Thesis Proposal

Institute for Software Research
Societal Computing

Computational Models of Interactional Stancetaking

Sumeet Kumar

Tuesday, Jan 22nd, 2019 - 1 pm
Newell Simon Hall 4305

People express their opinions on blogs and other social media platforms. Automated ways to categorize views of people in such user-generated corpora is of immense value. My thesis aims to develop computational models to learn to predict the opinion of users who are active on social media platforms like Twitter.

In particular, I explore machine learning approaches for stance learning, which involves learning people's opinion about a topic of interest. Most existing studies on stance learning take a simplistic view that assumes a `sentence' (like a Tweet) holds a perspective that is independent of the context and the author. This approach to stance learning ignores the complex activities and interactions of social media users. According to one definition of stance, it is a public act of evaluating objects, positioning subjects, and aligning with other subjects. In the same spirit, I approach stance learning in the broader context of social action wherein authors take a stance to position themselves on topics of interests, thereby aligning with other stance takers. This approach, therefore, brings a new direction to the stance learning problem which is grounded in social theory and is more amenable to analyzing conversations on social media.

In this research, I plan to develop models for different interactions and then combine them to improve the stance prediction accuracy. First, to predict stance from text, I propose a weakly supervised learning method that reduces the cost of collecting data to train models for new topics. In this approach, I use specific hashtags found on social media platforms that carry stance information. Using these hashtags as noisy labels, I build a classifier that only needs labeled data for validation, yet it achieves accuracy comparable to the state-of-the-art on a human-labeled stance dataset. Second, I plan to extend the text-based stance learning to use multiple posts from a user. Using different social media posts of users is expected to improve the stance classification accuracy. Third, I plan to extract different types of networks on Twitter (e.g. follower network, mentions network), and use graph-based algorithms to align users based on the similarity in the networks. Fourth, I want to use conversation threads to learn the authors' position about an issue based on their posts and the replies to these posts. On Twitter, users can reveal their alignment with other users based on the stance they take in replying to other posts. Such interactions are useful in additionally finding disagreement among users since most other interactions, e.g. liking and retweeting only reflect alignment. After developing independent models for different types of interactions, I intend to propose a joint model that combines the various sources of stance information. To combine different modalities of information (e.g., users' follower graphs and users' likes), I plan to embed features from all information sources into one continuous representation space. Finally, I would like to apply these models to measure community polarization in a real world example.

**Committee: Dr. Kathleen M. Carley (Chair), Dr. Tom Mitchell,
Dr. Louis-Philippe Morency, Dr. Huan Liu (Arizona State)**

institute for SOFTWARE RESEARCH

Carnegie Mellon University
School of Computer Science