Clueless: Explorations in the unsupervised, knowledge-lean extraction of lexical-semantic information

Lillian Lee
Cornell University
http://www.cs.cornell.edu/home/llee

Joint work with:
Cristian Danescu-Niculescu-Mizil, Rick Ducott, Bo Pang, and Mark Yatskar

CoNLL 2010

Introduction:

Preview of two extraction problems

Some unifying themes

The first problem: a one-slide preview

You know she'll give a talk

You know she'll give a talk today

The first problem: a one-slide preview

You know she'll give a talk

$\not\Downarrow$

You know she'll give a talk today

The first problem: a one-slide preview

You know she'll give a talk

$\not\Downarrow$ $\Uparrow$

You know she'll give a talk today

The first problem: a one-slide preview

$$\text{You know } \underbrace{\text{she'll give a talk}}_{\textit{“set”}}$$

$$\not\Downarrow \Uparrow$$

$$\text{You know } \underbrace{\text{she'll give a talk today}}_{\textit{“subset”}}$$

The first problem: a one-slide preview

You know $\underbrace{\text{she'll give a talk}}_{\text{``\textit{set}''}}$

$\not\Downarrow \Uparrow$

You know $\underbrace{\text{she'll give a talk today}}_{\text{``\textit{subset}''}}$

You doubt she'll give a talk

You doubt she'll give a talk today

The first problem: a one-slide preview

You know $\underbrace{\text{she'll give a talk}}_{\textit{“set”}}$

$\not\Downarrow$ $\Uparrow$

You know $\underbrace{\text{she'll give a talk today}}_{\textit{“subset”}}$

You doubt she'll give a talk

$\Downarrow$

You doubt she'll give a talk today

The first problem: a one-slide preview

You know $\underbrace{\text{she'll give a talk}}_{\textit{“set”}}$

$\not\Downarrow$ $\Uparrow$

You know $\underbrace{\text{she'll give a talk today}}_{\textit{“subset”}}$

You doubt she'll give a talk

$\Downarrow$ $\not\Uparrow$

You doubt she'll give a talk today

The first problem: a one-slide preview

You know $\underbrace{\text{she'll give a talk}}_{\textit{“set”}}$

$\not\Downarrow$ $\Uparrow$

You know $\underbrace{\text{she'll give a talk today}}_{\textit{“subset”}}$

You doubt she'll give a talk

$\Downarrow$ $\not\Uparrow$

You doubt she'll give a talk today

What terms are downward entailing — allow reasoning “from sets to subsets” in their arguments, like doubt does?

An important, challenging, understudied problem ... more details to come!

The second problem: a one-slide preview

How can we learn lexical-level simplifications, like ...

indigenous $\rightarrow$ native
stands for $\rightarrow$ is the same as

...so as to automatically simplify text?

This is *not* sentence compression/summarization [e.g., Knight & Marcu '02, C. Lin '03, Turner & Charniak '05, Clarke & Lapata '06]

It *complements* syntactic simplification (usually a small set of rules like "change passive voice to active") [e.g., Chandrasekar & Srinivas '97, Siddharthan et al. '04, Vickrey & Koller '08]

A unified approach

We face two hard lexical-semantics problems.
  downward entailing operators; lexical simplifications

Initially, we are "clueless": no annotated data, and no* examples.

Get a clue from an interesting source. Problem solved...
  debate about "NPIs" in linguistics; socially-authored media

A unified approach

We face two hard lexical-semantics problems.
downward entailing operators; lexical simplifications

Initially, we are "clueless": no annotated data, and no* examples.

Get a clue from an interesting source. Problem solved...
debate about "NPIs" in linguistics; socially-authored media

... Oh, wait, the clues turn out to be very noisy.

Find a simple, knowledge-lean (resource-light) way to overcome the noise.

A note on presentation style

"Find a **simple** knowledge-lean way to overcome the noise."

We'll focus on simple descriptions of the main ideas.

A note on presentation style

"Find a **simple** knowledge-lean way to overcome the noise."

We'll focus on simple descriptions of the main ideas.

- You might start working on these problems because you're sure you can do better. Please do! (Our data is online.)
- Stuart Shieber says:
  "Convince [people] that *your solution is trivial* ...

A note on presentation style

"Find a **simple** knowledge-lean way to overcome the noise."

We'll focus on simple descriptions of the main ideas.

- You might start working on these problems because you're sure you can do better. Please do! (Our data is online.)
- Stuart Shieber says:
  "Convince [people] that *your solution is trivial* ...
  The advantage of [them] thinking your solution is trivial or obvious is that it necessarily comes along with the notion that *you are correct*."

