Opinion Dynamics in Social Networks

A Thesis
Submitted to the Faculty
in partial fulfillment of the requirements for the
degree of

Doctor of Philosophy

by

Qi Gu

Thayer School of Engineering
Dartmouth College
Hanover, New Hampshire

April 2014

Examining Committee:

Chairman ____________________
Eugene Santos, Jr.

Member ____________________
George Cybenko

Member ____________________
Mark Borsuk

Member ____________________
Hien Nguyen

____
F. Jon Kull
Dean of Graduate Studies

ABSTRACT

Opinion dynamics is a complex procedure that entails a cognitive process when it deals with how a person integrates influential opinions to form revised opinion. Early research on opinion formation and social influence can be traced back to the eighteenth century. The original research focus was to study the conditions for people to aggregate information and reach consensus. Recently, due to the rise of the World Wide Web more and more studies tend to model opinion dynamics in large-scale social networks via computational methods. Among those works, non-Bayesian rule-of-thumb learning models keep gaining popularity due to their simplicity and computational efficiency. Unlike many non-Bayesian methods that treat individual opinions on various issues as independent beliefs but overlook the connections between knowledge fragments, we leverage from Bayesian approaches to consider opinions as a product inferred from one's knowledge-based system, where new knowledge fragments are acquired through social interaction and learning experiences. We study how an individual evaluates and adopts such knowledge fragments from others sources, both visible and invisible, on the basis of the findings from well-established social theories.

A computational framework was developed to model opinion dynamics, in which we applied a probabilistic model named Bayesian Knowledge Bases to represent an individual's knowledge base. Opinion dynamics is studied by modeling opinion formation as a process of knowledge fusion, learning the impact metric that estimates the reliability of knowledge fragments, and identifying influential sources whose impact patterns are hidden. The contributions of this work can be summarized as 1) the development of a domain-independent computational method to model opinion formation by emphasizing the dependencies between knowledge pieces, 2) the capability to model different aspects of opinion dynamics in one entire system, 3) the intuitiveness in representing opinions such that the intents behind the opinion change can be readily captured, 4) the ability to characterize the influences in a social community by realizing and enriching theories of social commu-

nication, and 5) the flexibility of application on detecting and tracking hidden influential sources.

Acknowledgements

First, I would like to thank my advisor Prof. Eugene Santos, Jr. for providing me with invaluable mentorship, guidance, and support. More than anyone else he has influenced my view of the research process, and instilled in me the importance of aiming to produce high-quality research with the potential for impact. He fostered my interest in Opinion Dynamics and provided me an inspiring environment with all the support, freedom, patience, and encouragement I needed to achieve the work presented in this dissertation.

I would like to thank all members from the Distributed Information and Intelligence Analysis Group. I count myself lucky to have been surrounded by a number of outstanding researchers: Deqing Li, Fei Yu, Jacob Russell, Jeremy E. Thompson, Keumjoo Kim, et al.. In particular, I offer a special thank to Deqing, Fei, Jeremy and Jacob for their time and support in different stages of my PhD, especially when I reached bottlenecks in the research. Over the years we have shared countless hours working together in the lab. They are not only good colleagues, but good friends. I would also like to thank Ms. Louise A. Cullen for helping me carefully proofread the thesis.

I feel fortunate to have had the opportunities to work closely with Prof. John Korah and Prof. Hien Nguyen in our collaborating projects Cross-boarder and Metalearning. From them, I have learned many research skills from writing to presentation to designing an experiment. I would also like to thank my committee members Prof. George Cybenko, who taught me modern modeling technologies through his classes, Prof. Mark Borsuk, who gave me lectures and instructions on Bayesian approaches and modeling, and Prof. Hien Nguyen, who provided me with valuable comments and advice on my thesis. I thoroughly enjoyed the opportunities to collaborate with these exceptional researchers.

Most importantly, I am definitely lucky to have the support from my family. I especially thank my parents in China for the high standards they raised me with. Living on the other side of the Earth complicates many things for a family and they have always stood by me. I also would like to thank my boyfriend Mu Lin for his love and support irrespective of the unpredictable nature of my work schedule and the uncertainty it brought. I do not think I could have finished all this without them.

Contents

1 Introduction **1**
1.1 A Framework for modeling opinion dynamics 6
1.1.1 Motivations 6
1.1.2 Framework 7
1.2 Literature review of opinion dynamics modeling using computational methods 12
1.3 Contributions 18
1.4 Outlines 19
1.5 List of Acronyms 19

2 Backgrounds **21**
2.1 Social Influence 21
2.1.1 Relationship 23
2.1.2 Contact Frequency 23
2.1.3 Opinion Similarity 23
2.2 Bayesian Knowledge Bases 24
2.2.1 Framework of Bayesian Knowledge Bases 26
2.2.2 BKB Fusion 29
2.2.3 BKB Tuning 29
2.3 Social Network 33

3 Opinion Formation **36**
3.1 Opinion as a Reasoning Product 36
3.2 Computational Approaches for Opinion Representation 37
3.3 Model of Opinion Formation 41
3.3.1 Mathematical Formulation 42
3.3.2 Discussion 44
3.4 Nonlinearity in Opinion Formation 45

4 Social Influence Learning **51**
4.1 General Impact Metric 52
4.2 Parameter Estimation 54
4.3 Validation and Evaluation 55
4.3.1 Baseline 56
4.3.2 Experiments on Synthetic Data 57

4.3.3 Experiments on Real Social-network Data . . . 64
4.4 Summary . . . 72

5 Hidden Influential Source Tracking and Detection **74**
5.1 Overview . . . 74
5.2 Source Tracking and Detection Approach . . . 78
5.2.1 Mathematical Formulation . . . 78
5.2.2 Parameter Estimation . . . 79
5.3 Experiments . . . 83
5.3.1 Evaluation on Synthetic Data . . . 84
5.3.2 Evaluation on Real Data . . . 100

6 Conclusions **114**
6.1 Accomplishment and Findings . . . 115
6.2 Limitations and Future Work . . . 117
6.2.1 Limitations . . . 117
6.2.2 Future works . . . 118

7 Appendices **121**
7.1 Appendix A . . . 121
7.2 Appendix B . . . 121
7.3 Appendix C . . . 124
7.4 Appendix D . . . 125

List of Tables

3.1 Correlation with significance 47
3.2 Percentage of positively correlated interested trends 50

4.1 Error comparison of two approaches in terms of different parameters 59
4.2 Average similarity with the true factor weights 61
4.3 Selected Political Survey Questions and Possible Answers 67
4.4 Predefined probability values based on the survey answers 68
4.5 Average error comparison between two models 70
4.6 Factor weights learned from two methods 70
4.7 Average error comparison between two contact networks 71
4.8 Factor weights learned by the proposed approach from two communication means 71
4.9 Performance comparison with different initial settings 71
4.10 Average Error using Different factor sets (F = Friendship, C= Contact, OD = Opinion Difference) 72

5.1 Statistics of the distance between the true and learned probability distribution. . . . 91
5.2 p-value of paired sample t-test using different algorithms 91
5.3 Learning error of FFM v.s. two baselines with different diversity rate d 93
5.4 Learning error differences of FFM vs two baselines with different complexity rate r. 94
5.5 Learning error differences of FFM vs two baselines with different completeness rate q. 95
5.7 Two-level factorial design 97
5.8 Statistics of the difference between each paired model 106
5.9 Major events happened during H1N1 pandemic. 111

7.1 Distance measure between the true and learned probability distribution using different algorithm 124

List of Figures

1.1 Prior opinion formation and evolution 9
1.2 Opinion evolution in a social environment 11
1.3 Complete framework of opinion dynamics 13

2.1 Example of a BKB fragment 25
2.2 Example of BKB fusion 30
2.3 Example of a small contact network 35

3.1 Sample BKB fragment from an intent framework modeling the real scenario 39
3.2 Comparison of opinion formation process 40
3.3 Partial residuals plot 49

4.1 Social influence learning from Gu et al. [1] 55
4.2 Comparison of two methods with respect to different parameters 60
4.3 Mean RSE difference between two models 63
4.4 Convergence rate comparison between two methods given different connection degree. 66
4.5 Learned causal dependencies between four political issues 69

5.1 Multiple-source detection algorithm 83
5.2 Comparison of impact tracking for FFM and two baselines with increasing T 88
5.3 Cross-correlation analysis with T = 10,50 and 100 90
5.4 The aggregated performance for tracking multiple-hidden-source impact trend 98
5.5 Performances for new hidden source detection 99
5.6 Cross-correlation on dataset 1: Adults 104
5.7 Cross-correlation on dataset 2: CoverType 105
5.8 Cross-correlation on dataset 3: Color-based segmentation 105
5.9 Example BKB fragment with belief and action 107
5.10 Cultural fragment for middle-aged population 109
5.11 Event fragment 110
5.12 Posterior probability of illegal migration to US 111
5.13 Detection results on event impact 113

7.1 BKB fusion from Santos et al.[2] 122

Chapter 1

Introduction

Opinions are subjective perspectives that evolve over time. Unlike phenotypic characteristics that are mainly decided by our genes, opinions or beliefs[1] are generally acquired through learning experience [4]. Some of this learning results from the observation via one's own actions and experiences. Much of the learning, however, takes place through the process of social interaction. As summarized by Acemoglu and Ozdaglar (2011) [5],

> *Learning is social because any given individual observes the behavior of or receives information through communication with a small subset of society –those we may want refer to as her social network, consisting of her friends, her coworkers and peers, her distant and close family members, and a certain group of leaders that she listens to and respects.*

In fact, the process of interpersonal influence that affects individuals' beliefs and opinions lies at the heart of socialization, identity and decision-making. Such process also serves as a guide to how people change their behaviors by taking into account the attitudes and opinions of others with whom they interact [6]. Evidence of reliance on others' opin-

[1] From Schwitzgebel [3], belief is the fact or proposition that an individual holds to be true. Opinions, in contrast, include both personal beliefs and attitudes or judgments that are not founded on proof or certainty. In our work, we assume that all opinions are supported by some facts or beliefs.

ions in one's decision can be found in many fields, from advertising with social cues [7] [8] to the diffusion of political views [9]. The results from Bindel et al.[10] show that when the influence from a social leader is strong enough, people may shift their opinions far from their internal opinions.

Overall, the importance of how opinions are formed and changed via social learning can never be overstated. Through the last few centuries the issue of whether a group of people with diverse information will be able to aggregate information and reach a consensus has been the focus of a great many efforts. In eighteenth century, Condorcet [11] observed that a large group of people with each holding some correlated information is sufficient for aggregation of information when everyone is providing true information. The work of Galton [12] discovered that when a group of relatively uninformed individuals average their knowledge they could guess the weight of an ox fairly accurately. An active line of recent work in economic theory studied the formation of consensus from the point of view of social networks. Friedkin and Johnsen [6] concluded that because of the changing social structure the outcomes of the groups could vary from the mean of the initial opinion of a group member. Golub and Jackson [13] suggested that opinions in a large society might only converge to the true values when the influence of the most influential agent diminishes as the society grows. Considering the fact that consensus in the real world is usually not reached [14], Hegselmann and Krause [15] extended the previous work by considering some non-consensus cases. Specifically, they investigated several models to understand when opinion formation will lead to consensus, polarization or fragmentation.

Despite these significant findings, to fully understand opinion dynamics through social processes is never a trivial job. Individual knowledge acquired via various channels is an essential ingredient in decision-making and opinion formation. It provides the foundation for predicting the outcomes of alternative courses of action [16]. However, not everyone holds the same base of knowledge. It is often the case that when information is hard to process or access due to incompleteness, one would shape their opinions by interacting with others who hold views on the given issues [17]. A simple illustration of this is that it is not hard for an individual to have opinions about the places they have not been to or the food they have not tasted, as decisions can be simply made based on the information they gained from others. On the other hand, other than self-directed learning where the knowledge is extended by the results of experimentation and systematically monitoring [18], a large portion of knowledge is gained in the process of social activities [19]. For example, a person may decide not to buy a Japanese car after being told by her friend that a political issue between Japan and her country have become severely strained, even though their conversation had nothing to do with car-buying. Even if two people do not exchange views on a specific problem, due to the causal relationships between knowledge fragments, the newly acquired information from the conversation could influence their opinions toward that problem. Lastly, the art of understanding personal opinions rests not only on the measurement of people's views but also on understanding the motivations behind those views. When people go to the polls, why do some of them vote for one candidate over another? Is it possible to change their views and by how much? Those questions cannot be simply answered by looking at the statistics of participants' responses. That is why the methods based on the traditional opinion surveys are insufficient as they fail to capture the motivations behind them [20].

Before the age of information the problems of opinion dynamics studied by conventional methods usually considered a small group of people due to limitations in communication. With the development of the Internet people now can exchange information and ideas more freely and frequently. The scale of community, therefore, can be extremely large such that the process of opinion dynamics can hardly be understood without a computational model. There are some significant and insightful studies trying to apply computational approaches to model opinion dynamics. One direction of these works focuses on discovering conditions and criterions for opinion formation [13] [21], while another direction lies on inferring the influence patterns from contact data [22] [23]. In recent years, the fashionable term of "Big Data" keeps gaining attention. Domain experts apply statistical methods to discover patterns in data. Nevertheless, how useful is the "Big Data" is a controversial topic. Its value comes from social patterns that can be derived by connecting pieces of data. If not used properly, however, Big data interpretation can be problematic as well. Aside from its ethical questions concerning the access of personal data, the dataset itself is not necessary to be objective nor representative. As Bollier [24] explains,

> *As a large mass of raw information, Big Data is not self-explanatory. And yet the specific methodologies for interpreting the data are open to all sorts of philosophical debate. Can the data represent an 'objective truth' or is any interpretation necessarily biased by some subjective filter or the way that data is 'cleaned'?*

Despite the success of FiveThirtyEight[2] that correctly predicted the winner of the 2012 US Presidential race in forty-nine of fifty states, even Silver himself admits that the data-driven prediction models are so highly domain dependent that the model, while performing well in solving one problem, may fail dramatically in another [25]. In fact, it is very dangerous to simply rely on a dataset without those limitations and bias being understood. Therefore we want to focus on the philosophy behind the problem, i.e. how an opinion is derived and how it is influenced by other sources.

To have a model address all of the above concerns is extremely challenging. However, we believe that there are some fundamental rules that guide the whole process of opinion dynamics. The overall objective of this thesis is to develop a framework to model opinion

[2] FiveThirtyEight is a polling aggregation website with a blog created by Nate Silver

dynamics based on the idea that the procedure of opinion formation involves a cognitive process when dealing with how a person integrates influential opinions to form a revised opinion. The whole framework can be broken down to three parts. Firstly, we present a computational opinion formation model (OFM) that quantifies individual's opinions in the process of knowledge fusion. The underlying hypothesis is that an opinion can be treated as a product inferred from one's knowledge base, where the knowledge base keeps growing and updating through self-directed learning or social interaction. The intuition is that there are connections between knowledge fragments. Modification of opinion on one issue can influence the decision made on another. Though the whole picture of how opinions are changed is still missing, we are affirmative that "influencing means" and "influencing source" are two main ingredients in the process of opinion dynamics. Moreover, individuals always selectively accept or disregard new information, where the amount of acceptance is decided by some mechanism(s) based on social influence. In the second part of our framework, we identify certain social factors that can potentially affect how individuals adopt others' ideas. We then show how to learn the "influencing means" on the basis of these factors from OFM. In the third part, we extend OFM to address the situation when influential sources are "invisible". In particular, we provide an approach that is capable of detecting hidden sources and capturing the evolution of their impact levels.

Evaluation of the framework is another challenging task due to the scope of the problem and the lack of ground truths in opinion formation. Alternatively, we evaluate the accuracy of the opinion formation model with respect to its learning capability of social factors and the robustness of capturing influencing sources with various inputs.

1.1 A Framework for modeling opinion dynamics

1.1.1 Motivations

As mentioned earlier, opinion formation is in essence a learning process. Over the past few decades prevalent computational models have been focusing on how opinion evolves in a social environment, in which the opinion is either simplified to an abstract number or studied with respect to a specific scenario. Does this imply that the conversation or experience regarding one particular issue has no impact on one's opinion on other issues? The answer should be no. Let us consider a scenario, say person A was thinking about going cycling in the afternoon. He consulted another person, X. We are not sure who X is, but we know two potential candidates: person B and person C. We also know that if A talked to B, then B would say "the weather forecast says it is going to rain in the afternoon." and if A talked to C, C would say "Biking is bad for your knees.". Regardless, A changed his mind and stayed at home. So whom did A talk to? If we parse the words from B and C literally, apparently C is more close to the topic "biking", as B is about another issue "weather". However, from the perspective of knowledge, the question of whether to do an outdoor activity is highly dependent on weather. This suggests that both B and C can be X. What if we know that C got his information from an unauthenticated website, which makes C's impact much smaller? What if we know that the reason person A stayed at home is because he knew there would be a football game? If this is the correct scenario, maybe there existed another source of information? In tho example, we can keep adding information and the answer can keep changing. As one can see from this scenario, opinion dynamics is actually a complex procedure that entails a cognitive process when dealing with how a person integrates influential opinions to form a revised opinion. We cannot separate an individual's opinion from his/her knowledge base [26]; we cannot overlook the metrics that quantify the influences between people [27], and we have to be aware of the hidden sources that may have individual's opinion influenced subconsciously [28].

All these findings motivate us to develop a comprehensive computational framework to overcome the following limitations of existing computational efforts:

1. Key components in opinion dynamics are modeled separately. As a consequence, the approach designed to address one aspect of the problem may not be applicable to another.

2. The cognitive process has not been well adopted to model opinion dynamics computationally. So most of the current works are either domain specific or fail to explain the findings from the well-established social theories.

The insight of this thesis is that we build our framework on the basis of the cognitive process involved in shaping individual opinions and learning nature of opinion dynamics: we treat opinion as a reasoning product from one's context knowledge base where new knowledge fragments are acquired through various types of learning experience. By doing so, we manage to combine multiple components in opinion dynamics to a coherent system, such that no matter what learning technique a component applies, they all share the same core representation.

1.1.2 Framework

Building a model of opinion dynamics that is consistent with existing social theories and capable of handling large problem scalability is challenging. However, we can always simplify and extract the core parts of a hard problem and tackle it using a divide-and-conquer strategy. We divide the whole framework into three parts. Each part corresponds to one key component in opinion formation proposed by Acemoglu and Ozdaglar [5],

- Prior opinions
- Method of information processing
- Information sources

Prior Opinion

Opinion dynamics is a continuous procedure featured by a sequence of varying opinions. Any method modeling opinion dynamics has to start with some type of initial opinions that can be personal or public. In our work, the opinion type is determined by the knowledge base from which an opinion is inferred, where the knowledge base can be elicitated from either a single person or a population. There is a difference between the personal knowledge base and group knowledge base. The knowledge base of a single person usually stores the structure of one's interrelated knowledge pieces. It can be used to infer one's attitudes or viewpoints regarding the issues related to a topic. In contrast, the knowledge base with respect to a group of subjects captures the aggregation of common knowledge or beliefs shared by the whole population. The knowledge base of this type can be learned from data or be modeled from statistics. In the past few years, we have constructed knowledge bases for modeling group intents in topics from cultural behaviors [29] to emergent behaviors [30].

Figure 1.1 shows the relationship between opinions and knowledge bases. The initial or prior opinion Opn_0 is reasoned from the initial knowledge base KB_0. Similarly, the opinion at each time step Opn_t is inferred from the corresponding knowledge base KB_t. KB_t is dependent (but not only dependent) on KB_{t-1}. As a note, Opn_t does not refer to a single opinion but covers the opinions with respect to multiple issues, where the richness of an Opn_t is determined by the complexity of the knowledge base at that particular time. Full details of the knowledge base representation and opinion reasoning process will be covered in Chapter 2.2 and Chapter 3.

Information Sources

One critical component in opinion dynamics is influential sources. By interacting with those sources, individuals can update their knowledge bases so as opinions in a certain

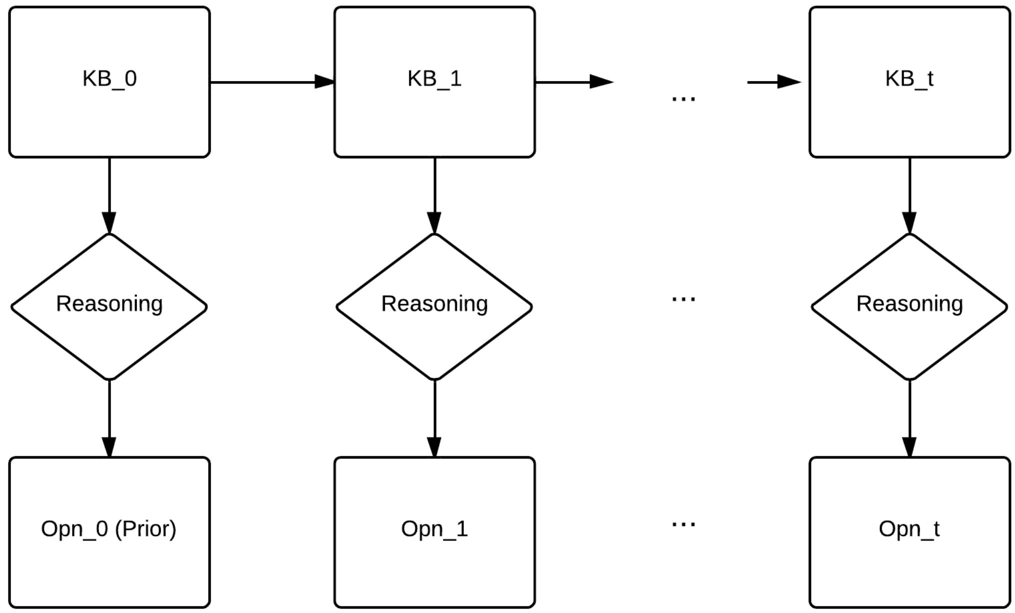

Figure 1.1: Prior opinion formation and evolution

way based on the received information. As we have discussed, some of the information is gained through visible channels, such as online social interaction, whereas others can be concealed due to various reasons. In this work, we assume that all communications take place in a subset of social community, usually referred to as a social network. Each node in the social network represents an actor. The information passed between actors is considered as knowledge fragments that have the same representation as the knowledge base. As a note, we have no restriction on the size of a fragment. A fragment can be as small as a single rule, e.g. "90% chance to rain", or be as large as a fraction of one's knowledge base. In order to model learning experience gained from information sources besides interpersonal communication, we allow actors to be media sources as well. Then the link from one news source to a person implies a media influence. To accommodate the dynamic nature of a social community, the structure of a social network/influential network does not have to be static, but can vary from time to time.

We extend Figure 1.1 to Figure 1.2 for the purpose of reflecting the opinion change in social networks. The knowledge base (KB) of an actor i is updated in a cycle. The initial opinions are inferred from her prior knowledge base KB_0. Then, at each time step t, she integrates the knowledge fragments sent from the actors/sources that she has interacted with to form KB_t in a certain fashion that will be covered shortly. Two types of sources exist in the social environment, namely visible sources and hidden sources. However, only the connections or interactions with visible sources can be tracked from the social networks, which makes it even harder to uncover the hidden sources that may have equal or even greater influences on shaping people's opinions. In Chapter 5, we will elaborate how we can leverage from the opinion formation process to detect *hidden influencing sources* meanwhile tracking the corresponding impact patterns by analyzing the opinion/behavior change over time. Preliminary results have been published in [31, 32].

Method of information processing

As we mentioned earlier, social learning lies at the heart of an individual's opinion formation and it is one of the most common means for people to gain knowledge. When new knowledge fragments are acquired through learning experience, instead of adopting them completely, it is a natural expectation that one should fuse the new fragments into his/her own knowledge base with some reliability/trust. Researches from social theories have shown that how people adopt the opinions of others depends on various factors, from relationship closeness [33] to personal and opinion similarity [34]. Although there is evidence that social factors correlate with level of reliability, how those factors interact still remains a mystery. In this thesis, to understand how social factors contribute to quantifying the reliance on each of the information source, we design a general impact metric by combining the influential factors in a linear fashion, where influential power associated with each factor can either be manually assigned or be learned from data. Linear models have been widely used in trust-based approaches due to its simplicity and intuitiveness [35, 8].

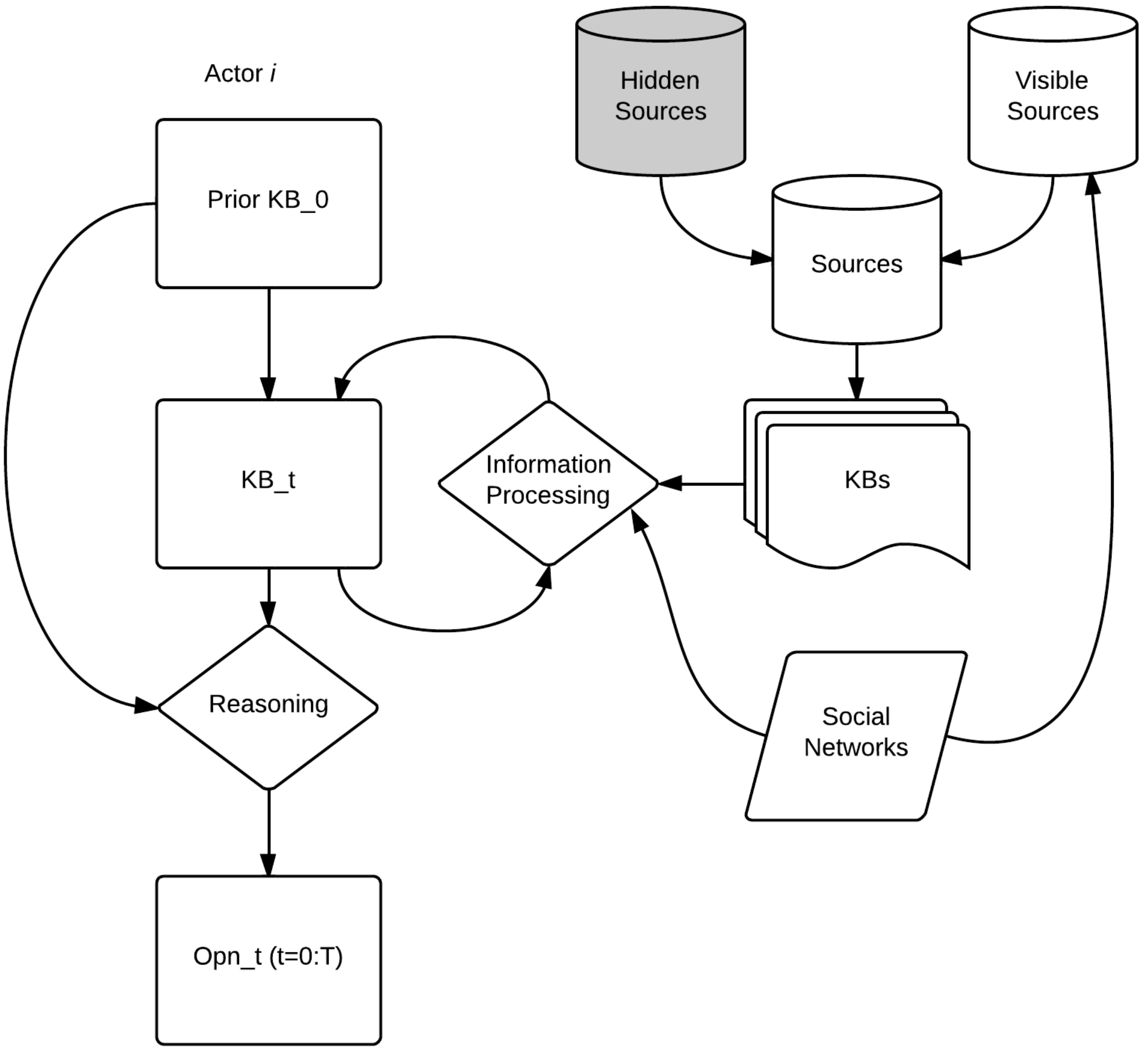

Figure 1.2: Opinion evolution in a social environment

After completing the information processing part, the whole framework of modeling opinion dynamics is depicted in Figure 1.3. The figure only shows three social factors that come into play to determine the impact metric: “friendship” based on the relationship network, “opinion similarity” based on the overall opinion distance and “contact frequency” based on the contact network. However, we will show that our model is not restricted to these factors. Newly discovered factors can be easily incorporated into the model. If we take a close look at the figure, “information processing” is actually the pivot of the whole system. It bridges opinion change with social influence factors as well as information sources. The impact from social factors and information sources will end up being reflected in the changing opinions. This makes it possible to fulfill the impact metric and capture hidden sources by computationally modeling the opinion change. In Chapter 4, we show how the weight of each influential factor can be efficiently learned via analyzing the opinion variation. Preliminary results have been published in [1].

1.2 Literature review of opinion dynamics modeling using computational methods

In this section, we provide an overview of the popular computational methods of opinion dynamics. Firstly, we discuss two traditional approaches to opinion dynamics modeling. One is a Bayesian approach based on Bayesian updating of beliefs, and the other one is a non-Bayesian approach based on reasonable “rules of thumb” on how people form their opinions on the basis of social influences. Then, we introduce some recent works that aim to study social influence and opinion dynamics using real data. We discuss their advantages and limitations and show how those previous researches inspired our work.

