means**
“Households that celebrate Halloween carve a scary or comical face on a pumpkin.”
- “You post too many photos of pumpkins.” “What do you **mean** by that?”
- The semanticist wants to know why *Kim believes that jack-o'-lanterns are scary* and *Kim believes that jack-o'-lanterns are scary and $1+1 = 2$* don't **mean** the same thing.

The conduit metaphor (Reddy 1979)

- “Your thoughts here don’t quite make it across.”
- “It is very difficult to put this concept into words.”
- “Never load a sentence with more thought than it can carry.”
- “The passage conveys a feeling of excitement.”
- “John says that he cannot find your idea anywhere in the passage.”
- “Get your insights down on paper at once.”
- “The concept made its way very quickly into the universities.”

Conduit metaphor, “major framework” (Reddy 1979: 290)

1. language functions like a conduit, transferring thoughts bodily from one person to another
2. in writing and speaking, people insert their thoughts or feelings in the words
3. words accomplish the transfer by containing the thoughts or feelings and conveying them to others
4. in listening or reading, people extract the thoughts and feelings once again from the words




FIG. 76. Trådtelefon.

Conduit metaphor, “minor framework” (Reddy 1979: 291)

1. thoughts and feelings are ejected by simply speaking them into an external "idea space"
2. thoughts and feelings are reified in this external space, so that they exist independently of any need for living human beings to think or feel them
3. these reified thoughts and feelings may, or may not, find their way back into the heads of living humans





FIG. 76. Trådtelefon.


Toolmaker's paradigm/radical subjectivism (Reddy 1979)

1. Each person has their own environment, with their own “indigenous materials”
2. Communication involves using language (and other signals) as clues to construct private representations from indigenous materials
3. Each interlocutor's representations will be scattered (different) unless they expend energy to coordinate them




Questions to consider

- In what ways do NLP/NLU technologies incorporate the conduit metaphor, and how might this affect their success?
- Given that machines have very different tools in their own “indigenous environments” how can they effectively join communication as understood in the toolmakers paradigm?




Questions to consider

- In what ways do NLP/NLU technologies incorporate the conduit metaphor, and how might this affect their success?
- Given that machines have very different tools in their own “indigenous environments” how can they effectively join communication as understood in the toolmakers paradigm?



Questions to consider

- In what ways do NLP/NLU technologies incorporate the conduit metaphor, and how might this affect their success?
- Given that machines have very different tools in their own “indigenous environments” how can they effectively join communication as understood in the toolmakers paradigm?



Clark 1996: *Using Language*

- Language use is a *joint action*: “Alan intends Barbara to recognize that he wants her to say whether or not she saw his dog run by on the sidewalk, and she is to see this in part by recognizing that intention. The remarkable thing about Alan’s intentions is that they involve Barbara’s thoughts about those very intentions.” (p.12)
- “Words and sentences are *types* of signals, linguistic units abstracted away from any occasion on which they might be used, stripped of all relation to particular speakers, listeners, times, and places.
- “To describe them is to describe the conventions for their use within speech communities.
- “But utterances are the actions of producing words, sentences, and other things on particular occasions by particular speakers for particular purposes.” (p.128)


Reconciling Clark and Reddy

- Through experience within our speech communities, we learn (and help create) shared linguistic conventions.
- These conventions support fairly consistent calculation of ***sentence meaning*** by different speakers in the same community.
- The sentence meaning of an utterance (together with its form) serves as a clue which a toolmaker-listener can use to construct his/her representation of the toolmaker-speaker's ***speaker meaning***

(See also Grice 1968)

Aspects of meaning

- Compositional semantics
- Lexical semantics
- Information structure (information packaging)
- Coreference
- Discourse structure
- ...



Compositional semantics

- Semantic dependencies
 - Who did what to whom
 - What kind of who did what kind of thing to what kind of whom, where, when, how and why
- Scopal operators
 - Quantifiers: *Every dog chased some cat*
 - Scopal modifiers: *The dog didn't chase every cat*
 - Clause embedding predicates: *A unicorn seems to be in the garden*
- Tense/aspect, sentential force, honorifics, evidentials, ...

Lexical semantics

- Predicate argument structure
 - *The cat chased the dog* != *The dog chased the cat*
 - But: *Kim resembles Sandy* = *Sandy resembles Kim*
- Qualia structure: Constitutive, Formal, Telic, Agentive (Pustejovsky 1995)
 - *Kim began the novel.*

Information structure (Lambrecht 1996)

- The way the content expressed by a sentence is structured with respect to the shared common ground in a discourse
 - *As for books, Kim only likes mystery novels.*
 - *It's Kim who gave the book to Sandy.*
 - *It's Sandy Kim gave the book to.*
 - *It's the book that Kim gave to Sandy.*

Coreference chains (e.g. Soon et al 2001)

- Sometimes called “anaphora resolution” because the problem is most obvious with pronouns: *They told her all about it.*
- But really a more general issue: NLU depends on knowing which phrases refer to the same entities (individuals, concepts, events)
- In the most general case, coreference chains include unexpressed arguments as well:
 - *The children promised Sandy to clean up after themselves.*
 - *Bake for 50 min, then remove from oven and cool thoroughly.*

Discourse Structure

- Rhetorical relations (e.g., Mann & Thompson 1988; Marcu 1997)
 - Relationships between clauses: Justification, Condition, Elaboration, Concession, etc.
 - Organizes clauses (within and without sentence boundaries) into a discourse tree.
- Dialog acts/Adjacency pairs (e.g., Shriberg et al 2004, Levinson 1983)
 - Relationships between adjacent turns by different speakers in dialog
 - Ex: Open-Ended Question + No Knowledge Answer


A Cute Example



- Victor Mair on Language Log (9/2/2013; <http://languagelog.ldc.upenn.edu/nll/?p=6606>) analyzes this comic from Zits (8/30/2013):
 - There seems to be no concise way to translate this use of *It wasn't* to Chinese.
- Compositional semantics: `neg(past(contextually-provided-pred(contextually-provided-ref)))`
- Reference resolution: `pred="be-song-stuck-in-head"`, `it="Puff"`
- Information structure: Contrast-focus

Aspects of meaning

- Compositional semantics
- Lexical semantics
- Information structure (information packaging)
- Coreference
- Discourse structure
- ...



