Better Together: Combining Language and Social Interactions in Predictive Models of Individual Behavior

PIs: Jennifer Neville and Dan Goldwasser

Overview

Low-level observational data in social networks often consists of **transactions**—actions and interactions observed over time, for individuals, pairs, and groups of users. While this observed data provides some information about people and their relationships over time, there is also an important unobserved aspect of the system that influences behavior (e.g., personal beliefs, relational influence, roles, communities). To improve understanding of this behavior, it is critical to use the observed data to make inferences about the unobserved components of the system. For example, to identify the formation of groups and predict their influence on personal beliefs/actions, we need to be able to reason about latent properties such as user roles, relationship types, community affiliations, etc.

The technical challenge to developing data-driven methods to learn models of individual and collective behavior lies in the dimensionality of the data and the complexity of the modeling task(s). The data is relational, temporal, and often spatial, and can include a large amount of high-dimensional information due to text (e.g., messages), images (e.g., photos in posts), or actions (e.g., reacting to posts). The modeling task also involves more than just learning to make independent predictions from static views of the data. Modeling relational behavior requires joint representations for learning and inference, and the discovery of hidden temporal patterns requires more complex dynamic, latent variable models that can optimize intermittent signals (i.e., reward) over time in structured data.

Current State of the Art

The interactions, social bonds, and relationships between people have been studied extensively in recent years. Broadly speaking, these works fall into two, almost completely disconnected, camps. The first, focusing on **social network analysis**, looks at the network structure and information flow on it as means of inferring knowledge about the network. For example, some works model the evolution of social network structure over time, and other works use the network structure to predict properties of links (e.g., tie strength, type of relationship). The second camp, focusing on **natural language analysis**, looks into tasks such as extracting social relationships from narrative text and analyzing the contents of the information flowing through the network. For example, some work extracts attributes of, and social relationships between, nodes by analyzing the textual communication between them. Other works use the social network to inform language analysis.

Both perspectives on social network analysis have resulted in a wide range of successful applications; however, they neglect to model the interactions between the social and linguistic representations and how they complement one another. One of the few exceptions was discussed in [7], which inferred sentiment links between nodes in a social network by jointly modeling the local output probabilities of a sentiment analyzer looking at the textual interactions between the nodes and the global network structure. While this work somewhat improved performance, inference is done over two independent representations, one capturing the linguistic information, and the other, the network structure.

Proposed Research

In this project, we aim to develop novel data-driven methods for learning predictive models using a structured machine learning framework—with a *joint representation* over relational and textual information. Given the wealth of social interaction and individual behavioral data available in online and electronic systems, there is a unique opportunity to learn models of individual/relationship dynamics from these data and discover latent patterns of individual beliefs, relational roles, and group norms that impact social behavior. We aim to develop robust statistical methods that can learn accurate, parsimonious models from observed sequences of social interactions. Specifically, we focus on analyzing relational discourse to identify content and structure patterns, including how they evolve and shift over time with respect to role/tie formation. Natural language interactions reveal important information about the speakers' identity, belief systems and their social relationships. This information can be captured by closely observing both the conversational structure (who-talks-to-whom), and its content (about what).

The key insight of our approach is derived from our extensive past work on developing ML methods for natural language processing, complex network, and relational domains. We have identified a common property of ML methods that are able to successfully expand the complexity of the model space without incurring the typical reduction in estimation effectiveness due to higher dimensionality of joint representations. In a wide variety of scenarios, we have observed that while it is necessary to increase dimensionality to improve model expressivity, estimation over the full range of model space is not necessary for accurate modeling observed data (i.e., large parts of the complex space have zero likelihood of occurrence). With this in mind we can use **inductive bias**, designed from background knowledge of social theories and/or desired algorithmic properties, to reduce the *effective* dimensionality for accurate estimation and inference in spite of an enormous full model space. For example, we have shown how to use generative models of network structure (without attributes) in a simple *proposal distribution*, to model a more complex target distribution of attribute networks with correlated nodal attributes [3, 6]. In our proposed work, we aim to apply these insights to develop models of **social trajectories**, in order to reason about individual actions and relational interactions *over time* and estimate summaries of the latent properties of people, dyads, and groups that impact behavioral decisions.