Bayesian approaches have been widely used for managing uncertainty, and more re-

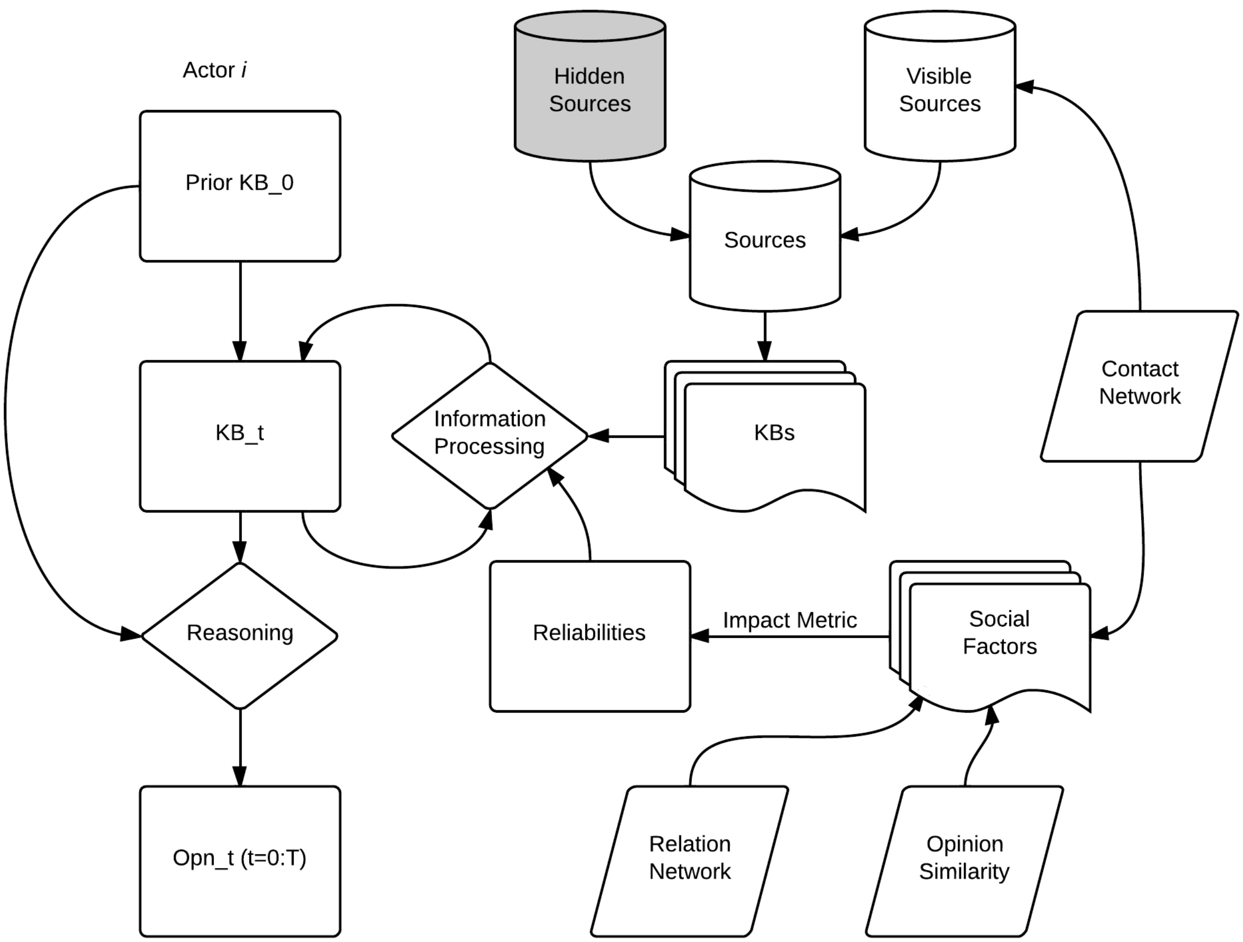

Figure 1.3: Complete framework of opinion dynamics

cently, for opinion formation. The Bayesian learning approaches assume that "individuals would update their beliefs optimally given an underlying model of the world" [5]. It was first proposed to explain uniformity in economic social behaviors [36]. The idea is that individuals should change their views to maximize their utility. Similar ideas have been applied to Bayesian learning from observations of past sequential actions [37] [38] and communication learning, in which individuals learn through communication other than the observation of others' actions [39] [40]. Despite the recognized success of these works, they may suffer from some common limitations. Considering a situation where an individual tries to shape an opinion about some social or political issues x, e.g. "whether the political view advocated by a political party should be trusted" after observing some evidence e. Then, Bayesian approaches postulate that the individual will update his / her prior opinion following Bayes rule:

$$P(x|e) = \frac{P(e|x) * P(x)}{P(e)}.$$

This equation implies that the individual has prior information on each possible value of x and also understands what observation may appear given a certain opinion. However, the requirement that the individual must have a complete set of priors is not innocuous. Especially when the possible values of x is large. Similarly, the assumption that the individual will have a reliable knowledge about the likelihood $p(e|x)$ can also be problematic when such information needs to be obtained from the behaviors/actions of others. Another type of limitation is that much of the literature has not tackled the issue of learning taking place in the context of a social network. The assumption that decisions are made without taking into account the social network or any contacts a person may have is also unrealistic.

As the assumption that individuals should have a reliable model of the world can be too restrictive and cause complexity issues, alternative approaches have been developed to overcome such limitations. Gilboa and Schmeidler [41], for example, suggest an approach for prior formulation based on “empirical frequency”. In this approach, individuals form their beliefs about a situation or make decisions based on their experiences in similar situations in the past. Bayesian Networks (BNs) [42] are another type of models that have been used in knowledge representation due to its sound theoretical foundations in probability theory combined with efficient reasoning. For example, Garg et al. [43] introduced a BN based divergence minimization framework to integrate opinions from different sources in order to solve the problem of standard opinion pooling. However, the incompleteness in people’s knowledge bases can become problematical to BNs, as they require a completely specified conditional probability table (CPT). BNs also require that information be topologically ordered which further restricts their general applicability to real-world situations.

The second class of opinion dynamics modeling is non-Bayesian approaches. How individuals form their opinions are determined by “rules of thumb”, which are either motivated by social theories or have been justified to be capable of producing desirable properties. The work provided by French [44] became he pioneer of this model family. In his work, individuals form their new opinions by averaging the opinions of other people with whom they have communicated directly. After that, several non-Bayesian approaches were developed in sociology and physics around the idea that “the opinions of each agent evolve dynamically over time as a function of their neighbors’ opinions”. Most works in this area are based on classical models of interacting particle systems motivated by statistical mechanics [45]. For instance, we have the first uses of the Ising model to describe the behavior of laborers in a strike [46]; the voter model, where only one neighbor affects the agent at a time [47] [48], and continuous opinion models like DeGroot [49].

The DeGroot Model is a simple but widely used model of belief and consensus formation over social networks. The DeGroot Model and its variations have been well exploited in the cooperative control literature for achieving collective behavior in a social environment [50, 51, 6]. The idea of this model is to replace the simple average function in French [44] with a weighted mean in order to assess opinion pooling of a dialogue among experts. Imagine that each person holds an opinion represented by a real number, which might be a political spectrum, or a probability that i assigns to a certain opinion. Then, at any time $t >= 0$, individuals in the defined network will update their opinions/beliefs represented by a vector $x(t)$ according to the following equation:

$$x(t+1) = A * x(t)),$$

where A is a stochastic weight matrix that reflects the extent of influences between individuals in a social network. Besides its simplicity and intuitiveness, another attractive feature is that the analysis of consensus is straightforward as well. The sufficient conditions to ensure the convergence can be easily derived from Markov Chain Theory. Despite the fact that the weighting matrix in DeGroot is stochastic and static, and it overlooks the situations when consensus is not reached, it still serves as a classic opinion formation model that can be easily modified to incorporate social theories.

After DeGroot [49], many extensions have been proposed to overcome its limitations. For example, Chatterjee and Seneta [52] modeled the phenomenon when the weights between two actors are modeled as a function of time. Krause [53] allowed each actor to only take opinions from the ones whose opinions differ from his/her own within a bounded confidence. However, many DeGroot-based models are still linear in the sense that the updated opinions are calculated as a linear combination of the adopted opinions. Of course, having a mathematical model does not mean that the updated opinion is always obvious and can be determined from the model. Wu and Huberman [17] developed a model to predict the evolution of a set of opinions by considering the structure of the social network. In their work, people update their opinion with a probability based on the opinion distribution in her neighbors. However, they assume that the influences from each neighbor are the same, which is hardly true in the real world. Most of these works are theoretical models where the strong assumptions attached to these methods, such as fixed social structure [6], equally distributed influence [17] or binary opinion dynamics [54, 55] restrict them to be only applicable to extremely simple scenarios.

Some recent efforts have been made on studying social influence using real data. For example, "Influence Model" has been applied to model influence within social systems. The influence model was first articulated in Basu and Choudhury [56] to describe the connections between Markov chains in terms of the "influence" between them. The main idea is that the distribution of an entity's state or an opinion on a given issue at time t is a combination of the pairwise transition probabilities from $t-1$ to t with some "influence" parameters. Pan et al. [22] extended it to model dynamical influence in human interaction, where the dynamical parameters that characterize transition probabilities and the influences between entities can be learned from data. However, the influence model stands on a hypothesis that every entity in a social community has the same transition probability distribution, the individual difference in shifts in the states/opinions, therefore is overlooked.

Another limitation of the existing methods is that most of them only consider single opinion dynamics but overlook the connections between knowledge fragments [15, 57, 8]. As we mentioned earlier, it is necessary to capture and model such relationships so that more insights on opinion dynamics can be derived via reasoning. All these works inspire us to develop a framework to model different aspects of opinion dynamics in one entire system. The foundation of the framework is that we treat opinions a product inferred from one's knowledge base, where the knowledge base keeps growing and updating via various type of social learning.

1.3 Contributions

The key contributions of this thesis include the following:

1. We provide a domain-independent computational model to study opinion dynamics.

   (a) The model is built on individual knowledge-based system to capture the intents behind opinion change.

   (b) The model considers individual difference in knowledge bases.

   (c) The model takes into account the dependencies between knowledge fragments.

   (d) The model characterizes the influences in a social community by realizing and enriching theories of social communication.

   (e) The model provides a flexible platform where new influential factors can be readily plugged in.

2. The results generated by the model are promising compared to the existing approaches.

   (a) The performances are consistently promising in both synthetic and real data.

   (b) The model is robust to different environmental settings.

   (c) The model is compatible with classical approaches under a certain conditions.

   (d) The model facilitates the validation of social influential factors.

   (e) The results can potentially be used to predict and regulate future opinion changes.

3. The model can be used to detect and track hidden sources.

   (a) The results show the capability of tracking influencing sources even when the impacts are small.

(b) Newly added hidden sources are detected more accurately than classical models.

(c) The results can be used to explain the role that external sources play in opinion and behavior change in complex scenarios.

1.4 Outlines

This thesis presents a framework for modeling opinion dynamics. The rest of the paper is organized as follows:

- Chapter 2 discusses background materials on social influence, Bayesian Knowledge Bases and a brief description on social network.
- Chapter 3 overviews some popular computational approaches for representing an opinion and formally introduces our opinion formation model.
- In Chapter 4, we introduce a general impact metric that estimates the reliability for each adopted knowledge fragment based on a set of social influential factors. We show how the impacts of each social factor can be efficiently learned based on the OFM derived in Chapter 3 .
- Chapter 5 extends the OFM and introduces a new modeling approach called a Finite Fusion Model for detecting and tracking hidden sources in a time-variant scenario. The goal is to analyze the behavior change over time by detecting hidden influencing sources and tracking the corresponding impact patterns.
- We conclude in Chapter 6 and describe a few limitations and potential future works.

1.5 List of Acronyms

- Knowledge Base KB

- Opinion Formation Model OFM
- conditional probability table CPT
- Bayesian Knowledge Bases BKBs
- Bayesian Networks BNs
- conditional probability rules CPR
- causal rule set CRS
- Linear Programming LP
- Relative Squared Error RSE
- Dynamic Bayesian Networks DBNs
- Hidden Markov Model HMM
- Expectation-maximization EM
- Finite Fusion Model FFM
- Sequential Quadratic Programming SQP
- Gaussian Mixture Model GMM
- Mixtures of Bayesian Networks MBNs
- Receiver Operating Characteristic ROC
- standard std
- Waikato Environment for Knowledge Analysis Weka

Chapter 2

Backgrounds

We provide background materials on social influence, Bayesian Knowledge Bases (BKBs) and social network in this chapter. We first introduce the primary observations and findings in social scientific studies that serve as the theoretical foundation of the social influence learning (Chapter 4). Then, we provide the formal definition of BKBs and its reasoning algorithms that are applied in this work to quantify each individual's opinion values (Chapter 3, 4). Lastly, we give a brief description on the social network that is employed to define a social influential environment.

2.1 Social Influence

Social influence lies at the heart of individual's opinion formation. It has been explored in many aspects. One branch of works is based on social network analysis. They consider social influence as an ability to affect the diffusion process. Works fitted in this category involve the study of diffusion patterns [58, 59], influential entities [60] and information flow tracking [61, 62]. Computer simulation of interpersonal influence processes has been employed to study the conditions leading to the massive information cascade. The findings from Watts and Dodds [63] suggest that the large cascades of information are due to a few critical individuals whose opinions can be easily influenced other than the influentials[1] . However, the analogue with the real world is not exact the same. Influence in real world is much more complex because the influential pattern concluded from one scenario may not apply to another; social behavior can be strategic; and people have a free will to choose whom they want to interact with.

In contrast, another angle in studying social influence channels its focus into the reasons behind the social influence. For instance, Granovetter [64] presented threshold models of collective behaviors to understand situations where outcomes are not intuitively consistent with the individual preference. They explained the paradoxical phenomenon as rioting and diffusion of innovation. Another topic of interest is social notions of trust. It defines a social network about with whom we should share information and from whom we should accept information. Trust-based models have been well developed to measure the magnitude of trust between people and to provide recommendations in social networks [27, 65, 66]. A common practice to measure the extent of trust in most of the popular approaches is via a simple scalar [67, 68]. However, social trust depends on a host of factors that can never be easily modeled computationally [27]. One reason is that it has various forms depending on the situational context of specific interactions. Moreover, it is often the case that one may believe someone's words, but refuses to change her mind. In this thesis, we study and model social trust from a different angle. Specifically, we only focus on how people view and adopt the opinions held by the people they have communicated with. The extent of people's willingness to adopt a new knowledge is characterized by a *reliability* value that can be estimated via some assessment mechanisms. Considering that how an individual adopts the beliefs of others may vary in different situations, we model reliability as a function of social factors. We start by considering three domain independent influential factors: relationship, contact frequency and opinion similarity, which have been shown to

[1] A minority people that have great influence on an exceptional number of their peers [63]

be highly related to opinion formation by communication scholars and social theories. As a note, our model is not restricted to these factors. Newly discovered factors can be easily incorporated into the model (Chapter 4).

2.1.1 Relationship

Relationship characterizes the closeness between people, is a strong indicator of how likely one may accept another's opinion. It forms the social environment, e.g. family or friend that people are willing to adjust their attitudes to conform to their group members [33]. For instance, people are inclined to take the opinions from their close friends more seriously than from random people. Rotter, in his work [69], conducted sociometric analysis to study interpersonal trust. The results show that trust, as measured sociometrically, is positively related to friendship.

2.1.2 Contact Frequency

The second factor we considered in the social influence study (Chapter 4) is the recent contact frequency. Without frequent contact, even two close friends may differ so much on their knowledge bases that one can hardly convince the other. Chong and Druckman [70] proposed a theory of framing to identify the key parameters that determine how public opinion is formed. One of their finding is that frequent communication allows one to deliver her opinions more credibly to the public.

2.1.3 Opinion Similarity

Similarity has been used as a metric to evaluate the relationship between people in social network analysis [71]. Instead of measuring the similarity of personal characteristics, such as age and location, we measure the opinion-wise similarity, because the former one can be highly correlated to the relationship. The work from Sniderman and Theriault [34] shows that individuals favored the idea that was consistent with their own values. Crandall et al. (2008) developed techniques for identifying and modeling the online interactions between social influence and selection. They addressed the problem via empirical analysis of Wikipedia data. Their results suggest that the similarity between two individuals leads to future interaction, and interaction leads to further similarity.

2.2 Bayesian Knowledge Bases

In this section, we introduce Bayesian Knowledge Bases (BKBs) - a rule-based probabilistic model that represents possible world states and their (causal) relationships using a directed graph. Instead of specifying the causal structure using conditional probability tables (CPT) as in BNs, it collects the conditional probability rules (CPR) in an “if-then" style. Figure 2.1 shows a graph structure of a BKB fragment, in which A is a random variable with possible instantiations (or “states” in BNs terminology) $\{a_1, a_2\}$. Each instantiation of a random variable is represented by an I-node, or “instantiation node", e.g. $A = a_1$, and the rule specifying the conditional probability of an I-node is encoded in an S-node, or “support node" with a certain weight. For example, q_7 corresponds to a CPR which can be interpreted as: if $A = a_2$ and $C = c_1$, then $B = b_2$ with a probability 0.5.

One benefit of this approach is that it allows for independence to be specified at the instantiation level instead of the random variable level. Also, it does not require the full table representation of the CPTs and allows for reasoning even with incompleteness. As the BKB shows in Figure 2.1, the dependency relationship at the variable level implies that variable B depends on both A and C. However, given the evidence of $A = a_1$, B becomes independent of C. This could happen in the real world when the role of a critical variable can dominate some local dependency relationships between variables. In a BN, in order to represent the probability distribution of B dependent on A and C, all CPT entries of

$P(B|A,C)$ are required to fill in, which could grow exponentially when the number of the parent nodes is large. BKBs in contrast, only capture the knowledge that is available and does not require a complete distribution specification.

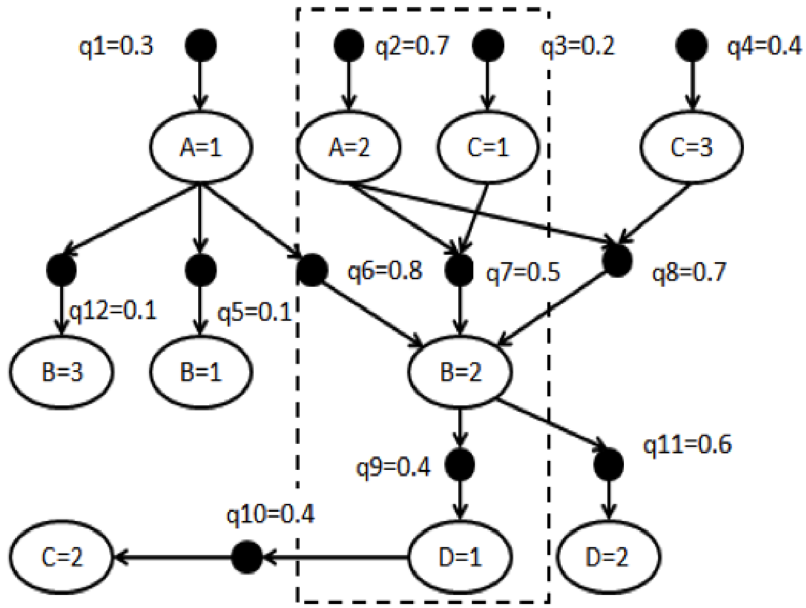

Figure 2.1: Example of a BKB fragment

The other feature of BKBs is that they also allow cyclic relationships among random variables. Imagine if the direction of some causal mechanism also depends on specific states of the variables. Santos et al. [2] gives an example of BKB modeling of a political election, in which the type of “race" may flip the causal direction between the belief of a piece of evidence and the voting action.

We first provide a formal definition of BKBs [72], followed by the BKB fusion algorithm [2] that allows us to merge multiple BKBs into one BKB without violating semantic and probabilistic soundness. Lastly, we describe our previous work: BKB tuning algorithm [73] which can be used to modify or correct a distribution without changing the structure.

2.2.1 Framework of Bayesian Knowledge Bases

The formal definition of the graphical representation of a BKB is given below:

Definition 1. (Santos and Dinh [72]) A *correlation-graph* is a directed graph $G = (I \cup S, E)$ in which $I \cap S = \phi, E \subset \{I \times S\} \cup \{S \times I\}$, and $\forall q \in S$, there exists a unique $\alpha \in I$ such that $(q, \alpha) \in E$. If there is a link from $q \in S$ to $\alpha \in I$, we say that $q\, supports\, \alpha$.

For each S-node q in a correlation graph G, we denote $Pred_G(q)$ as the set of I-nodes pointing to q, i.e. $Pred_G(q) = \{\alpha \in I | \alpha \to q \in E\}$ and $Desc_G(q)$ as the I-node supported by q, i.e. the α such that $q \to \alpha \in E$.

Definition 2. (Santos and Santos [74]) Two I-nodes, α_1 and α_2 are said to be *mutually exclusive* if there is an I-node $R = v_1$ in I_1 and an I-node $R = v_2$ for which $v_1 \neq v_2$.

Similarly, two sets of I-nodes I_1 and I_2 are mutually exclusive if there exists two I-nodes $\alpha_1 \in I_1$ and $\alpha_2 \in I_2$, such that α_1 and α_2 are mutually exclusive. For example, the sets of I-nodes $\{A = a_1, B = b_2\}$ and $\{A = a_2, B = b_2, C = c_1\}$ are mutually exclusive. A *state* is a set of I-nodes such that it contains no more than one instantiation of each random variable.

Definition 3. (Santos et al., [73]) A set of S-nodes R is said to be *complementary* if for all $q_1, q_2 \in R$, $Desc_G(q_1)$ and $Desc_G(q_2)$ are mutually exclusive, but $Pred_G(q_1)$ and $Pred_G(q_2)$ are not mutually exclusive. Variable v is said to be the *consequent variable* of R, if for any S-node $q \in R$, $Desc_G(q)$ is an instantiation of v.

We denote ρ_v as the set that contains all possible complementary sets of S-nodes w.r.t. variable v, such that for any complementary set $r \in \rho_v$, v is the consequent variable of r.

We also introduce Ψ_v to denote the subset of ρ_v that removes all complementary sets in ρ_v that are a subset of some other complementary set, i.e. $\Psi_v = \{r | r \in \rho_v \wedge \nexists r' \in \rho_v, r \subseteq r'\}$. For example in Figure 2.1, $\{q_5, q_6\}$ is a complementary set w.r.t. variable B and $\Psi_B = \{\{q_5, q_6, q_{12}\}, \{q_7\}, \{q_8\}\}$ Note that for a certain variable v in a Bayesian Network, the size of Ψ_v goes exponentially with the number of its parent variables. However, since BKBs only capture context-specific dependence rules, the size of Ψ_v can be considerably smaller in real-world cases.

Definition 4. A set of S-nodes B is said to be a *causal rule set* (CRS) for the random variable v if B contains all S-nodes pointing to the instantiations of v. As a note, each S-node only belongs to one CRS.

For each complementary set $r' \in \Psi_v, r'$ is a subset of B. For example, the CRS for variable B in Figure 1 is $\{q_5, q_6, q_7, q_8, q_{12}\}$.

Definition 5. A Bayesian knowledge base (BKB) is a tuple $K = (G, w)$ where $G = (I \cup S, E)$ is a correlation graph, and $w \colon S \to [0, 1]$ such that

1. $\forall q \in S, Pred_G(q)$ contains at most one instantiation of each random variable.

2. For distinct S-nodes $q_1, q_2 \in S$ that support the same I-node, $Pred_G(q_1)$ and $Pred_G(q_2)$ are mutually exclusive.

3. for any complementary set of S-nodes $R \subseteq S$, R is normalized: $\sum_{q \in R} w(q) \leq 1$ where $w(q)$ is a weight function that represents the conditional probability of $P(Desc_G(q) | Pred_G(q))$.

The intuition behind these three conditions is that each S-node can only support one I-node; two rules supporting the same I-node cannot be satisfied at the same time; and, to ensure normalization of the probability distribution, every complementary set of S-nodes should be normalized.

Definition 6. The CRS B w.r.t. the variable v is called normalized if for any complementary set $r \in \rho_v, r$ is normalized.

Theorem 1. Let B be the CRS of variable v, if for any complementary set $r' \in \Psi_v, r'$ is normalized, then B is also normalized.

Theorem 1 can be easily proved from the definition of Ψ_v, since any complementary set in ρ_v is a subset of some complementary set in Ψ_v.

As in BNs, reasoning with BKB is also based on the calculation of joint probabilities over the possible worlds. Here, a *world* is a subgraph of a BKB including at most one I-node of each random variable and the associated S-nodes (a world is referred to as an *inference* in [74]). For example in Figure 2.1, the dotted rectangle circles a world. As proved in Santos & Santos [74], the joint probability of a world τ is just the product of the weights of all S-nodes in τ, i.e.$P(\tau) = \Pi_{(q \in \tau)} w(q)$. The idea of a world plays an important role in two forms of reasoning with BKBs, belief revision (also called maximum a posteriori or MAP) [74, 42] and belief updating. In belief updating, the goal is to calculate the probability of a state of a random variable given some evidence, e.g. $P(Y = y | X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n)$. As shown in Rosen et al [75], the joint probability of a state $\theta = \{X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n\}$ is a summation of the probabilities of all possible worlds where θ is true, i.e. $P(\theta) = \sum_{(\tau \in I_\theta)} P(\tau)$, where I_θ represents the set of worlds containing θ. The worst-case complexity of reasoning with BKBs is NP hard. Some approximation algorithms, e.g. stochastic sampling methods, have been introduced to make the reasoning process more efficient [75]. However, given that BKBs only captures the knowledge available, for real-world problems [76, 77, 78], exact methods have been shown to be computationally feasible at least in terms of moderate size.

To better understand the role of each node in the probability calculation, an algorithm used to compute the *contribution* of node q to the probability of state θ, represented by $c_\theta(q)$, was presented in [76], in which the contribution is defined as the sum of the probabilities of the worlds including both θ and q.

Definition 7. $c_\theta(q) = \sum_{(\tau \in I_\theta \wedge q \in \tau)} P(\tau)$.

Here, q is not restricted to just S-nodes. If it is an I-node, then the contribution measures how much influence an event can have on a state.

2.2.2 BKB Fusion

One extension based on the BKB framework is BKB Fusion. Unlike BNs, multiple BKB fragments can be combined into a single valid BKB using the BKB fusion algorithm [2]. The idea behind this algorithm is to take the union of all input fragments by incorporating source nodes, indicating the source and reliability of the fragments. As such, all knowledge is preserved and source/contribution analyses can be conducted to determine the impact of various elements of source knowledge on reasoning results. Figure 2.2 shows an example of the knowledge fusion, where two BKB fragments from actor 1 and actor 2 are fused into one. A valuable property of knowledge fusion is that it considers the impacts from multiple sources when constructing explanations for any evidence observed. For detailed description of the BKB fusion algorithm, please refer to Appendix A.

2.2.3 BKB Tuning

Another useful extension named BKB Tuning was developed in one of our previous works Santos et al. [73]. In this work, we proposed a method to tune a BKB. Specifically, we

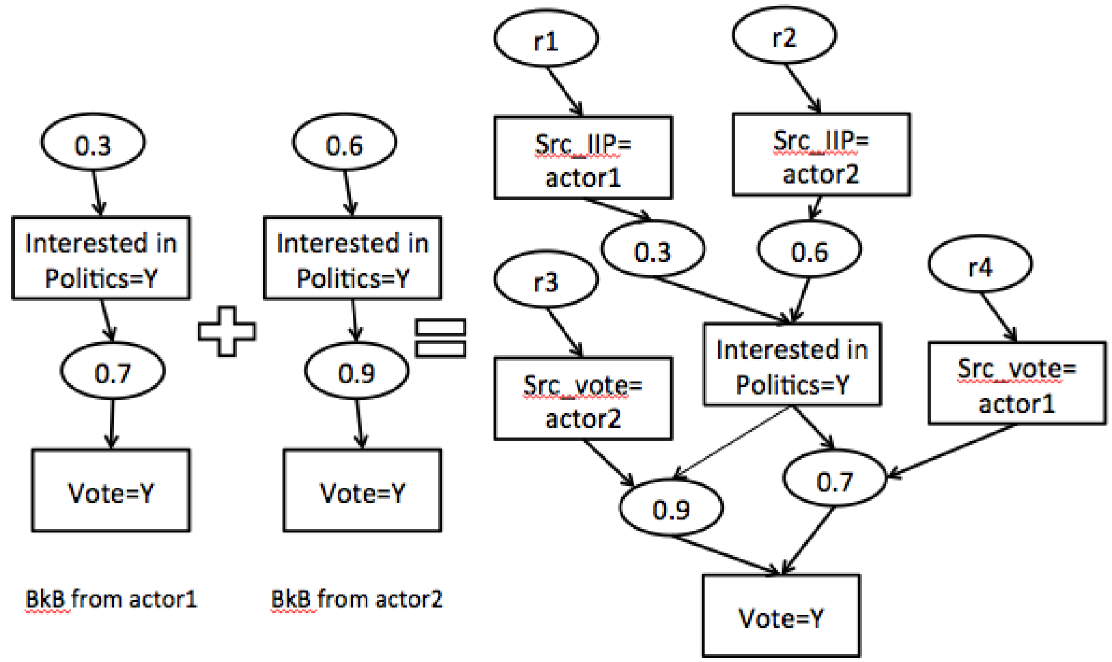

Figure 2.2: Example of BKB fusion

described how to explain and fix the conflict between the users' expectation and system answer by making a small change to the parameters. The optimal set of parameters is determined by evaluating their sensitivity, which minimizes the amount of change. Both single parameter and multiple parameters solutions are provided. We show that tuning a system via multiple parameter changes has the same computational complexity as single parameter changes. The problem of multiple parameter changes can be transformed into a Linear Programming (LP) problem and efficiently solved.

BKB Tuning can be applied to tune the marginal probabilities of a BKB. This is extremely useful to modify or correct a distribution with the whole structure being maintained. As BKB subsumes BN, the tuning method is applicable to BN problems. Moreover, tuning method can also be used for controlling reliabilities or weights of the knowledge bases from different sources in knowledge fusion.

The tuning problem is defined as: given a BKB, $K = (G, w)$ and some evidence e entered by the user, we determine how we can identify necessary changes to a set of S-nodes such that the tuned BKB can enable $P'(\mathbf{v} = v_1|e) > P'(\mathbf{v} = v_2|e)$, where P' is the distribution after we apply changes on the rules (as compared to the original distribution P). Before we move on to the tuning algorithm, we would like to first introduce the formal definition of *sensitivity* with respect to a single S-node and a CRS.

Definition 8. Given a BKB $K = (G, w)$, let x be the weight of S-node q, the sensitivity of x on $r(\mathbf{v}|e)$ is a partial derivative:

$$S(x|\mathbf{v}, e) = \frac{\partial r(\mathbf{v}|e)}{\partial x}$$

The ratio $r(\mathbf{v}|e)$ can be expressed in terms of x as:

$$r(\mathbf{v}|e) = \frac{P(\mathbf{v} = v_1|e)}{P(\mathbf{v} = v_2)|e} = \frac{P(\theta_1)}{P(\theta_2)} = \frac{\alpha x + a}{\beta x + b}$$

which is a fraction of two linear functions on x. Then the partial derivative of $r(v|e)$ on x turns to be:

$$\frac{\partial r(\mathbf{v}|e)}{\partial x} = \frac{ab - \alpha\beta}{(\beta x + b)^2} \tag{2.1}$$

The S-nodes with higher $|S(x|\mathbf{v}, e)|$ are more sensitive to the results, and thus more likely to correct the system with small change.

