

The Power of Context in Networks: Ideal Point Models with Social Interactions*

Mohammad T. Irfan and Tucker Gordon

Department of Computer Science, Bowdoin College, Brunswick, Maine 04011
mirfan@bowdoin.edu, tucker.gordon18@gmail.com

Abstract

Game theory has been widely used for modeling strategic behaviors in networked multiagent systems. However, the context within which these strategic behaviors take place has received limited attention. We present a model of strategic behavior in networks that incorporates the behavioral context, focusing on the contextual aspects of congressional voting. One salient predictive model in political science is the *ideal point model*, which assigns each senator and each bill a number on the real line of political spectrum. We extend the classical ideal point model with network-structured interactions among senators. In contrast to the ideal point model’s prediction of individual voting behavior, we predict *joint voting behaviors* in a game-theoretic fashion. The consideration of context allows our model to outperform previous models that solely focus on the networked interactions with no contextual parameters. We focus on two fundamental problems: learning the model using real-world data and computing *stable outcomes* of the model with a view to predicting joint voting behaviors and identifying most influential senators. We demonstrate the effectiveness of our model through experiments using data from the 114th U.S. Congress.

1 Introduction

Over the past several decades, game theory has been widely used to model and analyze strategic interactions among multiple agents. In particular, games with network-structured interactions, also known as graphical games [Kearns, 2007], have received a lot of attention in AI due to their succinct representation size and their applications to many real-world scenarios. Although one of the hallmarks of game theory is its reliable encoding of stability as a solution concept, most strategic interactions are extremely complex in practice, where context plays a significant role. As a result, any realization of the game-theoretic solution concept of stability

*This is an extended abstract of the original paper published at the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018) in Stockholm, Sweden.

very much depends on our attention to behavioral context, or lack thereof.

In this paper, we consider a networked multiagent system where agents take actions through strategic interactions with each other. An agent’s action depends not only on others’ actions in a network-structured fashion but also on the underlying behavioral context. Context being a domain-specific construct, we focus on voting behavior in Congress. In congressional voting, the agenda or issue on which a vote is taking place gives the context to the voting behavior of the members. Our framework can also be applied to other settings where there are (1) network-structured strategic interactions and (2) issues (or context) on which agents adopt binary actions.

There have been multidisciplinary quests into modeling and predicting how legislators vote, reflecting a multitude of factors in play. Perhaps a bill lines up closely with a legislator’s political agenda, or is popular amongst their constituents. Perhaps they feel pressure from their party or other legislators to confirm a nomination or to pass a bill. Despite such pressure, senators sometimes vote against the party line. One recent example is the failure of the so-called “skinny-repeal of Obamacare” in the U.S. Senate, where Senator McCain (R-AZ) gave a dramatic, last-minute *nay* vote. There are other examples where senators voted against the party line due to the bill not lining up with their particular legislative positions (e.g., pro-choice Republicans or anti-gun-control Democrats). Therefore, a model of legislative voting behavior should take into account a senator’s legislative position, the characteristics of the bills, *and* the network of influence among the senators. We present the first such model here.

Introduced by Davis *et al.* [1970], one popular approach to modeling legislative behavior is the *ideal point model*. In this model, each senator i has an ideal point $p_i \in \mathbb{R}$, representing the senator’s political leaning based on how liberal or conservative they are. Each bill l is characterized by its *polarity* $a_l \in \mathbb{R}$ and its *popularity* $b_l \in \mathbb{R}$, analogous to *discrimination* and *difficulty* in item-response theory [Hambleton and Swaminathan, 2013]. Typically, p_i and a_l are negative for liberal positions and positive for conservative, and the b_l parameter corresponds to the fraction of senators voting *yea*. The probability of senator i voting *yea* on a bill l is given by the following logistic function: $p(x_{i,l} = \text{yea} \mid p_i, a_l, b_l) = \sigma(p_i a_l + b_l)$.

There are numerous applications and extensions of this

model [Jenkins, 1999; Schickler, 2000; Canes-Wrone *et al.*, 2002; Clinton *et al.*, 2004; Gerrish, 2013]. For the most part, these models consider each legislator’s vote as an individual decision with respect to the issue being voted on. All elements of social interactions are typically lumped into the single popularity parameter b_l . In the real world, however, the voting behaviors of legislators are often interdependent (e.g., deal-making). As a result, each voting outcome can be thought of as a joint behavioral outcome, which naturally falls in the territory of game theory.