Annotating meanings

- Work in NLU is supported by corpora annotated for “meaning”
- Different projects have different ideas of what “meaning” is
- Dimensions of variation (not independent):
 - Machine annotation + human selection v. human annotation
 - Grammatically-constrained v. free-hand
 - Spanning v. partial
 - Sentence meaning v. speaker meaning



Meaning annotation example 1: Redwoods (Oepen et al 2004)

- Machine generated/grammatically constrained/spanning/sentence meaning

this sentence is easy to understand. [1 : 2]

TOP	h1						
INDEX	e3						
		<i>_this_q_dem<0:4></i>		<i>_easy_a_for<17:21></i>		<i>_understand_v_by<25:36></i>	
		LBL	<i>h4</i>	LBL	<i>h2</i>	LBL	<i>h11</i>
RELS	{	ARG0	<i>x6</i>	ARG0	<i>e3</i>	ARG0	<i>e12</i>
		RSTR	<i>h7</i>	ARG1	<i>h9</i>	ARG1	<i>i13</i>
		BODY	<i>h5</i>	ARG2	<i>i10</i>	ARG2	<i>x6</i>
HCONS	{	<i>h1=qh2</i>	<i>h7=qh8</i>	<i>h9=qh11</i>			
	}						

Meaning annotation example 2: Groningen Meaning Bank (Basile et al 2012)

