Research Statement

Victor Amelkin

Network processes have surrounded people for thousands of years. Heat diffusing through a complex material, warriors coordinating on the battlefield, commodities spreading through a network of merchants, people passing rumors via word of mouth, innovations proliferating through an organizational network—whenever we see a system of interacting entities, we deal with a network process. A special place is occupied by *the processes in social and economic networks*, as they are immediately connected to people's social and financial well-being. Most recently, with the emergence of massive-scale online social networks, and trading and collaboration platforms and tools, studying network processes has become feasible and important as never before.

The *central theme* of my research is **analysis, modeling, and control of dynamic processes in networks**, with an emphasis on online social networks, collaboration networks, and economic networks. In this statement, I delineate the scope of the inter-disciplinary field of dynamic network processes along three facets:

1. **Analysis**—creation of scalable methods for model-driven analysis of observed network processes.
2. **Modeling**—design and theoretical analysis of models for network processes rooted in first principles of social psychology and behavioral economics.
3. **Control**—development of theories and the accompanying scalable algorithms to guide network processes for optimizing and/or safeguarding these processes' performance.

Below, I discuss my contributions to and the plan for the future development of the dynamic network process cross-discipline along these three facets.

1 Model-driven Analysis of Network Processes

We often encounter a situation when we have no control over a network process, and only observe how its state evolves in time. For example, when user opinions spread through an online social network, we can periodically observe the user opinion distribution. Two general types of analyses useful in this setting are (1) *anomalous state detection* (e.g., detecting when the evolution of user opinions drastically deviates from the expected word-of-mouth dynamics, potentially indicating the start of a viral marketing campaign); and (2) *future state prediction* (predicting how user opinions will change in the future). Straightforward domain-oblivious approaches to towards performing these analyses—such as a domain-independent feature extraction followed by anomaly detection or extrapolation in the obtained vector space—are unlikely to be efficient. Indeed, to tell whether an observed change in the state of a network process is anomalous, we need a domain-specific definition of what is expected and unexpected in process' evolution. Thus, an effective process analysis should exploit domain-specific knowledge about the network process, captured in this process' *model*.

The key question to answer to enable the model-driven analysis of a network process' state dynamics is: *Given two states of a network process, what is the likelihood of the process' transitioning between these two states with respect to a given process model?* For example, assuming that the spread of opinions through a social network follows a fitted instance of the Independent Cascade model, we are after the likelihood of the transition between two observed user opinion distributions, where the pathways for the opinions' evolution are determined by the model. The exact computation of this likelihood—or the distance between two network process' states—is unfeasible, as it requires traversal of the process' entire exponential-size state space. A practically useful estimate, however, can be efficiently computed, as established in my past work.

Past Work [1],[2]: I have designed a general scalable method to estimate the (log-)likelihood of network process' transitioning between states in the context of opinion dynamics in social networks. This method—a distance measure for network process states termed Social Network Distance (SND)—computes the (log-)likelihood of the most likely transition between two observed user opinion distributions of an opinion spread process in a large online social network. SND reduces the state transition likelihood definition to that of a transportation problem or, alternatively, a well-known Earth Mover's Distance. The latter, roughly, computes the "cost" of optimal way to reshape one (opinion) distribution into another one, where the elementary transform—in the edit distance sense—is the transportation of an opinion between two users along the pathways define by the social network and the chosen opinion dynamics model. While the direct solution of the obtained transportation problem would have super-cubic time-complexity, SND exploits the special structure of the transportation problem and effectively uses a combination of radix-Fibonacci heap-based Dijkstra algorithm and bi-push min-cost network flow algorithm for unbalanced bipartite graphs, resulting in SND's computability in time linear in the network's size. In experiments with Twitter data, SND has shown to be effective at detecting controversial events—such as the introduction of "Obamacare" in the US in 2010 —that are known to have polarized the US society. SND has also shown to be effective at predicting future opinions of a small number of network users (however, see future work below).

Future Work: One open problem within the scope of the general network process analysis framework that I plan to study in the future is *how to efficiently explore the state space of an observed network process.* Consider a situation when, having observed a network process' dynamics, we need to predict how it will evolve in the near future. For example, having observed how the political opinions of a social network's users have evolved in the past, we want to predict how these opinions will change at the time of the upcoming elections. In other words, we want to predict the future state of a network process based on a series of its past states. While SND [1, 2] can estimate the likelihood of a specific change in a network process' state, here we are faced with the need to *search for the most plausible future state of the process.* One solution is as follows: extrapolate the series of distances between the adjacent past states of the process to estimate the expected distance to the process' future state, and, then, among all future state candidates take as the prediction the state the distance to which is closest to the made distance estimate. This distance-based search for the future state of a network process has been shown to work in practice [1], but its fundamental bottleneck is the need to enumerate a possibly exponential number of future state candidates. To solve the problem of efficient network process' state space traversal, we will need to (1) rely upon the domain-specifics structure of this space, and (2) exploit semi-metricity of SND, with the goal of reducing the size of the network process' state space.