Discovery of downward-entailing operators

Danescu-Niculescu-Mizil, Lee, and Ducott, '09

Danescu-Niculescu-Mizil and Lee, '10

More on downward-entailing operators (DEOs)

Recall: 'know' vs. 'doubt': You $\boxed{\text{doubt}}$ $\underbrace{\text{she'll give a talk}}_{\textit{``set''}}$

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ $\underbrace{\text{she'll give a talk today}}_{\textit{``subset''}}$

More on downward-entailing operators (DEOs)

Recall: 'know' vs. 'doubt': You $\boxed{\text{doubt}}$ $\underbrace{\text{she'll give a talk}}_{\textit{“set”}}$

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ $\underbrace{\text{she'll give a talk today}}_{\textit{“subset”}}$

'not' : I do $\boxed{\text{not}}$ <u>want food</u>

I do $\boxed{\text{not}}$ <u>want cheese</u>

More on downward-entailing operators (DEOs)

Recall: 'know' vs. 'doubt': You $\boxed{\text{doubt}}$ $\underbrace{\text{she'll give a talk}}_{\text{“}\textit{set}\text{”}}$

$\Downarrow\not\Uparrow$

You $\boxed{\text{doubt}}$ $\underbrace{\text{she'll give a talk today}}_{\text{“}\textit{subset}\text{”}}$

'not' ✓: I do $\boxed{\text{not}}$ $\underline{\text{want food}}$

$\Downarrow\not\Uparrow$

I do $\boxed{\text{not}}$ $\underline{\text{want cheese}}$

Which are downward-entailing operators (DEOs)?

‘doubt’, negations ✓:

You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk}}$

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk today}}$

Which are downward-entailing operators (DEOs)?

'doubt', negations ✓: You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk}}$

$\Downarrow\not\Uparrow$

You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk today}}$

'see' : Witnesses $\boxed{\text{saw}}$ $\underline{\text{a car}}$

Witnesses $\boxed{\text{saw}}$ $\underline{\text{a red car}}$

Which are downward-entailing operators (DEOs)?

‘doubt’, negations ✓: You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk}}$

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk today}}$

‘see’ ×: Witnesses $\boxed{\text{saw}}$ $\underline{\text{a car}}$

$\not\Downarrow \Uparrow$

Witnesses $\boxed{\text{saw}}$ $\underline{\text{a red car}}$

Which are downward-entailing operators (DEOs)?

‘doubt’, negations ✓: You $\boxed{\text{doubt}}$ she’ll give a talk

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ she’ll give a talk today

‘see’ ×: Witnesses $\boxed{\text{saw}}$ a car

$\not\Downarrow \Uparrow$

Witnesses $\boxed{\text{saw}}$ a red car

‘too weak to’ : She is $\boxed{\text{too weak to}}$ eat or drink

She is $\boxed{\text{too weak to}}$ eat

Which are downward-entailing operators (DEOs)?

'doubt', negations ✓: You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk}}$

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk today}}$

'see' ×: Witnesses $\boxed{\text{saw}}$ $\underline{\text{a car}}$

$\not\Downarrow \Uparrow$

Witnesses $\boxed{\text{saw}}$ $\underline{\text{a red car}}$

'too weak to' ✓: She is $\boxed{\text{too weak to}}$ $\underline{\text{eat or drink}}$

$\Downarrow \not\Uparrow$

She is $\boxed{\text{too weak to}}$ $\underline{\text{eat}}$

Which are downward-entailing operators (DEOs)?

‘doubt’, negations ✓: You $\boxed{\text{doubt}}$ <u>she’ll give a talk</u>

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ <u>she’ll give a talk today</u>

‘see’ ×: Witnesses $\boxed{\text{saw}}$ <u>a car</u>

$\not\Downarrow \Uparrow$

Witnesses $\boxed{\text{saw}}$ <u>a red car</u>

‘too weak to’ ✓: She is $\boxed{\text{too weak to}}$ <u>eat or drink</u>

$\Downarrow \not\Uparrow$

She is $\boxed{\text{too weak to}}$ <u>eat</u>

‘allow’ : One is $\boxed{\text{allowed}}$ to <u>use a credit card</u>

One is $\boxed{\text{allowed}}$ to <u>use Mastercard</u>

Which are downward-entailing operators (DEOs)?