- Machine + human generated/grammatically constrained/spanning/sentence meaning

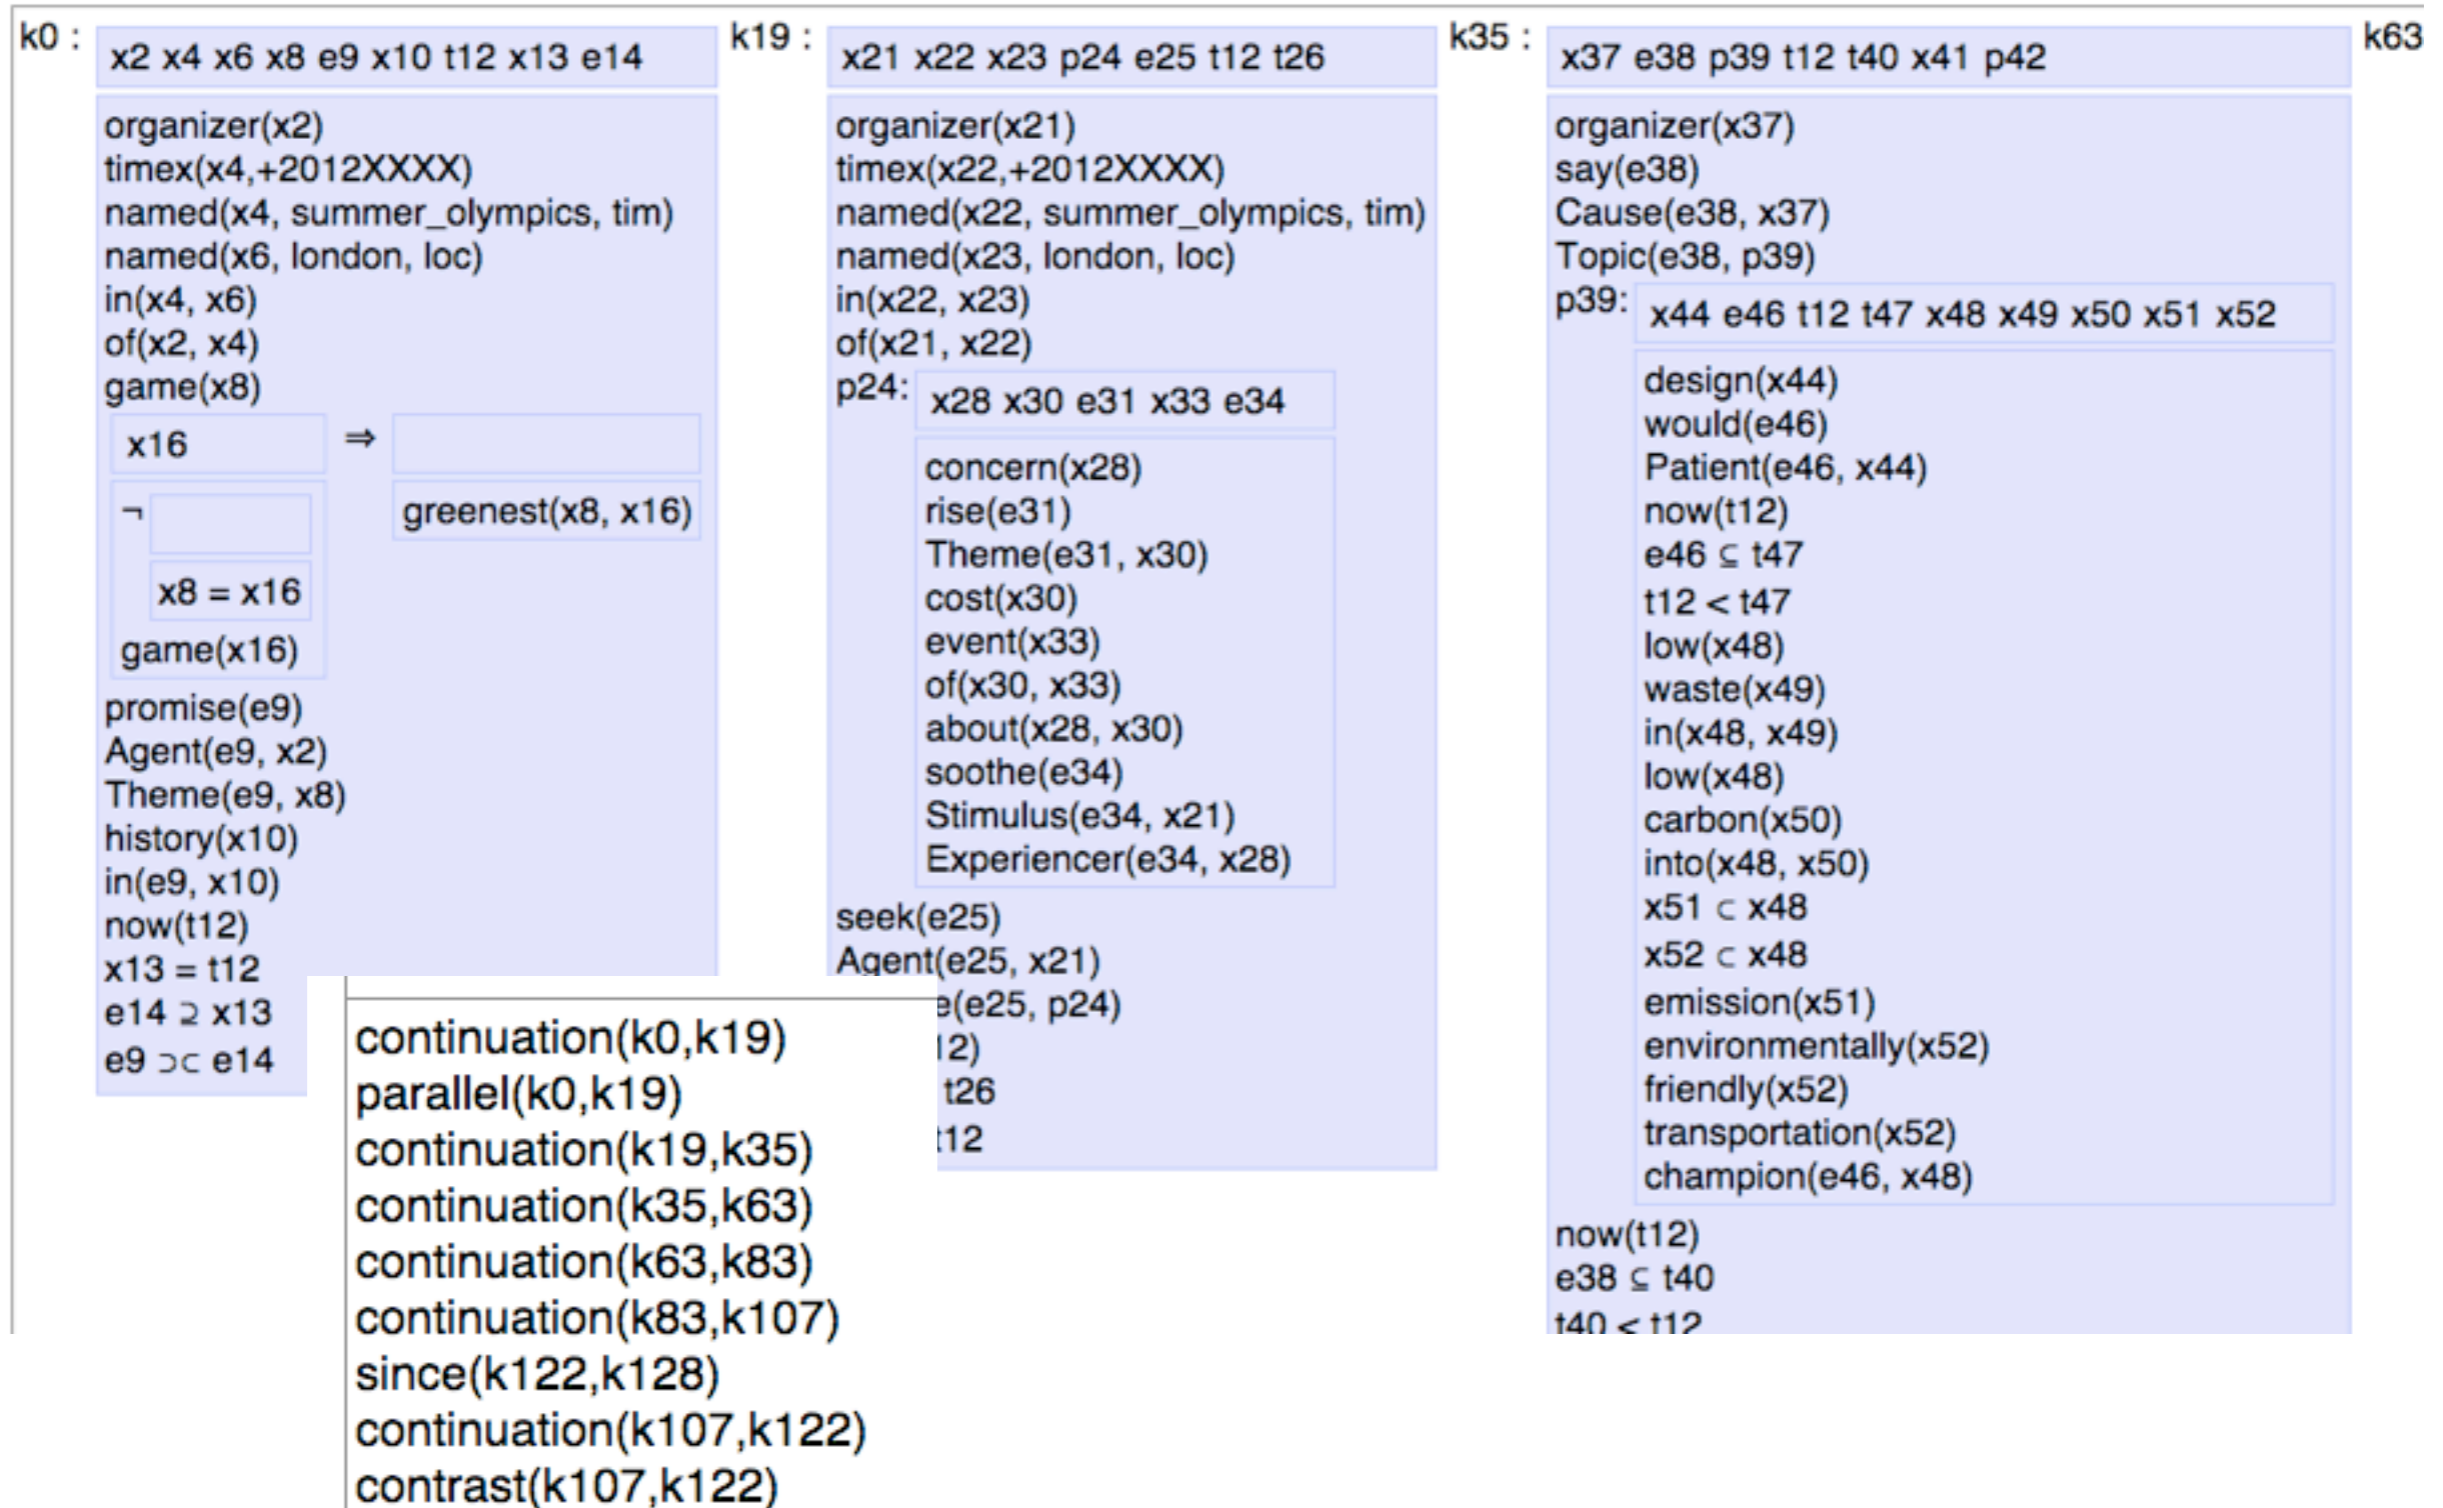
Organizers of the 2012 Summer Olympics in London have promised the "greenest games" in history and sought to soothe concerns about the rising cost of the event.