Definition 9. Given a BKB $K = (G, w)$, let $\vec{x} = [x_1, x_2 \ldots x_n]$ be the vector that denotes the weights of S-node in CRS $B = \{q_1, q_2 \ldots q_n\}$, the sensitivity of $\vec{x}$ on $r(\mathbf{v}|e)$ is a partial derivative:

$$S(\vec{x}|v, e) = \frac{\partial r(\mathbf{v}|e)}{\partial \vec{x}}$$

Similar to the way we derive $r(\mathbf{v}|e)$ as a function of the weight of a single S-node, $r(\mathbf{v}|e)$

on $\vec{x}$ can be given by:

$$r(\mathbf{v}|e) = \frac{P(\theta_1)}{P(\theta_2)} = \frac{\alpha_1 x_1 + \alpha_2 x_2 + \ldots + \alpha_n x_n + a}{\beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n + b}$$

Then the sensitivity of $\vec{x}$ on $r(\mathbf{v}|e)$ is a partial derivative vector:

$$\frac{\partial r(\mathbf{v}|e)}{\partial \vec{x}} = [\frac{\partial r(\mathbf{v}|e)}{\partial x_1}, \frac{\partial r(\mathbf{v}|e)}{\partial x_2}, \ldots, \frac{\partial r(\mathbf{v}|e)}{\partial x_n}] \tag{2.2}$$

The Euclidean norm of the partial derivative vector can be used to compute the sensitivity value. The CRS with higher sensitivity value will cause larger changes to the results.

Two types of solution are provided in [73]. The first one is to tune a single S-node. The idea is to transform the original problem into a set of conditions and find the solution set to satisfy the conditions. The conditions to correct the system for S-node **q** are listed as follows:

$$\begin{cases} \forall R \in \Psi_v(\mathbf{q}), \sum_{q \in R} w(q) + \mathbf{\Delta w} \leq 1 \\ w(\mathbf{q}) + \mathbf{\Delta w} > 0 \\ (\alpha_{\mathbf{q}} - \beta_{\mathbf{q}})\mathbf{\Delta w} > b - a + c_{\theta_2}(\mathbf{q}) - c_{\theta_1}(\mathbf{q}) \end{cases} \tag{2.3}$$

where

$$\alpha_{\mathbf{q}} = \frac{c_{\theta_1}(\mathbf{q})}{w(\mathbf{q})} \text{ and } \beta_{\mathbf{q}} = \frac{c_{\theta_2}(\mathbf{q})}{w(\mathbf{q})}$$

In Inequalities 2.3, $\Psi_v(\mathbf{q})$ collects the complementary sets in Ψ_v that contain **q**; $c_{\theta_1}(\mathbf{q})$ and $c_{\theta_2}(\mathbf{q})$ are contributions of **q** in θ_1 and θ_2; a and b are two constants independent of $w(q)$ Since $\mathbf{\Delta w}$ is the only unknown variable, the solution space can be easily found by solving for the equality condition. Then, which S-node **q** to tune is decided by the sensitivity analysis.

The second type of solution is to tune CRS. Combining with normalization constraints of Ψ_v to find the solution of parameter changes over the CRS that satisfies $P'(\mathbf{v} = v_1|e) - P'(\mathbf{v} = v_2|e) > 0$, the following conditions must hold:

$$\begin{cases} \forall R \in \Psi_v, \sum_{q \in R} w(q) + \sum_{q \in R} \mathbf{\Delta w}(q) \leq 1 \\ \forall q \in \mathbf{B}, w(q) + \mathbf{\Delta w}(q) > 0 \\ \sum_{i=1}^{n} (\alpha_i - \beta_i) \mathbf{\Delta w}(q_i) > b - a + \sum_{i=1}^{n} [c_{\theta_1}(q_i) - c_{\theta_2}(q_i)] = C_0 \end{cases} \tag{2.4}$$

where n is the number of S-nodes in **B**,

$$\alpha_i = \frac{c_{\theta_1}(q_i)}{w(q_i)} \quad \text{and} \quad \beta_i = \frac{c_{\theta_2}(q_i)}{w(q_i)}$$

Solving Inequalities 2.4 is treated as a LP problem using simplex algorithms, i.e. $Min \sum_{i}^{n} |\mathbf{\Delta w}(q_i)|$, such that the upper conditions hold. Similar to tuning one S-node, the CRS with the highest sensitivity to the results is selected for tuning. An example of how tuning algorithm is applied More details about the tuning algorithm and associated properties can be found in [73].

2.3 Social Network

Social networks play a fundamental role in the spread of information. Hot topics or rumors used to be propagated among people in a "word-of-mouth" form. The rise of World Wide Web accelerates the development of large-scale social networks and makes it possible to diffuse information faster. Large online social medias include Twitter and Facebook . In the last few decades, researchers from the social and behavioral science have made considerable efforts to analyze social networks. Much of these efforts can be attributed to the focus of relationships between social actors and the patterns of these relationships [6, 79]. As the goal of this thesis is not to develop new social network analysis techniques, we omit details of these techniques. A thorough review of social network analysis methods is provided by

Wasserman [80].

Recently, with the development of Internet, considerable attention has been devoted to analyzing information on online social medias. For example, Gruhl et al. [81] and Lin et al. [82] developed models to identify and track hot topics in online communities. Danescu-Niculescu-Mizil et al. [83] studied how people's attitudes are affected by online reviews. Their results suggest that users tend to consider the reviews that agree with the average to be more helpful.

The relationships encoded in social networks may be of many sorts from economic to interactional. In this work, we focus on direct communications from one individual to another, as direct communication is the most powerful and effective way for people to convey their ideas and knowledge. Communication has various forms as well: face-to-face conversations, email, phone, instant messaging, bluetooth file transfer, to name a few. No matter what form a commutation takes, they all serve as channels to pass knowledge fragments from one source to another. We assume that all communications take place in a subset of social community, usually referred to as a social network. Each node in the social network represents an agent. The information passed between actors is represented by BKB fragments. In order to model learning experience gained from information sources besides interpersonal communication, we allow agents to be media sources as well. Then the node with no incoming edges indicates a media influence. Apparently, the structure of such social network/influential network is not static, but varies in time. Figure 2.3 depicts an example contact network, where links in different colors represent different communication means.

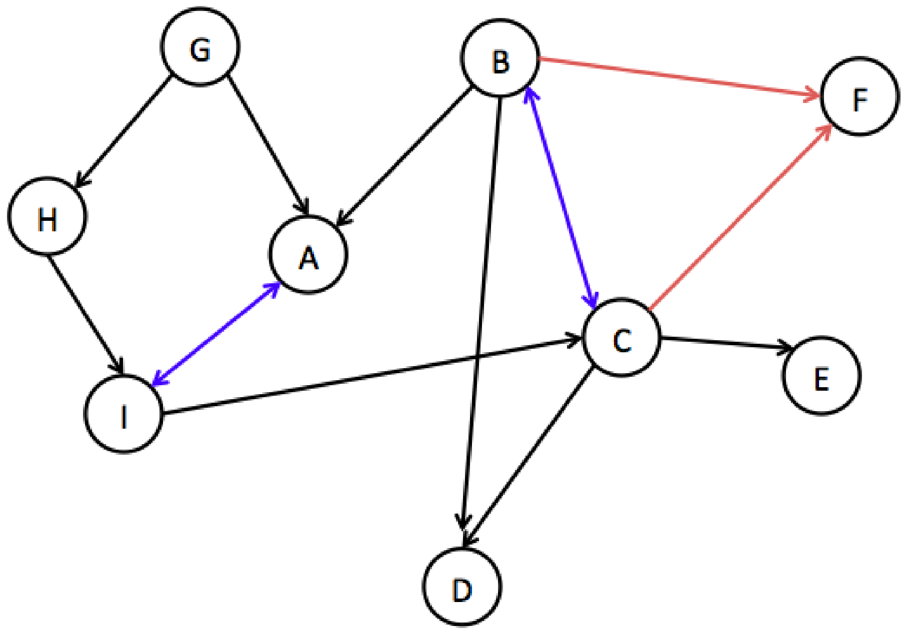

Figure 2.3: Example of a small contact network

Chapter 3

Opinion Formation

That social influences alter individual's opinions is a truism to which anyone can readily consent. However, people do not shape opinions on an issue for no reason. To understand opinion dynamics in a social community, it is crucial to study how opinions are formed and how to quantify and represent opinions.

3.1 Opinion as a Reasoning Product

According to Damer [84], opinions, supported by arguments, should be viewed as the conclusion resulting from certain process of rational reflection on evidences. Moreover, studies on cognitive processes have shown that arguments are inferred from one's knowledge bases [85, 86]. Thus, in this work, we assume that all opinions are accompanied by the reasons that support them. In fact, though everyone is entitled to have her opinions, supported opinions are considered more valuable when a person is evaluating a bunch of given opinions [84]. Traditional models mainly focus on the process of opinion formation where opinions are treated as independent values. Multiple opinions are integrated into one based on some weighting strategies [49, 6, 15]. Then the question becomes how to assign weights to each received opinion? Different weighting functions lead to different results, e.g. consensus or polarization [15]. In recent years, more and more works start evaluating the trust between people in social networks based on homophily effects such as the similarity in demographic features, e.g. age, gender, etc [27, 66]. The intuition is that individuals in homophilic relationships share common characteristics that make communication easier. However, to evaluate whether an opinion deserves acceptance, homophily only is not sufficient. For example, a person trusts someone's words does not necessarily mean that she will want to change her opinions. In addition, traditional methods usually simplify the problem by assuming that opinions are only formed via communication with others. However, what if an individual is temporarily isolated from others, can she change her opinions as a result of self-learning, e.g. through observations of her own experience? The answer is yes. The work from Gureckis and Douglas [87] suggested that self-directed learning is an important factor as it optimizes the education experience and allows people to focus on useful information. In a nutshell, in this work, opinions are treated as a reasoning product from one's context knowledge base where new knowledge fragments are acquired through various types of learning experience [5].

3.2 Computational Approaches for Opinion Representation

A knowledge base (KB) usually refers to an information repository that stores knowledge fragments. It is accompanied with a reasoning engine that enables effective retrieval of information and knowledge-based inference. Cognitive scientists have made efforts to structure individual knowledge-based system as interrelated argument pieces [26], from which one can form her opinions via reasoning. For example, Zukerman [85] proposed a graphical model to capture connections in context knowledge, where connections like causality and conditional dependencies play an important role in reasoning process. On the other hand, the knowledge-based system at any stage is necessarily associated with some degree of incompleteness and uncertainty. This results in a natural expectation that a knowledge-based system should exhibit enough flexibility and intuitiveness for capturing knowledge in order to make it easy and responsive for the knowledge engineers to build models/systems that they can still understand, maintain, validate, update, and so forth. Bayesian approaches, therefore, have been widely used for managing uncertainty, and more recently, for opinion formation. Among those, Bayesian Networks (BNs) [42] keep gaining attention in knowledge representation due to its sound theoretical foundations in probability theory combined with efficient reasoning. For example, Garg et al. [43] introduced a BN based divergence minimization framework to integrate opinions from different sources in order to solve the problem of standard opinion pooling. However, the incompleteness in people's knowledge bases can become problematical to BNs, as they require a completely specified CPT. BNs also require that information be topologically ordered which further restricts their general applicability to real-world situations. To overcome these limitations, we apply a probabilistic framework called Bayesian Knowledge Bases (BKBs) [74] to represent individual's knowledge base due to its flexibility on modeling incomplete information and allowance of cyclic on variable level. BKBs have been extensively used to represent individual's goals, intentions, opinions and cultural influences [76, 29]. Figure 3.1 shows a sample BKB fragment from an intent framework.

In sum, we consider opinions as a product inferred from one's knowledge-based system, where new knowledge pieces are acquired through social interaction and learning experiences. To illustrate the differences between our method and the existing approaches, e.g. DeGroot based models, we depict the comparison of our model of opinion formation with them in Figure 3.2. To the best of our knowledge, this is the first computational framework of opinions where opinion dynamics is modeled as an evolution of knowledge bases.

In this work, an individual fuses the knowledge fragments (represented by BKBs) received from n sources into her prior knowledge base. Each fragment will be assigned with

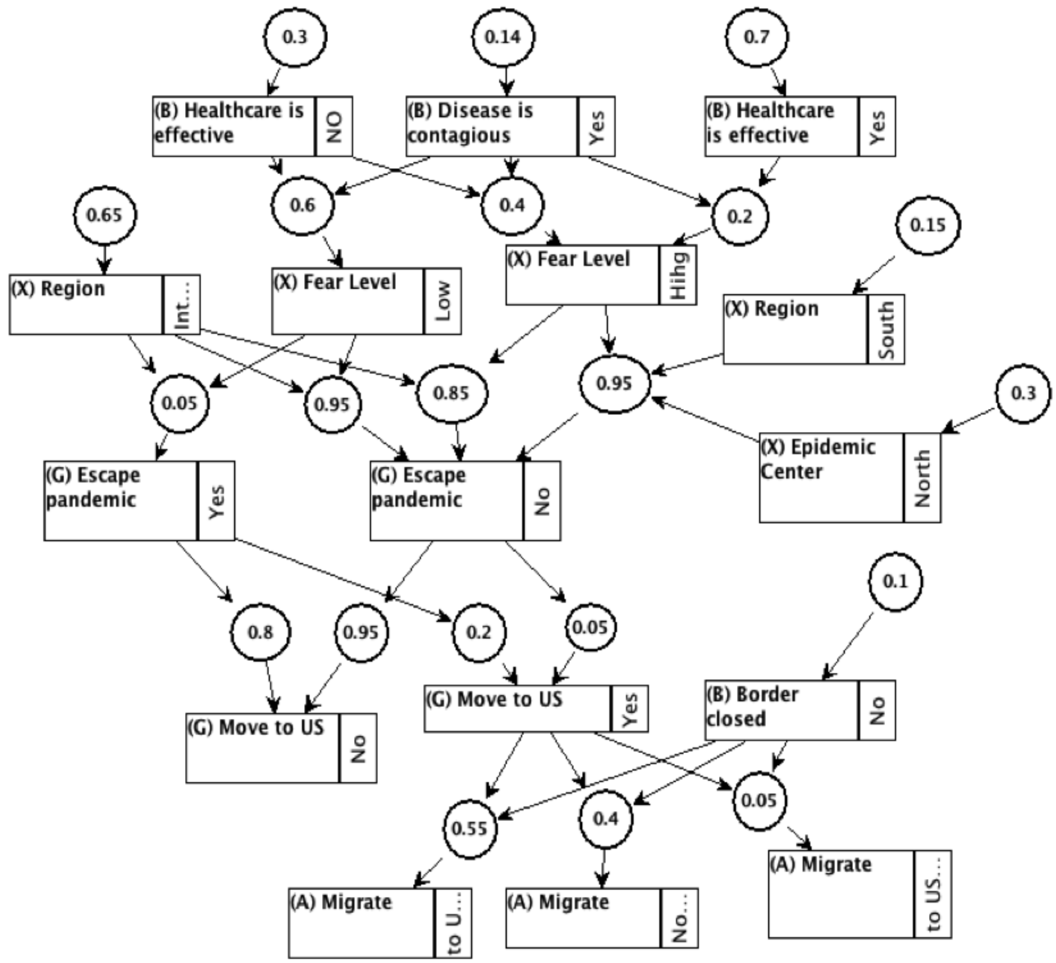

Figure 3.1: Sample BKB fragment from an intent framework modeling the real scenario

a reliability r_i depending on how much the person wants to adopt it. Then, new opinions are derived from the updated knowledge base. In contrast, traditional methods usually focus on quantifying opinion values on a certain issue first, and then integrate those opinion values directly to form a new opinion (as circled by the blue rectangle box). The limitations of such methods are threefold:

1. the dependencies and connections between knowledge fragments in a knowledge base are lost after the opinion acquisition process. In other words, they assume that different opinions on a given issue can be integrated independent of other knowledge fragments and whatever changes in other fragments will not affect the decision made on that given issue. Such assumption, unfortunately, is not true in the real world.

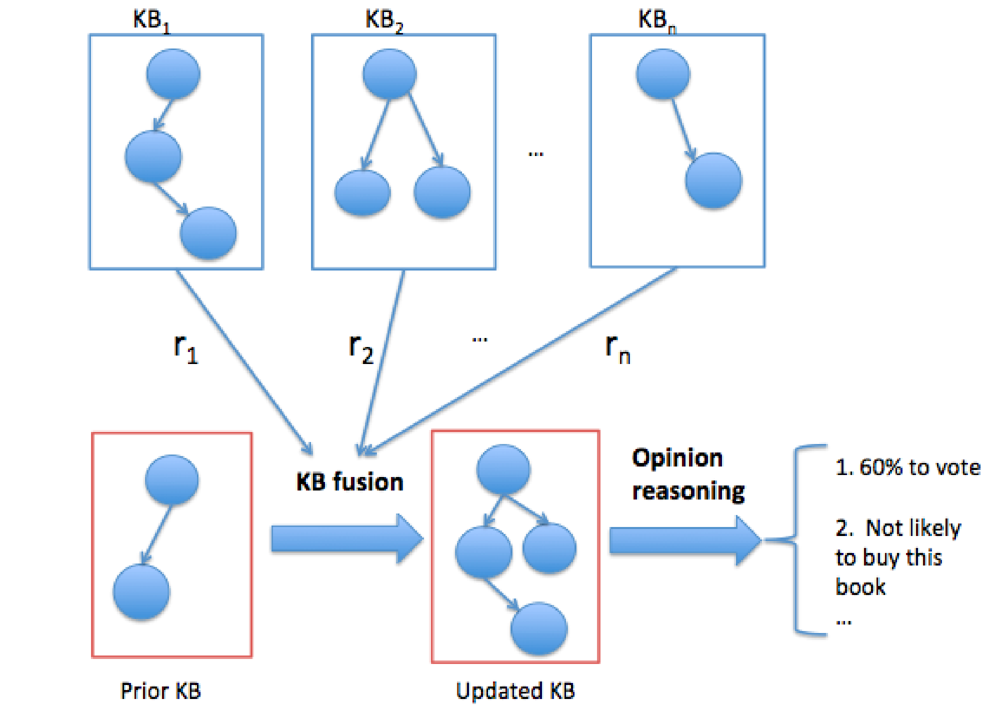

(a) The proposed opinion formation model

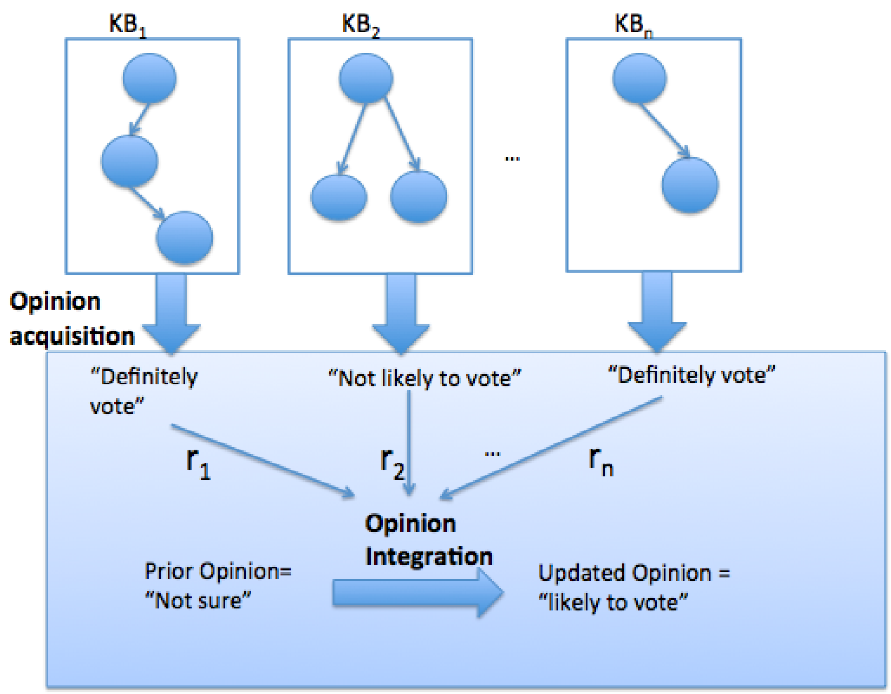

(b) Existing opinion formation models

Figure 3.2: Comparison of opinion formation process

2. As individual opinions are acquired in the first step, if modeling engineers want to model opinion formation in terms of different issues, they must have a clear picture of what opinions to collect in advance. However, the development of a knowledge-based system is fundamentally an incremental process. The domain knowledge engineers (subject matter experts) are inclined to start by selecting a subarea that they personally understand well, and then continuously introduce new information to the existing knowledge-based system as a basic cyclical process [88, 89]. Therefore, it is important to endow the model with some flexibility by making the knowledge integration process domain-independent, such that opinions on other topics/domains can be incrementally acquired from the updated knowledge base;

3. The third danger of the traditional design is that it is unable to explain where an opinion change comes from. As shown in Figure 3.2(b), by looking at the updated opinions only, it is impossible to decide if the change is due to the reliability variation or the knowledge base expansion/update. In contrast, as opinions in our model are reasoned from one's knowledge bases that is represented by a probabilistic model, we are capable of providing more insights on the reasons behind an opinion change from inference activities. In Chapter 5, we will present an approach to track and detect hidden sources by looking into the inconsistency in belief distribution.

3.3 Model of Opinion Formation

People do not like/dislike a thing for no reason. We treat all opinions as the reasoning results from ones' knowledge base. We assume that all knowledge fragments that serve as input to our model are valid BKBs.

From the concept of knowledge fusion, the individual knowledge base at time t ($t = [1 : T]$) can be viewed as an integration of the previous one and the knowledge fragments acquired through the interaction. A valuable property of knowledge fusion is that it considers the impacts from multiple sources when constructing explanations for any evidence observed. Thus, the opinions inferred from the fused BKB are not a linear combination of the original ones. The example in Appendix B shows that one's integrated opinion is actually a high-order polynomial function with the reliability of each knowledge fragment being indeterminate.

The structure of a knowledge base indicates the causal dependencies between different issues. For simplicity, we assume that everyone's initial BKBs share the same structure, but different distribution. The intuition is that people usually have common understandings on the causal relationships of basic issues. For instance, whether to bring a umbrella depends on the weather condition, not vice versa. However, we allow someone's knowledge base to be more complete than others. The trick is that if one's BKB misses one part of the information, we can always treat the corresponding conditional probabilities as 0, and the transformed BKBs are still valid. This is guaranteed by the power of BKBs in modeling incompleteness.

3.3.1 Mathematical Formulation

Given an actor 0 whose initial knowledge base is K_0^0 (K_t^i denotes the knowledge base for actor i at time step t), let p_t^0 represents the distribution of the BKB after fusing knowledge fragments from the people she contacted from $t-1$ to t, i.e. $K_t^0 = fusion(K_{t-1}^0, K_{t-1}^1, \ldots, K_{t-1}^n)$. Then, the opinion on a specific issue v, represented by the probability can be calculated from probability theory:

$$p_t^0(v) = \sum_j p_t^0(v|pa^j(v)) * p_t^0(pa^j(v)) \tag{3.1}$$

where $pa^j(v)$ denotes the jth parent set of v and $p_t^0(v|pa^j(v))$ is simply the weight of the jth S-node q_j pointing to v in K_t^0. Let $w_t(q)$ be the weight of the S-node q at time t. Then, Equation 3.1 becomes:

$$p_t^0(v) = \sum_j w_t(q_j) * p_t^0(pa^j(v)) \tag{3.2}$$

$$w_t(q_j) = p_t^0(v|pa^j(v)) = \frac{p_t^0(v, pa^j(v)))}{p_t^0(pa^j(v))}$$

For each S-node q_j, $pa^j(v)$ consists of all precedent I-nodes of q_j. So $pa^j(v)$ also contains the source node src_j indicating where a particular piece of knowledge comes from. Let $r_t(src_j)$ be the reliability of such a knowledge fragment at time t and let $\bar{pa}^j(v)$ be the collection of all precedent I-nodes of q_j, except src_j, i.e. $\bar{pa}^j(v) = pa^j(v) \backslash src_j$. Since the sources nodes are independent of any random variables in a fused BKB, we have

$$p_t^0(pa^j(v)) = r_t(src_j) * p_t^0(\bar{pa}^j(v)). \tag{3.3}$$

After combining Equation 3.2 and Equation 3.3, we have

$$p_t^0(v) = \sum_j w_t(q_j) * r_t(src_j) * p_t^0(\bar{pa}^j(v)). \tag{3.4}$$

Let r_t^i ($i = [0:n]$) be the reliabilities for K_{t-1}^i ($i = [0:n]$), r_t^0 is the normalized reliability for actor 0 herself. Then from Santos et al. [2], Equation 3.4 can be rewritten into

$$p_t^0(v) = \sum_i^n r_t^i \sum_{q_j \in \sigma_i} w_{t-1}(q_j) * p_t^0(\bar{pa}^j(v)), \tag{3.5}$$

where σ_i is the set of S-nodes pointing to v in K_{t-1}^i. Take the BKB fusion in Figure 2.2 as an example, let q_1 and q_2 be the S-node pointing to the I-node "Vote = Y" in actor 1 and 2's BKBs, respectively. Then in the fused BKB with new distribution p', the probabilities of $p'(\bar{pa}^1(v))$ and $p'(\bar{pa}^2(v))$ are identical and simply equal to $p'(\text{Interested in Politics} = Y)$. So following Equation 3.5, the probability of $p'(Vote = Y)$ can be calculated as:

$$\begin{aligned} p'(Vote = Y) &= r_4 * 0.7 * p'(\text{Interested in Politics} = Y) \\ + & \quad r_3 * 0.9 * p'(\text{Interested in Politics } = Y) \end{aligned}$$

Let $\rho_t^i = \sum_{q_j \in \sigma_i} w_{t-1} * p_t^0(\bar{pa}^j(v))$, then we simplified Equation 3.5 to:

$$p_t^0(v) = \sum_i^n r_t^i * \rho_t^i \tag{3.6}$$

where

$$\sum_i^n r_t^i = 1$$

In fact, ρ_t^i represents each actor's view on v in each of the input fragments given the distribution of v's the precedent variables being updated from p_{t-1}^0 to p_t^0.

So, what if $p_t^0(v)$ represents the prior probabilities of the variables with no precedent variables? From the fusion algorithm, the fusion of the prior probabilities is simply the weighted average of each input prior, i.e.

$$p_t^0(v) = \sum_i^n r_t^i * p_{t-1}^i(v) \tag{3.7}$$

3.3.2 Discussion

Equation 3.6 and 3.7 are key findings for our computational framework. The importance are threefold:

1. As ρ_t^i is calculated based on the fact that the distribution of v's parent variables being updated to p_t^0, it takes into account the opinion changes appeared on the issues that v is dependent on. Whatever happens on v's related issues will be reflected on the decision made on v. Moreover, if all variables are independent of each other, then as shown in Equation 3.7, our model will become equivalent to DeGroot-based models. This provides us the flexibility of modeling both simplified theoretical scenarios and complex real scenarios.

2. Though the marginal probability of $p_t^0(v)$ is in essence a high-order polynomial function of reliabilities (ρ_t^i is also dependent on the reliabilites), by following Equation 3.7, we manage to model $p_t^0(v)$ as a "weighted sum" of the reliability values, where the "weight" ρ_t^i can be easily calculated based on the inferencing results with respect to the distribution of v's precedent variables at time t. In the next Chapter, we will show how this feature enables us to transform the social influence learning problem into a regression problem and solve efficiently.

3. Similar to reasoning with BKB and BN, as $p_t^0(v)$ is independent of any descendent variables, the calculation of $p_t^0(v)$ can be performed in a top-down manner, where the prior probabilities are computed from Equation 3.7. The complexity of reasoning is same as BKB.

3.4 Nonlinearity in Opinion Formation

The main difference between DeGroot-based models and our model is that their opinion formation process is modeled linearly in the sense that the updated opinions are calculated as a linear combination of the adopted opinions, whereas our approach formulates opinion as a inferred product from a fused knowledge base that encodes the opinions from others in a nonlinear fashion (as shown in the example in Appendix B). In this section, we conduct analyses on the real data collected from Wikipedia[1] to evidently show the nonlinearity in opinion formation.

Wikipedia is a collectively authored encyclopedia where users can interact and discuss during the course of producing articles. To facilitate social interaction, each registered user is given a talk page that she and other users can edit in order to discuss updates to various articles on Wikipedia. We study how the interaction on Wiki's talk pages affects people's interests, where interests are extracted from users' edit history. Specifically, the interest on a particular topic is characterized by the edit frequency to the articles that belong to that topic. The intuition is that people tend to make more edits on the topics that they have strong interests. Similar representation can be found in Crandall et al.[23], in which they encode users' interests into a vector of edit activities.

The data[2] used in this study contains the complete edit history (all revisions, all pages) of all Wikipedia since its inception till January 2008. We construct the social network from Wiki talk pages as following: the nodes in the network represent wikipedia users and a directed edge from node i to node j represents that user i has edited a talk page of user j at least once. To reduce noise, we only select active users that have edit activities more than 10 times. The number of active users/nodes and edited pages are approximately 22000 (accounts for 0.9$\%$ of total users) and 440000, respectively. Anonymous edits are discarded, as they are recorded only by IP address that may reflect the activities of many users. Each wiki page may belong to one or more subcategories. We ignore the edits with no category information. We apply DBpedia Ontology 3.8[3] to map subcategories to 109 main topics, e.g. person, animal, sport, etc., as we only focus on users' interest on these main topics.

[1] http://en.wikipedia.org/wiki/Main_Page

[2] Datasets are available at http://snap.stanford.edu/data/wiki-Talk.html

[3] DBpedia Ontology is an open source cross-domain ontology. It is manually created based on the commonly used infoboxes in wikipedia. http://dbpedia.org/Ontology38

To study how a user's interests vary with the interests of others that he or she is connected to, we track each user's changing interests that are measured monthly from 2004 to 2008 January. For every two nodes (i, j) where there is an edge from i to j, we calculate the correlation between their interest trends with respect to each of the topics that they have discussed. Then we select 3 pairs with highest correlation to investigate the relationship between their interest changes. The idea is that if i and j are not correlated, then i must have little influence on j's interest/opinion. For each pair (i, j), to show that j's interest is not linearly correlated with i's interest, we perform a linear regression where j's interest value is the dependent variable and i's interest value is the independent variable. Given that j may be influenced by other neighbors, we depict the partial residual plot[4] [90] in Figure 3.3. The correlations with significance are presented in Table 3.1, ordered by correlation value.