Game-theoretic models, such as the *linear influence game* (LIG) model, have been proposed in the literature [Irfan and Ortiz, 2011; Irfan, 2013; Irfan and Ortiz, 2014]. In the LIG model, the actions of players (i.e., senators) influence those of others in a network-structured fashion, and every player acts in a way to maximize their payoffs with respect to others. For example, a Democratic senator may vote *yea* on a bill that many Republican senators oppose, and may vote *nay* if a close colleague opposes it.

While both the LIG and ideal point models are predictive models, each is limited in some capacity. The ideal point model is statistical and provides insight into the ideologies of individual legislators, but lacks any notion of strategic interactions. The LIG model is game-theoretic and captures the effects of social influence on votes, but does not take into account the characteristics of the issue being voted on. Furthermore, the statistical ideal point model gives individual predictions, while the the game-theoretic LIG model gives joint predictions. Previous studies to extend the ideal point model focused on the textual content of the bills [Gerrish, 2013]. Here, we extend the ideal point model in the direction of network-structured strategic interactions while considering the content of the bills via the subject tags. We are not aware of any study on ideal point model with social interactions.

2 Ideal Point Model With Social Interactions

We present a binary-action (+1 for *yea* vote and -1 for *nay*) graphical game [Kearns *et al.*, 2001] model with parametric payoff functions of quadratic forms. We use a weighted directed graph to represent the network among N agents (or senators). We use \mathcal{N}_i for the set of neighbors of node i .

Senator-Specific Parameters. Each node i is a senator and has two parameters: a threshold $t_i \in \mathbb{R}$ representing how “stubborn” Senator i is and an ideal point $p_i \in \mathbb{R}$ representing the senator’s ideology in the political spectrum. For each directed edge from j to i , $w_{i,j} \in \mathbb{R}$ denotes the influence weight that j has on i (this weight can be negative). The influence weights form a matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, with $\text{diag}(\mathbf{W}) = 0$.

Bill-Specific Parameters. Each bill l being voted on is parameterized by its polarity $a_l \in \mathbb{R}$ representing the bill’s position in the political spectrum.

Senators’ Choices of Actions. $x_i \in \{+1, -1\}$ denotes the vote of a senator i on a bill l .¹ Here, x_i depends on: (1) the votes of other senators and their influences on i , (2) the polarity a_l of the bill l , and (3) the threshold t_i . We combine

these three factors in the following *influence function*.

$$f_i(\mathbf{x}_{-i}, l) \equiv \sum_{j \in \mathcal{N}_i} w_{ij} x_j + (p_i \cdot a_l) - t_i \quad (1)$$

To give an intuition, the first term in (1) is collective, signifying how others’ voting decisions influence senator i ’s voting decision. In contrast, the second term is specific to an individual senator and a particular bill, capturing how well a senator’s ideal point matches the polarity of the bill in the political spectrum. The third term is purely an individualistic term modeling the level of “stubbornness” of a senator. Another way of viewing this influence function is to consider the b_l parameter of the classical ideal point model. It represents the fraction of senators voting *yea*, but does not give any further details. Our model breaks down b_l into a network-structured influence expression for each senator. The longer version of this paper provides a detailed discussion and justification [Irfan and Gordon, 2018].

Note that the senators’ choices of actions are interdependent on each other. Game theory naturally models such mutual dependencies. We next define some game-theoretic terminology in our context.

A vector of actions, one action per player (or senator), is called a *joint action*. A joint action \mathbf{x}^* is a *pure-strategy Nash equilibrium* (PSNE) if every player simultaneously plays their *best response* x_i^* w.r.t. others’ actions (i.e., votes “optimally”). We use PSNE to mathematically represent the *stable outcomes* of our model. Following is Senator i ’s (quadratic) *payoff function*, which is maximized by i ’s best response.

$$u_i(x_i, \mathbf{x}_{-i}, l) \equiv x_i f_i(\mathbf{x}_{-i}, l) = x_i (\mathbf{w}_{i,-i}^T \mathbf{x}_{-i} + p_i \cdot a_l - t_i).$$

To maximize $u_i(x_i, \mathbf{x}_{-i}, l)$, Senator i ’s best response x_i^* will be +1 if $f_i(\mathbf{x}_{-i}, l) \geq 0$. It will be -1 if $f_i(\mathbf{x}_{-i}, l) \leq 0$. An equality indicates indifference between +1 and -1.