Related Papers:

[1] **V. Amelkin**, P. Bogdanov, and A. K. Singh. A distance measure for the analysis of polar opinion dynamics in social networks. In *Proc. of IEEE Int. Conf. on Data Engineering (ICDE)*, pages 159–162, 2017. [pdf],[doi],[slides],[code].

[2] **V. Amelkin**, P. Bogdanov, and A. K. Singh. A distance measure for the analysis of polar opinion dynamics in social networks (Extended Paper). *In submission to ACM Trans. on Knowl. Discov. from Data (TKDD)*, 2017. [arxiv].

2 Modeling Network Processes

Having a theoretical model for a network process is essential, as not only it helps with the process' analysis, but also provides insight into the nature of the process. This prospect is particularly attractive for studying social network processes, since learning human nature, besides its fundamental value, immediately impacts our lives.

In 1954, Festinger in his social comparison theory stated that, as long as there is room for subjective judgment, humans evaluate their opinions by comparison with the opinions of others, thereby, establishing that the opinion formation process is inherently *a network process.* Over the following decades, sociologists and social psychologists—from Abelson to Friedkin—have been successfully applying mathematical tools to modeling opinion or attitude formation processes in the society. At first, the focus has primarily been on linear models, whose analysis is straightforward and can bring useful analytic results. Later, engineers and applied mathematicians proposed a range of non-linear opinion formation models—such as Deffuant and Hegselmann-Krause models. These models' behavior, however, has mostly been studied via simulation, and their complete theoretical analysis is yet to come.

In my work on modeling the opinion formation process in social networks, I generally focus on the design of non-linear models that (1) rely on first principles from sociology and social psychology; (2) incorporate multiple interacting opinion formation "forces"; and, at the same time, (3) allow for a comprehensive theoretical analysis. A few specific past and future social process modeling efforts of mine are discussed in what follows.

Past Work [3]: One fundamental model of opinion formation—stemming from Festinger's social comparison and cognitive dissonance theories—is DeGroot model stating that members of a social network—whose adjacency matrix' entries reflect the relative amounts of inter-personal trust—form their opinions via weighted averaging with the opinions of their network neighbors. One limitation of this model is that every person manifests the same opinion adoption behavior, while in reality people have different receptiveness to persuasion. Within the realm of linear models, this drawback is resolved by Friedkin-Johnsen model that allows the network's members to have different (though, constant) degrees of "stubbornness".

In [3], I have proposed a class of non-linear opinion formation models, having incorporated dynamic opinion-dependent susceptibility to persuasion into the opinion formation process[1]. This model is particularly suitable for the case of polar opinions, when people's shifting towards opinion extremes makes them less or more receptive to opinion change. For example, a person shifting towards the Republican ideology is harder to persuade to invert her or his political stance. Alternatively, in a society with strong social norms, people holding conservative opinions following the norms are harder to persuade, while extreme opinions are volatile. The main theoretical question here is *How will the user opinions evolve in a long term?* I have answered this question for the above defined non-linear model, having brought up the tools from non-smooth analysis, and used them together with min-max Lyapunov functions and the suitable version of LaSalle Invariance Principle. One qualitative finding of this analysis states that, as long as a user is at least to some extent susceptible to persuasion, that specific extent does not affect the long-term opinion adoption behavior of that user. Another finding is an analytical expression for the asymptotic user opinion distribution—dependent upon the structure of the network as well as the placement of stubborn agents in it, yet independent of the initial opinions of susceptible users.

Future Work: I plan to explore the design of sociologically plausible opinion dynamics models where social ties evolve in time based on the users' positions in the society. For example, a rising politician's opinion can gain weight even outside of his or her network neighborhood solely due to that person's growing social power. This idea has been explored in the past in the context of DeGroot–Friedkin model for the evolution of social influence networks proposed by the groups of Friedkin and Bullo, and recently revisited by the groups of Anderson and Başar. This model allows the social tie weights to occasionally[2] change based on the users' evolving centrality, but, since centrality is a global measure, the agents are assumed to observe the entire network.

I plan to design a non-linear DeGroot-type opinion formation model, where (1) *all social ties* in the network *continuously evolve* based upon the agents' social positions; and (2) the agents weight their social ties based on *locally available information*—the states and social positions of either immediate neighbors of those about whom one can learn by interacting with the neighbors. Such a model would be useful for applications to large-scale online social networks, which are inherently incomplete, and where users make decisions based on a small observed part of the network. The theoretical analysis of this model will combine the well-established dynamical system theory together with my recent theoretical analysis of how a social network's eigencentrality changes under link perturbation (the latter result is described in Sec. 3 as past work).

Additional future work ideas not covered in this statement include adaptation of Heider's structural balance theory to real-world incomplete online social networks, and the design of hierarchical opinion dynamics models.