With 2,012 days to go until the Games get under way, organizers said the design would champion low waste, low carbon emissions and environmentally friendly transportation. The Olympic Delivery Authority has promised to cut emissions 50 percent by generating energy on site and using renewable energy.

Prime Minister Tony Blair said London is farther ahead in preparations at this stage than any other previous Olympic host city.

The British government wants to have the Olympic budget finalized early this year, but costs have already risen substantially since London won the bid in July of 2005. A select committee report due Wednesday is expected to be highly critical of the government's financing of the Games.

Meaning annotation example 2: Groningen Meaning Bank (Basile et al 2012)



Meaning annotation example 5: Abstract Meaning Representation (Banarescu et al 2013)

- Human generated/grammar-independent/spanning/?sentence meaning
- About 14,000 people fled their homes at the weekend after a local tsunami warning was issued, the UN said on its Web site.

```
(s / say-01
  :ARG0 (g / organization
        :name (n / name
              :op1 "UN"))
  :ARG1 (f / flee-01
        :ARG0 (p / person
              :quant (a / about
                    :op1 14000))
        :ARG1 (h / home
              :poss p)
        :time (w / weekend)
        :time (a2 / after
              :op1 (w2 / warn-01
                    :ARG1 (t / tsunami)
                    :location (l / local))))
  :medium (s2 / site
          :poss g
          :mod (w3 / web)))
```

Meaning annotation example 3: GeoQuery (Tang & Mooney 2001)

- Human generated/grammar-independent/?spanning/speaker meaning
- `parse([which,us,city,has,the,highest,population,density,?],
answer(A,largest(B,(city(A),density(A,B))))).`

Meaning annotation example 4: AAWD (Morgan et al 2013)

- Human generated/grammar-independent/partial/speaker meaning

I think <authority-claim:social-expectations> in the minds of most people</authority-claim>, including the government, the word “war” and a formal declaration of war have come apart.

Meaning Representation Example 6: Sentiment Analysis (Wiebe et al 2005)

- Human generated/grammar-independent/partial/speaker meaning

(13) “The report is full of absurdities,” Xirao-Nima said. [“US Human Rights Report Defies Truth,” 2002-02-11, By Xiao Xin, Beijing China Daily, Beijing, China]

Objective speech event:

Text anchor: the entire sentence

Source: <writer>

Implicit: true

Direct subjective:

Text anchor: said

Source: <writer, Xirao-Nima>

Intensity: high

Expression intensity: neutral

Target: report

Attitude type: negative

Expressive subjective element:

Text anchor: full of absurdities

Source: <writer, Xirao-Nima>

Intensity: high

Attitude type: negative

Annotating meanings

- Work in NLU is supported by corpora annotated for “meaning”
- Different projects have different ideas of what “meaning” is
- Dimensions of variation (not independent):
 - Machine annotation + human selection v. human annotation
 - Grammatically-constrained v. free-hand
 - Spanning v. partial
 - Sentence meaning v. speaker meaning



Recovering sentence & speaker meaning: Semantic dependencies

- Semantic dependencies form a central part of sentence meaning
- Semantic dependencies are key to getting the most from sentence meaning when aiming to construct a representation of speaker meaning

Recovering sentence & speaker meaning: Semantic dependencies

- Semantic dependencies form a central part of sentence meaning
- Semantic dependencies are key to getting the most from sentence meaning when aiming to construct a representation of speaker meaning

先生によると男の子よりも女の子がポケモンがすきだ。

Recovering sentence & speaker meaning: Semantic dependencies

- Semantic dependencies form a central part of sentence meaning
- Semantic dependencies are key to getting the most from sentence meaning when aiming to construct a representation of speaker meaning

先生によると男の子よりも女の子がポケモンがすきだ。 [jpn]

Recovering sentence & speaker meaning: Semantic dependencies

- Semantic dependencies form a central part of sentence meaning
- Semantic dependencies are key to getting the most from sentence meaning when aiming to construct a representation of speaker meaning

先生によると男の子よりも女の子がポケモンがすきだ。 [jpn]

先生 によると 男の子 よりも 女の子 が ポケモン が すき だ。

Recovering sentence & speaker meaning: Semantic dependencies

- Semantic dependencies form a central part of sentence meaning
- Semantic dependencies are key to getting the most from sentence meaning when aiming to construct a representation of speaker meaning

先生によると男の子よりも女の子がポケモンがすきだ。 [jpn]

先生 によると 男の子 よりも 女の子 が ポケモン が すき だ。

teacher (.) boy (.) girl (.) Pokemon (.) like (.)

Recovering sentence & speaker meaning: Semantic dependencies

- Semantic dependencies form a central part of sentence meaning
- Semantic dependencies are key to getting the most from sentence meaning when aiming to construct a representation of speaker meaning

先生によると男の子よりも女の子がポケモンがすきだ。 [jpn]

先生 によると 男の子 よりも 女の子 が ポケモン が すき だ。

teacher (.) boy (.) girl (.) Pokemon (.) like (.)

teacher ACCORDING.TO boy THAN girl NOM Pokemon NOM like COP.PRES

Recovering sentence & speaker meaning: Semantic dependencies

- Semantic dependencies form a central part of sentence meaning
- Semantic dependencies are key to getting the most from sentence meaning when aiming to construct a representation of speaker meaning

先生によると男の子よりも女の子がポケモンがすきだ。 [jpn]

先生 によると 男の子 よりも 女の子 が ポケモン が すき だ。

teacher (.) boy (.) girl (.) Pokemon (.) like (.)

teacher ACCORDING.TO boy THAN girl NOM Pokemon NOM like COP.PRES

according-to(e4, e3, x6)

teachers(x6)

like(e3, x14, x17)

girls(x14)

pokemon(x17)

more-than(e22, x14, x23)

boys(x23)

Side-rant: Language-independence in NLP

Side-rant: Language-independence in NLP

- It is common for NLP researchers to assert that their methods “apply” to all (or many) languages, without testing that claim.