This completes the description of our model. As a historical backdrop, our influence function (1) is grounded in mathematical sociology as a threshold model of collective behavior [Granovetter, 1978]. Particularly within the CS community, Kempe *et al.* [2003]’s study of threshold models with the goal of identifying the most influential nodes in a network is the precursor to many studies on influence maximization [Even-Dar and Shapira, 2007; Kleinberg, 2007; Mossel and Roch, 2010; Chen *et al.*, 2009; Chen *et al.*, 2010]. In most threshold models, however, the influence weights are non-negative. Here, we do allow negative weights. Most importantly, in almost all of the previous threshold models, the context (e.g., the bills in our case) under which players choose actions has largely been ignored. Here, we incorporate context within threshold models. One subtle but crucial aspect of our model is that it moves away from the well-studied influence maximization setting to a strictly game-theoretic setting.

3 Learning the Model

As a synopsis, we wish to learn our game-theoretic model so that it captures as much of the observed roll call as possible in the form of PSNE. We can trivially capture *all* observed roll call data as PSNE by choosing an empty model, which is not desirable. Therefore, we want to capture as many of the roll

¹We use x_i instead of $x_{i,l}$, since l will be clear from context.

call instances as possible *while having as few PSNE as possible*. On the other spectrum, it may very well be impossible to capture *exactly* the observed data as PSNE. A natural way of approaching the middle ground between these two extremes is to use a generative mixture model [Murphy, 2012] that generates the observed data by picking a PSNE with probability q and a non-PSNE with probability $1 - q$. Obviously, we would like q to be as high as possible while avoiding overfitting. Honorio and Ortiz [2015] use this approach to learn threshold models from behavioral data, which we adapt here.

Let $\mathcal{NE}(\mathcal{G})$ be the set of PSNE of our game model \mathcal{G} . As noted above, we want to learn \mathcal{G} to maximize the “quality” of $\mathcal{NE}(\mathcal{G})$. Let $\pi(\mathcal{G})$ be the size of $\mathcal{NE}(\mathcal{G})$ relative to the total number of possible outcomes: $\pi(\mathcal{G}) \equiv |\mathcal{NE}(\mathcal{G})|/2^N$. If M is the number of observed bills, $q \cdot M$ is the number of bills that the generative model picks from the set of PSNE. In order to maximize the quality of our model’s $\mathcal{NE}(\mathcal{G})$, we want to maximize $q \cdot M$ relative to $\pi(\mathcal{G})$. This can be written as $\frac{q \cdot M}{|\mathcal{NE}(\mathcal{G})|/2^N}$. Since N and M are fixed, this expression is proportional to $\log_{10} \frac{q}{|\mathcal{NE}(\mathcal{G})|}$, which allows us to measure the quality of $\mathcal{NE}(\mathcal{G})$ while taking into account its size.

While we leave details to the longer version of the paper [Irfan and Gordon, 2018], it can be shown that the optimization task boils down to:

$$\min_{\mathbf{W}, \mathbf{t}, \mathbf{p}, \mathbf{a}} \frac{1}{M} \sum_l \log \left[1 + \sum_i e^{-x_i^{(l)} (\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} + (p_i \cdot a_l)) - t_i} \right] + \rho \|\mathbf{W}\|_1 + \rho' \|\mathbf{p} \cdot \mathbf{a}\|_1. \quad (2)$$

Here, we use an ℓ_1 -norm regularizer $\rho \|\mathbf{W}\|_1$ in order to bias the learning to sparse networks and to reduce overfitting. We add another ℓ_1 -norm regularizer $\rho' \|\mathbf{p} \cdot \mathbf{a}\|_1$ to penalize the ideal points and polarities. We use the ℓ_1 - projection method [Schmidt *et al.*, 2007] for optimization.

3.1 Issues with Learning

Anchoring Ideal Points. Traditionally, Democrats have negative ideal points and Republicans have positive ideal points. To avoid getting the opposite signs, we set anchors, which is common in the literature [Gerrish, 2013]. We used ideal points generated for the 114th Senate² to find the most extreme senators in each party for the 114th Senate. This meant giving Warren (D-MA) and Sanders (I-VT) ideal points of -4 , and Cruz (R-TX) and Shelby (R-AL) $+4$.