Related Papers:

[3] **V. Amelkin**, F. Bullo, and A. K. Singh. Polar opinion dynamics in social networks. *IEEE Transactions on Automatic Control*, 62, 2017. [pdf],[doi],[poster],[slides].

3 Control of Network Processes

A deep embedding of network processes in the social and economic spheres of human life makes it important to have control over these processes. *How to prevent malicious interventions in the process of opinion formation in online social networks? How to ensure that the coalition formation strategies among potential collaborators ensure the best performance of a financial system?* Below, I discuss how I address these questions in my past, ongoing, and future work.

Past Work [4]: Since the process of opinion formation is inherently a *network process*, it is lucrative for marketing and political technologists to tap into online social networks and shift the public opinion distribution towards business-imposed objectives. One widely studied example is influence maximization in viral marketing,

[1] The origins of the idea that susceptibility to persuasion is a function of the held opinion can be traced back at least to Abelson's 1964 work "Mathematical models of the distribution of attitudes under controversy", and has been periodically revisited since then. However, a theoretically analyzed network model had been absent prior to [3].

[2] The non-linear update of the interpersonal trust happens at discrete time steps, between which the opinion dynamics is linear.

where the goal is to strategically influence the opinions of select—desirably, influential—users, so that they efficiently distribute these "right" opinions through the network. The society would, however, benefit from safeguarding the opinion formation process from such external influence. To that end, in my recent work [4], I propose a novel problem of disabling external influence upon the opinion distribution in a social network, as well as an efficient method to solve it. I assume that the specific target for the attack is the so-called asymptotic consensus value—the sum of user opinions weighted by these users' eigenvector centralities. This value is the asymptotic limit for user opinions in DeGroot model, or, more generally, a value to which the opinions of all network users are attracted. I assume that the adversary maximizes the asymptotic consensus value by altering the opinions of some users. I, then, state DIVER—an NP-hard problem of disabling such external influence attempts via strategically adding a limited number of edges to the network. Relying on the theory of Markov chains, I provide perturbation analysis that shows how eigenvector centrality and, hence, DIVER's objective function change in response to an edge's addition to the network. The latter leads to a pseudo-linear-time heuristic for DIVER, that relies on efficient estimation of mean first passage times in a Markov chain. The obtained theoretical and algorithmic results can also be applied in other fields dealing with Markov processes.

Future Work: In the past, I have empirically confirmed [5] that there is a strong positive correlation between the collective performance of a network of agents and the way these agents collaborate. In the future, I plan to explore two related pending questions: (1) *What is the exact dependency between agents' collaboration and performance? Is it bi-directional?* and (2) *How to incentivize collaboration between the agents to optimize their collective performance?* I will study these questions in the context of a dynamic collaboration network of stock traders in a day-trading firm. Besides the natural assumption that *collaboration affects performance*, this modeling effort relies on another major and less obvious assumption—*performance affects collaboration*. Thus, traders are expected to pragmatically collaborate with others who have either manifested high performance or the past collaboration with whom has been fruitful. This assumption has been supported by a preliminary Granger causality analysis of real-world trading data, that has shown a statistically significant causal relationship between agents' global reputation and the average number of "collaboration offers" they receive.

One crucial aspect of the model will be the balance between focusing on collaboration with top-earning traders, traders with whom an agent has successfully collaborated in the past, and the remaining "unexplored" agents. While pure exploration (random collaboration) is clearly inferior, putting too much emphasis on exploitation (attempting to collaborate only with the top-performing traders) will likely also be suboptimal, as it can prevent discovery of high-performing, but not yet clearly visible collaboration alliances.

I plan to design a model of performance-driven collaboration among stock traders validated against real-world data, perform its theoretical analysis, and use the obtained insights to design an incentive scheme for the collaboration decision making to maximize the collective performance of the trading firm.

Related Papers:

[4] **V. Amelkin** and A. K. Singh. Disabling external influence in social networks via edge recommendation. In *submission to International Conference on Data Mining (SDM'18)*. SIAM. [arxiv], [supplement].

[5] **V. Amelkin**, O. Askarisichani, Y. J. Kim, T. W. Malone, and A. K. Singh. Dynamics of collective performance in collaboration networks. *Submitted to PLOS ONE (May, 2017)*. [slides].

Summary

As the world is becoming more linked, the social and economic network processes will clearly remain in the spotlight. Rapid development of online social network and collaboration tools on one side, and theoretical techniques on the other side allows the studies of network processes to enter a qualitatively new stage. The network process research—conducted along the three facets of analysis, modeling, and control—involves design, theoretical analysis, and optimization of sociologically and economically plausible models, as well as the development of practical scalable algorithms to make these models work in real-world applications. This research agenda requires bridging domain-specific areas, such as social psychology and behavioral economics, as well as the fundamental fields of combinatorial algorithm design, machine learning, linear algebra, and dynamical systems theory. In my work, spanning all these areas, I have contributed to and brought closer the formation of the analysis, modeling, and control of dynamic processes in networks as an independent cross-discipline.