Side-rant: Language-independence in NLP

- It is common for NLP researchers to assert that their methods “apply” to all (or many) languages, without testing that claim.
- Often, the claim is supported instead by the assertion that no (or minimal) language-specific resources were used.

Side-rant: Language-independence in NLP

- It is common for NLP researchers to assert that their methods “apply” to all (or many) languages, without testing that claim.
- Often, the claim is supported instead by the assertion that no (or minimal) language-specific resources were used.
 - “Look Ma, no linguistics!”

Side-rant: Language-independence in NLP


- It is common for NLP researchers to assert that their methods “apply” to all (or many) languages, without testing that claim.
- Often, the claim is supported instead by the assertion that no (or minimal) language-specific resources were used.
 - “Look Ma, no linguistics!”
- n -gram models, for example, are considered linguistics-free, but in fact are less useful in morphologically rich languages (Khudanpur, 2006, inter alia)

Side-rant: Language-independence in NLP

- It is common for NLP researchers to assert that their methods “apply” to all (or many) languages, without testing that claim.
- Often, the claim is supported instead by the assertion that no (or minimal) language-specific resources were used.
 - “Look Ma, no linguistics!”
- n -gram models, for example, are considered linguistics-free, but in fact are less useful in morphologically rich languages (Khudanpur, 2006, inter alia)
 - A little linguistics can easily explain why










Evaluating language independence (Bender 2011)

Number of languages/language pairs studied



ACL2008:HLT survey (long papers, monolingual systems)


ACL2008:HLT survey (long papers, monolingual systems)

- | | |
|---|---|
|  Germanic |  Slavic |
|  Romance |  Indic |
|  Semitic |  Chinese |
|  Japanese |  Turkish |
|  West Balkly | |



ACL2008:HLT survey (long papers, monolingual systems)

- Germanic
- Romance
- Semitic
- Japanese
- West Barkly

- Slavic
- Indic
- Chinese
- Turkish



ACL2008:HLT survey (long papers, monolingual systems)




English: 63%

EACL survey (monolingual systems)

EACL survey (monolingual systems)


- | | |
|--|---|
|  Germanic |  Slavic |
|  Romance |  Indic |
|  Semitic |  Chinese |
|  Turkish | |

EACL survey (monolingual systems)



EACL survey (monolingual systems)

- Germanic
- Romance
- Semitic
- Turkish
- Slavic
- Indic
- Chinese



Language pairs studied at ACL2008:HLT and EACL 2009 (Bender 2011)

L1, L2 genus	Studies		L1, L2 family	Studies	
	N	%		N	%
English, Romance	20	24.69	English, Indo-European	40	49.38
English, Germanic	8	9.88			
English, Slavic	6	7.41			
English, Romance + Germanic	3	3.70			
English, Greek	2	2.47			
English, Indic	1	1.23			
English, Chinese	18	22.22	English, Sino-Tibetan	18	22.22
English, Semitic	12	14.18	English, Afro-Asiatic	12	14.81
English, Finnic	4	4.94	English, Uralic	4	4.94
English, Southern Dravidian	2	2.47	English, Dravidian	2	2.47
English, Japanese	1	1.23	English, Japanese	1	1.23
English, Sundic	1	1.23	English, Austronesian	1	1.23
French, Romance	2	2.47	French, Indo-European	3	3.70
French, Germanic	1	1.23			
Total	81	100.00		81	100.00

TABLE 7 Language pairs studied in ACL 2008 and EACL 2009 papers by genus and family, exclusive of Davidov and Rappoport 2009

English is not a prototypical language

- 33 EACL 2009 papers neglected to state directly that English was the language under study
- Tendency more pronounced in papers on extracting meaning as opposed to those working with linguistic structure directly

Enter Linguistic Typology

- Languages vary within a finite and increasingly known range
- Practical point of language independence:
 - improve scalability of NLP to existing set of human languages
 - not to any possible language in the universe (human or extraterrestrial!)
- Linguistic Typology can tell us about the range
- Typological databases (e.g., WALS, Haspelmath et al 2008, <http://wals.info>) store typological information about many languages

Typological knowledge: An example (Dryer 2008)

- Expression of clausal negation:
 - Negative affix
 - Negative auxiliary verb
 - Negative particle
 - Negative word indeterminate between verb and particle
 - Variation between negative affixes and negative words
 - Two part negation (double marking in every negative clause)

Side-rant: Conclusion

- The best way to create language-independent systems is to include linguistic knowledge.
- Only by doing so can we avoid “overfitting” to our development languages.
- Typological knowledge is relatively inexpensive to incorporate, thanks to typologists and field linguists.
- We should ensure that the languages used in evaluations are representative of the language types and language families we are interested in.

Side-rant: Conclusion

- The best way to create language-independent systems is to include linguistic knowledge.
- Only by doing so can we avoid “overfitting” to our development languages.
- Typological knowledge is relatively inexpensive to incorporate, thanks to typologists and field linguists.
- We should ensure that the languages used in evaluations are representative of the language types and language families we are interested in.

</rant>

Morphosyntax is key to recovering semantic dependencies

- How do speakers know what semantic dependencies are implicit in the strings they encounter?

先生によると男の子よりも女の子がポケモンが好きだ。

teacher (.) boy (.) girl (.) Pokemon (.) like (.)

teacher ACCORDING.TO boy THAN girl NOM Pokemon NOM like COP.PRES [jpn]

according-to(e4, e3, x6)

teachers(x6)

like(e3, x14, x17)

girls(x14)

pokemon(x17)

more-than(e22, x14, x23)

boys(x23)



MORGAN & CLAYPOOL PUBLISHERS

Linguistic Fundamentals for Natural Language Processing

*100 Essentials from
Morphology and Syntax*

Emily M. Bender

***SYNTHESIS LECTURES ON
HUMAN LANGUAGE TECHNOLOGIES***

Graeme Hirst, *Series Editor*

A few sample “things”

- #1 Morphosyntax is the difference between a sentence and a bag of words
- #2 The morphosyntax of a language is the constraints that it places on how words can be combined both in form and in the resulting meaning
- #3 Languages use morphology and syntax to indicate who did what to whom, and make use of a range of strategies to do so
- #16 The notion ‘word’ can be contentious in many languages
- #29 Morphological features associated with verbs and adjectives (and sometimes nouns) can include information about tense, aspect and mood