‘doubt’, negations ✓: You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk}}$

$\Downarrow \not\Uparrow$

You $\boxed{\text{doubt}}$ $\underline{\text{she'll give a talk today}}$

‘see’ ×: Witnesses $\boxed{\text{saw}}$ $\underline{\text{a car}}$

$\not\Downarrow \Uparrow$

Witnesses $\boxed{\text{saw}}$ $\underline{\text{a red car}}$

‘too weak to’ ✓: She is $\boxed{\text{too weak to}}$ $\underline{\text{eat or drink}}$

$\Downarrow \not\Uparrow$

She is $\boxed{\text{too weak to}}$ $\underline{\text{eat}}$

‘allow’ ✓: One is $\boxed{\text{allowed}}$ to $\underline{\text{use a credit card}}$

$\Downarrow \not\Uparrow$

One is $\boxed{\text{allowed}}$ to $\underline{\text{use Mastercard}}$

Why discover downward-entailing operators (DEOs)?

'doubt', 'not', 'too weak to', 'allow', and many more
('without', 'reluctant to', 'ban', 'regardless', 'rarely', 'bar from', ...)

DEOs are key to understanding the implications of sentences [van der Wouden '97, van Benthem '86, Hoeksema '86, Dowty '94, Sànchez Valencia '91, MacCartney and Manning '07]

▷ Important for textual inference, QA , summarization, ...

Downward inferences induce greater cognitive load [Geurts et al. '05]

▷ lists of DEOs useful for natural language generation

Current systems only have lists of a small number of manually* collected DEOs (mostly negations) [Nairn et al. '06, MacCartney and Manning '08, Christodoulopoulos '08; Bar-Haim et al. '08]

Challenges to discovering DEOs

So why aren't there large lists of downward-entailing operators?
Because we don't have a clue how to automatically identify them.

DEOs exhibit great diversity (not just verbs, not just "stuff that feels negative", etc.)
Reminder: 'doubt', 'not', 'too weak to', 'allow', and many more ('without', 'reluctant to', 'ban', 'regardless', 'rarely', 'bar from', ...)

The relevant information is "not available in or deducible from any public lexical database" [Nairn, Condoravdi and Karttunen '06]

There are no DEO-annotated corpora to learn from

Linguistics to the rescue

New concept: Negative Polarity Items (NPIs) — words that *tend* to occur only in negative contents.

'any' (= at all):

$\textit{They do } \underbrace{\textit{not}}_{\text{negative}} \textit{ have any drugs}$ vs. $*$*They do have any drugs*

Also yet, ever, have a clue, etc.

But other things license NPIs too, e.g., 'I doubt they have a clue'
[Linebarger '87, von Fintel '99, Giannakidou '02]

Linguistics to the rescue

New concept: Negative Polarity Items (NPIs) — words that *tend* to occur only in negative contents.

'any' (= at all):

They do $\underbrace{\textit{not}}_{\text{negative}}$ *have any drugs* vs. $*$*They do have any drugs*

Also yet, ever, have a clue, etc.

But other things license NPIs too, e.g., 'I doubt they have a clue' [Linebarger '87, von Fintel '99, Giannakidou '02]

Ladusaw ['80]: NPIs appear only in the scope of DEOs!

DEO discovery: precision-at-k on English newswire

Ladusaw ['80]: ⇒ "extract as DEOs words frequently co-occurring with NPIs." (many details suppressed)

Note: can't measure recall — there's no complete list of DEOs. (Which is the point of our work.)

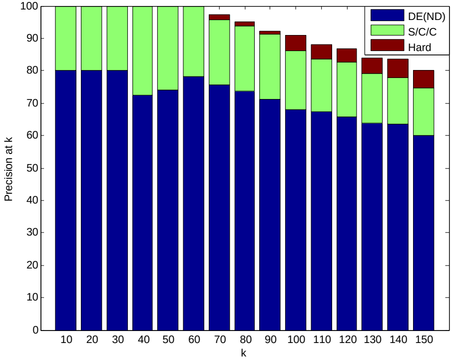

Oh, wait, what about all the non-English languages?

"extract as DEOs words frequently co-occurring with NPIs."

There's an NPI list for English ... but not for any* other language.

Oh, wait, what about all the non-English languages?

"extract as DEOs words frequently co-occurring with NPIs."

