Natural Language Processing for Analyzing Disaster Recovery Trends Expressed in Large Text Corpora

Lucy H. Lin, Scott B. Miles, Noah A. Smith
University of Washington

19 October 2018

Introduction

Problem: empirical data describing disaster recovery is scarce

$\Rightarrow$ constrains pre-/post-event recovery planning

Available data source: **text corpora** (e.g., news articles)

Manual analysis of large text corpora is slow...

$\Rightarrow$ + natural language processing

Introduction

Proposition query:
"Dealing with authorities is causing stress and anxiety."

query corpus

Matched sentences:
"Unfamiliar bureaucratic systems are causing the majority of the stress."

"Those in charge of recovery are making moves to appease the growing anger among homeowners."

aggregate

Frequency across time:

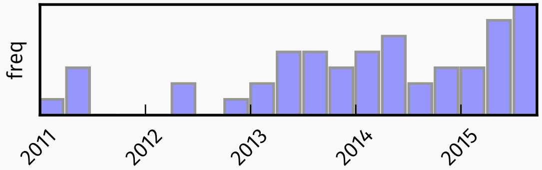

Outline

1. Introduction
2. Case study: 2010-2011 Canterbury earthquake disaster
3. NLP method for semantic matching
4. User evaluation
5. Qualitative/quantitative output
6. Conclusions

Outline

2010–2011 Canterbury earthquake disaster: timeline

September 2010:

- Epicenter: 35km west of Christchurch
- Moderate damage

February 2011:

- Epicenter: 10km southeast of Christchurch
- Extremely high ground acceleration
- 185 deaths, thousands of felt aftershocks

2010–2011 Canterbury earthquake disaster: impacts

Damages:

- Estimated $40 billion
- Housing: 100k houses in need of repairs
- Water, utilities, road infrastructure: extensive damage

Recovery groups:

- Government: CERA, SCIRT (sunset after 5 years)
- Community: Regenerate Christchurch

Recovery still ongoing: public development projects, residential rezoning

2010–2011 Canterbury earthquake disaster: text data

Corpus: 982 NZ news articles (2010–2015) post-earthquakes

- `stuff.co.nz`, `nzherald.co.nz`

Proposition queries: 20 queries, covering

- Community wellbeing
- Infrastructure
- Decision-making

e.g.: "The council should have consulted residents before making decisions."

Outline

1. Introduction
2. Case study: 2010-2011 Canterbury earthquake disaster
3. NLP method for semantic matching
4. User evaluation
5. Qualitative/quantitative output
6. Conclusions

Semantic matching

Goal: find sentences with similar meaning to the query.

- Needs to be more powerful than word/phrase-level matching.
- Related to information retrieval, but want *all matches*.