Top correlated pairs	Correlation	p-value
1 (Geography)	0.8178	1.958E-05
2 (Technology)	0.7788	5.952E-06
3 (Sport)	0.6980	1.719E-04

Table 3.1: Correlation with significance

As shown in Table 3.1, all three pairs are significantly correlated. To understand how often does an individual's interest variation correlates with the ones they have interacted with, we calculated the percentage of the investigated pairs whose correlation is significantly greater than a threshold (with the p-value < 0.01). The percentage with different threshold is depicted in Table 3.2. 68.9% of the test cases are shown to be positively correlated with each other, which supports our hypothesis that social interaction will affect

[4]partial residual is defined as the linear component of the partial regression between the response variable Y and independent variable X plus the least-squares errors, i.e. $e_i^j = e_i + a_j X_{ij}$. The partial residuals can then be plotted versus X_j, where a_j is the slope of the simple regression of e_i^j on X_j.

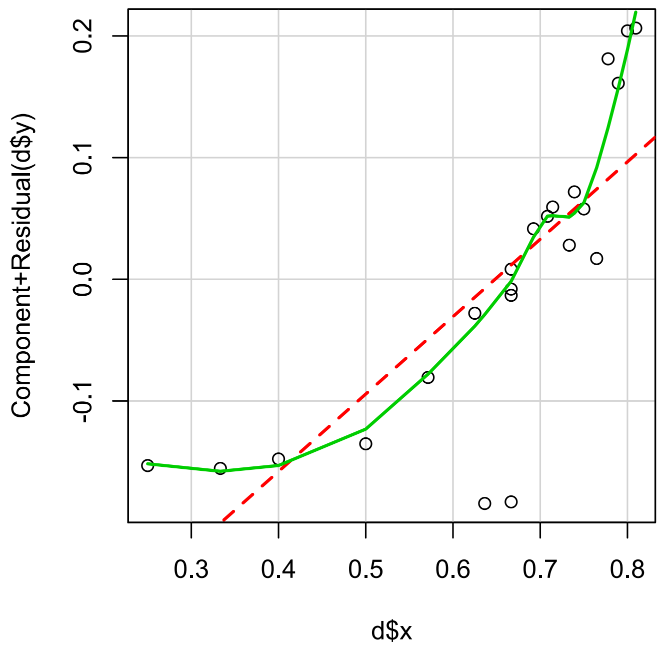

(a) UserY_id :247981, UserX_id :28107; Topic: Geography

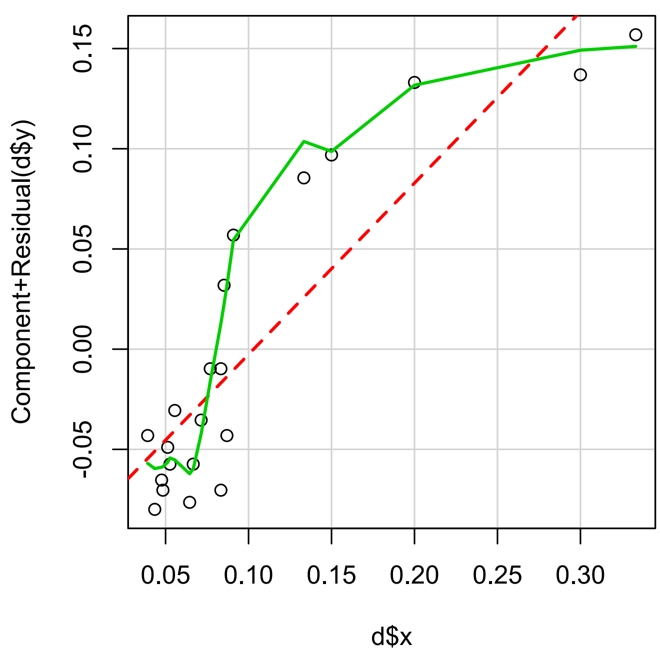

(b) UserY_id :149, UserX_id :5103; Topic: Technology

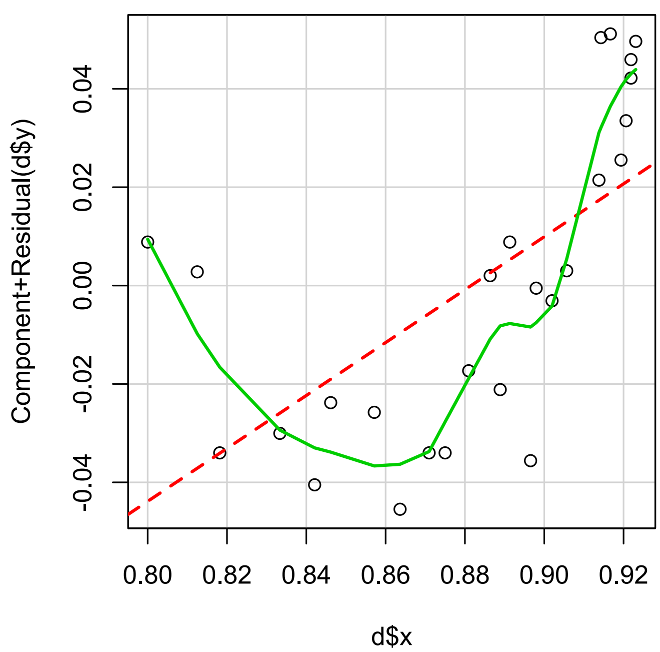

(c) UserY_id :1591, UserX_id :147410; Topic: Sport

Figure 3.3: Partial residuals plot

individuals' interests/opinions over time. Moreover, our finding is consistent with the results from Crandall et al. [23]. They examined the interplay between social interaction and interest similarity and show that in Wikipedia, people's interests will become more similar through communication.

The partial residual plot in Figure 3.3 shows that even when two interest trends are highly correlated, the relationship is not linear. In fact, if one's interest/opinion with respect to topic x at time step t is a weighted average of his/her neighbors' interests at previous time step, then we should expect to see:

$$P_t(x) \in [min(P_{t-1}(x)), max(P_{t-1}(x))], \tag{3.8}$$

where $min(P_{t-1}(x))$ and $max(P_{t-1}(x))$ are the min and max values of his/her neighbors' interests at time $t-1$. However, in Wikipedia, about 45% of the test cases actually violate the above condition. Among those, 15% of the violation is due to the fact that none of his/her neighbors have shown interests on x. This could be attributed to many factors: incomplete information, hidden influential sources or emergent behavior. In sum, people do not change their opinions as a linear integration of others'. The opinion formation process in real world is a complex process that needs to be addressed by a flexible platform that is capable of modeling different situations in one entire system.

Correlation threshold	Percentage of investigated pairs
> 0.5	16.5
> 0.4	26.8
> 0.2	60.3
>0	68.9

Table 3.2: Percentage of positively correlated interested trends

Chapter 4

Social Influence Learning

Social impact lies at the heart of individual's opinion formation. People obtain new information and update their knowledge bases through various learning. Among those, social learning is one of the most common means for people to gain knowledge. As suggested by Acemoglu and Ozdaglar [5], "social learning" takes place where people update their opinions by observing others' behaviors and experiences, by communicating with others about their beliefs, or by absorbing information from media sources. In the process of social learning, a person will neither fully adopt nor completely disregard the opinions of other people, but takes into account the opinions of others to a certain extent in forming her own opinion. However, how an individual adopts the opinions of others may vary in different situations. For instance, people are inclined to take the opinions from their close friends more seriously than from random people.

In modeling influence in social networks, efforts have been made on characterizing the influence of each individual on information diffusion based on the social network structure and interpersonal interaction [91, 60]. However, different from these approaches whose main goals are to target the influential people that can persuade a great number of people in their society, we emphasize on finding the rules of how one adopts another's opinion.

In this chapter, on top of the opinion formation model, we introduce a general impact metric that estimates the reliability for each adopted knowledge fragment based on a set of social influential factors. The individual opinion at each time step is computed from Equation 3.6 and 3.7. We show that though opinions from different people are not linearly fused, the relationship between the opinion change and influential factors can be transformed into a regression problem and be solved efficiently.

To understand the role of social influence in the formation of one's opinions, we start by considering three domain independent influential factors: friendship, recent contact frequency and opinion similarity that have been shown to have influence on how one takes other's opinions in social communication theories [33, 70, 34] (Details about social influence are discussed in Chapter 2). However, we will show that our model is not restricted to these factors. Newly discovered factors can be easily incorporated into the model.

4.1 General Impact Metric

Though there is evidence that reliability is correlated with these social factors, how do they interact and contribute to a reliability value is needed to be addressed.To model and learn the impact of each social factor computationally, we introduce a linear impact metric that combines these factors in a coherent scheme. Linear models have been widely used in trust-based approaches due to its simplicity and intuitiveness [35, 8].

At time t, let $X_t(i,j)$ denote the relationship closeness between actor i and j from actor i's view, scaled from 0 to 1, where 1 is the closest and 0 means totally unknown of each other. Let $Y_t(i,j)$ denotes the normalized contact frequency, i.e. number of contacts with j divided by the total number of contacts with all others from time $t-1$ to t and let

$Z_t(i,j)$ denotes the opinion difference on the target issue, i.e. $|p^i_{t-1}(v) - p^j_{t-1}(v)|$. Then, the reliability $r_t(i,j)$ that actor i assigns to the knowledge fragment sent by actor j, can be defined as:

$$r_t(i,j) = \alpha X_t(i,j) + \beta Y_t(i,j) + \gamma Z_t(i,j) + C \tag{4.1}$$

where α, β, γ and C are unknown weights/coefficients that need to be learned.

Now, let us reconsider Equation 3.6 and 3.7:

$$\begin{aligned} p^0_t(v) &= \sum_i^n r^i_t * \rho^i_t \qquad \text{(if } v \text{ has precedent variables)} \\ p^0_t(v) &= \sum_i^n r^i_t * p^i_{t-1}(v) \qquad \text{(otherwise)} \end{aligned}$$

We replace r^i_t with Equation 4.1. Then, Equation 3.6 turns into

$$\begin{aligned} p^0_t(v) &= \sum_{i=0}^n [\alpha X_t(0,i) + \beta Y_t(0,i) + \gamma Z_t(0,i) + C] * \rho^i_t \\ &= (1 - \sum_{i=1}^n [\alpha X_t(0,i) + \beta Y_t(0,i) + \gamma Z_t(0,i) + C]) * \rho^0_t \\ &+ \sum_{i=1}^n [\alpha X_t(0,i) + \beta Y_t(0,i) + \gamma Z_t(0,i) + C] * \rho^i_t. \end{aligned} \tag{4.2}$$

Rearranging the terms, we get

$$
\begin{aligned}
p_t^0(v) - \rho_t^0 &= \alpha \sum_{i=1} X_t(0,i)(\rho_t^i - \rho_t^0) \\
&+ \beta \sum_{i=1} Y_t(0,i)(\rho_t^i - \rho_t^0) \\
&+ \gamma \sum_{i=1} Z_t(0,i)(\rho_t^i - \rho_t^0) \\
&+ C \sum_{i=1} (\rho_t^i - \rho_t^0)
\end{aligned}
\tag{4.3}
$$

Similarly, Equation 3.7 now becomes:

$$
\begin{aligned}
p_t^0(v) - p_t^0(v) &= \alpha \sum_{i=1} X_t(0,i)(p_{t-1}^i(v) - p_{t-1}^0(v)) \\
&+ \beta \sum_{i=1} Y_t(0,i)(p_{t-1}^i(v) - p_{t-1}^0(v)) \\
&+ \gamma \sum_{i=1} Z_t(0,i)(p_{t-1}^i(v) - p_{t-1}^0(v)) \\
&+ C \sum_{i=1} (p_{t-1}^i(v) - p_{t-1}^0(v))
\end{aligned}
\tag{4.4}
$$

4.2 Parameter Estimation

Let n be the number of people in the social community under consideration. To model opinion dynamics, we look into the opinions over a number of time periods. Let $P = \{p_0^i, p_1^i, p_2^i, \ldots, p_T^i\}(i = [1:n])$ be a sequence of belief distributions generated over T time periods. Then, given the friendship network X and the contact network Y over time[1], the goal is to learn the factor weights α, β, γ and C from opinion change. Note, considering that the measure is taken from each actor's view, neither the friendship network nor the contact network needs to be symmetric.

Now that we have shown how to model the opinion change as a function of influential

[1] Opinion similarity network Z_t can be constructed from the opinion values at each time step.

$[\alpha, \beta, \gamma, c]$ = Influential-factor-weights (P, X, Y)

A $\leftarrow \emptyset$; H $\leftarrow \emptyset$
k $\leftarrow$ 1
for every instantiation/state $v; t = 1 : T$; actor$i = 1 : n$
    if v has no precedent variable
        $U(k) \leftarrow p_t^i(v) - p_{t-1}^i(v)$
        $d_{ij} \leftarrow p_{t-1}^j(v) - p_{t-1}^i(v) \quad (j = 1 : n, j \neq i)$
    else
        $U(k) \leftarrow p_t^i(v) - p_t^i$
        $d_{ij} \leftarrow p_t^j - p_t^i \quad (j = 1 : n, j \neq i)$
    end, if
    $A(k, :) \leftarrow [\sum_j X_t(i,j) * d_{ij}, \sum_j Y_t(i,j) * d_{ij},$
    $\sum_j |p_{t-1}^i(v) - p_{t-1}^j(v)| * d_{ij}, \sum_j d_{ij}]$
k $\leftarrow$ k + 1
end for
$[\alpha, \beta, \gamma, C]$ = regression (U, A)

Figure 4.1: Social influence learning from Gu et al. [1]

factors (Equation 4.3 and 4.4) for one particular issue/I-node v. We apply the rule to all possible I-nodes. As α, β, γ and C are only unknown parameters, we perform a linear regression with $p_t^0(v) - \rho_t^0$ being response variable ($p_t^0(v) - p_t^0(v)$ for Equation 4.4) and α, β, γ and C being independent variables. The complete algorithm can be described more formally in Figure 4.1:

4.3 Validation and Evaluation

In what follows, we present results of experiments that were carried out on both simulated data and a real world dataset. We start by introducing the baseline we compared with in the experiments.

4.3.1 Baseline

The DeGroot model is one of the classical opinion formation models, in which one forms her opinions by averaging over the opinions acquired through communication [49]. Specifically, they represent the opinion on a given issue by a real number. At time $t+1$, person i adjusts her opinion on issue v by taking a weighted mean of others' opinions. So if we represent the opinion value by the probability of her support for that issue, then the process can be formulated as:

$$p_t^i(v) = \sum_j^n r(i,j) * p_{t-1}^j(v)$$

where

$$\sum_j^n r(i,j) = 1$$

Other variations of this model include the Friedkin-Johnsen model [6], which assumes that a person would adhere to her initial opinion $p_0^i(v)$ to a certain extent. Though there are some other models in which the weight matrix may vary depending on time or opinion itself [52], their opinion formation process is still linear.

To compare with our work, we modify the DeGroot model by assigning the weight between agent i and j with the reliability value $r_t(i,j)$ generated by the impact metric introduced in the last section. Then, the influential factor weights can be transformed into a regression in a similar way to Equation 4.4. In fact, if there are no dependencies between different variables/issues, then our approach will devolve into the classical linear model.

4.3.2 Experiments on Synthetic Data

Data synthesis

To evaluate the effectiveness of our method, we first randomly generate an n-by-n friendship network X. Then the contact networks $Y = \{Y_1, Y_2, \ldots, Y_T\}$ for every time step are simulated based on the friendship network but with a certain deviation. In this experiment, the deviation is set to be 0.2 at maximum. The assumption here is that close people tend to be in contact more frequently. We also simulate each individual's initial opinion distribution $\{p_0^1, p_0^2, p_0^3, \ldots, p_0^n\}$ from m ($m \gg n$) predefined BKBs. At time $t(t = 1 : T)$, we fuse one's previous BKB K_{t-1}^i with the BKB fragments from the people she contacted, where the contacting information is gained from the Y_t. The reliability for each fragment is calculated using the general impact metric, where the factor weights are predefined for simulation, e.g. $[\alpha, \beta, \gamma, C] = [0.4, 0.2, -0.2, 0.2]$.

Evaluation of Social Influence Learning

In the first experiment, we test the performance of our approach on learning the influential factors through opinion dynamic modeling. Since the problem can be solved via regression, we use Relative Squared Error (RSE) to measure the error rate of the regression model. RSE works well when the errors of models are measured in different units. Let u_i, u_i' and $\bar{u}_i$ be the real opinion change (simulated opinion values), the predicted change and mean of the real values, then RSE can be defined as:

$$RSE = \frac{\sum_i (u_i' - u_i)^2}{\sum_i (\bar{u}_i - u_i)^2}$$

Also, to evaluate the prediction ability, we apply a 5-fold cross-validation on the regression.

With regards to the regression error We run multiple experiments to test our performance in terms of different social network size, initial knowledge bases and network connection degree. In particular, we first increase the number of nodes in the friendship network and contact network while maintaining other settings the same. We run each experiment for 10 times. Table 4.1(a) shows the average RSE using our approach and the baseline respectively. Then, we maintain the community size to 50 and vary the complexity of the predefined BKBs from which we generate the initial opinion distribution by increasing the number of the variables. By doing so, we test whether our method is capable of handling complex knowledge bases. The results are presented in Table 4.1(b). Last but not least, we measure the network connection degree as a summarized degree centrality. The connection degree is formally defined as:

$$C_D = \frac{\sum_{i=1}^{n} deg(i)}{(n^2 - n)/2},$$

where $deg(i)$ calculates the number of links that node i has. If the contact network is fully connected, then C_D equals to 1. The community size is set to 50 and the number of variables is set to 20. We show the regression error in terms of C_D in Table 4.1(c).

From Tables 4.1 (a), (b) and (c) we can see that the error of using our method is consistently smaller than the baseline. In addition, when there are more variables in the BKB, our advantages appear to be more significant (Table 4.1 (b)). We plot the error differences in terms of the number of variables in Figure 4.2 (a). The main reason is that a complex knowledge system is always accompanied by strong dependencies between knowledge fragments. The baseline model, however, fails to capture such dependencies, as it assumes that people's opinions on different issues are independent of each other. Moreover, we depict how error evolves as connection degree increases in Figure 4.2 (b). As shown in the figure, the error decreases as the connection degree gets closer to 1. Actually, a higher connection degree will lower the impact from each of the neighbor. So the errors from the inaccurate factor weights learning will be averaged out after fusing in a large portion of neighbors' opinions.

Network Size	Proposed	Baseline
10	0.0998	0.1347
20	0.0508	0.1019
30	0.1162	0.1358
40	0.0674	0.1509
50	0.0532	0.0840
60	0.1197	0.1684
70	0.0714	0.1361
80	0.0802	0.1389
90	0.0765	0.1359
100	0.1106	0.1390
Proposed vs Baseline	p-value	2.8576E-05

(a) network size

Number of Variables	Proposed	Baseline
5	0.0613	0.1322
10	0.0814	0.1311
15	0.1177	0.1485
20	0.0361	0.0739
25	0.0329	0.1497
30	0.0573	0.1326
35	0.1034	0.1834
40	0.0742	0.2123
45	0.0995	0.3219
50	0.0501	0.2501
Proposed vs Baseline	p-value	8.9047E-04

(b) number of variables

Connection degree	Proposed	Baseline
0.1	0.0594	0.1187
0.2	0.0637	0.1462
0.3	0.0679	0.1358
0.4	0.0572	0.0807
0.5	0.0388	0.0529
0.6	0.0356	0.0391
0.7	0.0194	0.0344
0.8	0.0201	0.0289
0.9	0.0118	0.1121
1	0.0056	0.0057
Proposed vs Baseline	p-value	0.0100

(c) connection degree

Table 4.1: Error comparison of two approaches in terms of different parameters

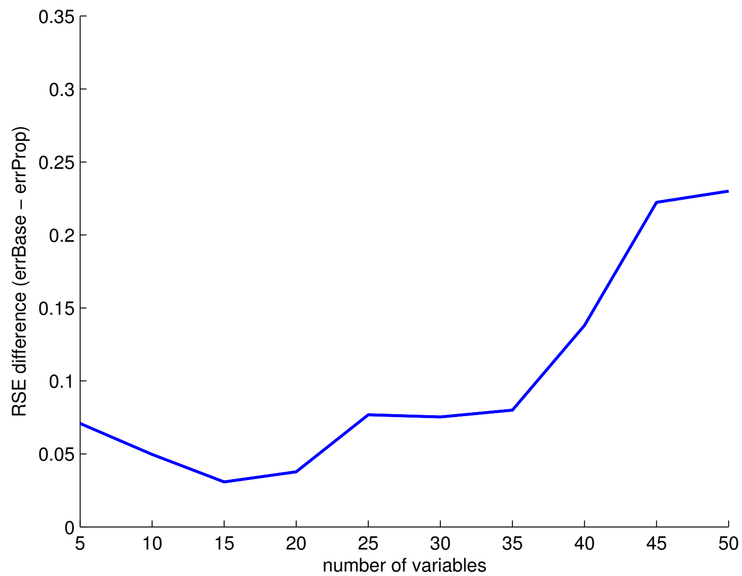

(a) the complexity of the initial knowledge bases

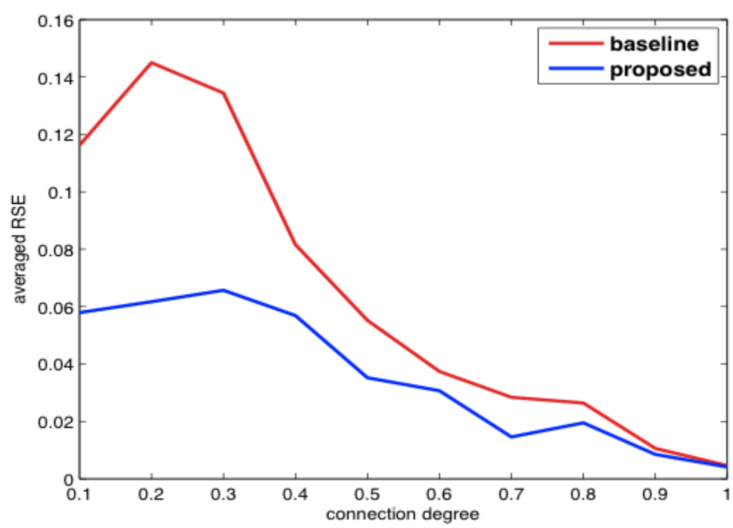

(b) connection degree

Figure 4.2: Comparison of two methods with respect to different parameters

With regards to the coefficient similarity Another way to evaluate our learning performance is to directly measure the similarity between the factor weights learned from two methods against the true predefined weights. As the factor weights are represented by a vector, we apply cosine similarity to measure the accuracy, i.e.

$$sim = \frac{v_1 \cdot v_2}{||v_1||||v_2||}.$$

Table 4.2 reports the average cosine similarity with standard divination to show our learning capability. As we can see, our approach outperforms the baseline model, which gives us confidence on using our results as a guide to control or predict future opinion dynamics.

Cosine Similarity	Baseline	Our Model
Mean	0.5867	0.7438
Std	0.0417	0.0684

Table 4.2: Average similarity with the true factor weights

Once we have the factor weights learned from the opinion trend, we can apply them to study and even predict the future opinion dynamics. We conduct another experiment to evaluate how close the opinion distribution trend simulated using the learned factor weights are to the true trend. Details about this experiment can be found in Appendix C.

Convergence Analysis

As we discussed earlier, if there is no dependency between different issues or knowledge pieces, then the way we transform the problem will become the same as the way we deal with the baseline. However, it is not true in the real world, especially when the knowledge bases evolve over time. To better understand the connections between our method and the baseline and to answer whether our method is compatible with the classical model, we examine under what situation will our method produce similar results as the baseline.

Lemma 1 If everyone's initial BKBs share the same conditional probabilities, i.e. $p_0^i(v|pa(v)) = p_0^j(v|pa(v))$, then our approach becomes equivalent to the baseline.

Proof, from Rosen et al. [75], for actor i, the probability on a given issue v can be calculated as the summation of the probabilities of all possible inferences where v is true. So at time t, we have:

$$p_t^i(v) = \sum_{\tau \in I_v} p_t^i(\tau) = \sum_{\tau \in I_v} \prod_{q \in \tau} w_t^i(q)$$

where I_v represents the set of inferences containing v, $w(q)$ is the weight of the S-node q. Then, the aggregated opinion at time $t+1$ using baseline method will be:

$$p_{t+1}^i(v)_{baseline} = \sum_{j=1}^{n} \gamma_j * p_t^j(v) = \sum_{j=1}^{n} \gamma_j \sum_{\tau \in I_v^j} \prod_{q \in \tau} w_t^j(q)$$

Considering that our opinion formation method is based on BKB fusion, for every I-node in an inference, the weight of its supporting S-node in the fused BKB can be viewed as the average over the weights of the corresponding S-nodes in the input BKB fragments, i.e. ${w'}_t^i(q) = \sum_{j=1}^{n} r_j * w_t^j(q)$. So using our method, $p_{t+1}^i(v)$becomes

$$p_{t+1}^i(v)_{proposed} = \sum_{\tau \in I_v} \prod_{q \in \tau} \sum_{j=1}^{n} \gamma_j * w_t^j(q)$$

Given the condition that $w_t^i(q) = w_t^j(q)$, we can now derive:

$$p_{t+1}^i(v)_{baseline} = \sum_{\tau \in I_v^j} \prod_{q \in \tau} w_t^j(q) = p_{t+1}^i(v)_{proposed}$$

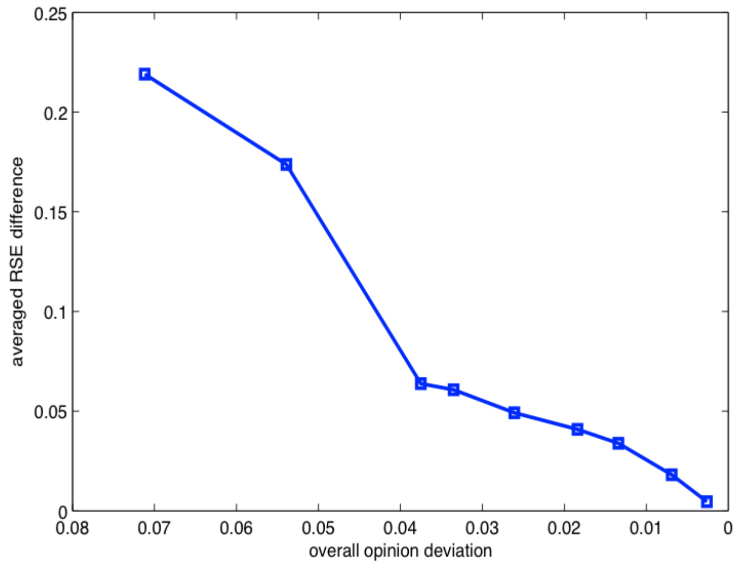

Figure 4.3: Mean RSE difference between two models

To illustrate Lemma 1, we run an experiment by decreasing the deviations between the conditional probabilities in each BKB. Figure 4.3 shows the mean error difference of using two methods, where the error difference is defined as *proposed_rse-baseline_rse*. As we can see, the differences between two methods shrink when the distributions of different initial BKBs get closer.

So, what can we gain from lemma 1? In fact, in a closed social environment, if there is no new information coming in, then both of our method and the baseline will lead to opinion convergence, where the convergence rate highly depends on the social connection degree. From Lemma 1, if people's opinions are close to each other, e.g. converge to certain values, then the advantage of our methods will diminish. (Note, our performance will never be worse than the baseline.) We test the convergence rate by measuring the normalized square of interpersonal opinion difference at each time step t, i.e.

$$OD_t = \sum_{i,j} \sum_{v} (p_t^i(v) - p_t^j(v))^2 / n^2 V,$$

where V is the total number of states/issues in the BKB. We assume that everyone's initial BKBs have different distribution. Then, to examine how the social connection degree affects the convergence rate, we choose three different connection degrees: $C_D = 0.2, 0.5$ and 0.8 and the results are plotted in Figures 4.4. As shown in the figure, the convergence rate increases with the connection degree. This finding also explains why in Figure 4.2 (c), the performance difference decreases when C_D gets larger. Additionally, the opinions converge to different values when the degree is low, e.g. 0.3. The fact is that different opinion groups are formed when people only communicate with the people in a subgroup. Moreover, on closer examination of the convergence trend, we see that the opinions formed using our method converge slightly faster than the baseline. The reason is that we do not form opinions linearly but considers the relationships between variables. So the distribution characteristics of higher-level variables could be propagated to the lower-level variables, which speeds up the convergence.

In reality though, changes in people's opinions are gradual [92]. Plus, people keep absorbing new information from external sources and keep making new friends. So, it is very unlikely at any point that people hold the same opinions, which makes the conditions described in Lemma 1 hardly to satisfy. This finding is valuable. It makes our approach more applicable to real-world problems.