A few sample “things”

- #45 Syntax provides scaffolding for semantic composition
- #49 There is no one universal set of parts of speech, even among the major categories
- #54 The number of semantic arguments provided for by a head is a fundamental lexical property
- #55 In many (perhaps all) languages, (some) arguments can be left unexpressed

A few sample “things”

- #68 There is no agreed upon universal set of semantic roles, even for one language; nonetheless, arguments can be roughly categorized semantically
- #69 Arguments can also be categorized syntactically, though again there may not be universal syntactic argument types
- #73 Syntactic and semantic arguments aren't the same, though they often stand in regular relations to each other
- #77 Identifying the grammatical function of a constituent can help us understand its semantic role with respect to the head
- #83 There are a variety of syntactic phenomena which obscure the relationship between syntactic and semantic arguments

#94 Long-distance dependencies separate arguments/adjuncts from their associated heads

- (173) a. Kim saw the movie.
b. What did Kim see?
c. What did Sandy claim everyone hoped Lee would believe Kim saw?
- (174) a. Kim read the book in the library.
b. This is the library in which Kim read the book.
c. This is the library in which no one believes anyone could imagine Kim read the book.
- (175) a. I don't think Kim eats eggs. Kim likes to eat BAGELS.
b. I don't think Kim eats eggs. BAGELS, Kim likes to eat.
c. I don't think Kim eats eggs. BAGELS, I seem to recall Sandy saying Pat had mentioned Kim likes to eat.

(Bender 2013:116-117)

- (177) Коѡ кѡде мислиш, че е отишѡл?
Koj kŭde misli-ř, ře e otiřŭl?
who.NOM.M.SG where think-2SG.PRES that be.PRES.3SG go.PST.PTCP.M.SG
'Who do you think went where?' [bul] (Rudin, 1988, 450)

#88 Many (all?) languages have semantically empty words which serve as syntactic glue

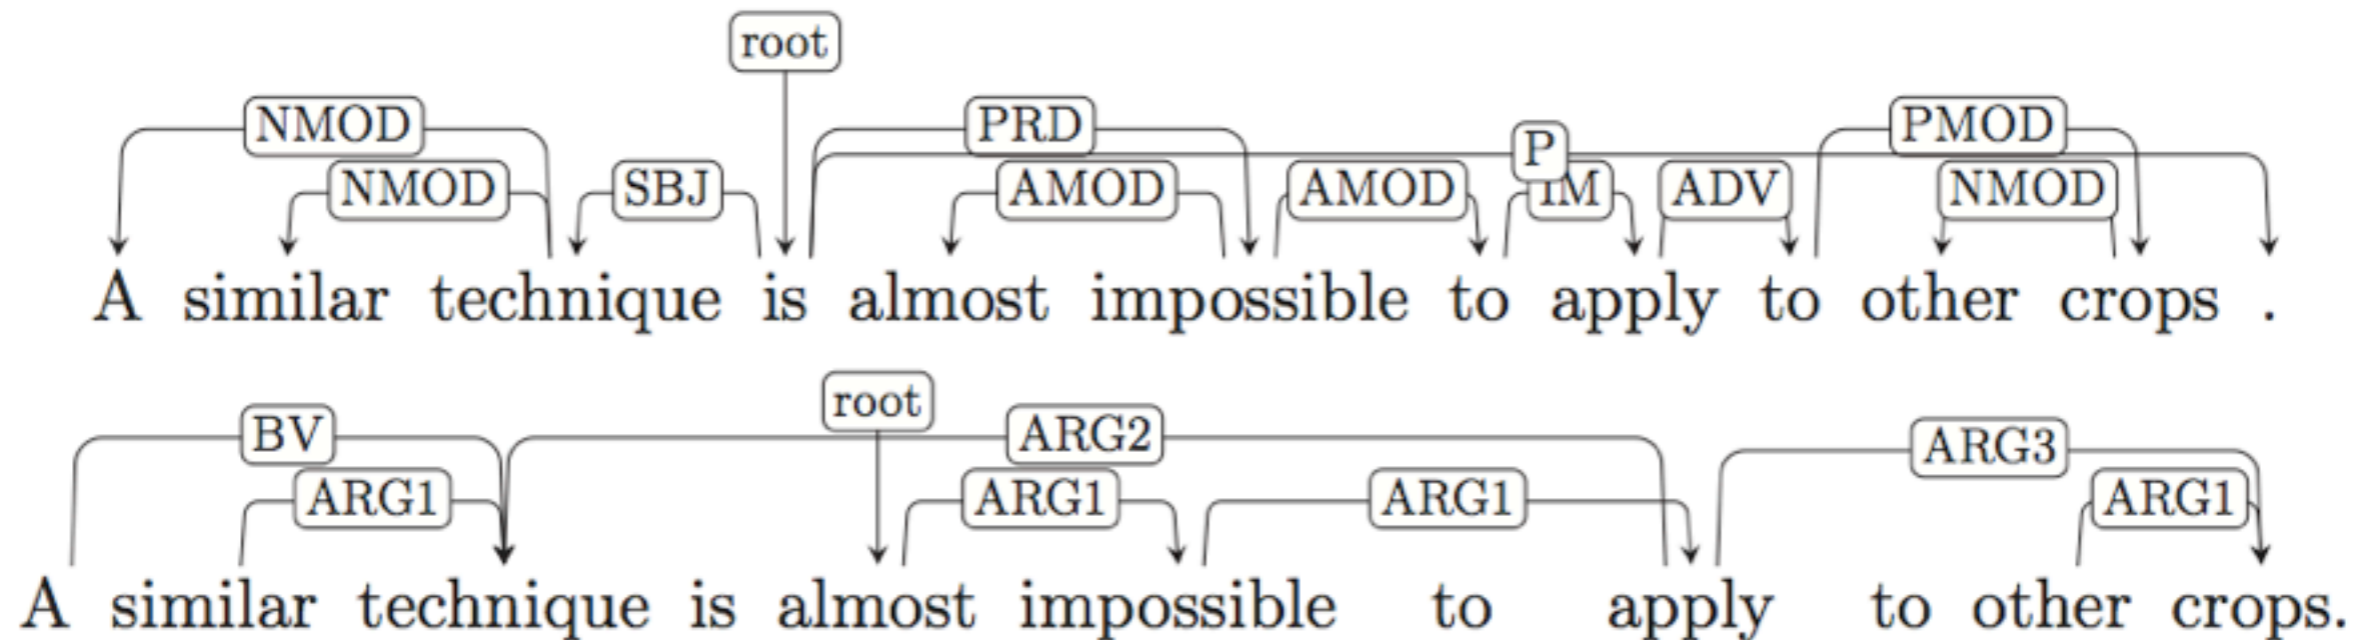


Figure 8.1: Syntactic (CoNLL 2008, top) and semantic (ERG, bottom) dependency structures