Semantic matching: method overview

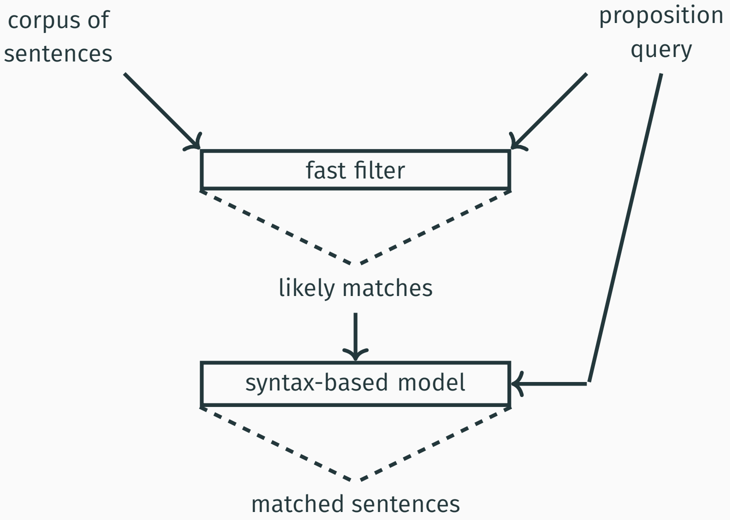

Semantic matching: method overview

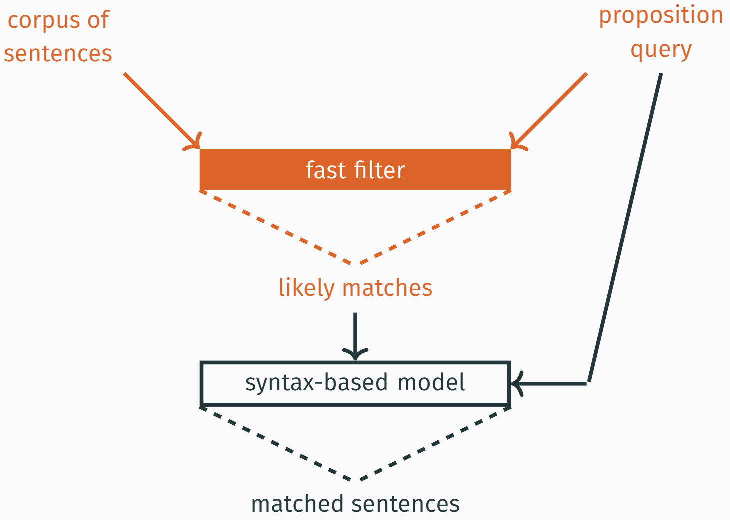

Semantic matching: fast filter

Goal: quickly filter out unlikely matches.

Word vector based comparison between two sentences:

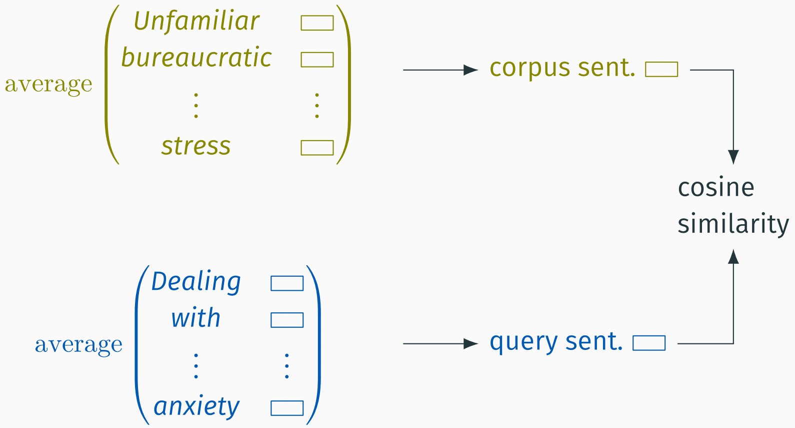

Semantic matching: method overview

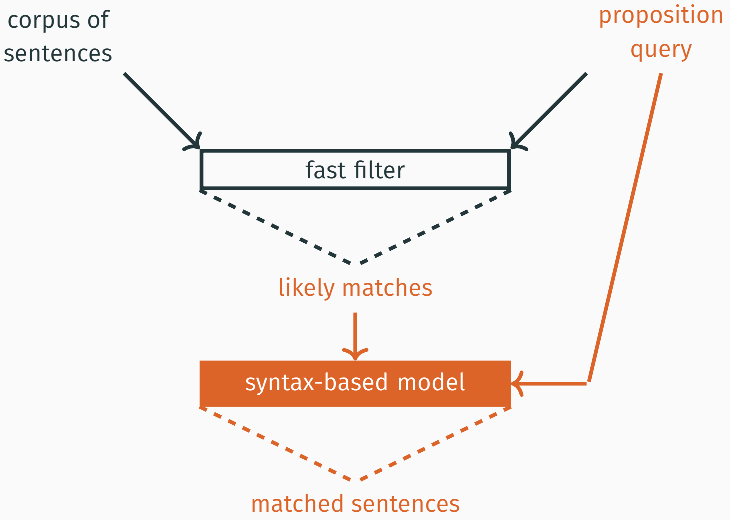

Semantic matching: syntax-based model

Finer-grained matching: take word order/syntax into account.

Intuition: transformation between sentences is indicative of their relationship.

