

Research statement

Modern natural language processing (NLP) offers a diversity of powerful computational and statistical methods [2, 3, 11] which can help answer vital social questions. For instance, NLP can help find social dynamics which lead to war, identify regulators formerly employed in the industries they now regulate or learn how many civilians are killed by police in the United States.

However, because understanding conclusions from statistical language systems currently requires years of specialized study, many contemporary NLP methods are inaccessible to the journalists, policy makers and social scientists who can make best use of their output. Thus, in my research, I create tools which help non-technical users understand, interpret and make decisions based on statistical language processing. The tools I create, developed in close collaboration with end users such as reporters and editors, synthesize many different NLP methods and make use of a wide range of different machine learning techniques. Many of my projects involve user studies, which examine how well the systems I build actually work.



Figure 1: A “concept map” summary of economic forces in 1990s Russia

My current research is focused on creating “concept map summaries”: interpretable, interactive visualizations which summarize complicated relationships between actors in large corpora (fig. 1). Many current text analysis methods [9] involve finding important terms in documents: concept map summarization can help explain how these terms are linked. Creating such summaries requires developing custom statistical language methods. These include:


- Defining *relational summarization*. Traditional work on summarization synthesizes important information in a document or collection of documents. But describing the connections between, say, “Boris Yeltsin” and “free market reforms” requires summarizing their relationship across a corpus. My research will formalize, model, analyze and evaluate relational summarization, a new natural language processing task.
- *Extracting relevant concepts from the corpus*. Decades of research in NLP focuses on extracting people, places and organizations (“named entities” or “NER”) from documents. However, these methods cannot find important phrases like “free market reforms” which should be included in a concept map. In order to identify these important phrases, I worked with a team to create software called **phrasemachine** which we presented at the “**Text as data**” [6] conference. Our open-source implementation, available in **python** and **R**, has gained attention on *GitHub* and found use in social science projects around the country! (See <https://github.com/slanglab/phrasemachine> for a growing list of users and projects).
- *Creating edge labels*. Relationships between actors and concepts are often expressed in long sentences, which must be shortened for display on a concept map. I would like to adapt recurrent neural models for shortening sentences (a process called “sentence compression”) for the new goal of creating edge labels.

I hope that these custom NLP methods will help make concept maps a ubiquitous tool for exploring and analyzing text. However, I also understand how language systems can still be tremendously useful, even if they fall short of “Strong AI” [4]. For instance, over the past several years, I have worked as part of a team to develop a system which can identify instances of civilians killed by police from news text, presented this year at **EMNLP** [8]. While our group’s statistical model does not achieve perfect precision or perfect recall, journalists at Fatal Encounters¹ currently use our work as part of a specialized search interface which helps


¹www.fatalencounters.org

track and document such killings. Similarly, I expect that when journalists begin to use concept maps, the output of the NLP system will only serve as a helpful starting place. Journalists will want to add, modify and investigate [10] new vertexes and edges — guided by suggestions from a computer which directs their attention to relevant passages from an underlying corpus.

This suggests one important way that natural language processing algorithms can become more interpretable for non-specialists: by helping users read. I explored this idea with my “word-as-pixel” visualization at the **Workshop on Human Interpretability in Machine Learning at ICML 2016** [5], which helped users find and interpret similar passages from decades of U.S. State of the Union speeches using a latent variable model (fig. 2 a). Similarly, another project, ROOKIE, used statistical and computational methods to offer “linked views” of a corpus (fig. 2 b), guiding a user’s reading with NLP. (In a linked view interface, the same information is shown from multiple “perspectives”: when one view changes, the other views update to reflect the change). I developed ROOKIE’s “linked views” in close collaboration with reporters and editors, funded with a **Knight foundation prototype grant**. A rigorous evaluation demonstrated that the system helped users find the answer to a complex historical research question 37% faster than a traditional interface [7]. In the future, I hope to use many of the ideas from ROOKIE to create similar text analytics systems which help non-specialists interpret complex NLP methods, such as latent variable models over dependency trees [1].



(a) Visualizing an entire corpus during inference
slanglab.cs.umass.edu/topic-animator/



(b) The ROOKIE system
slanglab.cs.umass.edu/Rookie/

Figure 2: Two systems which help non-specialists interpret and make use of statistical language processing

Modern natural language processing has tremendous power to help journalists, policy makers and social scientists. I hope that by developing new statistical tools, my research can help put these methods into the hands of those who can best use them.

References

- [1] David Bamman and Noah Smith. Open extraction of fine-grained political statements. In *EMNLP*, 2015.
- [2] Allison Chaney, Hanna Wallach, Matthew Connelly, and David Blei. Detecting and characterizing events. In *EMNLP*, 2016.
- [3] Jacob Eisenstein, Brendan O’Connor, Noah A Smith, and Eric P Xing. A latent variable model for geographic lexical variation. In *EMNLP*, 2010.
- [4] Spence Green, Jeffrey Heer, and Christopher D Manning. Natural language translation at the intersection of AI and HCI. *ACM Queue*, 13(6), 2015.
- [5] Abram Handler, Su Lin Blodgett, and Brendan O’Connor. Visualizing textual models with in-text and word-as-pixel highlighting. In *Workshop on Human Interpretability in Machine Learning, ICML*, 2016.
- [6] Abram Handler, Matthew J. Denny, Hanna Wallach, and Brendan O’Connor. Bag of what?: Simple noun phrase extraction for text analysis. In *Text as Data*, 2016.
- [7] Abram Handler and Brendan O’Connor. Rookie: A unique approach for exploring news archives. In *Workshop on Data Science + Journalism at KDD 2017*, Halifax, Canada, 2017.
- [8] Katherine Keith, Abram Handler, Michael Pinkham, Cara Magliozzi, Joshua McDuffie, and Brendan O’Connor. Identifying civilians killed by police with distantly supervised entity-event extraction. In *EMNLP*, 2017.
- [9] Brendan T O’Connor. *Statistical Text Analysis for Social Science*. PhD thesis, Carnegie Mellon University, 2014.
- [10] Jonathan Stray. Network analysis in journalism: Practices and possibilities. In *Workshop on Data Science + Journalism at KDD*, 2017.
- [11] John Wilkerson, David Smith, and Nicholas Stramp. Tracing the flow of policy ideas in legislatures: A text reuse approach. *American Journal of Political Science*, 59(4), 2015.