There's an NPI list for English ... but not for any* other language.
So... iteratively *co-learn* downward-entailing operators and "NPIs", using **one** seed NPI — the translation of 'any'

Test case: Romanian

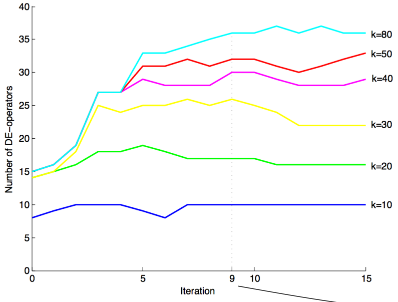

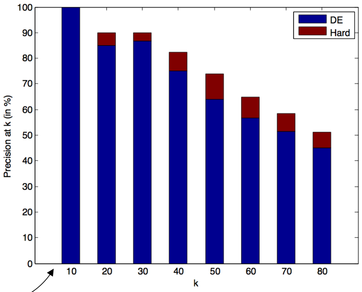

Some surprises

Co-learning DEOs and NPIs isn't *that* straightforward!

- Learning NPIs (even from DEOs) has previously proven hard [Hoekseman '97, Lichte and Söhn '07]
- Hubs and authorities [Kleinberg '98] was not successful

... but we learn "pseudo-NPIs" (English example: 'allegations')

In English, co-learning iterations basically don't alter performance!

- This seems to relate to some linguistics results regarding cross-linguistic variation in the behavior of indefinite pronouns, like 'any' [Haspelmath '01]

DEO discovery: Summary and contributions

The first method for learning downward-entailing operators

▷ complex semantic effect captured from raw text + **one** seed

Inspiration: linguistic insights about NPIs as DEO cues

▷ but, can operate effectively on languages without extensive NPI lists (= everything[1] but English)

Our findings regarding "pseudo-NPIs" and empirical cross-linguistic performance may contribute back to current research in linguistics.

I'm super-excited about this synergy!

[1] Actually, there's a few *noisy* non-English NPI lists, but pseudo-NPIs outperform them!

Now that that's over, you might be thinking ...

"Gee, what a marvelously clear explanation!
If only we could automatically make everything easier to understand!"

"Gee, that was kind of complicated.
If only we could automatically make things easier to understand."

Unsupervised extraction of lexical simplifications from Wikipedia

Yatskar, Pang, Danescu-Niculescu-Mizil, and Lee, '10

Why discover lexical-level simplifications?

Examples: indigenous $\rightarrow$ native, classified as $\rightarrow$ called, stands for $\rightarrow$ is the same as

Make more texts accessible to larger audiences

Eventual goal: a style dial for documents

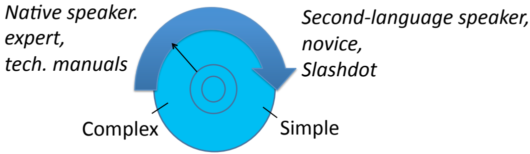

Near-term application: suggest simplifications to readers or authors

Contrasts with prior work on lexical simplification

We want a method for extracting lexical simplifications ...
that is not domain specific,
and doesn't need pre-compiled resources or annotated corpora.

Other work focuses on:
... syntactic simplification [e.g., Chandrasekar & Srinivas '97, Siddharthan et al. '04, Vickrey & Koller '08], identifying simple vs. non-simple documents [Napoles and Dredze '10], or monolingual sentence alignment [Barzilay and Elhadad '03, Nleken and Shieber '06]

... bio-medical text [Elhadad and Sutaria '07, Deléger and Zweigenbaum '09]

... using a thesaurus and word frequency [Devlin and Tait '98, Scarton et al '10]

Getting a clue

The *Simple English Wikipedia*: an independently-maintained "spin-off" of Wikipedia.

Roughly 18,000 editors have produced more than 60,000 articles written in simple English.

Oh, wait, it's not that simple...

Just treat Simple English Wikipedia as a translation (or directional paraphrase) of the "complex" (regular) Wikipedia?

But they aren't parallel: articles are written independently.

Oh, wait, it's not that simple...

Just treat Simple English Wikipedia as a translation (or directional paraphrase) of the "complex" (regular) Wikipedia?

But they aren't parallel: articles are written independently.

Simple English Wikipedia is a living corpus with rich metadata — treat edits as instances of simplifications?

But many edits aren't simplifications.

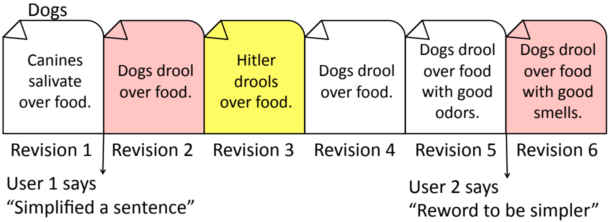

Two filtering approaches

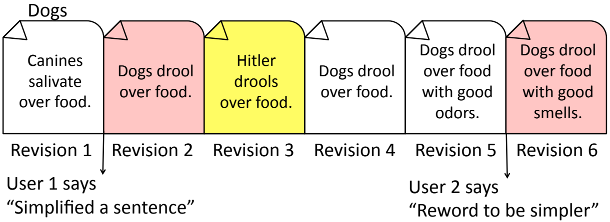

Both methods *first* create candidates by

1. sentence-aligning consecutive revisions using tf-idf [Nelken and Shieber '06]
2. then identifying differing segments

$$\text{“Canines salivate”} \rightarrow \text{“Dogs drool”};$$
$$\text{“Dogs drool”} \rightarrow \text{“Hitler drools”}$$

The methods differ in what they do next...