4.3.3 Experiments on Real Social-network Data

To better understand the opinion dynamics in the real world, it is crucial to test our approach on a real dataset. In particular, we conduct our experiment on the adoption of political opinions through social communication. People's prospects on political issues, especially voting behaviors can be influenced by various information sources, e.g. media and social interaction. Madan et al. [9] studied how opinions about political candidates and parties, and voting behavior, spread through a social network using mobile sensing

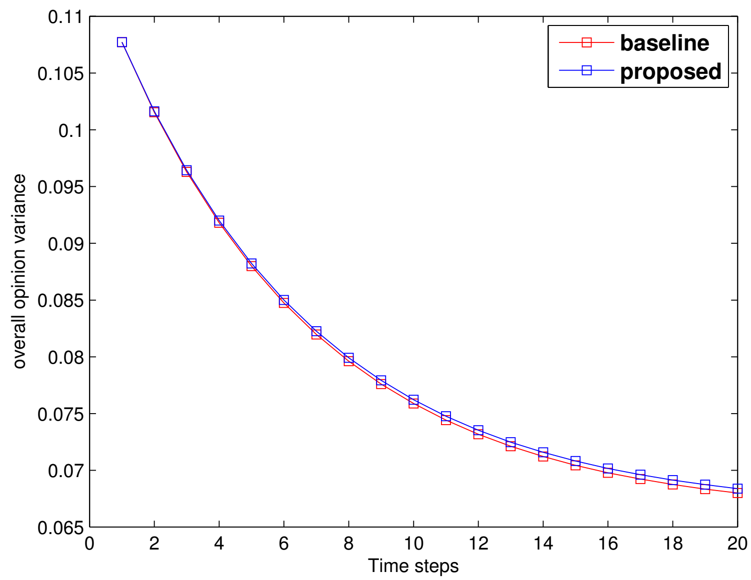
baseline
proposed
overall opinion variance
Time steps
0.11
0.105
0.1
0.095
0.09
0.085
0.08
0.075
0.07
0.065
0
2
4
6
8
10
12
14
16
18
20

(a) $C_D = 0.2$

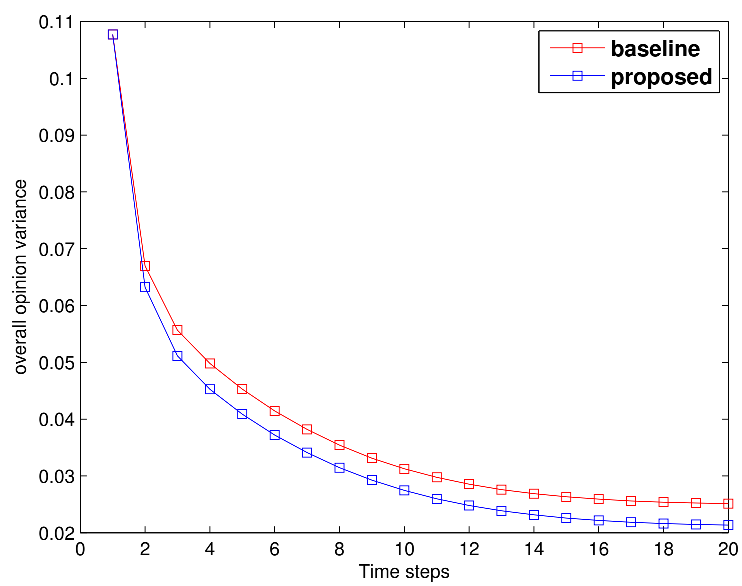
baseline
proposed
overall opinion variance
Time steps
0.11
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0
2
4
6
8
10
12
14
16
18
20

(b) $C_D = 0.5$

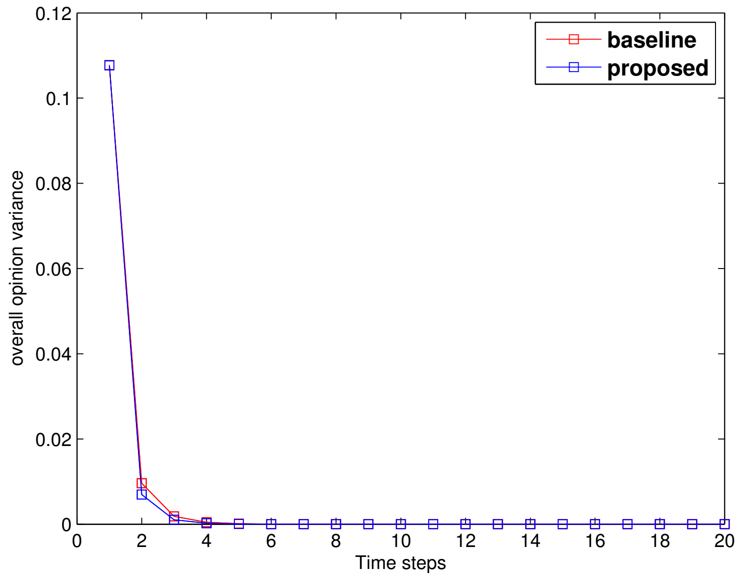

(c) $C_D = 0.8$

Figure 4.4: Convergence rate comparison between two methods given different connection degree.

data. The study was conducted among approximately 70 undergraduate students living in the same dormitory. They built up a mobile sensing platform to collect the communication data varies from Bluetooth signal transmission to phone call. The political opinions were captured using three monthly surveys before and after the presidential election in 2008. More descriptions about the data can be found in Madan et al. [9].

To model opinion changes on different issues, we select four questions from the survey. The questions and the possible responses are listed in Table 4.3. Then, we use Bayes Net Toolbox[2] to capture the causal dependencies between variables. The data we used for structure learning contains survey results over three months. The assumption is that the dependencies between those four common political issues are less likely to change. The learned causal structure is depicted in Figure 4.5. As we expected, the candidate that a person decides to vote for is highly dependent on her preferred party and how likely a student

[2]BNT is a toolbox for learning Bayesian network's parameters and structure. Here we only use BNT to learn the causal structure. https://code.google.com/p/bnt/

is going to vote is dependent on how much interest he/she have in politics. Next, we convert each of the questions into a binary state variable. For each student, her knowledge base is built based on the learned causal structure. Then, the probabilities/opinions for every variable/issue at each time step can be estimated from her survey answers. For example, if one is very interested in politics, then her probability for "interested in politics = Yes" is assigned to 0.9. Also, if one decides to vote for McCain, then her probability for "vote for Obama = Yes" will become very low. The extracted I-nodes/states and example probability values are listed in Table 4.4. In a later experiment, we will show that our results are not sensitive to the predefined values.

Question	Possible Answers
Interested in politics	[Very, Somewhat, Not-at-all]
Preferred political party	[Republican, Independent, Democrat]
Likelihood to vote	[Not-vote, Not-sure, Vote]
Who to vote	[Probably John McCain, Definitely John McCain, Undecided, Probably Barack Obama, Definitely Barack Obama]

Table 4.3: Selected Political Survey Questions and Possible Answers

Now that we have everyone's opinion distribution over time, the next step is to generate the friendship network X and contact network Y. Aside from the communication data and political view data, Madan et al. [9] also provides a survey of the relationship that a subject indicates she has with another, e.g. close friends or political discussant. We quantify the friendship using a single value between 0 (stranger) and 1 (extremely close). As for the contact network, we start with modeling only face-to-face contact using *proximity* data[3], where the weight between two nodes is calculated as the the normalized contact frequency in a certain period. To reduce the noise from the situations that two students are physically

[3]Proximity data captures bluetooth transmission between cell phones. It indicates that two person are within 10 meters at the time of the record [9]

I-nodes	Original Answers	
Interested in politics = Yes	Very	0.9
	Somewhat	0.4
	Not-at-all	0.1
Preferred political party is Republican = Yes	Republican	0.8
	Independent	0.5
	Democrat	0.2
Vote = Yes	Not-vote	0.1
	Not-sure	0.5
	Vote	0.9
Vote for Obama = Yes	Probably John McCain	0.3
	Definitely John McCain	0.1
	Undecided	0.5
	Probably Obama	0.7
	Definitely Barack Obama	0.9

Table 4.4: Predefined probability values based on the survey answers

close to each other but on different floors, only the records with the probability of on the same floor being greater than a threshold ζ are counted.

We test our approach on the real data and compare the results with the baseline. Table 4.5 shows the average error from cross-validation. As we can see, the error produced from our method is significantly smaller than the baseline. Moreover, we analyze the learned factor weights between two methods. As plotted in Table 4.6, the value of -0.0381 from our method suggests a negative effect of the opinion difference factor, which fits well with our expectation that people are more willing to take opinions that are similar to theirs. In contrast, the baseline model produces a positive coefficient for opinion difference, but negative weights for friendship and the contact frequency, which fails to agree with the social theory findings. Also, the regression coefficients learned using the baseline model are

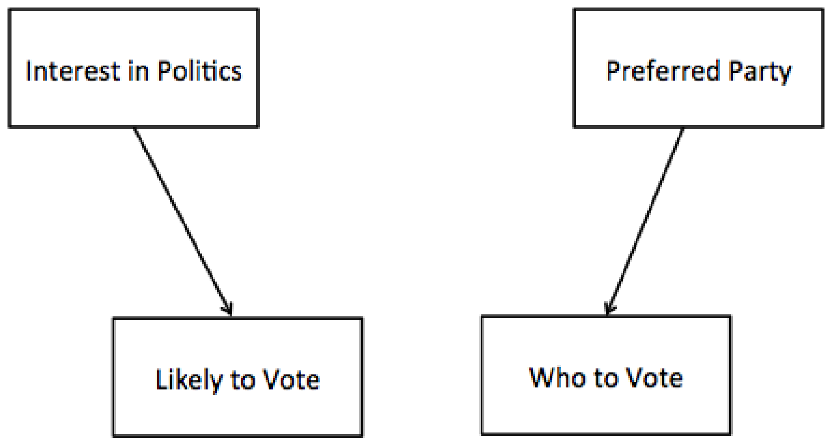

Figure 4.5: Learned causal dependencies between four political issues

not statistically significant, which suggests that it is not appropriate to use those coefficients to interpret the relationship between opinion change and social factors. Besides, in this particular case, the influence of the friendship appears to be the highest , whereas the impact from frequent contact is less significant. This indicates three things,

1. For the undergraduate students, their political views are more likely to be affected by their close friends.

2. Two persons might be in proximity very often without communicating to each other or even realizing the presence of each other.

3. Their opinions could be influenced by other communication means other than face-to-face.

So, our next task is to apply our model on contact networks that are formed by alternative communication means. In particular, we compare with the contact networks built from *call* data[4] and *SMS* data[5], respectively. Both networks are constructed in the same

[4]This data is collected based on the call log containing when and who has called whom.
[5]SMS data captures when and who has sent short messages to whom.

	Mean Error	Std
Proposed	0.5432	0.0478
Baseline	0.6467	0.0698
Significance (p-value)	Proposed vs Baseline	0.0021

Table 4.5: Average error comparison between two models

	α		β		γ		C	
	Weight	p-value	Weight	p-value	Weight	p-value	Weight	p-value
Proposed	0.0768	0.0453	0.0134	0.109	-0.0381	0.0028	0.1234	3.41E-04
Baseline	-0.1389	0.9576	-0.0001	0.6234	0.0142	0.2930	0.1138	3.82E-04

Table 4.6: Factor weights learned from two methods

way as proximity. The results in terms of regression error and learned coefficients are presented in Table 4.7 and 4.8. As we can see, the results of voice call (both error-wise and learned coefficient-wise) are slighter better than the proximity, but the impact of frequent phone call is still less than the friendship. Voice call is a strong indicator of the relationship between two person. However, like proximity, phone call only accounts for a fraction of one's daily contact. Many people often socialize with non-friends who may never call each other. The impact from SMS contact is the lowest among the three communication means. The reason is that comparing to voice call, SMS is a more convenient and easier way for people to communicate, especially when they are not friends at all. The statistics from Time mobility poll[6] suggests that 32% of all respondents would rather communicate by text than phone. Therefore, a frequent contact with someone via SMS may not suggest a strong social bond. Moreover, it is less likely for people to share their thoughts or get influenced through SMS than face-to-face chat or phone call. Psychologist Sherry Turkle believes that having a conversation with another person is an important process that teaches people to think, reason and self-reflect[93].

[6]http://www.qualcomm.com/media/documents/time-mobility-poll-cooperation-qualcomm/

	Mean Error	Std
Call network	0.4632	0.0415
SMS network	0.6128	0.0501

Table 4.7: Average error comparison between two contact networks

	α		β		γ		C	
	Weight	p-value	Weight	p-value	Weight	p-value	Weight	p-value
Call network	0.0801	0.0198	0.0277	0.0943	-0.0312	0.0009	0.1467	7.19E-05
SMS network	0.0674	0.0581	0.0127	0.1590	-0.0416	0.0015	0.1278	1.41E-04

Table 4.8: Factor weights learned by the proposed approach from two communication means

Last but not least, we test whether our approach is sensitive to the predefined conditional probability values. We maintain the order of the answers but vary its value. For example, the probability of "very interested in politics" is always higher than "not at all". We randomly generate the predefined values and run experiment for 100 times. The mean and standard deviation of the error are presented in Table 4.9. As we can see, the result of our model is close to the one in Table 4.5, which suggests that no matter how we quantify the survey answer, as long as the assigned probabilities can distinguish different levels of likelihood, our approach is robust to variation in initial settings.

	Mean Error	Std
Proposed	0.5718	0.0577
Baseline	0.6830	0.0897

Table 4.9: Performance comparison with different initial settings

There can be tons of factors that determine how an individual takes others' opinions.

In this paper, we only consider three factors that we believe are playing an important role in opinion dynamics. However, we show that our method is not restricted to those three factors. Table 4.10, shows the results of using only two factors out of three. The error of using factors “Friendship” and “Opinion Difference” appears to be smaller than the other two combinations, which is consistent with the previous finding that contact frequency is a weaker indicator of social influence. Moreover, it is not hard to see that our method outperforms baseline in all factor combinations. Thus, even when new factors are available, our method can be easily extended to accommodate new influential factors.

	[F,OD]	[F,C]	[C,OD]
Proposed	0.4936	0.5901	0.6130
Baseline	0.6532	0.6203	0.7509
Proposed vs Baseline (p-value)	3.404E-05	0.0025	6.325E-06

Table 4.10: Average Error using Different factor sets (F = Friendship, C= Contact, OD = Opinion Difference)

4.4 Summary

Though how the knowledge bases are expanded and updated through social interaction is a complex procedure, there may exist a core mechanism that guides people to adopt others' opinions or to absorb new information. The goal of this chapter is to uncover the rules behind opinion dynamics by starting with several influential factors that could affect the selection of opinions. We illustrated how to formulate the problem using a mathematical model and solve the factor weights via regression. We conducted multiple experiments on synthetic data. The results validate the effectiveness of our approach over the baseline. A thorough analysis and discussion is provided on the compatibility of our approach. Another main contribution of this part of work is that we tested our approach on modeling the political opinion change using a real world dataset. The error produced by our approach is significantly smaller than the baseline. Besides that, the factor weights learned from our approach are also consistent with social theory.

Chapter 5

Hidden Influential Source Tracking and Detection

5.1 Overview

As we emphasized in a couple of previous chapters, a person's beliefs are key elements for inferring the meaning of opinions held by individuals and groups and for predictions of future behaviors. These elements/perceptions, however, can be dynamic and affected by external sources over time through social interaction or exposure to information, e.g. media messages [94]. Meanwhile, such influencing sources may have qualitatively different effects depending on how likely an individual adopts the opinions of the sources, which vary in different situations as well [95]. For example, in the context of socialization of children, a child who has a strong bond with his family is inclined to take parental attitudes and actions with full trust. In contrast, people only selectively accept the arguments and views supported by online news sources, e.g. consumer review sites. Opinions adopted with different reliabilities will differ in terms of their qualitative characteristics, and affect a person's subsequent behavior. Moreover, the impact trend of each influencing sources may follow a certain pattern that can be used to make predictions. For instance, after an event breaks out, its social effect typically decays as time goes by [96, 76]. So a sudden rise of the impact value is either indicative of a new event or a sign of anomaly. Therefore, it is critical to understand the role that external sources play in opinion and behavior change at each time period, such that we can provide more insights and explanations on the observed changing-behavior and further answer the questions like: Why did the level of illegal migration from Mexico to the US increase sharply after April 29, 2009? What determinate factors make people want to escape from Mexico? Will that factor (continue to) cause panic?

However, the characteristics of external sources that affect people's opinions are rarely open to the public due to a variety of reasons – private communications, what one randomly reads or hears, and implicit social hierarchies, to name a few. Likewise, it is impossible to track how people view and adopt the opinions held by each of the sources they have interacted with. Such information can be concealed subconsciously when the influence is subtle or the reliability is not quantifiable, whereas sometimes people will intentionally conceal this information. For example, terrorists tend to protect criminal organizations by hiding their connections with the group. Recently, statistical-based studies on detecting influencing sources that cause behavior change has begun to emerge, particularly in the area of event detection and anomaly detection [97, 98]. However, even these advanced methods fail to distinguish whether one's changing-behavior is caused by the new influencing sources or the evolution with respect to the actual impact of some existing/internal sources. In a complex real-world scenario, it becomes a significant challenge to have a model that is compatible with state-of-the-art belief representation approaches, and still be able to detect hidden sources and capture the evolution of their impact levels.

Anomaly detection has been applied to detect the presence of any observations or patterns that are different from the normal behavior of the data. Approaches based on Bayesian

Networks include detecting anomalies in network intrusion detection [99] and disease outbreak detection [100]. The typical approach in BN-based anomaly detection is to compute the likelihood of each record in the dataset and report records with unusually low likelihoods as potential anomalies. Different from these approaches whose main goals are to achieve early detection and identify anomalous change in terms of a probability distribution [97], we focus on detecting the reasons behind the behavior change. Moreover, many statistics-based anomaly detection methods only focus on detecting events whose patterns are anomalous enough to be distinguishable from normal data. Furthermore, they overlook the situation when certain external opinion sources that have subtle influences at present, but may cause a butterfly effect later, as triggered by other events. It happens in the real world when some less substantial events become the key clue for analyzing the future behavior change. We show that our work overcomes the above limitations by being able to track influencing sources even when the impacts are small.

There are some other techniques that attempt to handle changing belief networks. Methods based on learning Dynamic Bayesian Networks (DBNs) [101] have provided mechanisms for identifying conditional dependencies in time-series data, such as for reconstructing transcriptional regulatory networks from gene expression data [102] and speech recognition using Hidden Markov Model (HMM) [103]. Nevertheless, most DBN implementations assume for the sake of efficiency that the Markov property holds for the domain they represent, which restricts knowledge engineering by requiring that the probability distribution of variables at time t depends solely on the single snapshot at time $t-1$. Thus, for real world cases when the future outcomes are highly dependent on the hidden factors whose prior information is not identified, we need another model that can easily express such abstract temporal relationships. Process Query System (PQS) [104] is an advanced tracking system designed to determine which processes produced which events. Unfortunately, its detection strategy is based on the observable events generated by hidden states, which may not be available in our case.

Mixture models have been used in modeling opinions of populations. For instance, Hill and Kriesi [98] apply a Finite Mixture Model to support their theory of opinion-changing behavior, where the attitude of each of the group is represented by a distribution and the mixed distribution is described by a weighted aggregation of n different distributions. However, the Expectation-maximization (EM) [105] based mixture decomposition methods show propensity to identify local optima [106], which makes it also sensitive to initial guesses. In addition, the separation of parameter estimation and component identification increases the probability of converging to boundary values when the number of model components exceeds the true one [107]. These considerations led us to develop a variant mixture model that is suitable for our problem of detecting hidden belief sources by taking advantage of time-varying information, as well as loosening the requirement of a predefined number of sources.

In this chapter, we introduce a new modeling approach called a Finite Fusion Model (FFM) for detecting and tracking hidden sources in a time-variant scenario that consists of a sequence of opinions encoded in BKBs. Specifically, we build FFM on the basis of the opinion formation model introduced in Chapter 2, in which the formation of individual opinion at each time period is treated as a process of aggregating opinion/information from different sources. Then, FFM leverages BKB fusion to model the integrated belief distribution by taking into consideration the impact of hidden sources. The latent parameters that characterize the distribution of the underlying hidden sources and the corresponding impact weights are learned via a constrained optimization problem.

5.2 Source Tracking and Detection Approach

The goal of this chapter is to analyze the behavior change over time by detecting hidden influencing sources and tracking the corresponding impact patterns. However, without a sound theoretical foundation, the methods developed will simply be ad hoc. Before formally introducing the model, we first explain several key observations that motivate the model:

<u>Observation 1</u>. Studies on social influence have shown that one's opinions will be affected by external ideas through social interaction [108]. In many situations, a person will not accept these external ideas in total but only adopts the pieces that fit into her own situation [4].

<u>Observation 2</u>. The deeply ingrained opinion is the foundation of one's behavior and should not change dramatically within a short time [109]. This results in a natural expectation that the impacts from other sources will sequentially affect one's opinions.

5.2.1 Mathematical Formulation

Motivated by Observation 1, the individual opinion at time t, ($t = [1 : T]$) can be viewed as an integration of the previous opinion and certain opinions held by hidden sources that contain both new sources and existing sources, where the integration of the opinions from existing sources can be viewed as a reinforcement of their relative reliabilities/impacts. Moreover, from the opinion formation model introduced in Chapter 3, individual opinions can be treated as a reasoning product from one's knowledge base that keeps updating by fusing new knowledge fragments. So, in this study, we assume that the initial knowledge bases and influential knowledge fragments that serve as input to our algorithm are represented by valid BKBs, where the influential knowledge fragments are acquired from hidden sources. We also assume that all beliefs and opinions that serve as input to our algorithm are valid BKBs, where the influential knowledge fragments are acquired from hidden sources.

Let $H = \{h_1, h_2, \ldots, h_m\}$ be the collection of all possible hidden fragments that could potentially affect a person's opinion across the time. The individual belief distribution p_t, encoded in the fused BKB K_t, at each time step t is defined as a *finite fusion* of the previous belief distribution p_{t-1} and m hidden fragments:

$$p_t = fusion(p_{t-1}, H, w_t) \tag{5.1}$$

where $w_t = [w_{t,1}, w_{t,2}, \ldots, w_{t,m}]$ is the impact vector representing the impact value of each hidden source at time t. Note, $w_{t,i} = 0$ means that the source h_i has no effect at time t.

Then, for each of the I-node/state v in K_t, Equation 3.5 can be rewritten into:

$$p_t(v) = \frac{\sum_{q_j \in \sigma_0} p_{t-1}(q_j) * p_t(pa^j(v)) + \sum_{i=1}^{m} w_{t,i} \sum_{q_j \in \sigma_i} h_i(q_j) * p_t(pa^j(v))}{Z_t} \tag{5.2}$$

$$Z_t = 1 + \sum_{i=1}^{m} w_{t,i}$$

where σ_i is the set of S-nodes pointing to v in h_i and σ_0 is the set of S-nodes pointing to v in K_{t-1}. $h_i(q_j)$ characterizes the probability distribution for v conditioned upon v's parent $pa^j(v)$, i.e. $h(q_j) = \frac{h(v, pa^j(v))}{h(pa^j(v))}$. Z_t is a normalizer so that the weights of all sources for a given random variable do not exceed 1.

5.2.2 Parameter Estimation

Given a sequence of belief distributions $\{p_0, p_1, p_2, \ldots, p_T\}$ generated over T time periods, the goal is to learn the probability distribution for each of the hidden sources $h_i (i = 1 : m)$, as well as its time varying impact $w_{t,i} (t = 1 : T)$ with no prior knowledge.

Single-source Tracking

We start by solving a simpler problem when the behavior change across time is caused by the evolution with respect to the impact of a single source h. We denote its impact values across the entire time sequence by an impact vector $\vec{w} = [w_1, w_2, \ldots, w_T]$. Then, Equation (7) can be simplified as:

$$p_t(v) = \frac{\sum_{q_j \in \sigma_i} p_{t-1}(q_j) * p_t(pa^j(v)) + w_t \sum_{q_j \in \sigma_i} h_i(q_j) * p_t(pa^j(v))}{1 + w_t}, \forall t \in [1, \ldots, T] \tag{5.3}$$

As shown in Equation (8), the opinion value $p_t(v)$ depends on the distribution of v's parent variables at time t. So what if $p_t(v)$ represents the prior probabilities of the variables with no precedent variables? Similar to how Equation 3.7 is derived , for $p_t(v)$ that represents the prior probability of the variable with no precedent variables, the above equation simply becomes:

$$p_t(v) = \frac{p_{t-1}(v) + w_t h(v)}{1 + w_t}, \forall t \in [1, \ldots, T] \tag{5.4}$$

To fully represent the joint probability distribution of hidden source h, we consider all possible I-nodes in K_t. Let k be the number of states and let $p_i^t = p_t(v_i)$, then we have

$$p_i^t = \frac{\alpha_i^t + w_t \sum_{q_j \in \sigma_i} h(q_j) * \beta_{ij}^t}{1 + w_t}, \forall i \in [1, \ldots, k] \tag{5.5}$$

where

$$\alpha_i^t = \begin{cases} \sum_{q_j \in \sigma_i} p_{t-1}(q_j) * p_t(\bar{pa}^j(v)) & \text{if } v \text{ has precedent variables} \\ p_t - 1(v) & \text{Otherwise} \end{cases} \tag{5.6}$$

and

$$\beta_{ij}^t = \begin{cases} p_t(\bar{pa}^j(v)) & \text{if } v \text{ has precedent variables} \\ 1 & \text{Otherwise} \end{cases} \tag{5.7}$$

Note, p_i^t, α_i^t and β_{ij}^t are known values that can be efficiently inferred from the input belief trend using the stochastic sampling methods introduced in [75]. The time complexity is $O(X)$, where X is the number of random variables.

Additionally, as each of the hidden sources is still a valid BKB, let R be a complementary set of S-nodes (Definition for complementary set is introduced in Chapter 2), then the following property must be satisfied:

$$\sum_{q \in R} h(q) \leq 1 \tag{5.8}$$

Now, to leverage the time-varying knowledge, we propagate our modeling at a single time step t to the entire series. The objective is to find the probability distribution for the hidden source h that can best fit with the entire belief sequence. Let r be the total number of complementary sets of S-nodes, we estimate parameters $\vec{w}$ and h via the following constrained optimization problem:

$$\begin{aligned}[w^*, h] = \arg\min_{\vec{w},h} \quad & \textstyle\sum_{t=1}^{T} \quad \sum_i \exp^{[(1+w_t)p_i^t - a_i^t - w_t \sum q_{j \in \sigma_i} h(q_j) * \beta_{ij}^t]^2} \\ s.t. \quad & \textstyle\sum_{q \in R_i} \quad h_(q) \leq 1, \forall i \in [1, \dots r]\end{aligned}$$

We apply Sequential Quadratic Programming (SQP) algorithm [110] to do the optimization, as the linear algebra routines it uses are more efficient in both memory usage and speed than the active-set routines.

Multiple Hidden Sources Detection

Next, we extend the problem by allowing a series of hidden sources affecting individual opinions. According to Observation 2, we consider a simplified situation when there is at most one piece of new information fused into the previous knowledge base at each time step. In other words, two hidden sources will not have effect at the same time. This could happen when people adopt the attitude from the source whose opinion is most convincing in a particular field of knowledge. Without loss of generality, let $h_1, h_2, \ldots, h_m$ be a series of hidden sources. Thus the task is to address the following problems: 1) detect when a new source h_i gets fused in, denoted by s_i; 2) learn its probability distribution; and, 3) learn how its impact value varies after time s_i.

The last two problems can be solved in the same way as the single source tracking as long as we know the moment a source shows effect. In fact, if there is no new source fused in at time t, then the distribution of hidden source h_t learned at time t should be similar to h_{t-1}, i.e.

$$p_t \approx \frac{p_{t-1} + w_t h_{t-1}}{Z_t},$$

given some w_t. Otherwise, there is no way to transform p_{t-1} into p_t by simply varying the impact value of h_{t-1}. In our previous work, Santos et al. [73] proposed a tuning algorithm to adjust a knowledge-based system represented by a BKB such that the tuned BKB can lead to desired behavior or distribution with minimal change. We apply the idea of tuning in our work to calculate the minimum change required to tune p_{t-1} into p_t with respect to the source nodes of h_{t-1}. Large change (greater than a threshold δ) indicates that a new belief distribution has been fused into the current one. Additionally, instead of using just one belief point p_t to learn the distribution of the newly detected hidden source h_t, we leverage the time-varying information by using a subsequence of belief distributions starting from time t to strengthen our estimation.

The detection algorithm can be described more formally as follows: The input is a sequence of belief distributions $\{p_0, p_1, p_2, \ldots, p_T\}$. The output consists of a set of detected hidden sources distribution $H = \{h_1, h_2, \ldots, h_m\}$, the time varying impact for each of the source $W_{t,i}(i = 1 : m, t = 1 : T)$ and the moment that the source shows effect $S = \{s_1, s_2, \ldots, s_m\}$.

Multiple-source-detection ($p_0,\ p_1,\ p_2,\ \ldots\ p_r$)**:**

H $\leftarrow \emptyset$
i $\leftarrow$ 0
$s[i] \leftarrow 1$
for t = 1: T
    if i > 0
        $d = BKB\_tuning(p_{t-1}, p_t, H[i-1])$
        if $d > \delta$
            $S[i] \leftarrow t$
        else $i \leftarrow i - 1$
    end, if
    subsequence = $[p_{s(i)}, \ldots, p_t]$;
    [h, w] = single-source-detection (subsequence);
    $H[i] \leftarrow h$:
    for $j = S[i] : i$
        $W[i][j] \leftarrow w[j - S(i)]$;
    end, for
    $i \leftarrow i + 1$
end, for
return [H, W, S].

Figure 5.1: Multiple-source detection algorithm

5.3 Experiments

In what follows, we present results of experiments that were carried out on both simulated data and a real world scenario. We start by introducing two baselines we compared with in the experiments.

GMM: The first baseline is Gaussian Mixture Model (GMM) [106], one of the most statistically mature methods for mixture model clustering. GMM is a parametric probability density function represented as a weighted sum of Gaussian component densities. To compare with our work, we treat each component as an external opinion source and the mixture coefficient/weight of each component as the impact. Parameters are estimated from training data using the iterative EM algorithm [105]. We randomly choose the initial value and run the learning process 10 times to get the average result.

MBNs: As a general Gaussian density does not characterize the causal relationship between variables, to better represent the individual belief information, we apply Mixtures of Bayesian Networks (MBNs)[111] in our experiment as the second baseline. MBNs generalize BNs and several other important classes of models including mixtures of multivariant-Gaussian distributions. Each mixture component in MBNs is a BN encoding a conditional-Gaussian distribution, in which the variables may include both discrete and continuous variables. Considering that the causal relationship in human belief systems is less likely to change, we focus on learning the probability distribution of the component BNs and their mixing weight. Since at any time step, there will be only one source fused into the previous belief distribution, we set the number of components for both GMM and MBNs to 2.

5.3.1 Evaluation on Synthetic Data

Evaluation of our approach is challenging, as the data consisting of a human's knowledge bases and individual attitudes is usually inaccessible. Statistically designed sample surveys have helped researchers to measure human opinions on a wide range of issues. Nevertheless, they still suffer from the massive human effort required in the survey and the methodology biases associated with the survey techniques. The work from Wiseman [112], suggests that responses collected from a public opinion polling are dependent on the method used to collect the data. If not designed carefully, respondents may provide socially desired responses even when they hold different views. Moreover, none of the existing datasets contains a time-varying knowledge base influenced by the hidden sources whose characteristics have been well captured. All of these limitations drive us to conduct our evaluation with synthetic data. One of the main advantages of using synthetic data is that it can be used to discover unexpected situations and certain quantifiable conditions that may not be readily apparent or even available in real data. Although to enumerate all variants of human knowledge bases is not realistic, we still try to cover as many as possible by tuning the factors that may affect a knowledge base. We apply randomness in all sampling processes to eliminate the chance of a model favouring one pattern over ethers. Eventually, we want to demonstrate that our method can achieve reasonable results regardless of how a knowledge base is constructed.

Data Synthesis and Experimental Setup

To evaluate the effectiveness of our method, we simulate a person's belief sequence $\{p_0, p_1, p_2, \ldots, p_T\}$ from a set of predefined BKBs. For each experiment, we generate one's initial belief p_0 and a series of hidden sources $H = \{h_1, h_2, \ldots, h_m\}$ based on a belief template p_{temp} but with different conditional probability distribution, where p_{temp} is represented by a BKB. A person's belief p_t at time $t(t = 1 : T)$, then is the fusion of p_{t-1} and some hidden source $h_i(h_i \in H)$ with a randomly assigned hidden weight/impact w_t. To compare with the two baselines, at each time step t, we sample 1M records from the belief distribution p_t as the input dataset, denoted as $Data_t$ for both GMM and MBNs.

