

My research aims to draw insights and develop algorithms for a deeper language understanding that involves the meaning of texts, the intention and behaviors of people in online media. Specifically, I have focused on three areas: (1) analyzing semantics for intent understanding, (2) extracting and modeling user roles & relations, and (3) facilitating social interaction. My work focuses on *better understanding the user needs and developing new technologies to enhance online interaction* – for example, predicting what kind of support (informational or emotional support) cancer patients need and recommending the right caregivers to help them.

Analyzing Semantics for Intent Understanding. The recent popularity of social media provides opportunities to conduct large-scale analysis on user behaviors, in areas from social science, health, psychology, economics, etc. Understanding user behaviors requires algorithms to extract what they do, who they are, what they like, and also a deeper language analysis of the meaning of words and the structure of texts[16]. My research try to address these issues. In one such work [10], I identified several semantic structures that people commonly use to express humor, including incongruity structure, ambiguity, interpersonal effect and phonetic style. In addition to automatically identifying humorous texts, I also designed a maximal meaning decrement method to extract anchors that trigger humor in a given sentence, which correlate quite well with human judgments. In one study to understand racism, I developed natural language models to identify personally experienced racial discrimination in Reddit threads, and which type of support (e.g. commiseration, encouragement, empathy, etc.) are provided in subsequent replies, based on which I then built a system to automatically generate supportive responses for any specific instance of racism experienced [4]. In another study [13, 9], I quantified how students express confusion and how such emotion affects their course participation experience. We also identified around 60 types of annoying behaviors that people dislike on Twitter, including bragging, interrupting others, being arrogant, talking too much, etc [1]. I also examined what makes a persuasive loan request in online crowd-funding websites by designing methods to recognize different persuasion strategies such as using authority, showing liking, emphasizing shared identity [7]. These analyses revealed nuanced ways in which user in online communities express their intentions and interact with one another, and hold the potential to facilitate a variety of downstream applications such as recommending friends/content, detecting trolls, supporting minorities, etc.

Extracting Roles & Relations. Millions of volunteers participate in online production communities, exchange their expertise and ideas, and collaborate to produce complex artifacts. Matching individual volunteers to appropriate roles benefits the volunteers by clarifying how they should behave to participate effectively and also benefits the community overall by encouraging members to contribute in ways that best use their skills and interests and coordinate complex activities. To extract volunteers who occupy different editing roles in Wikipedia [6, 5], I developed a fine-grained taxonomy of edit types and distinguish edits contextually in terms of the object being edited (e.g. body copy, reference, template, etc.), functionally in terms of the edit operation (e.g. insert, delete, modify) and semantically in terms of edit intention (e.g. copy editing, verifiability, elaboration, etc.). Via a graphical model, I found a set of editor roles in Wikipedia such as Fact Checker, Copy Editor, etc., which are not directly reflected by prior methods. I use a similar approach to identify elemental

actions such as seeking informational/emotional support and receiving informational/emotional support, in online cancer support groups [8]. We then examined how exchange of such elemental actions influences people's commitment and how the influence changes with their experience in the community, and found that support exchange was generally associated with increased commitment with its impact increasing with members' tenure. I also examined how people occupying different roles collaborate to achieve a successful teamwork. In a team-based MOOC context, I developed a lightly supervised approach to identify behavioral profiles that described the emergent conversational roles team members take up, in order to achieve a successful group grade for their team project [15]. Such identification of effective role-based behavior profiles can enable work towards supporting effective team formation in subsequent studies.

Facilitating Social Interaction. Some online communities [15] fail to provide the social environment that is friendly enough for members' sustained engagement. For example, although Massive Open Online Courses (MOOCs) enable thousands of students to register for courses to learn at their convenience with no monetary cost, they suffer from high rates of attrition due to the rare interaction within courses. To encourage more social interaction on this platform, I first investigated the reasons behind this extremely high dropout rate [14, 2, 12]. These analyses reveal several significant predictors in better understanding students' experience of participation as they struggle and then ultimately dropout of a course. As a further step, I designed a recommender system to suggest appropriate answerers to students' posted questions in the discussion forums. our system provides a generalized question recommendation framework which also considers load balancing and expertise matching constraints [3, 11]. This enables the recommendation to be reciprocal and attractive to both the asker and the helper side, while existing techniques are inadequate in modeling. I then deployed this system into an edX MOOC that leads to real world benefits. Similarly, to increase user engagement, we are building another intervention in the discussion forum of Cancer Survivor Network, aiming at recommending useful threads for cancer patients and connecting them with similar others.

