

Data-Driven Dialogue Systems for Social Agents

Kevin K. Bowden, Shereen Oraby, Amita Misra, Jiaqi Wu, and Stephanie Lukin

Natural Language and Dialogue Systems Lab
University of California, Santa Cruz
{kkbowden, soraby, amisra2, jwu64, slukin}@ucsc.edu

Abstract In order to build dialogue systems to tackle the ambitious task of holding social conversations, we argue that we need a data-driven approach that includes insight into human conversational “chit-chat”, and which incorporates different natural language processing modules. Our strategy is to analyze and index large corpora of social media data, including Twitter conversations, online debates, dialogues between friends, and blog posts, and then to couple this data retrieval with modules that perform tasks such as sentiment and style analysis, topic modeling, and summarization. We aim for personal assistants that can learn more nuanced human language, and to grow from task-oriented agents to more personable social bots.

1 From Task-Oriented Agents to Social Bots

Devices like the Amazon Echo and Google Home have entered our homes to perform task-oriented functions, such as looking up today’s headlines and setting reminders [1, 2]. As these devices evolve, we have begun to expect social conversation, where the device must learn to personalize and produce natural language style.

Social conversation is not explicitly goal-driven in the same way as task-oriented dialogue. Many dialogue systems in both the written and spoken medium have been developed for task-oriented agents with an explicit goal of restaurant information retrieval, booking a flight, diagnosing an IT issue, or providing automotive customer support [6, 16, 23, 24, 4, 21]. These tasks often revolve around question answering, with little “chit-chat”. Templates are often used for generation and state tracking, but since they are optimized for the task at hand, the conversation can either become stale, or maintaining a conversation requires the intractable task of manually authoring many different social interactions that can be used in a particular context.

We argue that a social agent should be spontaneous, and allow for human-friendly conversations that do not follow a perfectly-defined trajectory. In order to build such a conversational dialogue system, we exploit the abundance of human-human social

media conversations, and develop methods informed by natural language processing modules that model, analyze, and generate utterances that better suit the context.

2 Data-Driven Models of Human Language

A myriad of social media data has led to the development of new techniques for language understanding from open domain conversations, and many corpora are available for building data-driven dialogue systems [17, 18]. While there are differences between how people speak in person and in an online text-based environment, the social agents we build should not be limited in their language; they should be exposed to many different styles and vocabularies. Online conversations can be repurposed in new dialogues, but only if they can be properly indexed or adapted to the context. Data retrieval algorithms have been successfully employed to co-construct an unfolding narrative between the user and computer [20], and re-use existing conversations [5]. Other approaches train on such conversations to analyze sequence and word patterns, but lack detailed annotations and analysis, such as emotion and humor [19, 22, 8]. The large Ubuntu Dialogue Corpus [9] with over 7 million utterances is large enough to train neural network models [7, 10].

We argue that combining data-driven retrieval with modules for sentiment analysis and style, topic analysis, summarization, paraphrasing, and rephrasing will allow for more human-like social conversation. This requires that data be indexed based on domain and requirement, and then retrieve candidate utterances based on dialogue state and context. Likewise, in order to avoid stale and repetitive utterances, we can alter and repurpose the candidate utterances; for example, we can use paraphrase or summarization to create new ways of saying the same thing, or to select utterance candidates according to the desired sentiment [12, 13]. The style of an utterance can be altered based on requirements; introducing elements of sarcasm, or aspects of factual and emotional argumentation styles [15, 14]. Changes in the perceived speaker personality can also make more personable conversations [11]. Even utterances from monologic texts can be leveraged by converting the content to dialogic flow, and performing stylistic transformations [3].

Of course, while many data sources may be of interest for indexing knowledge for a dialogue system, annotations are not always available or easy to obtain. By using machine learning models designed to classify different classes of interest, such as sentiment, sarcasm, and topic, data can be bootstrapped to greatly increase the amount of data available for indexing and utterance selection [15].

There is no shortage of human generated dialogues, but the challenge is to analyze and harness them appropriately for social-dialogue generation. We aim to combine data-driven methods to repurpose existing social media dialogues, in addition to a suite of tools for sentiment analysis, topic identification, summarization, paraphrase, and rephrasing, to develop a socially-adept agent that can carry on a natural conversation.