Semantic matching: syntax-based model

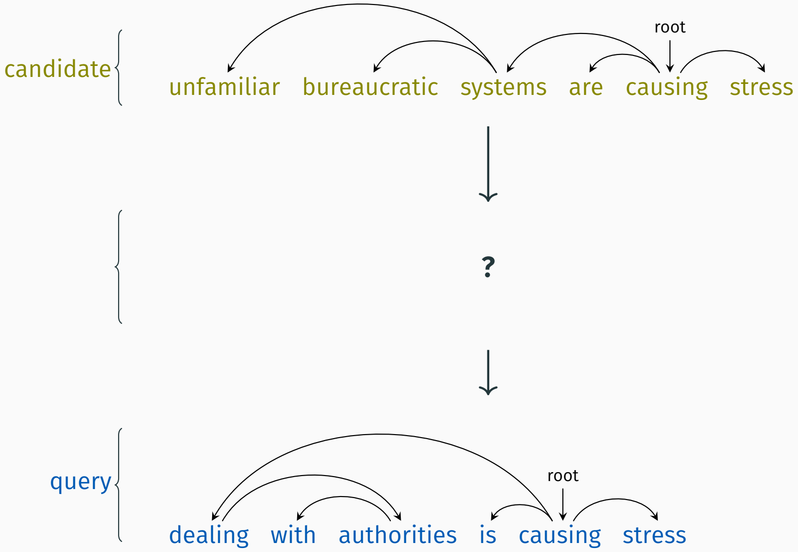

Semantic matching: syntax-based model

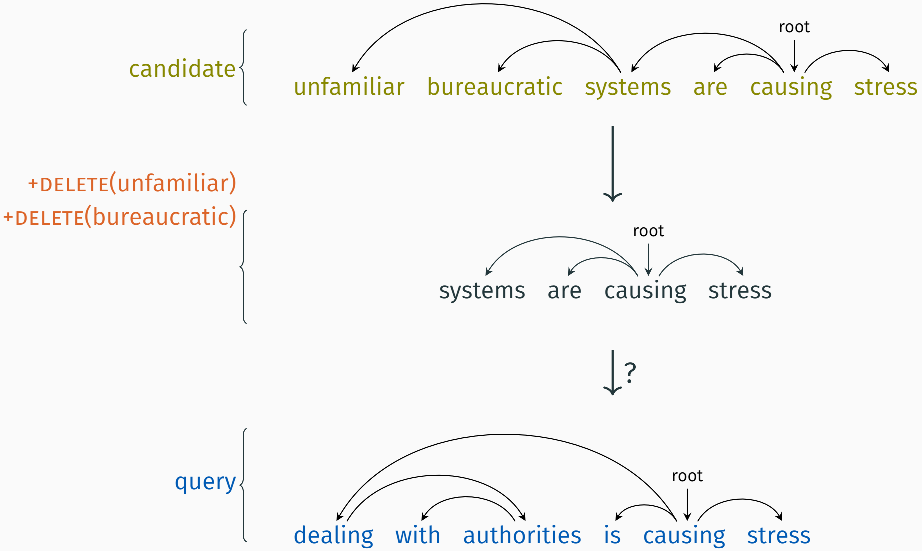

Semantic matching: syntax-based model

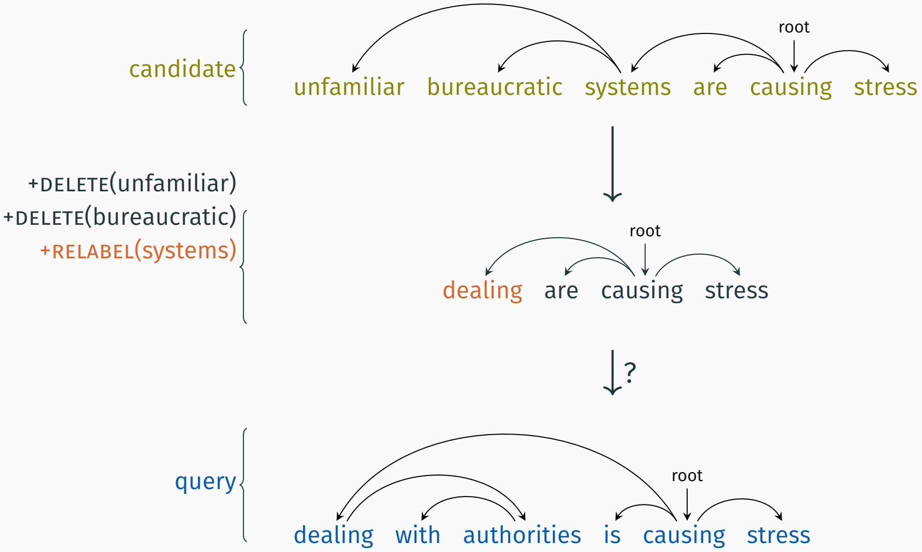

Semantic matching: syntax-based model

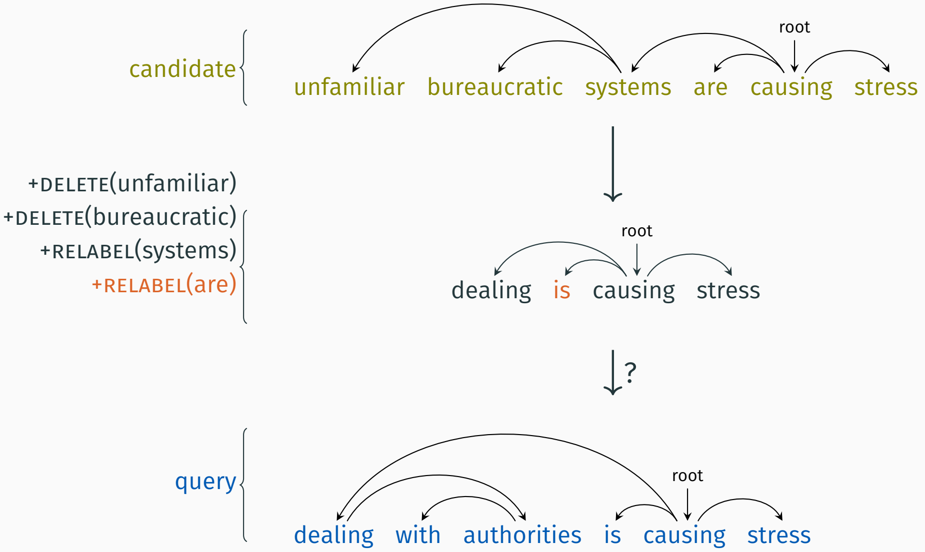

Semantic matching: syntax-based model

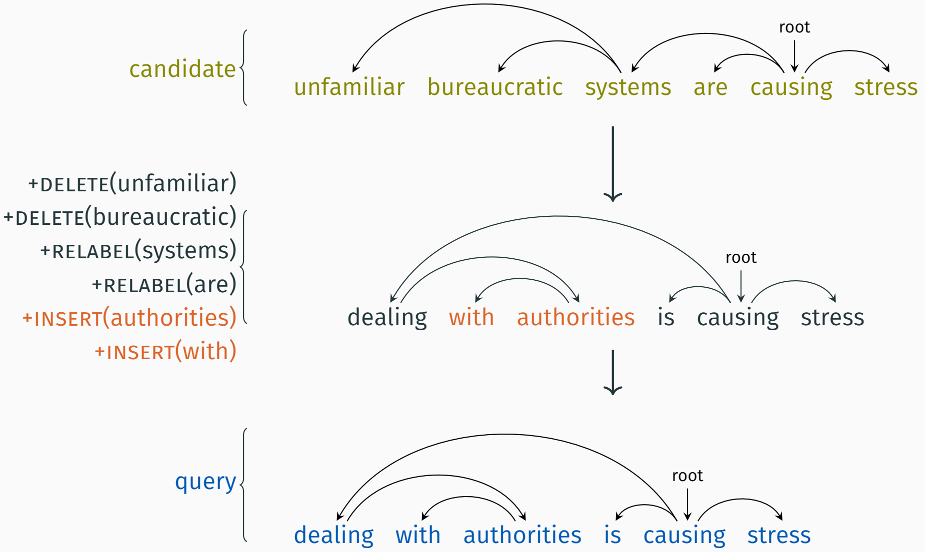

Semantic matching: method overview

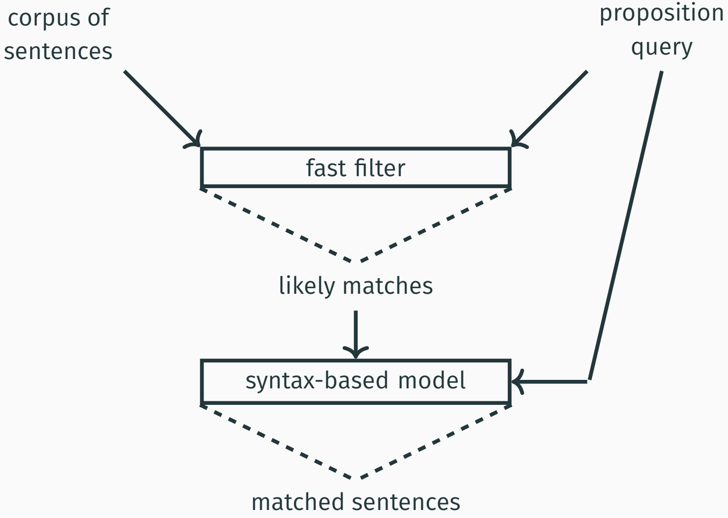

Outline

User evaluation

Questions:

- How good are the sentences matched by our method?
- Do potential users think this kind of tool will be helpful?

User study: 20 emergency managers

User evaluation: output quality

Rated output from 20 proposition queries:

- Different method variants
- Different parts of method:
  - Not selected by filter
  - Selected by filter, but not part of final output
  - Top-scoring output from filter
  - Method output (from syntax-based model)
- 1-5 scale (Krippendorf's $\alpha = 0.784$)

User evaluation: example

Query: **There is a shortage of construction workers.**

"The quarterly report for Canterbury included analysis on Greater Christchurch Value of Work projections."