(Bender 2013:89, adapted from Ivanova et al 2012:7)


These complications are frequent enough to matter (Rimell et al 2009, see also Bender et al 2011)


Construction	WSJ	Brown	Overall
Obj rel clause	2.3	1.1	1.4
Obj reduced rel	2.7	2.8	2.8
Sbj rel clause	10.1	5.7	7.4
Free rel	2.6	0.9	1.3
RNR	2.2	0.9	1.2
Sbj embedded	2.0	0.3	0.4

Table 2: Frequency of constructions in the PTB (percentage of sentences).

And not captured terribly well by statistical parsers
(Bender et al 2011)

Phenomenon	enju	xle	c&j	c&c	stanford	mst	rasp	average
vger-A	0.85	0.79	0.89	0.93	0.91	0.84	0.06	0.75
vger-B	0.91	0.67	0.88	0.91	0.89	0.78	0.64	0.81
vpart-A	0.88	0.76	0.88	0.81	0.85	0.75	0.84	0.82
vpart-B	0.79	0.62	0.80	0.61	0.74	0.50	0.57	0.66
control-A	0.72	0.75	0.86	0.86	0.84	0.92	0.85	0.83
control-B	0.85	0.75	0.62	0.75	0.59	0.00	0.09	0.52
argadj-A	0.78	0.61	0.67	0.78	0.71	0.81	0.02	0.63
argadj-B	0.55	0.30	0.40	0.45	0.28	0.50	0.20	0.38
barerel-A	0.21	0.68	0.75	0.68	0.34	0.70	0.55	0.56
barerel-B	0.58	0.67	0.55	0.62	0.22	0.10	0.08	0.40
rnr-A	0.78	0.61	0.75	0.70	0.80	0.75	0.48	0.70
rnr-B	0.08	0.31	0.08	0.30	0.06	0.08	0.03	0.14
tough-A	0.80	0.51	0.42	0.20	0.78	0.85	0.70	0.61
tough-B	0.65	0.59	0.01	0.01	0.01	0.01	0.01	0.18
ned-A	0.71	0.65	0.64	0.71	0.62	0.66	0.56	0.65
ned-B	0.05	0.45	0.06	0.05	0.09	0.03	0.02	0.11
itexpl-A	0.50	0.70	0.31	0.26	0.30	0.30	0.15	0.36
absol-A	0.66	0.45	0.86	0.45	0.60	0.70	0.47	0.60
absol-B	0.04	0.07	0.20	0.22	0.21	0.18	0.01	0.13
average	0.60	0.58	0.56	0.54	0.52	0.50	0.33	





Interim summary (nearly there)

- Sentence meaning is important to speaker meaning
- Morphosyntax is important to sentence meaning, and also non-trivial.
- Knowledge of morphosyntax is important
 - For feature design, especially when using parser output
 - For designing parsers
- For a basic understanding of morphosyntax: Bender 2013
- As for parsers: This is a problem that calls for reusable solutions

DELPH-IN: Hand-crafted, linguistic grammars and associated software (www.delph-in.net)



- Flagship grammar: the English Resource Grammar (Flickinger 2000, 2011)
 - 20+ person years of development
 - 80%+ coverage on (well-edited) text from new domains
 - Grammatical framework: HPSG (Pollard & Sag 1994) + Minimal Recursion Semantics (Copestake et al 2005)
 - Bi-directional (analysis and generation)
- Application-ready
 - Fast processing (Callmeier 2002)
 - Stochastic parse selection (Toutanova et al 2002)
- Open-source: moin.delph-in.net/ErgProcessing

ERG sample output (analysis direction)



- Kim is reclusive and seems to be difficult to contact.

```
TOP    h1
INDEX  e3

      |proper_q<0:3>| |named<0:3>| |_reclusive_a_1<7:16>| |_and_c<17:20>| |_seem_v_to<21:26>|
      |LBL      h4| |LBL      h8| |LBL      h9| |LBL      h2| |LBL      h12|
      |ARG0     x6| |ARG0     x6| |ARG0     e10| |ARG0     e3| |ARG0     e11|
      |RSTR     h5| |CARG  Kim| |ARG1     x6| |L-HNDL   h9| |ARG1     i13|
      |BODY     h7| |ARG1     x6| |ARG2     h14| |L-INDEX  e10| |ARG2     h14|
RELS  { | | | | | }
      |_difficult_a_for<33:42>| |_contact_v_1<46:54>|
      |LBL      h15| |LBL      h19|
      |ARG0     e17| |ARG0     e20|
      |ARG1     h16| |ARG1     i21|
      |ARG2     i18| |ARG2     x6|

HCONS { h1=qh2, h5=qh8, h14=qh15, h16=qh19 }
```


ERG sample output (analysis direction)



- Kim is reclusive and seems to be difficult to contact.

```
TOP    h1
INDEX  e3
```

	<i>proper_q</i> <0:3>	<i>named</i> <0:3>	<i>_reclusive_a_1</i> <7:16>	<i>_and_c</i> <17:20>	<i>_seem_v_to</i> <21:26>
	LBL h4	LBL h8	LBL h9	LBL h2	LBL h12
	ARG0 x6	ARG0 x6	ARG0 e10	ARG0 e3	ARG0 e11
	RSTR h5	CARG Kim	ARG1	L-HNDL h9	ARG1 i13
	BODY h7			L-INDEX e10	ARG2 h14
RELS {				R-HNDL h12	
				R-INDEX e11	}
	<i>_difficult_a_for</i> <33:42>	<i>_contact_v_1</i> <46:54>			
	LBL h15	LBL h19			
	ARG0 e17	ARG0 e20			
	ARG1 h16	ARG1 i21			
	ARG2 i18	ARG2 x6			

```
HCONS { h1=qh2, h5=qh8, h14=qh15, h16=qh19 }
```