In order to compare with GMM and MBNs in terms of the ability to detect new sources, we extend them with a statistics-based detection method. Similar to the way we derive the multiple-source-detection algorithm, if there is no new source fused in at time t, then the hidden distribution h_{t-1} learned at time $t-1$ using two baselines should be similar to h_t. We borrowed the idea from [97] to calculate the likelihood ratio statistic:

$$F(t) = \frac{P(Data_t|p_t)}{P(Data_t|p'_t)} \tag{5.9}$$

where

$$p_t = \frac{p_{t-1} + w_t h_t}{Z_t} \quad and \quad p'_t = \frac{p_{t-1} + w_t h_{t-1}}{Z_t}$$

It is not hard to find that a larger $F(t)$ is achieved when p'_t differs from p_t, which indicates that a new source has been fused in at time t. We calculate $F(t)$ for all time steps and treat the moments whose likelihood ratio is greater than a threshold ζ as detection results.

Evaluation of single-source tracking

In the first experiment, we test the performance of our approach on tracking one single hidden source h, i.e. $p_t = func(p_{t-1}, h, w_t)$. Both p_0 and h are generated based on a 10-variable p_{temp} (as shown in Figure 3.1). We choose three different numbers of time steps: $T = 10, 50$ and 100 to examine how the length of belief sequence affects our tracking performance. Figure 5.2 plots the impact trend detected using FFM, GMM and MBNs respectively in terms of the number of time steps. To quantitatively measure the consistency of the impact trend against the true values, we employ a cross-correlation score, which measures similarity of two time series as a function of time-lag applied to one of them [113]. As defined in [82], the normalized cross-correlation function between two series $x = \{x_1, x_2, \ldots, x_n\}$ and $y = \{y_1, y_2, \ldots, y_n\}$ is:

$$\rho_{xy}(k) = \frac{E[(x_t - \mu_x)((y_{t+k} - \mu_y)]}{\sigma_x \sigma_y} \quad k = 1-n, \ldots, 0, \ldots, n-1.$$

We run multiple experiments to test our performance on different initial belief p_0 and hidden source h. In particular, we gradually increase the size of the belief template p_{temp} in terms of the number of variables from 5, 10 to 20. Then, for each p_{temp}, we run experiments on 10 randomly picked hidden source h. Figure 5.3 depicts the aggregated cross-correlation curves, i.e. $\rho_{xy}(k) = \frac{1}{30}\sum_{i=1}^{30} \rho_{xy}^{i}(k)$, between each method with the true trend for the three time series.

From Figure 5.2 and Figure 5.3 we can see that FFM always has the best performance. The trend captured using FFM is pretty consistent with the true trend. MBNs outperform GMM but not as well as FFM. The reason is that when some of the variables are highly correlated, GMM is likely to converge to a local solution. MBNs on the other hand, characterize the causal relationship between variables and overcome this restriction. However, both MBNs and GMM suffer from the drawbacks of EM-based decomposition methods, e.g. the requirement of a good initial guess. Moreover, both MBNs and GMM fail to track the impact value when it is rather small. The reason is that general mixture models usually need large datasets. So when the mixing coefficient of a component distribution is small, then the data corresponding to that distribution is not enough to learn its parameters accurately. In contrast, as shown in Figure 5.2, even when the hidden impact values are very small, our detection results are still accurate. This fact enables us to detect less substantial influencing sources. Last but not least, the performance of FFM increases with the number time steps, which indicates that our method is capable of improving detection performance by leveraging time-varying knowledge. As FFM does not learn the impacts of hidden sources for each individual time step, but optimizes through the entire altogether, the overall runtime is much smaller than MBNs and GMM.

Moreover, we compare the distribution of a hidden source learned during the detection process with the true one. Chan and Darwiche [114] proposed a distance measure between

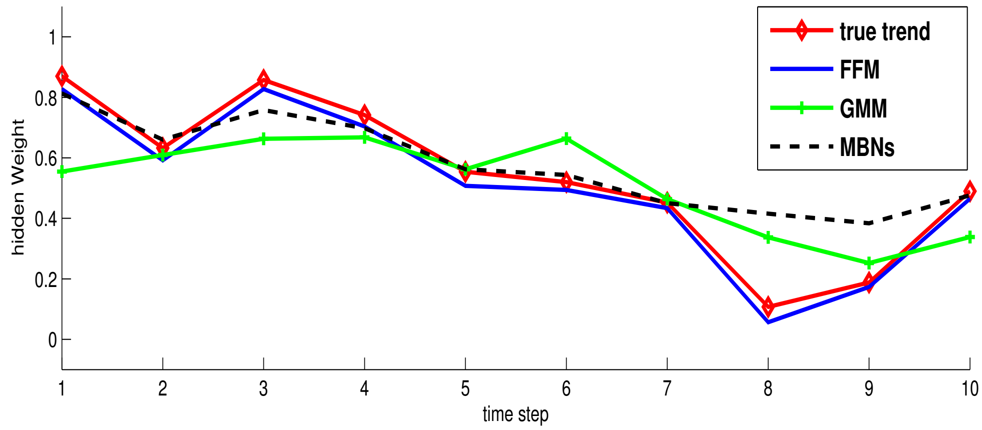

(a) T=10

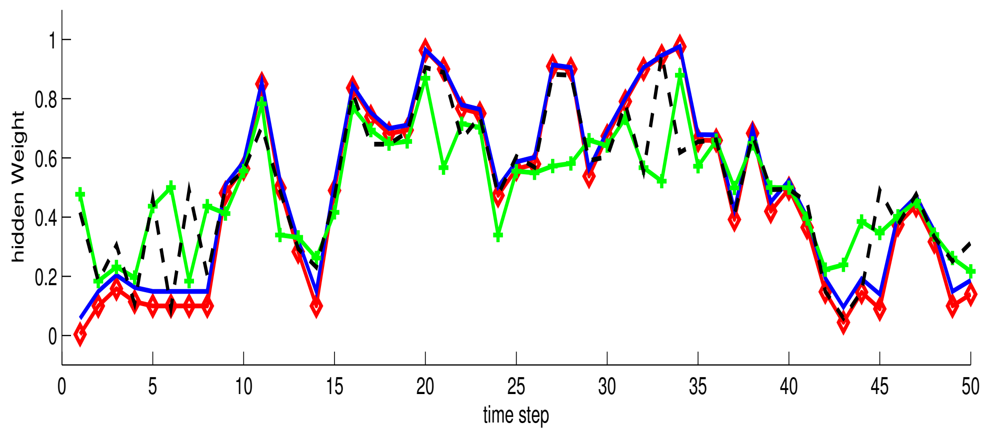

(b) T=50

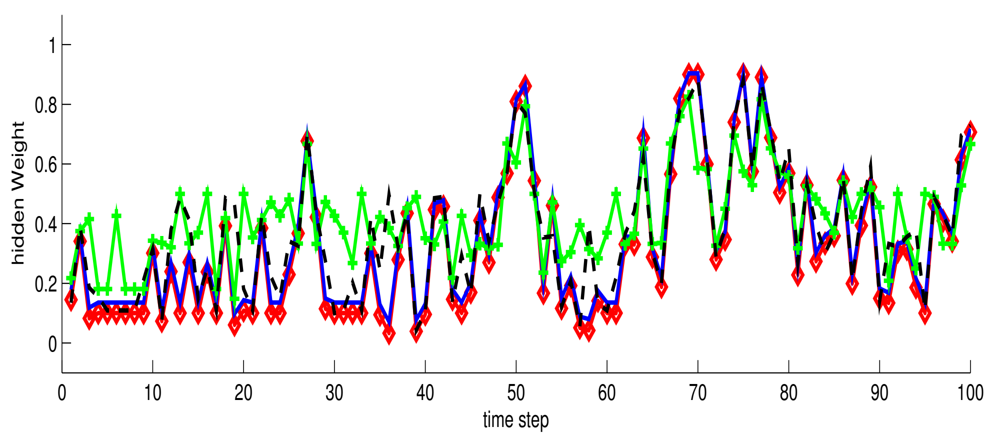

(c) T=100

Figure 5.2: Comparison of impact tracking for FFM and two baselines with increasing T

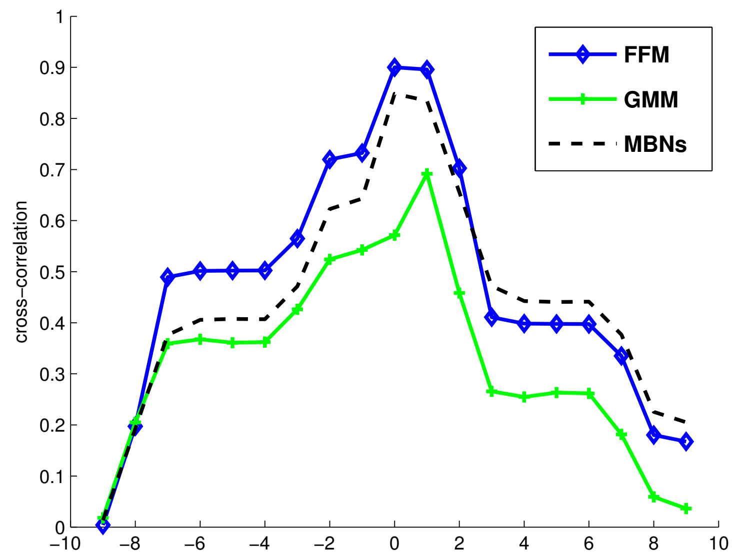

(a) T=10

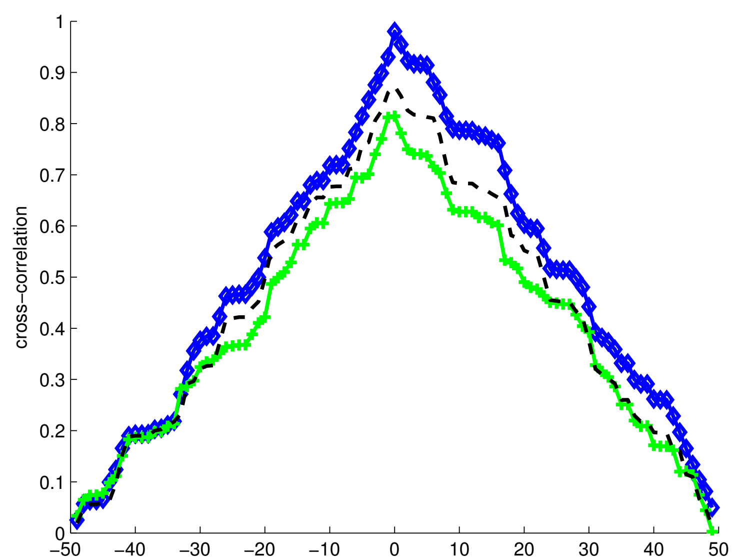

(b) T=50

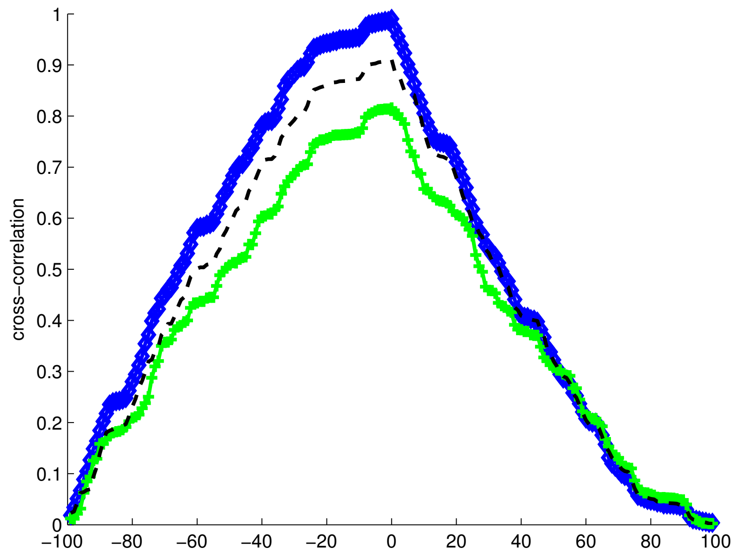

(c) T=100

Figure 5.3: Cross-correlation analysis with T = 10,50 and 100

two probability distributions, where the distance is defined as:

$$D(p, p') = \ln \max_{w} \frac{P'(w)}{P(w)} - \ln \min_{w} \frac{P'(w)}{P(w)} \tag{5.10}$$

We apply this metric in our evaluation due to its ability to bound belief changes comparable to KL-divergence. Table 5.1 shows the statistics of the distance between the learned distribution and the true one over total 30 runs. We apply Student's t-test to examine the significance of the improvements of using FFM. The results provided in Table 5.2 suggest that the distribution of the hidden source we learned is significantly better than both GMM and MBNs.

10 steps	Avg	Min	Max	Std
FFM	0.5618	0.4125	0.6587	0.0636
MBNs	2.1324	1.6724	2.3678	0.1999
GMM	2.3145	2.0135	2.5431	0.1647

50 steps	Avg	Min	Max	Std
FFM	0.5137	0.4074	0.6109	0.0689
MBNs	1.8934	1.6538	2.1546	0.1470
GMM	1.9723	1.8439	2.2768	0.1187

100 steps	Avg	Min	Max	Std
FFM	0.4742	0.4192	0.5308	0.0407
MBNs	1.9819	1.6938	2.2456	0.1547
GMM	2.1687	1.9437	2.3645	0.1845

Table 5.1: Statistics of the distance between the true and learned probability distribution.

p-value	10 steps	50 steps	100 stepts
FFM vs MBNs	4.1381E-5	1.6289E-6	5.8265E-7
FFM vs GMM	3.5809E-6	1.7764E-6	6.7676E-7

Table 5.2: p-value of paired sample *t*-test using different algorithms

Sensitivity Analysis

Since our testbed is simulated based on the predefined knowledge bases, e.g. initial belief p_0 and hidden source h, it is critical to understand how the diversity in the belief distribution, the complexity of the knowledge bases and the completeness of the hidden influential fragments affect FFM's performance. With limited real data on individual's knowledge bases, we conducted experiments on the synthetic data to test the robustness of our model with respect to these parameters.

Diversity of the Hidden Fragment Distribution

As the characteristics of the hidden source is captured by analyzing the distribution change of a person's knowledge base between each time step, a follow-up question to ask is what if the distribution of the hidden source is similar to one's initial belief? In this section, we evaluate how the diversity in the conditional probability distribution of the hidden fragment affects the test results. In particular, we assume that p_0 and h share the same BKB structure, but have different conditional probability distributions. To control the diversity level, we define a diversity rate d ranging from 0 to 1 where $d = 0$ means that the hidden distribution h is identical to p_0 and $d = 1$ means h is totally opposite of p_0. For example, if the probability of $p(A = 1)$ equals to 0.3 in p_0, then in h, $p'(A = 1) = 0.7$. Basically, for every single rule x, d controls the difference between $p(x)$ in p_0 and $p'(x)$ in h. We vary the diversity rate r from 0.1 to 0.9, where 0.1 is marked as min and 0.9 is marked as max. We measure and compare the learning errors of FFM and two baselines for each different d, where the error is defined as the averaged difference between the learned impact and the true value. To be consistent and comparable with other experiment settings, we fix the number of variables to 20, T to 50. The structure of h and p_0 is randomly generated for each run. We run experiments 10 times for each d and present the averaged results along with the standard deviation in Table 5.3. As shown in the table, the error for FFM is consistently smaller than two baseline models. The p-value indicates that the outperformance is statistically significant. In addition, the error gets larger when the diversity rate is low. The reason is that when the distribution of the hidden source is very similar to the initial belief, it becomes difficult to accurately detect the hidden impact, as the varied opinion at each time step is insensitive to the value of impact.

Complexity of the Knowledge Bases

Besides the impact from the diversity in the conditional probability distribution, the structure of the initial belief and the hidden source is another factor that may affect a model's tracking power. One main factor that contributes to the BKB structure is the causal dependencies between variables, which decide the complexity of a knowledge base. We start with the simplest knowledge base consisting of only a set of variables, but with no dependencies between them. Then, we continuously increase the complexity of the knowledge base by adding more dependencies. This process actually mimics the way people acquire

Diversity Rate d	FFM		GMM		MBNs	
	Avg Err	Std	Avg Err	Std	Avg Err	Std
Max	0.0195	0.0152	0.1620	0.0256	0.0982	0.0194
0.5 - 0.75	0.0374	0.0128	0.1728	0.0182	0.1514	0.0276
0.25 - 0.5	0.0428	0.0210	0.1845	0.0166	0.1566	0.0139
Min	0.1019	0.0101	0.2329	0.0198	0.1837	0.0089

(a) Error with standard deviation

Significance (p-value)	FFM vs GMM	FFM vs MBNs
Max	1.9047E-09	2.2929E-12
0.5 - 0.75	9.8032E-09	3.2728E-07
0.25 - 0.5	2.9403E-07	2.6374E-07
Min	5.0498E-08	2.9412E-08

(b) Significance (p-value)

Table 5.3: Learning error of FFM v.s. two baselines with different diversity rate d

and store new knowledge in their brain. When people are new to a field, it is easy to pick up some new entity names than figuring out their relationships. Specifically, we define a dependency rate r ranges from 0 to 1. $r = 0$ means that all variables are independent of each other and $r = 1$ means that the dependencies are at maximum. For example, given n variables $a_1, a_2, \ldots, a_n$, without loss of generality, let $a_1, a_2, \ldots, a_n$ be their dependency order. i.e. a_i will never be dependent on the variables after it. So $r = 1$, if and only if for all $i = 1 : n - 1$, a_i is a parent of $a_j (j = i + 1 \ldots n)$. We fix the number of variables to 20, T to 50 but only tune the dependency rate r. For each possible link from a_i to a_j, we randomly generate a value x from uniform distribution $[0, 1]$ and add a link from a_i to a_j only if $x <= r$. Similar to the previous experiment, we run our experiments 10 times for each r.

To show how the advantage of FFM over two baseline models changes when r increases, we measure and compare their learning errors. The error differences between the models are shown in Table 5.4. The result shows that no matter how a knowledge base is constructed, our method consistently produces a better result. Actually, the advantage of FFM becomes more obvious when the structure gets more complicated in the beginning, then tends to converge after a certain value. Moreover, MBNs generally outperform GMM. But when r is in the range of $[0.25, 0.5]$, GMM produces slighter smaller error than MBNs (highlighted in bold), which suggests that the advantage over MBNs converges faster than GMM. Also, when there is little dependency, i.e. $r = min$, the differences between models is subtle. Because most of the variables are independent of each other, FFM becomes similar to a mixture model.

Complexity rate r	e_GMM - e_FFM		e_MBNs - e_FFM		FFM vs GMM	FMM vs MBNs
	Avg Err	Std	Avg Err	Std	Significance (p-value)	Significance (p-value)
Max	0.309	0.039	0.292	0.042	8.23E-09	6.45E-09
0.5 - 0.75	0.295	0.037	0.286	0.035	2.47E-10	8.34E-09
0.25 - 0.5	**0.262**	0.031	**0.264**	0.053	1.22E-10	5.37E-06
Min	0.023	0.022	0.027	0.024	0.009	0.006142

Table 5.4: Learning error differences of FFM vs two baselines with different complexity rate r.

Completeness of the Hidden Sources

In order to be compatible with the baseline models, all of our previous experiments assume that both of the initial BKB and the hidden influential fragments share the same knowledge structure. This assumption generally does not hold in the real world, as people usually transform the ideas or opinions from external sources into small knowledge fragments and then fuse them into their own knowledge bases. Therefore, in contrast to the initial knowledge base, hidden sources at any stage are necessarily associated with some degree of incompleteness. In this experiment, we fix the initial belief p_0, but change the hidden fragment from simple ones to complex ones in a similar way to the previous section. Specifically, we apply a completeness rate q from 0 to 1. We select the variables and edges in p_0 with a probability q to construct h. The minimum structure in h is just a simple rule (one S-node pointing to one I-node) and the maximum situation $(q = 1)$ is when h has the same structure as p_0. Similar to the previous two experiments, we fix the number of variables to 20, T to 50, r of p_0 to 0.2. The results from Table 5.5 suggest that FFM performs consistently better than the baseline models, especially when the hidden sources are highly incomplete comparing to the initial knowledge base. The reason is that when the hidden source contains only a few rules (small q value), the overall distribution variation at each time step will be too subtle for mixture models to learn the mixture weight. In contrast, the essence of FFM guarantees that even when the knowledge fragment is incomplete, it can still capture the impact change by focusing on the changed marginal probabilities.

Complexity rate q	e_GMM - e_FFM		e_MBNs - e_FFM		FFM vs GMM	FMM vs MBNs
	Avg Err	Std	Avg Err	Std	Significance (p-value)	Significance (p-value)
Max	0.073	0.026	0.086	0.031	1.08E-05	1.12E-05
0.5 - 0.75	0.103	0.023	0.142	0.029	3.46E-07	2.14E-07
0.25 - 0.5	0.217	0.078	0.250	0.098	1.10E-05	2.05E-05
Min	0.300	0.115	0.314	0.162	1.72E-05	1.73E-04

Table 5.5: Learning error differences of FFM vs two baselines with different completeness rate q.

Two Level Factorial Design

To identify the main effects of the three factors and study how they interact to affect the results, we apply a two level factorial design to analyze the effects of all three parameters: diversity rate d, complexity rate r, and completeness rate q on the error differences between FFM and two baseline models. Factorial Design [115] has been widely used due to its statistical power to effect estimates. It provides contrasts of averages and show how results vary with the factors. Each of the three factors is studied at two levels: min (coded as -) and max (coded as +). We set min and max values to 0.1 and 0.9, respectively, as extreme values like 0 and 1 may cause noise to the result. The factorial design results including main effects and interaction effects are presented in Table 5.6. A couple of observations can be found from the table:

1. The effect of q with low complexity rate is much smaller than with high complexity (Table 5.6 (a), run(1, 5) vs run(3, 7)). The reason is that when the structure gets complicated, the incompleteness will occur not only on the variable base but also on the causal dependency base. Both of the incompleteness will harm the performance of the baseline models. This observation also indicates that the effect of q is dependent on the value of r, which is proved to be true in Table 5.6 (b), where the interaction between r and q is rather high comparing to the other two pairs of interactions.

2. When q is small, MBNs are more sensitive to r changes than GMMs. The reason is that the advantage of MBNs rests on the assumption that the probability distribution of a belief is specified completely. So, when the assumption cannot be satisfied, increasing the complexity of a BKB can only enlarge its learning error.

3. Compared to the other two factors, diversity has a notable positive effect, which is consistent with the experiment results in the previous section. Additionally, its effect is somewhat independent of the others as the interaction effects with r and q are both

small. This finding is not surprising, as the effect of diversity is mainly determined by the learning capability other than the structure.

Run number	D (diversity)	R (complexity)	Q (completeness)	Err_GMM - err_FFM	Err_MBNs - err_FFM
1	-	-	-	0.125	0.099
2	+	-	-	0.184	0.134
3	-	+	-	0.185	0.178
4	+	+	-	0.254	0.229
5	-	-	+	0.121	0.091
6	+	-	+	0.182	0.135
7	-	+	+	0.175	0.164
8	+	+	+	0.210	0.175

(a) average error difference with different factor values

Average		Effects	
		Err_GMM - err_FFM	Err_MBNs - err_FFM
Main effects	Diversity d	0.056	0.03525
	Complexity r	0.053	0.07175
	Completeness q	-0.015	-0.01875
Two-factor interaction	d*r	-0.004	-0.00425
	d*r	-0.008	-0.00775
	R*q	-0.012	-0.01525
Three-factor interaction	D*r*q	-0.009	-0.01225
Standard error		0.0048	0.0029

(a) calculated effects and standard errors for three factors: d, r, and q

Table 5.7: Two-level factorial design

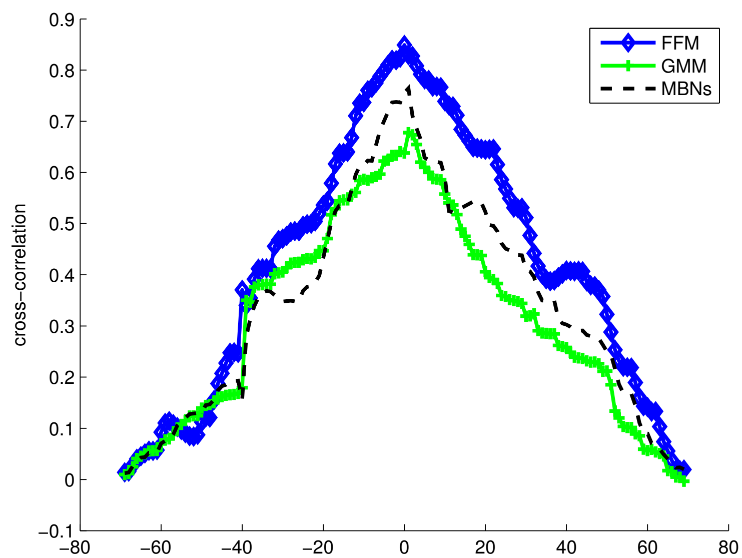

Figure 5.4: The aggregated performance for tracking multiple-hidden-source impact trend

Multiple-source Detection

Next, we examine the ability of our method to detect and track a series of hidden sources. We set the number of time steps T to 70. For each experiment, we randomly select 7 hidden sources from H and the belief sequence of a person is simulated similar to the first experiment except that after every 10 steps, we fuse a different hidden source into the previous belief distribution p_{t-1}. The impact weight for each step is still assigned randomly. We apply the detection algorithm described in the last section to capture the time that a new hidden source shows effect and track its time-varying impact pattern at the same time. We run experiment for 30 times in terms of different p_{temp} (similar to single-source tracking), and Figure 5.4 reports the average cross-correlation curves to show our tracking capability. As we can see, our approach is robust on tracking multiple hidden source impact trends.

To measure the accuracy of detecting new hidden sources, we apply an Receiver Operating Characteristic (ROC) curve to examine each method's tradeoff between its false positive rate (proportion of the time steps that are falsely detected as new source activation

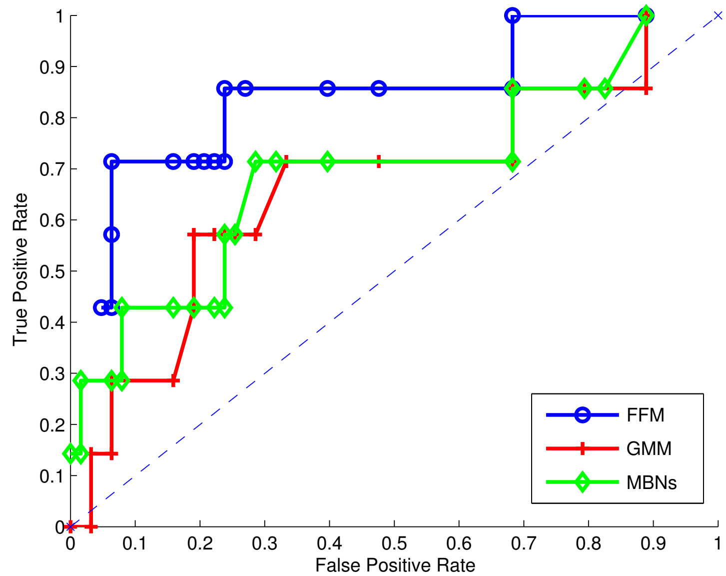

Figure 5.5: Performances for new hidden source detection

time) and true positive rate in terms of different threshold values . A higher curve or a larger area under the curve indicates a better detection performance. Figure 5.5 shows the ROC curves for each of the methods respectively.

We can see that, FFM performs better than the other two baselines. The reason is that the detection processes of GMM and MBNs depend only on the previous step, whereas FFM takes advantage of the belief points with no new sources fused in to refine the previous detected hidden source. The best tradeoff between true positive and false positive for FFM is achieved when threshold ζ is set to 0.2. By taking a close look at our results, we find that some miss-detected cases occur when the distribution of new hidden source is similar to the previous one. Thus, our method treats them as the same one as their subtle difference can hardly been captured without further information.

5.3.2 Evaluation on Real Data

The results from our synthetic data suggest that when the variation of a knowledge base's distribution is caused by the fusion of a hidden knowledge fragment, our algorithm is capable of capturing the evolution of its impact and characterizing its distribution. Moreover, our algorithm is insensitive to the complexity of the knowledge bases and has no requirement for the completeness of the influential knowledge fragments. To illustrate the applicability of FFM to real world problems, it is necessary to conduct experiments on a dataset collected in real scenarios. To the best of our knowledge, there is no other work trying to solve the exact same problem as we do. Wong et al [100] applies a BN-based anomaly detection approach on real emergency department data to test the performance of detection for disease outbreaks. However, they aimed to achieve early detection on the distribution variation without distinguishing the influential sources. Such a data set, unfortunately, is not suitable for our purpose as there is no ground truth on the characteristics of the sources.

Another stream of work that attempt to handle changing belief networks are built on top of DBNs, e.g. HMMs. They exploit time-series data to solve problems like speech recognition and behavior analysis. However, DBN models require a strong assumption that the observations at each time step must be dependent on one or more finite states. This may not be the case in our scenario, where the distribution of variable v at the next state is dependent on both the previous state and the hidden source whose prior information is not available. Moreover, the focus of our approach is not to predict how a person's knowledge base looks like in the future, but to discover and characterize the hidden influential sources based on the historical data, such that we can better analyze and explain this process. Considering the difficulties of finding the changing knowledge bases by fusing other knowledge fragments in the real world, we decided to evaluate our model's performance on real data in two steps. Firstly, we want to show that even when the knowledge fragments are integrated in a linear mixture fashion, our method can still get comparable performance as the two baseline mixture models. Secondly, we illustrate one specific real scenario to show how our method can be applied to epidemic spread analysis.

Mixture Data

To evaluate FFM's tracking ability on the data that is compatible with our model, but also captures the characteristics of particular problems, we decide to simulate our time-varying belief distributions based on three real-world datasets: Adults, Cover type and Color-based image segmentation. All datasets are collected from the UCI machine learning repository [116]. The original use case for these datasets was classification. We leverage the data, however to evaluate each algorithm's capability at learning mixing components. This idea was borrowed from Thiesson et al [111] who applied real classification dataset in the evaluation of MBNs. For a detailed description of each dataset, please refer to Appendix D.