To sum up, my research bridges social computing and natural language processing, by developing natural language understanding techniques to solve problems in computational social science, in a variety of applications. In the future, I would like to continue social role identification in online communities especially in online health support groups. By analyzing social roles and expertise levels, we can model *how different role configurations affect the group success* and identify strategies that lead to successful outcomes for online groups. We then can take advantages of our identified roles to provide guidance to these unsuccessful groups. We are also able to identify *trolls, spammers, lurkers and other social roles* via our deep text analysis techniques. My work about detecting personally experiences of racial discrimination sheds lights on how to monitor and track different kinds of discrimination issues, protests and social movements in social media, e.g. who discriminates against whom, what triggers a conflict, etc.

References

- [1] William Yang Wang and Diyi Yang. That's so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets.

- In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2557–2563, Lisbon, Portugal, September 2015. Association for Computational Linguistics.
- [2] Miaomiao Wen, Diyi Yang, and Carolyn Penstein Rosé. Linguistic reflections of student engagement in massive open online courses. In *Proceedings of the International Conference on Weblogs and Social Media*, 2014.
 - [3] Diyi Yang, David Adamson, and Carolyn Penstein Rosé. Question recommendation with constraints for massive open online courses. In *Proceedings of the 8th ACM Conference on Recommender Systems, RecSys '14*, pages 49–56, New York, NY, USA, 2014. ACM.
 - [4] Diyi Yang and Scott Counts. Detecting and supporting victims of racism in social media. In *Under Review*, 2016.
 - [5] Diyi Yang, Aaron Halfaker, Eduard Hovy, and Robert Kraut. Identifying semantic edit intentions from revisions in wikipedia. In *EMNLP*, September 2017.
 - [6] Diyi Yang, Aaron Halfaker, Robert Kraut, and Eduard Hovy. Who does what: Role identification in wikipedia. In *Proceedings of the International Conference on Weblogs and Social Media*, 2016.
 - [7] Diyi Yang and Robert Kraut. Persuading teammates to give: Systematic versus heuristic cues for soliciting loans. In *CSCW*, 2018.
 - [8] Diyi Yang, Robert Kraut, and John Levine. Commitment of newcomers and old-timers to online health support communities. In *CHI*, 2017.
 - [9] Diyi Yang, Robert Kraut, and Carolyn Rose. Exploring the effect of confusion in massive open online courses. In *Journal of Educational Data Mining*, 2016.
 - [10] Diyi Yang, Alon Lavie, Chris Dyer, and Eduard Hovy. Humor recognition and humor anchor extraction. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2367–2376, Lisbon, Portugal, September 2015. Association for Computational Linguistics.
 - [11] Diyi Yang, Jingbo Shang, and Carolyn Penstein Rosé. Constrained question recommendation in moocs via submodularity. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM '14*, pages 1987–1990, New York, NY, USA, 2014. ACM.
 - [12] Diyi Yang, Tanmay Sinha, David Adamson, and Carolyn Penstein Rose. Turn on, tune in, drop out: Anticipating student dropouts in massive open online courses. In *Proceedings of the 2013 NIPS Data-Driven Education Workshop*, volume 10, page 13, 2013.
 - [13] Diyi Yang, Miaomiao Wen, Iris Howley, Robert Kraut, and Carolyn Rose. Exploring the effect of confusion in discussion forums of massive open online courses. In *Proceedings of the Second (2015) ACM Conference on Learning @ Scale, L@S '15*, pages 121–130, New York, NY, USA, 2015. ACM.
 - [14] Diyi Yang, Miaomiao Wen, Abhimanu Kumar, Eric P Xing, and Carolyn Penstein Rose. Towards an integration of text and graph clustering methods as a lens for studying social interaction in moocs. *The International Review of Research in Open and Distributed Learning*, 15(5), 2014.
 - [15] Diyi Yang, Miaomiao Wen, and Carolyn Rose. Weakly supervised role identification in teamwork interactions. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1671–1680, Beijing, China, July 2015. Association for Computational Linguistics.
 - [16] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016.