(**1**: completely unrelated to the query)

User evaluation: example

Query: **There is a shortage of construction workers.**

"The construction sector's workload was expected to peak in December."

(**3**: related to but does not adequately express the query)

User evaluation: example

Query: **There is a shortage of construction workers.**

“Greater Christchurch’s labour supply for the rebuild was tight and was likely to remain that way.”

(**5**: expresses the query in its entirety)

User evaluation: results

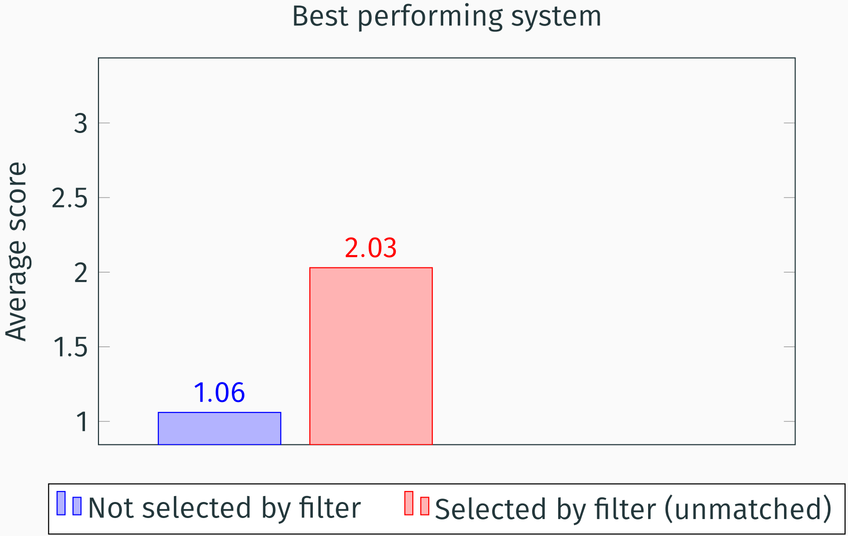

User evaluation: results

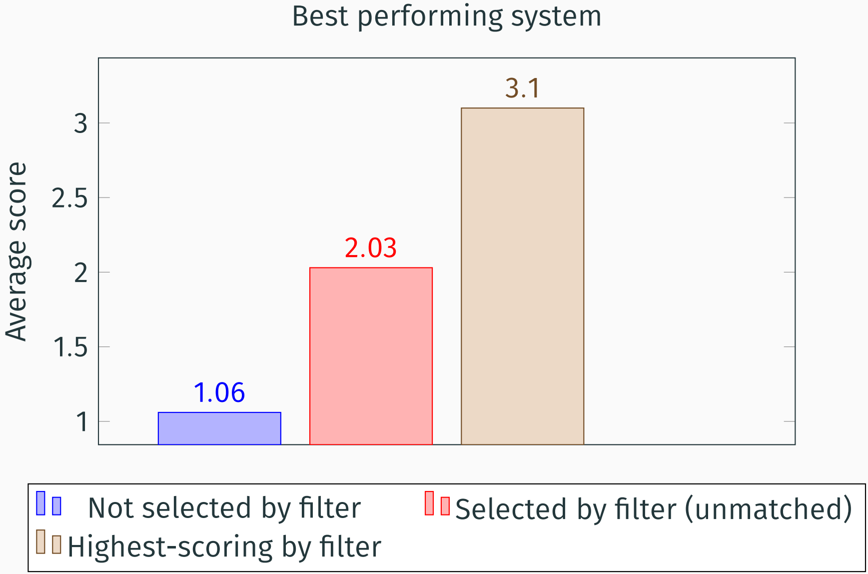

User evaluation: results

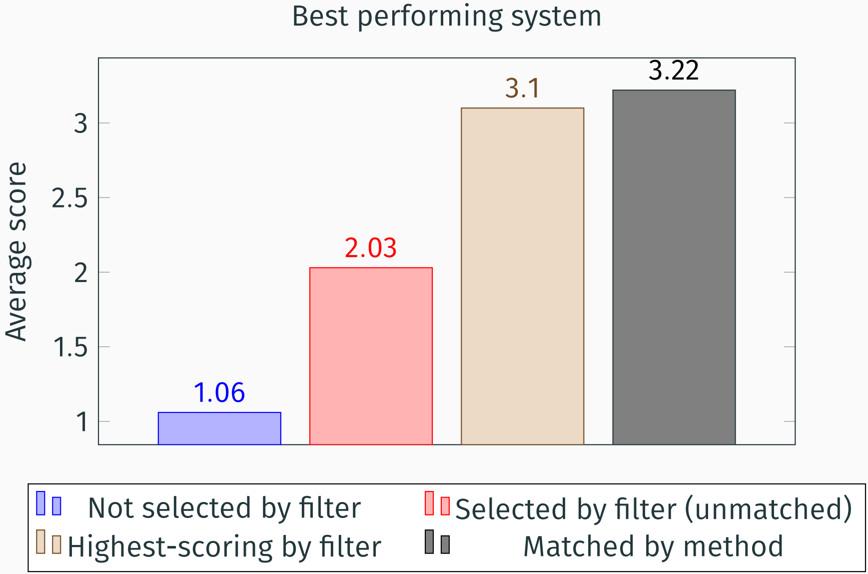