Determining Validation Bill Polarities. We cross-validated our model parameters using held-out roll call data from the same Congress. However, when our model is given a new bill l' , its polarity $a_{l'}$ is unknown and has not been learned as l' was not part of the training set. We solved this problem using the subject codes provided by Congress.gov. These hand-entered codes are terms that describe the bills’ contents (e.g. “Child safety and welfare”, “Drug Enforcement Administration”, or “Health care coverage and access”) and are known to be a good approximation to topic modeling [Gerrish, 2013]. Each bill usually contains multiple codes, with the median bill in the 114th Senate

²<https://voteviewblog.com/2015/08/16/alpha-nominate-applied-to-the-114th-senate>

containing 60 subject codes out of a total of 737. For a new bill l' , we approximate its polarity by finding the closest bill (in terms of Euclidean distance) among the bills already seen by the model.

3.2 Model Selection

The interplay between the regularization parameters ρ and ρ' can lead to a mixture of an ideal point model and a social interaction model, or cause the model to strongly shift towards one or the other. Ideally, we want ρ and ρ' to produce a network that is sparse enough to avoid overfitting but complex enough to represent the observed data while maintaining reasonable ideal points. We trained the model using Session I voting data and rigorously performed cross-validation using Session II data. We tracked four different measures: the number of edges, the proportion of training votes that are PSNE, errors in best responses, and the Euclidean distance between the learned and the average ideal points. Leaving all details, including plots, to the longer paper [Irfan and Gordon, 2018], we settled on $\rho = 0.00225$ and $\rho' = 0.0004$. This produced a graph of 998 edges, and a validation error of 16.0%.

4 Equilibria Computation

Our model gives joint predictions of senators’ voting behaviors in the form of PSNE. As a result, the learned model is not immediately useful without computing the PSNE. Using complexity results on LIGs, we show that this problem is provably hard.

Theorem 1. *It is NP-complete to decide whether there exists a PSNE; whether there exists a PSNE consistent with a designated set of senators voting yea; and whether there exists a PSNE with at least a given number of k senators voting yea.*

Theorem 2. *It is #P-complete to count the number of PSNE, even when the underlying network is bipartite or star.*

To compute *all* PSNE of our model, we applied a backtracking search algorithm [Irfan and Ortiz, 2014] that systematically explores the search space of 2^{100} due to 100 senators and binary actions. Whenever a node is assigned an action, we use the **NashProp** algorithm [Ortiz and Kearns, 2003; Irfan and Ortiz, 2014] to propagate its action locally and see if it leads to any contradiction. Perhaps not surprisingly, non-zero values of the contextual parameters \mathbf{p} and \mathbf{a} significantly improve the runtime compared to when either of them is 0. This is because the backtracking search finds contradictions more quickly when \mathbf{p} and \mathbf{a} are non-zero.

When run on a high performance computing grid with $a_l = 0$, the algorithm took 9 minutes to compute all PSNE in the 100-player game. When $a_l \approx -3.22$ (one of the largest negative polarities), it took only 0.49 seconds.

5 Experiments: The 114th Senate

We ran our model using roll call and bill-specific data from the 114th U.S. Senate (2015-2017). We considered all voting instances that had an associated bill (including amendments), and ignored all other types of votes (e.g. confirmations).

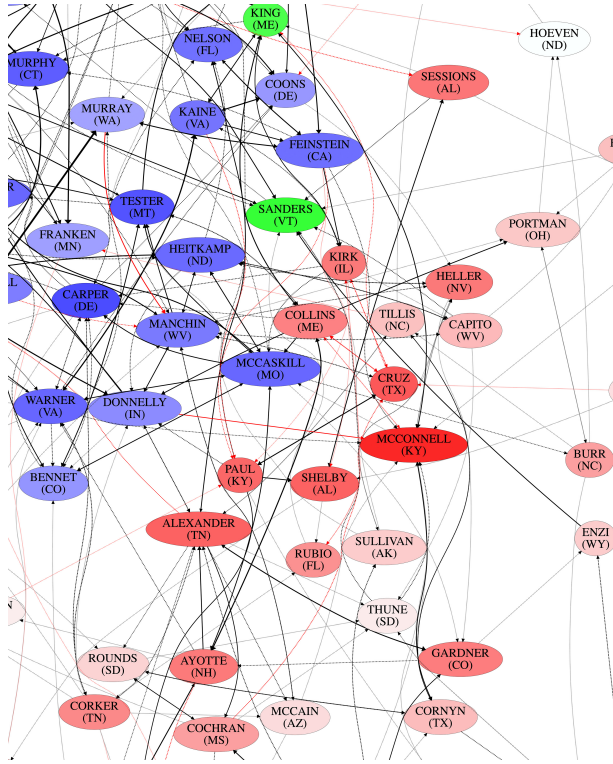


Figure 1: A partial view of the learned influence network. Only the strongest 30% incoming and outgoing edges for each node have been selected for visualization. Each node represents a senator, with colors to indicate party and shades to indicate threshold level t_i (darker means higher threshold). Edges represent one’s influence on another, with black edges representing positive influence and red edges negative; thickness represents magnitude of influence.