First filtering approach: Comment-based filtering

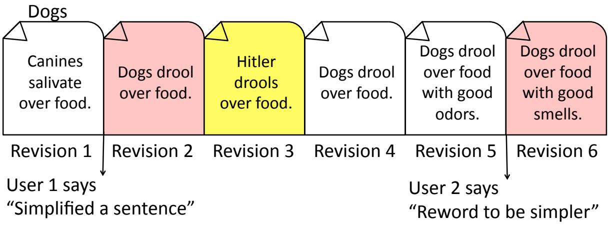

Only consider revisions accompanied by a *user comment* containing the substring "simpl".
(The candidate simplifications are then ranked by pointwise mutual information).

This is noisy annotation: comments correspond to the *whole document*, which can contain multiple revisions

Alternate filtering approach: Edit mixture model

Distinguish different types of operations o_i: simplification, fix (of grammar, content, etc.), spam, no change

The change $A \rightarrow a$ can come about via different operations on A, with different operations having different results:

$$P(a \mid A) = \sum_{o_i} P(o_i \mid A) P(a \mid A, o_i)$$

Alternate filtering approach: Edit mixture model

Distinguish different types of operations o_i: simplification, fix (of grammar, content, etc.), spam, no change

The change $A \rightarrow a$ can come about via different operations on A, with different operations having different results:

$$P(a \mid A) = \sum_{o_i} P(o_i \mid A) P(a \mid A, o_i)$$

We use various simplifying assumptions to estimate these parameters from the *two* Wikipedias.
$\triangleright$ Most important: Regular Wikipedia contains only "fix" edits.

$\Rightarrow$ filter out $A \rightarrow a$ from Simple English Wikipedia if it's frequent in regular Wikipedia

Results

Data: 38,000 Simple and regular Wikipedia articles

Method	Prec@100	# of pairs
Human	86%	2000
Edit Model	77%	1079
Comment Method	66%	2970
Frequent	17%	-
Random	17%	-

Top 100 pairs from each method were manually annotated

Manually assembled dictionary: SpList (by a SimpleWiki author)

Edit and Comment produce correct pairs not found in SpList (71% and 62%)

Results: some examples

Some correct instances:

Comment method	Edit model
voyage $\rightarrow$ trip	indigenous $\rightarrow$ native
legend $\rightarrow$ story	classified as $\rightarrow$ called
disbanded $\rightarrow$ broke up	discussed $\rightarrow$ talked about

Some incorrect instances:

Comment method	Edit model
could $\rightarrow$ can	counting $\rightarrow$ recounting
the $\rightarrow$ a	mistakes $\rightarrow$ members

Lexical simplification: contributions and future directions

We can learn lexical-level simplifications from Simple Wikipedia, **if** we figure out filter non-simplifications out.

Future possibilities: try bootstrapping from metadata, comparing against thesaurus-based approaches, more sophisticated modeling, context-sensitive rewriting, etc.

Socially-authored media allow us to “observe” humans at work, and learn from them!

I am super-excited about the possibilities this offers!

Conclusion: A unified approach

We faced two hard lexical-semantics problems.

downward entailing operators; lexical simplifications

Initially, we were "clueless": no annotated data, and no* examples.

Got a clue from an interesting source. Problem solved...

debate about "NPIs" in linguistics; socially-authored media

... Oh, wait, the clues turned out to be very noisy.

Found a simple, knowledge-lean way to overcome the noise.

Conclusion: A unified approach

We faced two hard lexical-semantics problems.
downward entailing operators; lexical simplifications

Initially, we were "clueless": no annotated data, and no* examples.

Got a clue from an interesting source. Problem solved...
debate about "NPIs" in linguistics; socially-authored media

... Oh, wait, the clues turned out to be very noisy.

Found a simple, knowledge-lean way to overcome the noise.

Data online at http://www.cs.cornell.edu/home/llee/data
I encourage you to look at these problems —
in fact, I bet you can improve on our work —
and to start with, I welcome your questions. THANKS!