GMMs have been employed to perform tasks such as real-time color-based tracking and segmentation as they are able to smooth over gaps resulting from sparse sample data and provide tighter constraints in assigning object membership to color-space regions. Considering that many approaches based on GMMs have been focusing on cases where the observed variables are continuous, MBNs, a more generalized mixture model, aim at being compatible with both discrete and continuous variables.

Time-series Data Preparation In order to simulate a time-varying belief distribution where the changes are caused by mixing instances from a hidden distribution, we select the instances from each group/category in a certain proportion that is determined by an impact value. As we assume that no prior information about the impact trend is available, we set the impact weight at each time step to be independent of each other. In our experiment, the impact value for the hidden source at each time step is obtained from a uniform distribution over the range $[0, 1]$. As a note, the impact distribution is not restricted to uniform. For each generated impact series, we apply it on all three data sets, such that we can also compare the performances resulted from different feature characteristics.

To test our model's ability on single-source tracking, we select two classes from the original datasets to mimic the initial belief distribution p_0 and the hidden source distribution h. For instance, in the Adult dataset, people with annual income $> 50K$ and income $<= 50K$ exhibit two distributions (see Appendix D for the specific classes labelled as $class\_1$ and $class\_2$). Without loss of generality, let the data from $class\_1$: D_0 and the data from $class\_2$: D_h represent p_0 and h respectively. Then, at time $t(t = 1 : T)$, we mix D_h with D_{t-1} in a certain proportion to form D_t, where the proportion is a normalization of the impact value at t. Again, take the Adult data as an example. Assume a person belongs to the $> 50K$ group and at t_1, he or she is influenced by another person from the $<= 50K$ group with a normalized impact value equals to 0.2. Then, to generate N learning samples for that time step, we will randomly select $0.8N$ instances from $> 50K$ and $0.2N$ instances from $<= 50K$ to form the mixture data.

Model Learning To apply FFM on a series of mixture data, we first learn the causal relationships between variables (except the category attribute) for both $class\_1$ and $class\_2$. As every BN can be transformed into an equivalent BKB, we employ Bayes Net toolbox to learn the initial structures. Then, the marginal probabilities of $p_t(v)$ and $p_t(pa^j(v))$ in Equation 5.3 can be estimated from the mixture data D_t using maximum likelihood approach. When we need to discretize the continuous features (currently, BKBs only accept nominal features), we use the Discretize filter from Weka[1] [117] on the default setting. As MBNs and GMMs are designed for learning the components from mixture data, we apply them directly to the mixture data at each time step to learn the impact.

[1]Weka (Waikato Environment for Knowledge Analysis) is a popular suite of machine learning software written in Java

Results We gradually increase the number of time steps from 5, 10, 20 to 50 and run the experiments on each data set 10 times. The average results for the cross-correlation analysis are shown in Figure 5.6-5.8. As one can observe from the figures, with only a limited number of time steps, GMM produces slightly better results than FFM (Figure 5.8 (a)). But when the time series is long enough, e.g. over 10 steps. FMM will catch up by leveraging the time-varying knowledge. Also, GMM's performance is significantly worse than FFM on categorical features (Table 5.8, data 1 and data 3), but not on numerical features (Figure 5.8 (c) and (d)). Note that GMM presumes that all the data points are generated from a mixture of Gaussian distributions. This assumption may not hold for data sets 1 and 2. We apply Student's *t*-test to examine the statistical significance of the differences between the two models. The results are presented in Table 5.8, the p-value is in bold if FFM performs worse than the control model. As we can see, for most of the cases where FFM performs worse than the baseline models, the p-value indicates statistical insignificance. The exception happens at the 5-step experiment in data 3, which suggests that when we are limited to only a few time steps, FFM could underperform the baselines significantly due to the reason we mentioned earlier. After all, FFM is not designed to learn the components of a mixed distribution, but to characterize the hidden influential source from analyses of the differences between the consecutive time steps.

Data 1

p-value	FFM vs GMM	FFM vs MBNs
5 steps	1.75E-05	**0.204**
10 steps	8.41E-04	0.106
20 steps	3.70E-06	0.003
50 steps	1.48E-06	7.84E-07

A Real Scenario Case Study

Now, we show how to apply our method to detect and track events that happened during the H1N1 pandemic in Mexico.

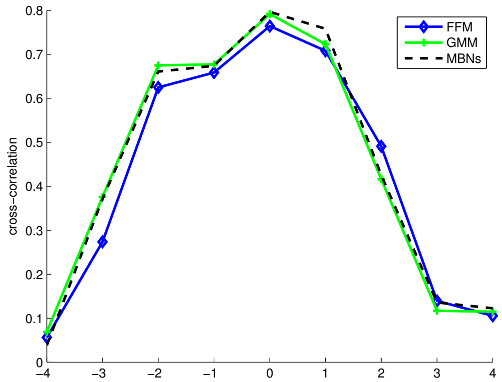

(a) Data1: 5 steps

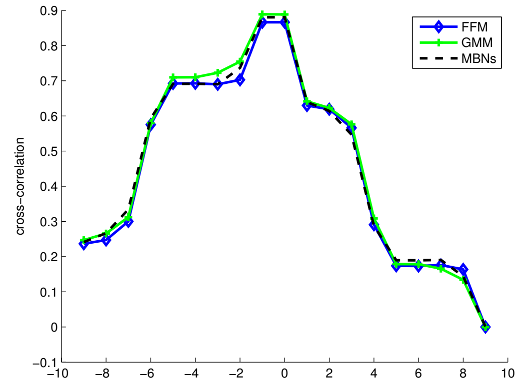

(b) Data1: 10 steps

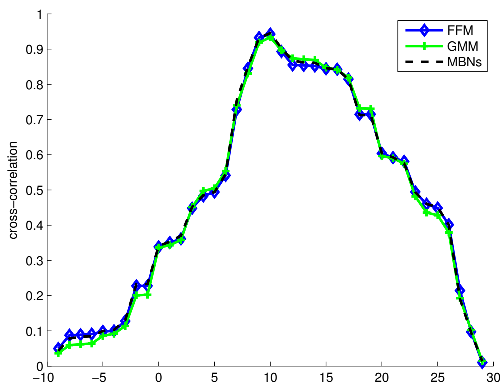

(c) Data1: 20 steps

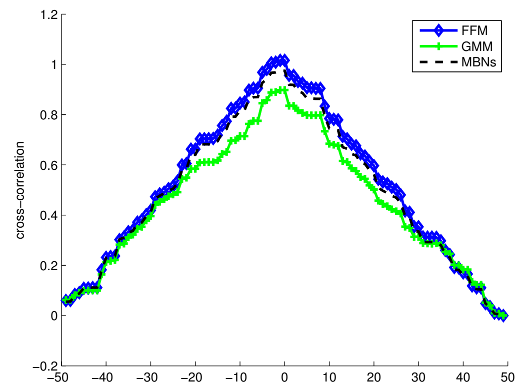

(d) Data1: 50 steps

Figure 5.6: Cross-correlation on dataset 1: Adults

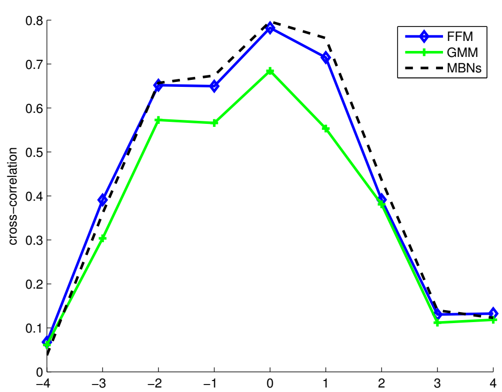

(a) Data2: 5 steps

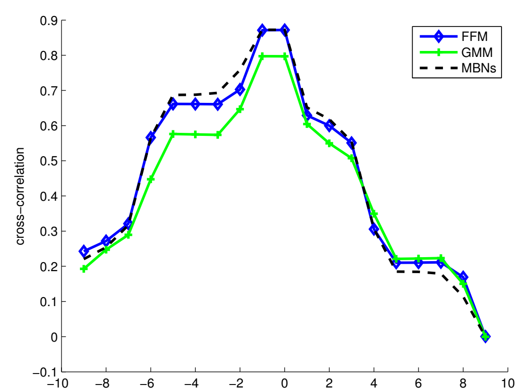

(b) Data2: 10 steps

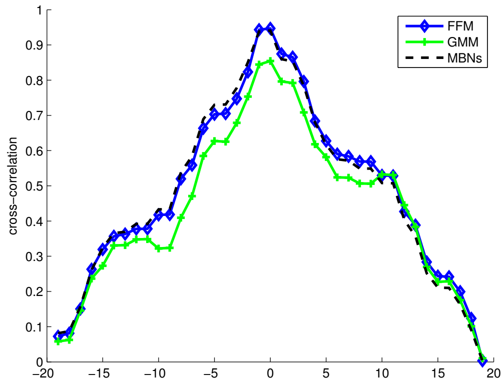

(c) Data2: 20 steps

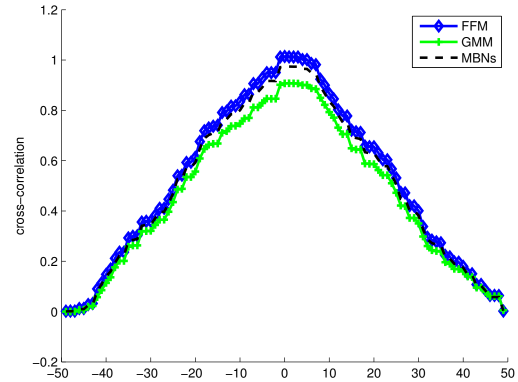

(d) Data2: 50 steps

Figure 5.7: Cross-correlation on dataset 2: CoverType

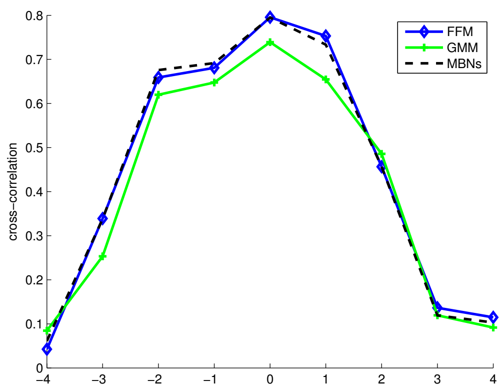

(a) Data3: 5 steps

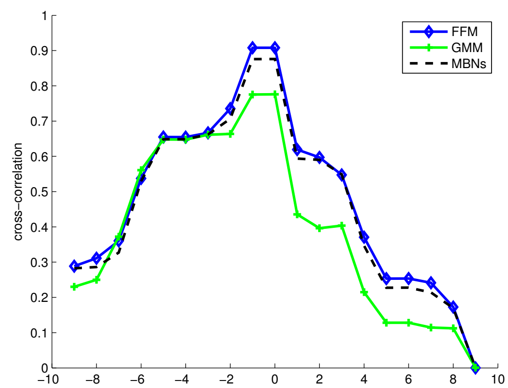

(b) Data3: 10 steps

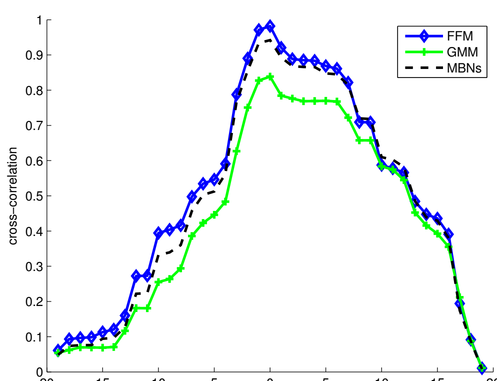

(c) Data3: 20 steps

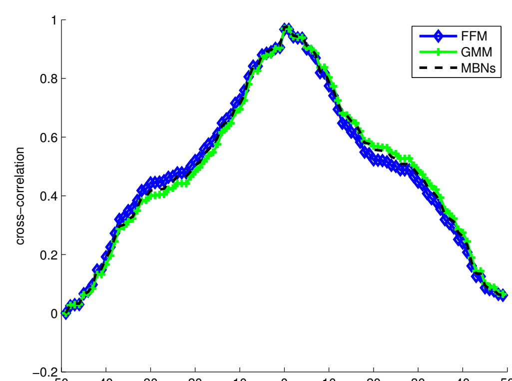

(d) Data3: 50 steps

Figure 5.8: Cross-correlation on dataset 3: Color-based segmentation

Data 2

p-value	FFM vs GMM	FFM vs MBNs
5 steps	6.87E-04	**0.195**
10 steps	8.51E-05	0.166
20 steps	1.03E-05	0.106
50 steps	8.77E-06	0.056

Data 3

p-value	FFM vs GMM	FFM vs MBNs
5 steps	**0.008**	**0.061**
10 steps	**0.140**	**0.151**
20 steps	0.221	0.176
50 steps	0.324	0.261

Table 5.8: Statistics of the difference between each paired model

Background In one of our previous works, Santos et al. [29] conducted a Cross-Border Epidemic Spread project to study why and under what circumstances would people be driven to cross the border both legally and illegally with respect to epidemic spread. In order to understand such human behavior as well as capturing actors' intent, we modeled actors' beliefs and subsequent actions by incorporating their background, history, education and experience with related issues, i.e. culture elements in a knowledge-based system. As described in Santos et al. [29]: Culture affects actors' beliefs, goals and decisions, and thus is critical to properly represent the entirety of influences affecting actors' decisions. Including cultural influences in the model allows analysts the power and freedom to fully represent their intuitions, and to explore the model for explanations to emergent behaviors. As the culture information is inherently uncertain and incomplete, we represented the intent of actors or entities using their cultural traits modeled probabilistically in BKBs. Each cultural fragment represents a certain aspect of an actor's goals, intentions, opinions and cultural influence. Then, multiple fragments can be fused into a single BKB to represent the actor's overall behavior.

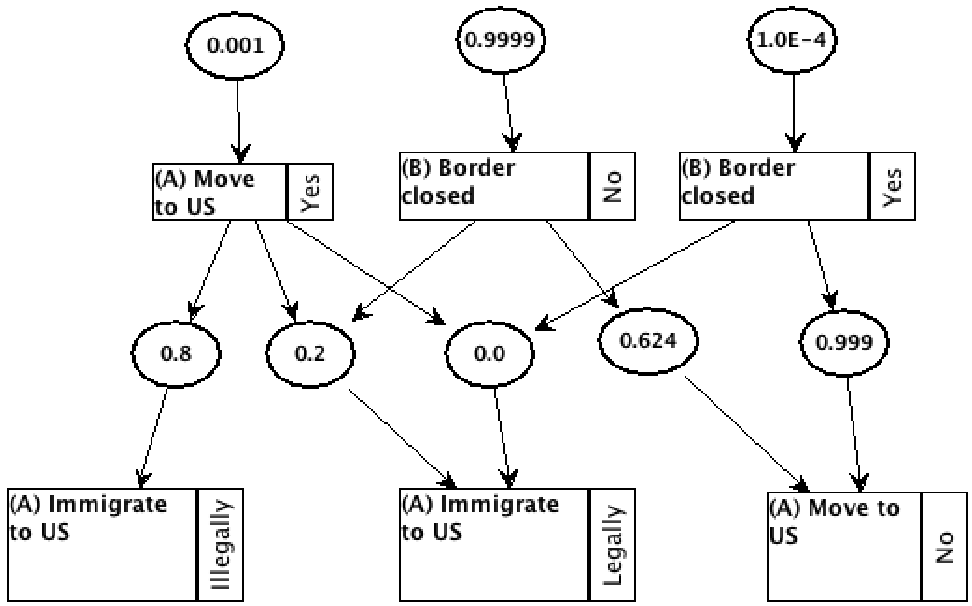

Figure 5.9: Example BKB fragment with belief and action

To distinguish an actor's action from his belief, we treated "Actions" as another type of variables (denoted as 'A') in the BKBs to represent the possible strategy or actions adopted by the actor or entity to achieve their goals. All "Beliefs" variables (denoted as 'B') can influence other beliefs and actions, but can only be influence by other beliefs. In contrast, actions can only influence other actions, but not beliefs. Figure 5.9 shows a representative BKB fragment with the hierarchy of beliefs and actions.

The original goal of this project was to show the effectiveness of our ability to explain outcomes produced by the model. We built our intent-driven model based on the simulated pandemic scenario in Mexico. As the initial groundwork for the simulation, we chose to model the Mexican population and focus on changes in their border crossing behaviors and intent. By building BKB cultural fragments to represent various sectors of the Mexican populace, we constructed the model fed by inputs from a set timeline representing various other events and actions/decisions taken by other actors in the scenario, such as the Mexican and US governments, and international organizations such as the European Union (EU) and the WHO. In total, a team of researchers used information from dozens of newspaper articles, news websites, government and non-governmental organizations, open sources and general knowledge, to build 23 BKB fragments and 22 random variables. The conditional probabilities assigned in the BKBs represent the subjective view of the researcher, informed by the data sources. However, when statistical information was available, it was used to generate the probabilities.

The Mexican populace was modeled by two main demographic details: age and location. Age was broken down into three groups that corresponded well with reported vulnerability to H1N1[2]: young (0-17), middle-aged (18-64), and old (65+). These breakouts were relevant because there was a widespread belief that the middle-aged group were disproportionately vulnerable, which would likely affect their decisions and behavior. Consequently, we expect the behavior for the middle-aged group to be markedly different from that of the other two age groups. A cultural fragment for the middle- aged group (given in Figure 5.10) represents the commonly held view that H1N1 was more lethal for this group[3]. Note that the probability values represent the intuition of the SME that built the BKB, based on his/her expertise, experience and available facts. We also implemented a high-level regional model, dividing the Mexican state into three geographic areas: the north, the south, and the interior. We subsequently modeled the H1N1 epicenter as located in the interior, as the earliest outbreaks were reported in the Federal District of Mexico and San Luis Potosi[4], both in central Mexico. To incorporate diverse behaviors of the Mexican population, we built BKB fragments to represent cultural influences of demographic examples of typical Mexicans, such as middle-aged and interior, or elderly and north. At the same time, we fused in fragments representing individual details, such as region of residence in relation to the epicenter of the outbreak. All of these cultural and individual details can influence

[2] http://www.who.int/csr/disease/swineflu/notes/h1n1_vaccine_20090713/en/ index.html
[3] http://www.suite101.com/content/who-needs-an-h1n1-swine-flu-shot-a187205
[4] http://www.who.int/csr/don/2009_04_24/en/index.html

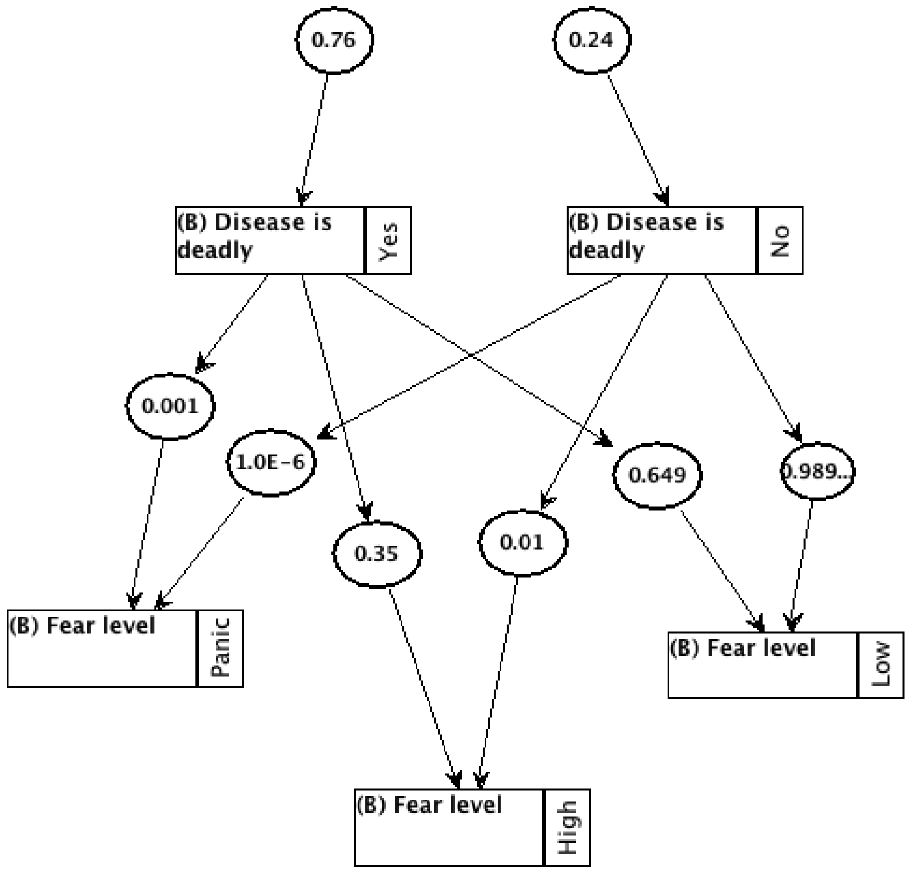

Figure 5.10: Cultural fragment for middle-aged population

decision making.

Once a representative individual is created by fusing relevant cultural fragments, we fused in BKB fragments representing events from the scenario timeline (shown in Table 5.9). An example of an "event fragment" (given in Figure 5.11) represents the event E6 in the timeline when the availability of the vaccine led to a change in belief that healthcare is more effective. There are 16 time steps in the simulated scenario from April 12th to December 5th in 2009. The fused BKB at each time step is a fusion of cultural fragments and event BKBs with the reliabilities that are predefined by the researchers.

Of primary interest are the changes in the migration patterns. Figure 5.12 shows peo-

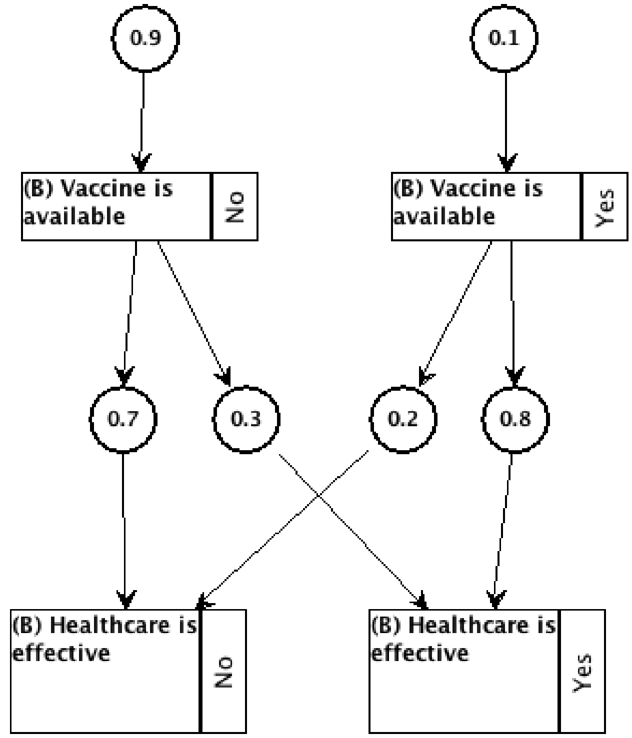

Figure 5.11: Event fragment

ple's intent variation to migrate to US illegally, which is represented by the probability of variable "illegal migration to US" inferred from the fused BKBs at each time step.

Event Detection using FFM The whole intent system is constructed through the fusion of cultural BKB fragments that are created based on sources such as demographic information and news articles. When a major event occurs, the intent system will fuse the event BKB and update its probability distribution adaptively to reflect an individual/group's belief change caused by the event. As the occurrence of an event will change people's beliefs and subsequence behaviors, the characteristics of these events and their impact patterns are key to analyzing people's reactions. We apply FFM on a series of belief distribution generated from the fused BKBs in [29] to detect the implicit events without any foreknowledge. Furthermore, we show how our approach can be used to analyze behavior of populations during H1N1 and their propensity to escape to neighboring countries, e.g. illegal migration

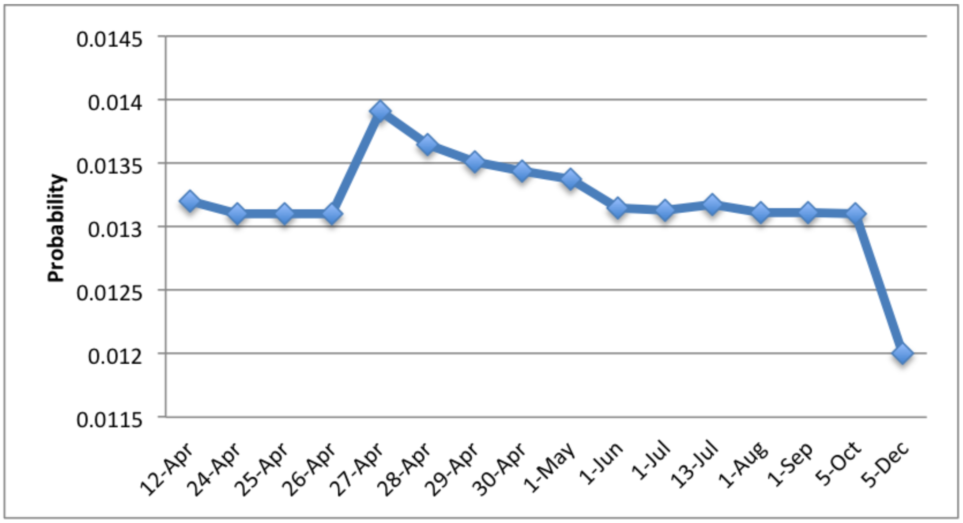

Figure 5.12: Posterior probability of illegal migration to US

to US.

	Date	Event Description
E1	4/24	WHO sends experts to Mexico
E2	4/27	EU advises European not to travel to Mexico
E3	4/29	WHO raises the pandemic level from 4 to 5
E4	5/1	Government shut down most parts of the country
E5	7/13	Businesses and government have reopened
E6	10/5	More vaccine is available

Table 5.9: Major events happened during H1N1 pandemic.

Again, let us look at the probability variation of "illegal migration to US" during pandemic period. As plotted in Figure 5.12, after H1N1 outbreak was detected on April 12th, there is a slight decline since April 24th, followed by a peak on 27th. The panic drops slowly till July 13st, then it increases a little bit on July 13th and starts to fall till the end.

So, what happened during this period? What's the cause behind the fluctuation? To answer these questions, we apply FFM to detect and track the implicit events. Figure 5.13

displays our results, where five out of six events (comparing to the timeline in Table 5.9) are successfully detected. As we can see, the behavior fluctuation that happened at each time step is not necessarily caused by the breakout of a new event, but mostly from the variation of the event impact. The only event that we fail to detect is E3. Actually, both "EU advises European not to travel to Mexico" and "WHO raises the pandemic level" made people believe that H1N1 is contagious and even deadly, which increases the desire to escape the pandemic. Furthermore, E3 happened only two days after E2, so our method treats them as the same event as they affect people's belief in a very similar way.

We also analyze the learned distribution of each event so as to gain more insights. The distribution of the first event we detected suggests an increase in the probability of "believe healthcare is effective", which becomes the main reason that lowers the fear level and migration behavior. This observation matches perfectly with the fact that "WHO sends experts to Mexico on April 24th" helps control the panic. Likewise, the breakout of the event on July 13th causes a temporary increase on people's belief regarding the contagious nature of the disease due to reopened business, which encourages migration behavior for a short period.

Moreover, on closer examination of the impact trend, we see that the impact of each event keeps declining after breakout. This explains why the probability of migration has an apparent decline on April 24th, but slows down on April 25th. Actually, event impact has been modeled as a function of days after an event occurs [76], i.e. $w_{ij} = \frac{1}{j-i+1}$. Our results fit with this pattern as well. As shown in Figure 5.13 the influence of event 3 decays close to a small value after 30 days. Thus, it is reasonable to believe that the moderate increase of the tendency to migration on July 13th is mainly caused by the event 4 individually as it has been over a month since all previous events.

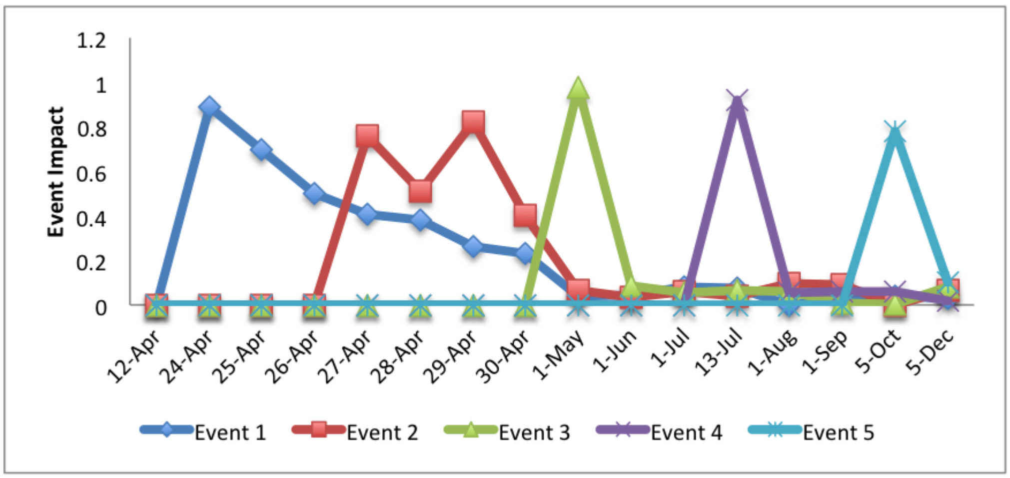

Figure 5.13: Detection results on event impact

Chapter 6

Conclusions

We have developed a framework for modeling opinion dynamics based on human knowledge bases and reasoning patterns. Different from the mainstream opinion formation methods that treat individual opinions on various issues as independent beliefs, we consider opinions as a product inferred from one's knowledge-based system, where new knowledge fragments are acquired through social interaction and learning experiences. Our underlying opinion formation model provides a computational foundation for quantifying individual opinions in the process of knowledge fusion. It subsumes existing opinion formation models in the sense that it is capable of capturing the connections between knowledge fragments, but is also compatible with prevalent social theories of opinion formation. After describing the opinion formation model, we introduced a general impact metric to measure the reliabilities of the influential knowledge fragments. We characterized the social influence on the basis of a set of domain-independent social factors. We also ensure that the impact metric exhibits enough flexibility that newly discovered influential factors can be readily incorporated into the model. Lastly, on top of the opinion formation model, we derived a Finite Fusion Model (FFM) to detect hidden influential sources and capture the evolution of their impact levels. To the best of our knowledge, this is the first computational framework of opinions where opinion dynamics is modeled as an evolution of knowledge bases. The primary advantage of such design is that we managed to model multiple components in opinion dynamics in a coherent system, such that no matter what learning techniques a component applies, they all share the same core representation. In the rest of this chapter, we will elaborate the main achievements in each component and discuss limitations and future works in the last section.