Learned Ideal Points. With very few exceptions, our model assigns negative ideal points to Democrats and positive ideal points to Republicans. This is an encouraging result as it demonstrates our model’s ability to capture ideal point data in a manner similar to traditional ideal point models. The average ideal point for Democrats and Independents is -1.49 with a standard deviation of 0.76 , and for Republicans 2.47 with a standard deviation of 0.79 . Details are left to the longer paper [Irfan and Gordon, 2018].

Influence Network. Fig. 1 gives a visualization of the learned network. Note that the structure of this network will not change for different bills.

Comparison with Ideal Point Models. Clinton *et al.* [2004] report an accuracy of 89.9% on individual votes in the 106th U.S. House of Representatives. However, that is training accuracy without any held-out data. The challenge with measuring test accuracy with held-out data is that the polarity a_l of a bill l in the held-out data is unknown. We describe how we mitigate this challenge using the subject codes of bills in Section 3.1. Even if we apply the same procedure to the classical ideal point model, it will perform poorly with respect to our desirable quality measure $\log_{10} \frac{q}{|\mathcal{N}\mathcal{E}(\mathcal{G})|}$, because it is designed to give *individual* predictions, not *joint* predictions.

More specifically, the network structure is basically empty in the classical ideal point model, which would lead to a staggering number of PSNE (computationally infeasible). This would exponentially bring down the above quality measure.

Comparison with LIG Model. We compare and contrast our model with the linear influence game (LIG) model [Irfan and Ortiz, 2011; Irfan, 2013; Irfan and Ortiz, 2014], which is the state-of-the-art game-theoretic model on influence in Congress to our knowledge. Since the LIG model does not have any notion of bill context while ours does, in order to compare ours with LIG we look into three distinct types of bills in the political spectrum: bills that are very much left-leaning (very negative polarity), bills that are in the middle (polarity close to 0), and bills that are very much right-leaning (very positive polarity). Overall, our model produces a much narrower band of PSNE as its stable-outcome prediction than the LIG model. In particular, the prediction of our model is remarkably precise and high-quality if a bill falls in either end of the political spectrum. Details on the quality of PSNE can be found in the longer paper [Irfan and Gordon, 2018].

Case Studies. For a right-leaning bill we picked one of the most controversial bills of the 114th Senate, the Keystone XL Pipeline bill [Erickson and Lazarus, 2014], for which $a_l = 1.426$. The bill instigated protests from environmental groups and the public and ultimately passed by a 62–36 vote. Our model produced one single PSNE for this bill, which correctly predicted 91 votes. In contrast, the LIG model [Irfan and Ortiz, 2014] produced 287,400 possible PSNE, where the median number of correctly predicted votes was 50. We obtained similar results for other types of bills (left leaning and centrist). Details, including figures, are left to the longer paper [Irfan and Gordon, 2018].

Most Influential Senators. It is natural to ask who the most influential senators are for passing a bill. The implication is that when the most influential senators group up and vote *yea* on a bill, then everybody else would be influenced to vote *yea* (note the lack of diffusion or dynamics and the contrast with influence maximization). Given the set of all PSNE, we can compute a small set of most influential senators by applying the well-known greedy set-cover algorithm with a provable log-factor approximation guarantee, as applied to LIGs [Irfan and Ortiz, 2014]. In our model, a set of most influential senators consists of only four senators: {Carper (D-DE), Schumer (D-NY), Cruz (R, TX), Peters (D, MI)}. In contrast, due to its wide band of PSNE, the LIG model gives a much bigger set of nine most influential senators. This not only highlights an improvement over LIGs but also makes our framework applicable to the minimal targeted intervention problem where we are interested in targeting as few individuals as possible to achieve a desirable social outcome.

Concluding Remarks

We have presented a model of strategic behavior grounded in threshold models. Unlike the existing literature on threshold models, ours takes into account the context under which the behaviors take place. This leads to significant improvements over previous models, both in terms of the number of PSNE as well as the quality of these PSNE.