User evaluation: is this interesting?

Other feedback:

- 17/20 respondents interested in measuring ideas in news/other text corpora

User evaluation: round two

Follow-up study:

- Participant-supplied queries (18)
- 7 return participants
- Replicated findings of first user study

Outline

Recovery trends: example #1

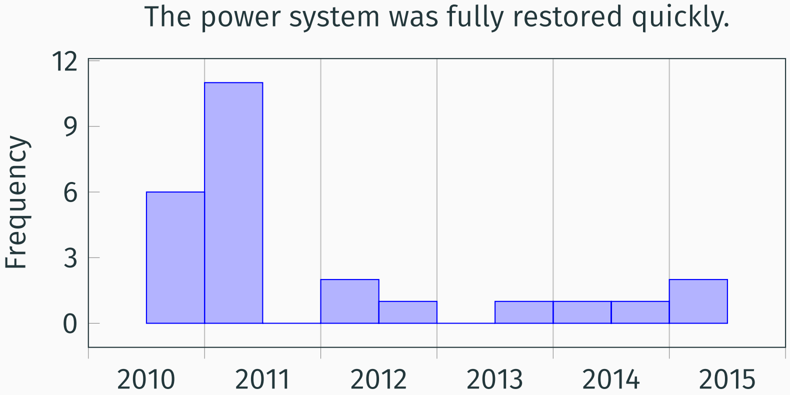

Recovery trends: example #1

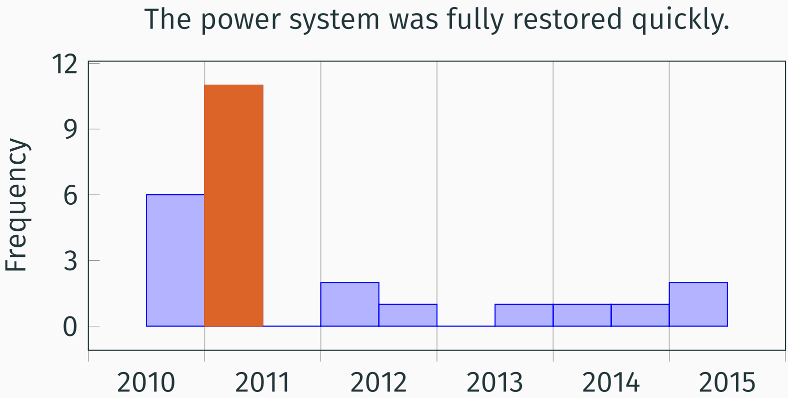

"Orion Energy CEO Roger Sutton says most of the west of Christchurch now has fully restored power."

Recovery trends: example #1

The power system was fully restored quickly.

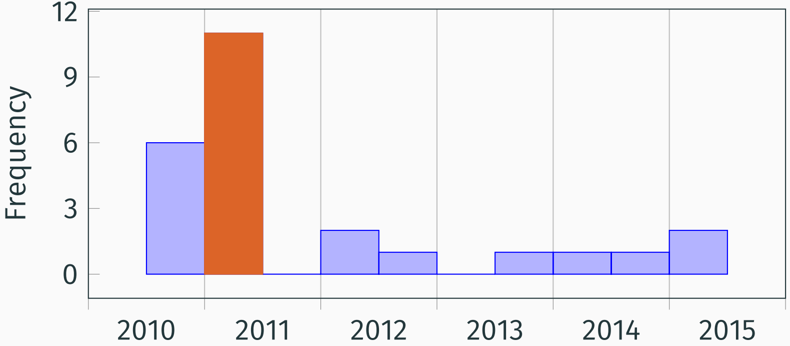

"He had no water but power had been restored in his area."

