

As of 2021, an estimated 85% of US adults own a smartphone, many of which are powered by natural language processing systems such as autocomplete. Many modern NLP systems are developed under the assumption that people use the same type of language regardless of the social situation. Autocomplete systems often generate output based only on the input text without explicitly considering the social situation of the text, such as whether the writer is communicating with a friend or a professional colleague. NLP systems may implicitly encode some social cues such as politeness, but they often lack explicit controls that writers may want to leverage for specific social goals, e.g. “provide more information to non-experts about key concepts.” As NLP systems become more integrated into daily life for more diverse situations, it will become increasingly important to develop methods for *personalization*, in order to help the systems address complicated social needs. I am interested in building human-centered NLP systems that can adapt to diverse social situations, with the goal of improving productivity among writers and addressing social “gaps” in model utility. I have extensive experience in applying and developing NLP systems that utilize methods such as word embeddings, information extraction, and language identification, and I look forward in my next position to pushing NLP technology to the limit in order to make it more socially useful.

Prior experience

Traditional research in sociolinguistics, the study of how society impacts language use, has often focused on small-scale studies of spoken language, which can struggle to capture the full scale of language variation for more sparse phenomena such as loanwords. In contrast, computational sociolinguistics research has leveraged the high volume, variety, and velocity of social media data to answer difficult sociolinguistic questions. My PhD thesis in this area investigated how language changes over time online and how multilingual people choose between languages. I used a combination of natural language processing techniques, including word embeddings, parsing, and causal inference to identify long-term effects of language change in online communities, including during disaster events. Furthermore, my work uncovered a strong connection between the language choice of social media members and their intended

audiences: multilingual speakers with smaller audiences often moderated their choice of language and of loanwords, often to accommodate existing norms among smaller audiences. This line of research has provided both theoretical insight into sociolinguistics as well as practical takeaways for NLP systems, such as the need to adapt to diverse social expectations among system users.

Current research

As a postdoctoral fellow at University of Michigan, I have worked to augment existing natural language processing methods with sociolinguistic insight, for the purpose of providing writers with better feedback about their writing. Typical writing feedback systems such as Grammarly have focused heavily on language-internal phenomena such as typos and non-standard grammar, but writing well also requires knowing who will be reading the text and how familiar they are with the subject area. To that end, I am developing a system to generate personalized clarification questions for writers seeking advice, where the system's output is based partly on the writer's intended audience. In a similar vein as my thesis work, I have leveraged the rich social signals available from online communities to train such reader-aware models. Automatic and human evaluation has shown that personalized text generation models do well in cases where when the questions are highly specific to different reader groups, e.g. where an expert reader asks a very different question from a novice. I plan to continue this work with more human-centered testing by providing writers with personalized feedback in the form of questions and in other forms, including keyword extraction and span selection.

Outside of model development, my research at University of Michigan has included projects about identifying bias in NLP systems and assessing the impact of a writing intervention on COVID-related health behavior. I am mentoring several students in projects that relate to detecting racial bias in writing with both experimental and observational data, including associations between racial categories and stereotypes. While not my main focus, these other projects address critical societal issues via computational social science and machine learning.

The bias-related work has highlighted the need for socially-aware NLP systems, because NLP systems should adapt to diverse social groups without reproducing harmful biases.

Research agenda

As a research scientist, I hope to continue work into personalization for the sake of human-centered NLP. Socially-aware NLP systems are not only important for improved performance, they are also critical to addressing the diversity of needs that people have for specific systems. I am interested in helping build NLP systems to preempt these kinds of miscommunications and improve writers' understanding of their target audience.

From a methods perspective, I would like to build on the growing understanding of language model personalization to develop controllable generation systems, for the sake of helping system users get more fine-tuned output to suit their particular social needs. For example, automated email response generation often suffers from generic outputs (e.g. “Thanks”) and will benefit from a stronger representation of the social situation in which the text is generated. Such NLP systems will first require diverse data sets to learn about social norms and expectations, and my prior work has prepared me for the task of collecting and curating large-scale text data sets for use in machine learning. Furthermore, training personalized systems will require developing more complicated model architectures to incorporate social information, such as via user-level embeddings, and I have developed the necessary skills to build such models through my current work in reader-aware models. Finally, personalized NLP models should be evaluated not only in their aggregate performance but also in more “extreme” social cases, such as measuring performance on text that is highly specific to a particular social setting. An NLP system that performs well in general but fails badly in a particular social setting should not be considered robust at all. My current work in personalized generation has prepared me to develop these real-world oriented evaluation strategies, which can highlight the strengths of personalized text generation.

Service and outreach

During my postdoctoral fellowship, I have worked to promote NLP research among a wide range of students. I have mentored five undergraduate and Master's students in independent research projects, with the goal of helping them develop quantitative skills, such as machine learning, and build their own research interests. For each student, I have worked with them to find a project that matches their personal interests and skill set, then adjusted my expectations during the mentorship to match their own goals. My overall goal has been to expose each student to a wide variety of research activities, including literature review, data analysis, and model development, in order to help them decide if research is right for them. I hope to continue this “full stack” style of mentorship in my next position.

Outside of academic outreach, I have organized an ongoing open-source tutorial series to encourage the adoption of new NLP methods among computational social science researchers. The series encourages interactive exploration of methods including dynamic word embeddings and information extraction, for students who have an introductory understanding of NLP. In our initial sessions, we have noted consistent engagement among participants with the code and the concepts, as well as a wide variety of social science backgrounds ranging from political science to economics. In my next position, I intend to continue collaborating with domain experts to develop and disseminate explanations of new NLP methods. Methods such as text-based causal inference can be transformative to social science research, and it's important for domain experts to share knowledge even outside of formal teaching settings.

Extra ideas

In a related line of work, I am also interested in developing collaborations with subject matter experts in computational social science, to better understand the full potential of NLP models for social impact. My prior work in computational social science has relied mainly on interpretable feature-based machine learning, but it remains to be seen whether more powerful approaches, such as transformer models, can be creatively applied to address complex social science problems. While some NLP models such as transformers can be more difficult to interpret, they can also capture complex constructs such as issue framing that can unlock new questions for computational social scientists, e.g. how specific sentences and phrases contribute to the framing of a particular issue. I am particularly interested in turning these more complex models toward the identification of bias in text, which can provide downstream benefits such as filtering biased data for training NLP systems.