References

- [Canes-Wrone *et al.*, 2002] Brandice Canes-Wrone, David W Brady, and John F Cogan. Out of step, out of office: Electoral accountability and house members' voting. *American Political Science Review*, 96(01):127–140, 2002.
- [Chen *et al.*, 2009] Wei Chen, Yajun Wang, and Siyu Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 199–208. ACM, 2009.
- [Chen *et al.*, 2010] Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In *2010 IEEE international conference on data mining*, pages 88–97. IEEE, 2010.
- [Clinton *et al.*, 2004] Joshua Clinton, Simon Jackman, and Douglas Rivers. The statistical analysis of roll call data. *American Political Science Review*, 98(02):355–370, 2004.
- [Davis *et al.*, 1970] Otto A. Davis, Melvin J. Hinich, and Peter C. Ordeshook. An expository development of a mathematical model of the electoral process. *American Political Science Review*, 64(2):426–448, 1970.
- [Erickson and Lazarus, 2014] Peter Erickson and Michael Lazarus. Impact of the keystone xl pipeline on global oil markets and greenhouse gas emissions. *Nature Climate Change*, 4(9):778–781, 2014.
- [Even-Dar and Shapira, 2007] Eyal Even-Dar and Asaf Shapira. A note on maximizing the spread of influence in social networks. In *Internet and Network Economics (WINE)*, 2007.
- [Gerrish, 2013] Sean M Gerrish. Applications of latent variable models in modeling influence and decision making. Technical report, Princeton University, 2013.
- [Granovetter, 1978] Mark Granovetter. Threshold models of collective behavior. *The American Journal of Sociology*, 83(6):1420–1443, 1978.
- [Hambleton and Swaminathan, 2013] Ronald K Hambleton and Hariharan Swaminathan. *Item response theory: Principles and applications*. Springer Science & Business Media, 2013.
- [Honorio and Ortiz, 2015] Jean Honorio and Luis E Ortiz. Learning the structure and parameters of large-population graphical games from behavioral data. *Journal of Machine Learning Research*, 16:1157–1210, 2015.
- [Irfan and Gordon, 2018] Mohammad T Irfan and Tucker Gordon. The power of context in networks: Ideal point models with social interactions. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 910–918. International Foundation for Autonomous Agents and Multiagent Systems, 2018.
- [Irfan and Ortiz, 2011] Mohammad T Irfan and Luis E Ortiz. A game-theoretic approach to influence in networks. In *AAAI Conference on Artificial Intelligence*, 2011.
- [Irfan and Ortiz, 2014] Mohammad T Irfan and Luis E Ortiz. On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach. *Artificial Intelligence*, 215:79–119, 2014.
- [Irfan, 2013] Mohammad Tanvir Irfan. *Causal Strategic Inference in Social and Economic Networks*. PhD thesis, Stony Brook University, Department of Computer Science, August 2013. http://www.bowdoin.edu/~mirfan/papers/Mohammad_Tanvir_Irfan_Dissertation.pdf.
- [Jenkins, 1999] Jeffery A Jenkins. Examining the bonding effects of party: A comparative analysis of roll-call voting in the us and confederate houses. *American Journal of Political Science*, pages 1144–1165, 1999.
- [Kearns *et al.*, 2001] Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. Graphical models for game theory. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pages 253–260, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- [Kearns, 2007] Michael Kearns. Graphical games. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 7, pages 159–180. Cambridge University Press, 2007.
- [Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003.
- [Kleinberg, 2007] Jon Kleinberg. Cascading behavior in networks: Algorithmic and economic issues. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 24, pages 613–632. Cambridge University Press, 2007.
- [Mossel and Roch, 2010] Elchanan Mossel and Sebastien Roch. Submodularity of influence in social networks: From local to global. *SIAM Journal on Computing*, 39(6):2176–2188, 2010.
- [Murphy, 2012] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [Ortiz and Kearns, 2003] Luis E. Ortiz and Michael Kearns. Nash propagation for loopy graphical games. In Suzanna Becker Becker, Sebastian Thrun Thrun, and Klaus Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 817–824, 2003.
- [Schickler, 2000] Eric Schickler. Institutional change in the house of representatives, 1867–1998: a test of partisan and ideological power balance models. *American Political Science Review*, 94(02):269–288, 2000.
- [Schmidt *et al.*, 2007] Mark Schmidt, Glenn Fung, and Romer Rosales. Fast optimization methods for l1 regularization: A comparative study and two new approaches. *Proceedings of the 18th European Conference on Machine Learning, ECML '07*, pages 286–297, 2007.