Alternatively: ERG-derived treebanks (with semantic dependencies) (Oepen et al 2004, Ivanova et al 2012)



-
- moin.delph-in.net/ErgTreebanks
 - Compute minimal discriminants (Carter 1997) over the parse forest provided by the ERG
 - Annotators accept or reject discriminants, until one analysis is left, and then accept or reject the analysis => supports consistent deployment of complex annotations
 - Treebank stores selected analyses and discriminants; treebank can be kept consistent with current grammar version with minimal manual effort
 - Genres include travel dialogues (Oepen et al 2004), wikipedia articles on computational linguistics (Ytrestøl et al 2009), and WSJ (Flickinger et al 2012)
 - 1,150,000 tokens of annotated text; available in CoNLL-08 bi-lexical semantic dependency format (Ivanova et al 2012)

This talk in a nutshell



- Conventional sentence meaning is only one clue to speaker meaning.
- Understanding the relationship between the two is critical to creating NLU applications.
- Morphosyntax is critical to extracting sentence meaning.
- Creating cross-linguistically portable systems requires understanding cross-linguistic variation in morphosyntax.

References

- Banarescu, Laura, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, 178–186, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Basile, Valerio, Johan Bos, Kilian Evang, and Noortje Venhuizen. 2012. Developing a large semantically annotated corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)*, 3196–3200, Istanbul, Turkey.
- Bender, Emily M. 2011. On achieving and evaluating language independence in NLP. *Linguistic Issues in Language Technology* 6:1–26.
- Bender, Emily M. 2013. *Linguistic Fundamentals for Natural Language Processing: 100 Essentials from Morphology and Syntax*. Morgan & Claypool.
- Bender, Emily M., Dan Flickinger, Stephan Oepen, and Yi Zhang. 2011. Parser evaluation over local and non-local deep dependencies in a large corpus. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, 397–408, Edinburgh, Scotland, UK., July. Association for Computational Linguistics.
- Callmeier, Ulrich. 2002. Preprocessing and encoding techniques in pet. In S. Oepen, D. Flickinger, J. Tsujii, and H. Uszkoreit (Eds.), *Collaborative Language Engineering. A Case Study in Efficient Grammar-based Processing*. Stanford, CA: CSLI Publications.

- Carter, David. 1997. The TreeBanker. A tool for supervised training of parsed corpora. In *Proceedings of the Workshop on Computational Environments for Grammar Development and Linguistic Engineering*, 9–15, Madrid, Spain.
- Clark, Herbert H. 1996. *Using Language*. Cambridge: Cambridge University Press.
- Copestake, Ann, Dan Flickinger, Carl Pollard, and Ivan A. Sag. 2005. Minimal recursion semantics: An introduction. *Research on Language & Computation* 3(4):281–332.
- Flickinger, Dan. 2000. On building a more efficient grammar by exploiting types. *Natural Language Engineering* 6 (1) (Special Issue on Efficient Processing with HPSG):15–28.
- Flickinger, Dan. 2011. Accuracy v. robustness in grammar engineering. In E. M. Bender and J. E. Arnold (Eds.), *Language from a Cognitive Perspective: Grammar, Usage and Processing*, 31–50. Stanford, CA: CSLI Publications.
- Grice, H. P. 1968. Utterer's meaning, sentence-meaning, and word-meaning. *Foundations of Language* 4(3):225–242.
- Ivanova, Angelina, Stephan Oepen, Lilja Øvrelid, and Dan Flickinger. 2012. Who did what to whom? a contrastive study of syntacto-semantic dependencies. In *Proceedings of the Sixth Linguistic Annotation Workshop*, 2–11, Jeju, Republic of Korea, July. Association for Computational Linguistics.

- Khudanpur, Sanjeev P. 2006. Multilingual language modeling. In T. Schutlz and K. Kirchhoff (Eds.), *Multilingual Speech Processing*, 169–205. Burlington, MA: Elsevier.
- Lambrecht, Knud. 1996. *Information Structure and Sentence Form: Topic, Focus, and the Mental Representations of Discourse Referents*. Cambridge, UK: Cambridge University Press.
- Levinson, Stephen C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- Mann, William C., and Sandra A. Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text* 8(3):243–281.
- Marcu, Daniel. 1997. The rhetorical parsing of unrestricted natural language texts. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, 96–103, Madrid, Spain, July. Association for Computational Linguistics.
- Morgan, Jonathan T., Meghan Oxley, Emily M. Bender, Liyi Zhu, Varya Gracheva, and Mark Zachry. 2013. Are we there yet?: The development of a corpus annotated for social acts in multilingual online discourse. *Dialogue & Discourse* 4:1–33.
- Oepen, Stephan, Daniel Flickinger, Kristina Toutanova, and Christopher D. Manning. 2004. LinGO Redwoods. A rich and dynamic treebank for HPSG. *Journal of Research on Language and Computation* 2(4):575–596.

- Pollard, Carl, and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. Chicago, IL and Stanford, CA: The University of Chicago Press and CSLI Publications.
- Pustejovsky, James. 1995. *The Generative Lexicon*. MIT Press.
- Reddy, Michael J. 1979. The conduit metaphor: A case of frame conflict in our language about language. In *Metaphor and Thought*, 164–201.
- Shriberg, Elizabeth, Raj Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey. 2004. The icsi meeting recorder dialog act (mrda) corpus. In M. Strube and C. Sidner (Eds.), *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue*, 97–100, Cambridge, Massachusetts, USA, April 30 - May 1. Association for Computational Linguistics.
- Soon, Wee Meng, Hwee Tou Ng, and Daniel Chung Yong Lim. 2001. A machine learning approach to coreference resolution of noun phrases. *Computational Linguistics* 27(4):521–554.
- Tang, Lappoon R., and Raymond J. Mooney. 2001. Using multiple clause constructors in inductive logic programming for semantic parsing. In *Proceedings of the 12th European Conference on Machine Learning*, 466–477.
- Toutanova, Kristina, Chris Manning, and Stephan Oepen. 2002. Parse ranking for a rich HPSG grammar. In *Proceedings of The First Workshop on Treebanks and Linguistic Theories (TLT2002)*, Sozopol, Bulgaria.
- Wiebe, Janyce, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation* 39:165–210.