Recovery trends: example #1

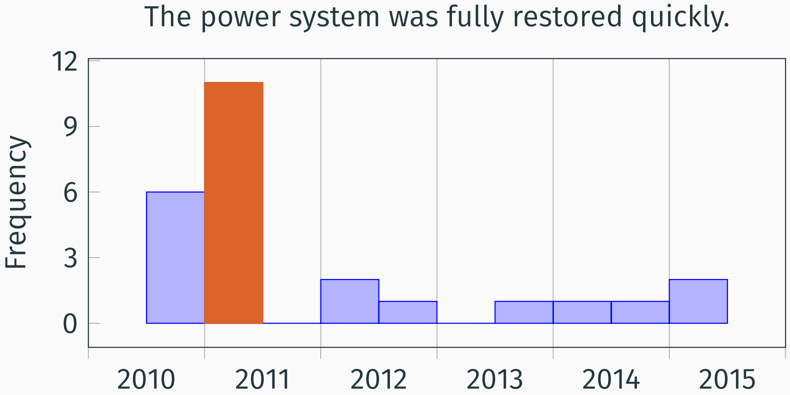

"TV3 reports that power has now been restored to 60 per cent of Christchurch."

Recovery trends: example #1

The power system was fully restored quickly.

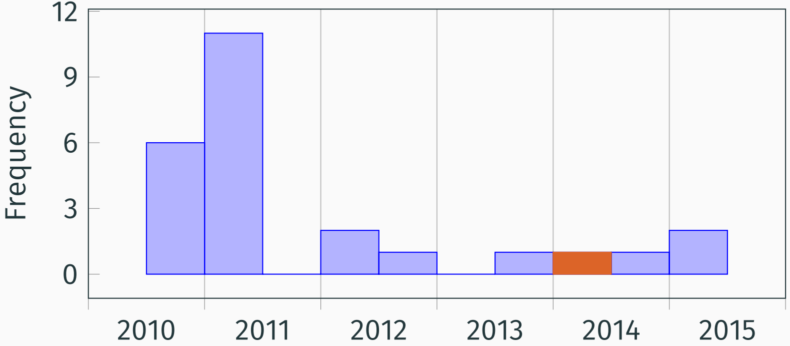

“It had been unable to access the electricity network to restore power and the situation could remain for the next few days.”

Recovery trends: example #2

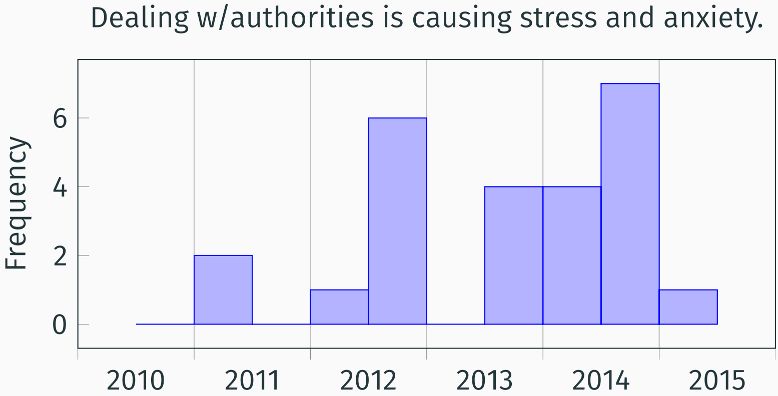

Recovery trends: example #2

Dealing w/authorities is causing stress and anxiety.

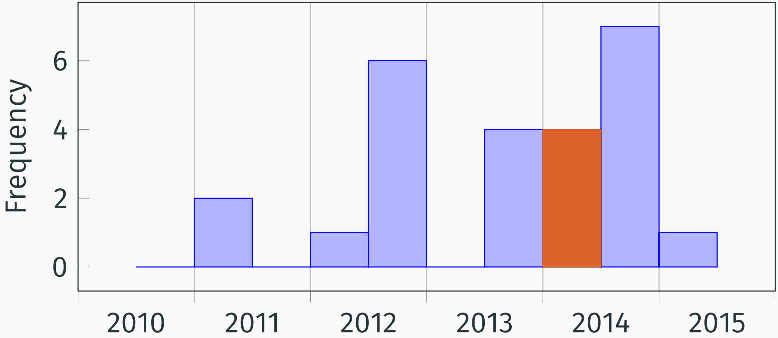

"The initial trauma may be over but [...] Christchurch residents will endure at least six years of 'man-made' stressors as the region battles bureaucracy." (5)

Recovery trends: example #2

Dealing w/authorities is causing stress and anxiety.