References

1. Amazon Echo. <https://www.amazon.com/Amazon-Echo-Bluetooth-Speaker-with-WiFi-Alexa/dp/B00X4WHP5E>.
2. Google Home Features. <https://madeby.google.com/home/features/>.
3. Kevin K. Bowden, Grace I. Lin, Lena I. Reed, Jean E. Fox Tree, and Marilyn A. Walker. M2D: monolog to dialog generation for conversational story telling. In *Interactive Storytelling - 9th International Conference on Interactive Digital Storytelling, ICIDS 2016, Los Angeles, CA, USA, November 15-18, 2016, Proceedings*, pages 12–24, 2016.
4. Matthew Henderson, Blaise Thomson, and Jason Williams. The second dialog state tracking challenge. In *Proceedings of SIGDIAL*. ACL Association for Computational Linguistics, June 2014.
5. Ryuichiro Higashinaka, Kenji Imamura, Toyomi Meguro, Chiaki Miyazaki, Nozomi Kobayashi, Hiroaki Sugiyama, Toru Hirano, Toshiro Makino, and Yoshihiro Matsuo. Towards an open-domain conversational system fully based on natural language processing. In *COLING*. ACL, 2014.
6. Lynette Hirschman. Evaluating spoken language interaction: Experiences from the darpa spoken language program 1990–1995. *Spoken Language Discourse*. MIT Press, Cambridge, Mass, 2000.
7. Rudolf Kadlec, Martin Schmid, and Jan Kleindienst. Improved deep learning baselines for ubuntu corpus dialogs. *CoRR*, abs/1510.03753, 2015.
8. Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A Persona-Based Neural Conversation Model. *arXiv preprint arXiv:1603.06155*, 2016.
9. Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the SIGDIAL 2015 Conference, The 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 2-4 September 2015, Prague, Czech Republic*, pages 285–294, 2015.
10. Ryan Thomas Lowe, Nissan Pow, Iulian Vlad Serban, Laurent Charlin, Chia-Wei Liu, and Joelle Pineau. Training end-to-end dialogue systems with the ubuntu dialogue corpus. *D&D*, 8(1):31–65, 2017.
11. Stephanie Lukin, James Owen Ryan, and Marilyn A Walker. Automating direct speech variations in stories and games. In *Proceedings of the 3rd Workshop on Games and NLP*, October 2014.
12. Amita Misra, Pranav Anand, Jean E. Fox Tree, and Marilyn A. Walker. Using summarization to discover argument facets in online ideological dialog. In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, 2015.
13. Amita Misra, Brian Ecker, and Marilyn A. Walker. Measuring the similarity of sentential arguments in dialogue. In *Proceedings of the SIGDIAL 2016 Conference, The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 13-15 September 2016, Los Angeles, CA, USA*, 2016.
14. Shereen Oraby, Vrindavan Harrison, Lena Reed, Ernesto Hernandez, Ellen Riloff, and Marilyn A. Walker. Creating and characterizing a diverse corpus of sarcasm in dialogue. In *Proceedings of the SIGDIAL 2016 Conference, The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 13-15 September 2016, Los Angeles, CA, USA*, pages 31–41, 2016.
15. Shereen Oraby, Lena Reed, Ryan Compton, Ellen Riloff, Marilyn A. Walker, and Steve Whittaker. And that’s A fact: Distinguishing factual and emotional argumentation in online dialogue. In *Proceedings of the 2nd Workshop on Argumentation Mining, ArgMining@HLT-NAACL 2015, June 4, 2015, Denver, Colorado, USA*, pages 116–126, 2015.
16. Patti Price, Lynette Hirschman, Elizabeth Shriberg, and Elizabeth Wade. Subject-based evaluation measures for interactive spoken language systems. In *Proceedings of the workshop on Speech and Natural Language*, pages 34–39. Association for Computational Linguistics, 1992.

17. Alan Ritter, Colin Cherry, and William B. Dolan. Data-driven response generation in social media. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2011.
18. Iulian Vlad Serban, Ryan Lowe, Laurent Charlin, and Joelle Pineau. A survey of available corpora for building data-driven dialogue systems. *CoRR*, abs/1512.05742, 2015.
19. Alessandro Sordani, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. A Neural Network Approach to Context-Sensitive Generation of Conversational Responses. *arXiv preprint arXiv:1506.06714*, 2015.
20. Reid Swanson and Andrew S Gordon. Say Anything: A Massively Collaborative Open Domain Story Writing Companion. In *Interactive Storytelling*, pages 32–40. Springer, 2008.
21. Pirros Tsiakoulis, Milica Gasic, Matthew Henderson, Joaquin Plannels, Jorge Prombonas, Blaise Thomson, Kai Yu, Steve Young, and Eli Tzirkel. Statistical methods for building robust spoken dialogue systems in an automobile. In *4th International Conference on Applied Human Factors and Ergonomics*, 2012.
22. Oriol Vinyals and Quoc Le. A Neural Conversational Model. *arXiv preprint arXiv:1506.05869*, 2015.
23. Marilyn A Walker, Diane J Litman, Candace A Kamm, and Alicia Abella. Paradise: A framework for evaluating spoken dialogue agents. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*, pages 271–280. Association for Computational Linguistics, 1997.
24. Marilyn A Walker, Rebecca Passonneau, and Julie E Boland. Quantitative and qualitative evaluation of darpa communicator spoken dialogue systems. In *Proceedings of the 39th Annual Meeting on Association for Computational Linguistics*, pages 515–522. Association for Computational Linguistics, 2001.