We propose to develop methods for analyzing individual attitudes and interests, and the trajectory of their change over time. Based on textual content analysis we will identify the speakers attitudes and opinions, and look at their temporal patterns. Our goal is to capture similarities and differences in perspective (e.g., sharing political views, ability to unite behind a common goal, etc.) on a wide range of topics, for large population of users. This analysis can be quite subtle as often attitudes are not explicitly stated, but rather implied by the way topics are framed, the communication style and the response to real-world events as they occur. From a technical perspective, we will construct a holistic structured prediction model that moves beyond standard topic models (e.g., LDA) to combine multiple inter-connected sources of information, such as the conversational structure, social relationships, and relevant content indicators. These efforts will build on our recent work, which analyzed college student interactions in online courses, with the goal of modeling their engagement levels and social bonds formed between students [5, 1], and looked at political bias in social media, where we aimed to identify fine-grained stances on a wide range of political issues by observing these indicators [2].

In contrast with previous work, we aim to learn a **joint representation** over both linguistic and relational information, rather than treating the two independently. We follow the intuition that interactions in a social network can be fully captured only by taking into account both types of information together. To achieve this goal, we will embed the input social graph into a dense, continuous, low-dimensional vector space, capturing both network and linguistic similarities between nodes. Our approach aims to map both social and linguistic information into the same vector space, rather than embedding the two aspects into two independent spaces. The social graph, originally containing only quantitative properties of the interaction between nodes (e.g., number of messages exchanged between nodes), is extended to capture the contents of these interactions, by computing the textual similarity between the messages generated by each one of the nodes.

Impact Our proposed work will provide new data-driven methods to estimate models of complex social behavior from observed electronic/online signals. The available data is large-scale, multi-modal, streaming over time, but it provides a noisy indication of hidden personal and relational information. Machine learning methods that can accurately learn models from noisy low-level interaction data, account for the uncertainty, and extract parsimonious temporal patterns of social behavior and personal interests, combined with robust statistical methods for testing hypotheses—will improve predictive modeling for a range of domains, including e-commerce and online social network systems, as well as personalized healthcare and education.

Plan and Budget

Our preliminary work in this direction [4] makes the primary assumption that there is a latent space that influences the interactions we observe among people. Thus the aim is learn this latent representation from the observed data by embedding two aspects (individual attributes and textual content) into a single vector space. We have shown that the new representation can improve predictive performance in two social relation prediction tasks in real world datasets. Our plan is to: (1) formalize our initial approach, showing how to generally combine multiple sources of information into a single joint representation during learning, (2) develop principled methods to evaluate the quality of the learned (latent) representation, and (3) extend our initial model and apply them to a range of static and dynamic social prediction tasks, using both social and education data (in collaboration with Purdue Office of Institutional Research, Assessment, and Effectiveness).

Budget We request funding for one 50% time graduate research assistant for one year and $5,000 travel funds (for student travel to relevant conference/workshops).

References

[1] Snigdha Chaturvedi, Dan Goldwasser, and Hal Daumé III. Predicting instructor's intervention in mooc forums. In *ACL (1)*, pages 1501–1511, 2014.

[2] Dan Goldwasser and Hal Daumé III. " i object!" modeling latent pragmatic effects in courtroom dialogues. In *EACL*, pages 655–663, 2014.

[3] J. Pfeiffer III, S. Moreno, T. La Fond, J. Neville, and B. Gallagher. Attributed graph models: Modeling network structure with correlated attributes. In *Proceedings of the 23rd International World Wide Web Conference (WWW)*, 2014.

[4] Y. Lai, C. Li, D. Goldwasser, and J. Neville. Better together: Combining language and social interactions into a shared representation. In *Proceedings of the TextGraphs Workshop 2016, NAACL*, 2016.

[5] Arti Ramesh, Dan Goldwasser, Bert Huang, Hal Daumé III, and Lise Getoor. Learning latent engagement patterns of students in online courses. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 1272–1278. AAAI Press, 2014.

[6] P. Robles, S. Moreno, and J. Neville. Sampling of attributed networks from hierarchical generative models. In *Proceedings of the 22nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2016.

[7] Robert West, Hristo S Paskov, Jure Leskovec, and Christopher Potts. Exploiting social network structure for person-to-person sentiment analysis. *arXiv preprint arXiv:1409.2450*, 2014.