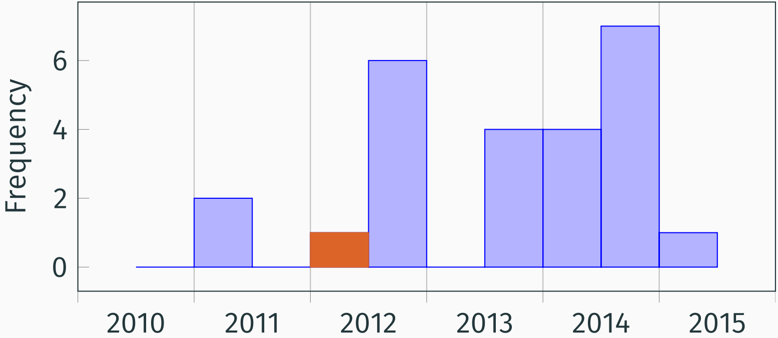

“Add to this the growing frustration **among** the new, youthful leaders of the community who emerged in the wake of the quakes.” (3)

Caveats

- Expected topics generally expressed in the output, but not necessarily relationships/quantities
    - Except some domain-specific entities, e.g., CERA & SCIRT
- Measurement plots best explored jointly w/text output (i.e., quantitative & qualitative)
- Small sample size (25 sentences per query)
- Reliance on sentences as unit of match

Outline

1. Introduction
2. Case study: 2010-2011 Canterbury earthquake disaster
3. NLP method for semantic matching
4. User evaluation
5. Qualitative/quantitative output
6. Conclusions

Conclusion

New NLP method to measure propositions in text corpora

- Potential applications: long-term recovery planning, exploratory research
- User study with participant interest in method
- Future work: richer models, further user engagement

Thanks!

Contact: `lucylin@cs.washington.edu`

Website: `homes.cs.washington.edu/~lucylin/research/semantic_matching.html`

Funding: National Science Foundation
(grant #1541025, graduate fellowship)

(more slides)

In relation to other NLP problems...

paraphrase (Dolan et al., 2004),
entailment (Dagan et al., 2006),
semantic similarity (Agirre et al., 2012)

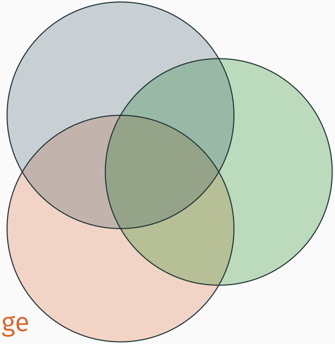

information retrieval,
passage retrieval for QA
(Tellex et al., 2003)

dynamics of language
across a corpus
(e.g., Blei & Lafferty, 2006)

Fast filter details

Pre-trained word vectors:

- word2vec (Mikolov et al., 2013), pre-trained on Google News
- paraphrastic word vectors (Wieting et al., 2015), based off the PPDB

Tree edit classifier details

Original model (Heilman and Smith, 2010):

- Extract 39 integer features from tree edit sequence: sequence length, counts of edit types
- Logistic regression (LR) $\rightarrow m(s_p, s)$

Tree edit classifier details

New variation: input tree edit sequence into a LSTM

Each operation in the sequence is vectorized as:

- One-hot encoding of the operation type
- Word vector Δ between the sentences pre- and post-operation
  - insert $\rightarrow$ word embedding of new word
  - relabel $\rightarrow$ difference between word embeddings
  - delete $\rightarrow$ negated word embedding of deleted word

Tree edit classifier details

Training: SNLI corpus (Bowman et al., 2015)

- 570k pairs of sentences
- labels: entailment, contradiction, neutral
- e.g.,: “A soccer game with multiple males playing.” *entails* “Some men are playing a sport.”

Mapping to our problem:

- $s \rightarrow$ premise, $s_p \rightarrow$ hypothesis
- match $\rightarrow$ entailment,
  non-match $\rightarrow$ contradiction/neutral

Disaster recovery queries

Residents are frustrated by the slow pace of recovery.

The repair programme is on schedule to be completed.

Money for repairs is running out.

The council should have consulted residents before making decisions.

Mental health rates have been rising.

Dealing with authorities is causing stress and anxiety.

Most eligible property owners have accepted insurance offers.

Confidence in Cera has been trending downwards.

Water quality declined after the earthquakes.

The power system was fully restored quickly.

Disaster recovery queries

Cera missed several recovery milestones.

Prices levelled off as more homes were fixed or rebuilt.

People are suffering because they've lost the intimacy of their relationships.

Coordination between rebuild groups has been problematic.

Few people said insurance companies had done a good job.

Having the art gallery back makes the city feel more whole.

Scirt has spent less money than predicted.

Disaster recovery queries

Traffic congestion was severe due to road repairs.

Some of the businesses forced out by the earthquake are returning.

Some of the burden on mental health services is caused by lack of housing.