6.1 Accomplishment and Findings

In this section, we present the contributions of our framework for modeling opinion dynamics. In particular, we list the accomplishments and findings we gained with respect to each chapter.

1. Opinion Formation (Chapter 3)

   (a) We provided a computational model to measure individual opinions from a sequence of changing knowledge bases. It takes into account the connections between knowledge fragments. So any change happened to the opinion regarding one issue will be reflected in the decision made on other related issues.

   (b) We identified the nonlinearity nature of opinion formation. To evidently show that individuals do not adopt the opinions from others in a linear fashion, we conducted analyses on the real data collected from Wikipedia and focused on how the interaction on Wiki's talk pages affects people's interests. The results show that one's interested topics will be influenced by the people they have communicated with. However, the particular updating process is complicated that can hardly be modeled as a linear integration of others.

2. Social Influence Learning (Chapter 4)

   (a) We extended the opinion formation model to study social influence by incorporating a general impact metric. The influence of each fused knowledge fragment is measured based on certain social factors. We started with several influential factors that have been shown to be related to process of opinion selection.

(b) With the sound mathematical formulation, we show that the impact of each social factor can be efficiently learned via regression methods.

(c) We conducted experiments on synthetic data with respect to different parameters. The results validate the effectiveness of our approach over the DeGroot model.

(d) We provided a thorough analysis and proved that the classical DeGroot model is essentially a special case of our model under certain conditions.

(e) We tested our approach on modeling the political opinion change using a real world dataset. The results of our approach are consistent with prevalent social theories. To the best of our knowledge, this is the first work to try and uncover the mechanism that guides the selection of opinions in the real world by modeling opinion change.

3. Hidden source tracking and detection (Chapter 5)

   (a) We presented a new approach to detect hidden sources of influence, as well as capture and characterize the patterns of their impact with regards to belief-changing trends.

   (b) We made several intuitive observations about opinion change, and propose a variant mixture model FFM that is specifically tailored to handle all the constraints.

   (c) The latent parameters that characterize the distribution of the underlying hidden sources and the corresponding impact weights are learned via a constrained optimization problem.

   (d) Experimental studies on synthetic datasets show that our approach outperforms the classic Gaussian Mixture Models and Mixture Bayesian Networks.

(e) We applied sensitivity analysis in terms of the diversity in the belief distribution, the complexity of the knowledge bases, and the completeness of the hidden influential fragments. The results validated the robustness of our approach.

(f) We applied our method on a real pandemic scenario. We demonstrated the effectiveness of our model in capturing the intentions behind the opinion change, and further showed that our approach can be used to better analyze and explain the behavior of populations.

6.2 Limitations and Future Work

In this section, we summarize the limitations of our current framework and discuss some possible directions that we can pursuit in the future.

6.2.1 Limitations

One limitation of this work exists in the accuracy of knowledge elicitation. Although individual knowledge bases can be constructed based on one's preferences and answers to various survey questions, to accurately extract how people feel without bias is hard in the first place. Like the real data we used in the Social Evolution experiment, political opinions and participations during presidential election are collected via 4 surveys over several months. Even with advanced survey techniques, manual retrievals of individual beliefs or attitudes are still time consuming and labor intensive. Moreover, when the amount of data to be processed is too large, the manual efforts may not scale and have difficulty producing consistent results. Recent years, automatic elicitation techniques have been developed. For example, Sentiment Analysis have been used to determine the attitude of a speaker or a writer with respect to some topics or the contextual polarity of a sentence/document. Nevertheless, these techniques rely on the assumption that people have conveyed their ideas explicitly in words and all expressions can be accurately interpreted through natural language processing. This apparently, can hardly be satisfied in a free online communication environment.

Another limitation of this work is that people may change their opinions via other channels that are not being monitored. Although we can detect hidden sources from the inconsistency in opinion change, more information is needed to understand how reliabilities are assigned to the knowledge fragments sent from those hidden sources.

6.2.2 Future works

In this work, we concluded that opinion formation is in essence a process of knowledge expansion and revision. Individuals selectively take into account the information they gained through social communication, and then revise their opinions via knowledge-based inference. Yet, there is still space for improvement on opinion formation, hidden source tracking and detection, BKB learning, and opinion propagation.

Analyses on opinion formation The convergence analysis illustrates the situations when opinions tend to converge to a stationary state. However, whether they will reach a consensus or disagreement leading to polarization or even more general opinion fragmentation still needs to be carefully studied. Additionally, only three social factors are studied in the real dataset due to the scope of the problem and the lack of ground truths in social communication. However, some online social networks, like Facebook[1] and Twitter allow users to provide demographic information. The likeness on demographic features may also affect whether a piece of information merits our acceptance.

[1]Facebook is an online social networking service, where users can post images, update status and exchange messages,with their friends. http://www.facebook.com

Hidden source tracking and detection The current FFM model assumes that there is at most one piece of new information (from hidden source) fused into the previous knowledge base at each time step. In reality, multiple sources may have influence on one's opinion at the same time. This happens when there is no single convincing source for a particular fragment and the final knowledge/belief system is formed by integrating all possible explanations. So our next step is to expand our approach by allowing multiple sources to affect the same part of the belief network.

BKB learning In this work, individual knowledge-based system is represented by BKBs. Currently, we do not have a methodology to learn a BKB fully automatically. Most of the BKBs in our research projects are constructed based on the statistics from authorized sources by our domain experts. However, since BKBs subsume BNs, any knowledge bases representable by BNs can be transformed into a BKB. To automatically build a BKB, one can always start with some small knowledge fragments and apply BKB fusion and BKB tuning to integrate them into a large knowledge base. Both BKB fusion and BKB tuning can be done automatically.

Opinion propagation Another possible direction of our future work is to study opinion propagation. Understanding of opinion propagation is necessary as it sheds light on how information flows and how the spread patterns can be affected. Different from existing information diffusion works whose main goal is to study how information reaches nodes, opinion diffusion takes one step further to discuss the possibility of adopting an idea. The problems that need to be addressed are twofold:

1. Social scientists have shown that besides network structure, the pattern of opinion diffusion may also depend on other factors. For example, Romero et al. [118] find that the way of opinion spread on different topics varies significantly and Buchan et al. [119] suggest that females scored slightly but consistently higher on scales of trust than male. Therefore, to fully understand each factor's role and how those

factors contribute to the opinion propagation, we plan to conduct analyses on the individual difference in both prior knowledge bases and reaction to other opinions.

2. The second task to tackle is to investigate which communicatee contributes most to one's opinion change. Such contribution relies not only on the reliabilities assigned to her neighbors but also on the information or knowledge pieces sent by them. Furthermore, if we can identify how the variation of one's reliability affects the opinion formation, we may find a clue of guiding and facilitating social policies. We plan to apply contribution analysis to measure the impact of each contributors and apply sensitivity analysis to quantify how changes in the reliability of a communicator influences one's opinion on a given issue.

Comprehensive Model In the current framework, hidden source detection and the social influence learning are modeled separately. Although they share the same method of information processing, further efforts are needed to integrate them into one comprehensive model, such that the parameters that characterize social factor weights and hidden source distributions can be learned at the same time. Moreover, the interactions with hidden sources indicate an implicit contact network that should be combined with the explicit contact network to enrich the social influence learning.

Chapter 7

Appendices

7.1 Appendix A

The method for fusing BKB fragments [2] can be described as follows. The input is a set of n BKB fragments, $\{K_1, K_2, \ldots, K_n\}$ where $K_i = (G_i, w_i, \sigma_i, r(\sigma_i))$. Each fragment K_i has an information source represented by σ_i, and $r(\sigma_i)$ is the reliability of source σ_i. The output is a new BKB $K' = (G', w')$ with $G' = (I' \cup S', E')$, that is the fusion of n input fragments. For an I-node a in some fragment, let R_a be the random variable of which a is an instantiation. The fusion algorithm is depicted in Figure 7.1

7.2 Appendix B

We show that the opinions inferred from a fused BKB are not a linear combination of the original ones. Take Figure 2.2 as an example, let K_1,K_2 be the knowledge bases for actor 1 and 2, respectively. After fusion, K_1 and K_2 are merged into a fused BKB K'. Originally,

Bayesian-Knowledge-Fusion $\{K_1, K_2, \ldots, K_n\}$:
Let $G' = (I \cup S', E')$ be an empty correlation graph

For all fragments K_i with $i \leftarrow 1$ to n
 For all S-nodes $q \in S_i$
  Let $\alpha \leftarrow Desc_{G_i}(q)$
  Let the source I-node for q be $s = (S_{R_a} = \sigma_i)$
  Add q, all nodes connected to q in G_i, and the corresponding edges to G'
  Add s to G' along with an S-node supporting it
 end for
end for
Let ρ be a normalizing constant
For all S-node q' supporting some source node s
 Let $w'(q') \leftarrow r(s)/\rho$
end for
return $K' = (G', w')$

Figure 7.1: BKB fusion from Santos et al.[2]

the probability of "Vote" in K_1 is:

$$\begin{aligned} p_1(\text{Vote = Y}) &= p_1(\text{Interested in politics = Y}) \\ &* \; p_1(\text{Vote = Y | Interest in politics = Y}) \\ &= 0.6 * 0.9 = 0.54. \end{aligned}$$

Similarly, the probability of "Vote" in K_2 is:

$$p_2(\text{Vote = Y}) = 0.3 * 0.7 = 0.21.$$

Then the probability of "Vote" in the fused BKB K' becomes:

$$\begin{aligned} p_1(\text{Vote = Y}) &= [r_1 * p_1(\text{Interested in politics = Y} \\ &+ \; r_2 * p_2(\text{Interested in politics = Y})] \\ &* \; [r_3 * p_1(\text{Vote = Y | Interest in politics = Y}) \\ &+ \; r_4 * p_1(\text{Vote = Y | Interest in politics = Y})] \\ &= (r_1 * 0.3 + r_2 * 0.6) * (r_3 * 0.9 + r_4 * 0.7) \end{aligned}$$

However, if we integrate the information by weighted averaging the probabilities of "Vote" in K_1 and K_2, then $p'(Vote = Y)$ should be equal to:

$$\begin{aligned} p'_{linear}(\text{Vote = Y}) &= r_3 * p_1(\text{Vote = Y}) \\ &+ \; r_4 * p_2(\text{Vote = Y}) \\ &= 0.54r_3 + 0.21r_4 \end{aligned}$$

Apparently, $p'(Vote = Y)$ does not equal to $p'_{linear}(Vote = Y)$ in most cases (some special cases are discussed in Chapter 4).

7.3 Appendix C

With the learned factor weights using two methods, we can recreate the opinion trend over time. The procedure is the same as the generation of the synthetic opinion distribution, except the predefined factor weights are replaced with the learned weights. Then we measure their probability distribution distance with the synthetic data using the metric in Equation 5.10 from [114].

We run our experiment for 10 times and the results are provided in Table 7.1. As shown in the figure, the mean distribution distance of our approach and the baseline model is [0.0513, 0.0671], with p-value equals to 0.0009. This result suggests that the recreated opinion trend from our approach is closer to the true synthetic trend.

Distance	Proposed	Baseline
1	0.0655	0.0927
2	0.0641	0.1029
3	0.0364	0.0679
4	0.0326	0.0310
5	0.0582	0.0840
6	0.0693	0.0737
7	0.0235	0.0270
8	0.0202	0.0235
9	0.0541	0.0628
10	0.0892	0.1052
Average	0.05131	0.06707

Table 7.1: Distance measure between the true and learned probability distribution using different algorithm

7.4 Appendix D

The descriptions of the three datasets we used in Chapter 5 are listed as follows:

1. Dataset 1: Adults[1] This data consists of 13 features (6 continuous and 7 categorical) extracted from two types of income group: >50K and <= 50K per year. The total number of instances is 48842. The original goal is to determine whether a person makes over 50K a year based on the consensus information. In our experiment, >50K and <=50K groups are labelled as class 1 and class 2.

2. Dataset2: Cover type[2] This data consists of 54 features (10 continues and 44 categorical). The original goal is to predict forest cover type from cartographic variables only (no remotely sensed data). The total number of cover type is 7 which is determined from US Forest Service (USFS) Region 2 Resource Information System (RIS) data.

3. Dataset3: color-based image segmentation[3] This data consists of 19 continuous features and 2310 instances in total. The instances were drawn randomly from a database of 7 outdoor images. The images were hand segmented to create a classification for every pixel. Each instance is a 3x3 region. In our experiment, we select 2 classes: window (class 1) and grass (class 2) to generate the mixture data.

[1] https://archive.ics.uci.edu/ml/datasets/Adult
[2] https://archive.ics.uci.edu/ml/datasets/Covertype
[3] https://archive.ics.uci.edu/ml/datasets/Image+Segmentation

Bibliography

[1] Q. Gu, E. Santos Jr., and E. E. Santos, "Modeling opinion dynamics in a social network," in *Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM International Joint Conferences on*, vol. 2. IEEE, 2013, pp. 9–16.

[2] E. Santos, J. Wilkinson, and E. Santos, "Fusing multiple bayesian knowledge sources," *International Journal of Approximate Reasoning*, vol. 52, no. 7, pp. 935–947, 2011.

[3] E. Schwitzgebel, "Belief," in *The Stanford Encyclopedia of Philosophy*, spring 2014 ed., E. N. Zalta, Ed., 2014.

[4] H. C. Kelman, "Processes of opinion change," *Public opinion quarterly*, vol. 25, no. 1, pp. 57–78, 1961.

[5] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, no. 1, pp. 3–49, 2011.

[6] N. E. Friedkin and E. C. Johnsen, "Social influence networks and opinion change," *Advances in Group Processes*, vol. 16, no. 1, pp. 1–29, 1999.

[7] E. Bakshy, D. Eckles, R. Yan, and I. Rosenn, "Social influence in social advertising: evidence from field experiments," in *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 2012, pp. 146–161.

[8] Y.-M. Li and Y.-L. Shiu, "A diffusion mechanism for social advertising over microblogs," *Decision Support Systems*, vol. 54, no. 1, pp. 9–22, 2012.

[9] A. Madan, K. Farrahi, D. Gatica-Perez, and A. S. Pentland, "Pervasive sensing to model political opinions in face-to-face networks," in *Pervasive Computing*. Springer, 2011, pp. 214–231.

[10] D. Bindel, J. Kleinberg, and S. Oren, "How bad is forming your own opinion?" in *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*. IEEE, 2011, pp. 57–66.

[11] M. J. A. N. de Caritat *et al.*, *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. L'imprimerie royale, 1785.

[12] F. Galton, "Vox populi (the wisdom of crowds)," *Nature*, vol. 75, 1907.

[13] B. Golub and M. O. Jackson, "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, pp. 112–149, 2010.

[14] D. Krackhardt, "A plunge into networks," *Science*, vol. 326, no. 5949, pp. 47–48, 2009.

[15] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *Journal of Artificial Societies and Social Simulation*, vol. 5, no. 3, 2002.

[16] M. C. Jensen and W. H. Meckling, "Specific and general knowledge, and organizational structure," 1990.

[17] F. Wu and B. A. Huberman, "Social structure and opinion formation," *arXiv preprint cond-mat/0407252*, 2004.

[18] R. Galman and A. L. Brown, "Early foundations of cognitive development," pp. 46–56, 1985.

[19] A. L. Brown and A. S. Palincsar, "Guided, cooperative learning and individual knowledge acquisition," *Knowing, learning, and instruction: Essays in honor of Robert Glaser*, pp. 393–451, 1989.

[20] D. Van Knippenberg and D. Daamen, "Providing information in public opinion surveys: Motivation and ability effects in the information-and-choice questionnaire," *International Journal of Public Opinion Research*, vol. 8, no. 1, pp. 70–82, 1996.

[21] M. O. Jackson, *Social and economic networks*. Princeton University Press, 2010.

[22] W. Pan, W. Dong, M. Cebrian, T. Kim, J. H. Fowler, and A. Pentland, "Modeling dynamical influence in human interaction: Using data to make better inferences about influence within social systems," *Signal Processing Magazine, IEEE*, vol. 29, no. 2, pp. 77–86, 2012.

[23] D. Crandall, D. Cosley, D. Huttenlocher, J. Kleinberg, and S. Suri, "Feedback effects between similarity and social influence in online communities," in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2008, pp. 160–168.

[24] D. Bollier and C. M. Firestone, *The promise and peril of big data*. Aspen Institute, Communications and Society Program Washington, DC, USA, 2010.

[25] N. Silver, *The signal and the noise: Why so many predictions fail-but some don't*. Penguin, 2012.

[26] P. Pasquier, I. Rahwan, F. Dignum, and L. Sonenberg, "Argumentation and persuasion in the cognitive coherence theory," *Frontiers in artificial intelligence and applications*, vol. 144, p. 223, 2006.

[27] J. Golbeck, "Trust and nuanced profile similarity in online social networks," *ACM Transactions on the Web (TWEB)*, vol. 3, no. 4, p. 12, 2009.

[28] A. G. Greenwald and M. R. Banaji, "Implicit social cognition: attitudes, self-esteem, and stereotypes." *Psychological review*, vol. 102, no. 1, p. 4, 1995.

[29] E. E. Santos, E. Santos, J. Korah, J. E. Thompson, K. Kim, R. George, Q. Gu, J. Jurmain, S. Subramanian, and J. T. Wilkinson, "Intent-driven behavioral modeling during cross-border epidemics," in *Privacy, security, risk and trust (passat), 2011 ieee third international conference on and 2011 ieee third international conference on social computing (socialcom)*. IEEE, 2011, pp. 748–755.

[30] E. E. Santos, E. Santos, J. Korah, J. E. Thompson, Q. Gu, K. J. Kim, D. Li, J. Russell, S. Subramanian, Y. Zhang *et al.*, "Modeling emergent border-crossing behaviors during pandemics," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2013, pp. 87 110Z–87 110Z.

[31] E. Santos Jr, Q. Gu, E. E. Santos, and J. Korah, "Finite belief fusion model for hidden source behavior change detection." in *KDIR*, 2012, pp. 17–24.

[32] E. Santos Jr., Q. Gu, E. E. Santos, and J. Korah, "Hidden source behavior change tracking and detection," in *IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology*, vol. 2, 2012, pp. 51–58.

[33] E. Noelle-Neumann, *The spiral of silence: Public opinion–Our social skin*. University of Chicago Press, 1993.

[34] P. M. Sniderman and S. M. Theriault, "The structure of political argument and the logic of issue framing," *Studies in public opinion: Attitudes, nonattitudes, measurement error, and change*, pp. 133–65, 2004.

[35] L. Xiong and L. Liu, "Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities," *Knowledge and Data Engineering, IEEE Transactions on*, vol. 16, no. 7, pp. 843–857, 2004.

[36] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, pp. 992–1026, 1992.

[37] I. H. Lee, "On the convergence of informational cascades," *Journal of Economic theory*, vol. 61, no. 2, pp. 395–411, 1993.

[38] X. Vives, "Learning from others: a welfare analysis," *Games and Economic Behavior*, vol. 20, no. 2, pp. 177–200, 1997.

[39] V. P. Crawford and J. Sobel, "Strategic information transmission," *Econometrica: Journal of the Econometric Society*, pp. 1431–1451, 1982.

[40] J. Morgan and P. C. Stocken, "Information aggregation in polls," *The American Economic Review*, vol. 98, no. 3, pp. 864–896, 2008.

[41] I. Gilboa and D. Schmeidler, *A theory of case-based decisions*. Cambridge University Press, 2001.

[42] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.

[43] A. Garg, T. Jayram, S. Vaithyanathan, and H. Zhu, "Generalized opinion pooling." in *AMAI*, 2004.

[44] J. R. French Jr, "A formal theory of social power." *Psychological review*, vol. 63, no. 3, p. 181, 1956.

[45] T. M. Liggett, "Particle systems," 1985.

[46] S. Galam, Y. Gefen, and Y. Shapir, "Sociophysics: A new approach of sociological collective behaviour. i. mean-behaviour description of a strike," *Journal of Mathematical Sociology*, vol. 9, no. 1, pp. 1–13, 1982.

[47] P. Clifford and A. Sudbury, "A model for spatial conflict," *Biometrika*, vol. 60, no. 3, pp. 581–588, 1973.

[48] R. A. Holley and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," *The annals of probability*, pp. 643–663, 1975.

[49] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[50] J. E. Cohen, J. Hajnal, and C. M. Newman, "Approaching consensus can be delicate when positions harden," *Stochastic Processes and their Applications*, vol. 22, no. 2, pp. 315–322, 1986.

[51] I. Lobel and A. Ozdaglar, "Distributed subgradient methods for convex optimization over random networks," *Automatic Control, IEEE Transactions on*, vol. 56, no. 6, pp. 1291–1306, 2011.

[52] S. Chatterjee and E. Seneta, "Towards consensus: some convergence theorems on repeated averaging," *Journal of Applied Probability*, pp. 89–97, 1977.

[53] U. Krause, "A discrete nonlinear and non-autonomous model of consensus formation," *Communications in difference equations*, pp. 227–236, 2000.

[54] J. Lyst, K. Kacperski, and F. Schweitzer, "Social impact models of opinion dynamics," *Annual reviews of computational physics*, vol. 9, pp. 253–273, 2002.

[55] R. Stocker, D. G. Green, and D. Newth, "Consensus and cohesion in simulated social networks," *Journal of Artificial Societies and Social Simulation*, vol. 4, no. 4, 2001.

[56] S. Basu, T. Choudhury, B. Clarkson, A. Pentland *et al.*, "Learning human interactions with the influence model." NIPS, 2001.

[57] C. J. Tessone, R. Toral, P. Amengual, H. S. Wio, and M. San Miguel, "Neighborhood models of minority opinion spreading," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 39, no. 4, pp. 535–544, 2004.

[58] M. Gomez Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2010, pp. 1019–1028.

[59] E. Adar and L. A. Adamic, "Tracking information epidemics in blogspace," in *Web intelligence, 2005. Proceedings. The 2005 IEEE/WIC/ACM international conference on*. IEEE, 2005, pp. 207–214.

[60] M. Kimura, K. Saito, R. Nakano, and H. Motoda, "Finding influential nodes in a social network from information diffusion data," in *Social Computing and Behavioral Modeling*. Springer, 2009, pp. 1–8.

[61] J. Leskovec, L. Backstrom, R. Kumar, and A. Tomkins, "Microscopic evolution of social networks," in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2008, pp. 462–470.

[62] D. Liben-Nowell and J. Kleinberg, "Tracing information flow on a global scale using internet chain-letter data," *Proceedings of the National Academy of Sciences*, vol. 105, no. 12, pp. 4633–4638, 2008.

[63] D. J. Watts and P. S. Dodds, "Influentials, networks, and public opinion formation," *Journal of consumer research*, vol. 34, no. 4, pp. 441–458, 2007.

[64] M. Granovetter and R. Soong, "Threshold models of diffusion and collective behavior," *Journal of Mathematical sociology*, vol. 9, no. 3, pp. 165–179, 1983.

[65] U. Kuter and J. Golbeck, "Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models," in *AAAI*, vol. 7, 2007, pp. 1377–1382.

[66] J. Golbeck and J. Hendler, "Filmtrust: Movie recommendations using trust in web-based social networks," in *Proceedings of the IEEE Consumer communications and networking conference*, vol. 96. Citeseer, 2006.

[67] D. Quercia, S. Hailes, and L. Capra, "Lightweight distributed trust propagation," in *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*. IEEE, 2007, pp. 282–291.

[68] Y. Katz and J. Golbeck, "Social network-based trust in prioritized default logic," in *AAAI*, vol. 6, 2006, pp. 1345–1350.

[69] J. B. Rotter, "A new scale for the measurement of interpersonal trust1," *Journal of personality*, vol. 35, no. 4, pp. 651–665, 1967.

[70] D. Chong and J. N. Druckman, "A theory of framing and opinion formation in competitive elite environments," *Journal of Communication*, vol. 57, no. 1, pp. 99–118, 2007.

[71] P. Singla and M. Richardson, "Yes, there is a correlation:-from social networks to personal behavior on the web," in *Proceedings of the 17th international conference on World Wide Web*. ACM, 2008, pp. 655–664.

[72] E. Santos Jr and H. Dinh, "On automatic knowledge validation for bayesian knowledge bases," *Data & Knowledge Engineering*, vol. 64, no. 1, pp. 218–241, 2008.

[73] E. Santos Jr, Q. Gu, and E. E. Santos, "Bayesian knowledge base tuning," *International Journal of Approximate Reasoning*, vol. 54, no. 8, pp. 1000–1012, 2013.

[74] E. Santos Jr and E. S. Santos, "A framework for building knowledge-bases under uncertainty," *Journal of Experimental & Theoretical Artificial Intelligence*, vol. 11, no. 2, pp. 265–286, 1999.

[75] T. Rosen, S. Shimony, and E. Santos, "Reasoning with bkbs–algorithms and complexity," *Annals of Mathematics and Artificial Intelligence*, vol. 40, no. 3, pp. 403–425, 2004.

[76] E. Santos, E. Santos, J. Wilkinson, and H. Xia, "On a framework for the prediction and explanation of changing opinions," in *Systems, Man and Cybernetics, 2009. SMC 2009. IEEE International Conference on*. IEEE, 2009, pp. 1446–1452.

[77] N. Pioch, J. Melhuish, A. Seidel, E. Santos Jr, D. Li, and M. Gorniak, "Adversarial intent modeling using embedded simulation and temporal bayesian knowledge bases," DTIC Document, Tech. Rep., 2009.

[78] E. Santos Jr, K. Kim, F. Yu, D. Li, E. Jacob, L. Katona, and J. Rosen, "A method to improve team performance in the or through intent inferencing," in *Proc. EUROSIM Congress on Modelling and Simulation*, vol. 2, 2010.

[79] E. W. Morrison, "Newcomers' relationships: The role of social network ties during socialization," *Academy of management Journal*, vol. 45, no. 6, pp. 1149–1160, 2002.

[80] S. Wasserman, *Social network analysis: Methods and applications*. Cambridge university press, 1994, vol. 8.

[81] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins, "Information diffusion through blogspace," in *Proceedings of the 13th international conference on World Wide Web*. ACM, 2004, pp. 491–501.

[82] C. X. Lin, B. Zhao, Q. Mei, and J. Han, "Pet: a statistical model for popular events tracking in social communities," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2010, pp. 929–938.

[83] C. Danescu-Niculescu-Mizil, G. Kossinets, J. Kleinberg, and L. Lee, "How opinions are received by online communities: a case study on amazon. com helpfulness votes," in *Proceedings of the 18th international conference on World wide web*. ACM, 2009, pp. 141–150.

[84] T. Damer, *Attacking faulty reasoning*. Cengage Learning, 2012.

[85] I. Zukerman, R. McConachy, and K. B. Korb, "Bayesian reasoning in an abductive mechanism for argument generation and analysis," in *AAAI/IAAI*, 1998, pp. 833–838.

[86] G. Carenini and J. D. Moore, "Generating and evaluating evaluative arguments," *Artificial Intelligence*, vol. 170, no. 11, pp. 925–952, 2006.

[87] T. M. Gureckis and D. B. Markant, "Self-directed learning a cognitive and computational perspective," *Perspectives on Psychological Science*, vol. 7, no. 5, pp. 464–481, 2012.

[88] A. J. Gonzalez and D. D. Dankel, *The engineering of knowledge-based systems: theory and practice*. Prentice Hall Englewood Cliffs, 1993.

[89] R. Knauf, S. Tsuruta, and A. J. Gonzalez, "Toward reducing human involvement in validation of knowledge-based systems," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 37, no. 1, pp. 120–131, 2007.

[90] N. R. Draper and H. Smith, "Applied regression analysis 2nd ed." 1981.

[91] V. Kandiah and D. L. Shepelyansky, "Pagerank model of opinion formation on social networks," *Physica A: Statistical Mechanics and its Applications*, vol. 391, no. 22, pp. 5779–5793, 2012.

[92] R. Huckfeldt and J. Sprague, "Discussant effects on vote choice: Intimacy, structure, and interdependence," *The Journal of Politics*, vol. 53, no. 01, pp. 122–158, 1991.

[93] S. Turkle, *Alone together: Why we expect more from technology and less from each other*. Basic books, 2012.

[94] M. MacKuen, "Exposure to information, belief integration, and individual responsiveness to agenda change," *The American Political Science Review*, pp. 372–391, 1984.

[95] R. B. Cialdini, *Influence: Science and practice*. Allyn and Bacon Boston, MA, 2001, vol. 4.

[96] D. Strang, "Adding social structure to diffusion models an event history framework," *Sociological Methods & Research*, vol. 19, no. 3, pp. 324–353, 1991.

[97] K. Das, J. Schneider, and D. B. Neill, "Anomaly pattern detection in categorical datasets," in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2008, pp. 169–176.

[98] J. L. Hill, "An extension and test of converseÕs Òblack-and-whiteÓ model of response stability," in *American Political Science Association*, vol. 95, no. 02. Cambridge Univ Press, 2001, pp. 397–413.

[99] P. Garcia-Teodoro, J. Diaz-Verdejo, G. Maciá-Fernández, and E. Vázquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges," *computers & security*, vol. 28, no. 1, pp. 18–28, 2009.

[100] W.-K. Wong, A. Moore, G. Cooper, and M. Wagner, "Bayesian network anomaly pattern detection for disease outbreaks," in *ICML*, 2003, pp. 808–815.

[101] T. Dean and K. Kanazawa, "A model for reasoning about persistence and causation," *Computational intelligence*, vol. 5, no. 2, pp. 142–150, 1989.

[102] J. W. Robinson and A. J. Hartemink, "Learning non-stationary dynamic bayesian networks," *The Journal of Machine Learning Research*, vol. 9999, pp. 3647–3680, 2010.

[103] M. Gales and S. Young, "The application of hidden markov models in speech recognition," *Foundations and Trends in Signal Processing*, vol. 1, no. 3, pp. 195–304, 2008.

[104] G. Cybenko, V. H. Berk, V. Crespi, G. Jiang, and R. S. Gray, "An overview of process query systems," in *Defense and Security*. International Society for Optics and Photonics, 2004, pp. 183–197.

[105] A. P. Dempster, N. M. Laird, D. B. Rubin *et al.*, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal statistical Society*, vol. 39, no. 1, pp. 1–38, 1977.

[106] G. McLachlan and D. Peel, *Finite mixture models*. John Wiley & Sons, 2004.

[107] M. A. Figueiredo and A. K. Jain, "Unsupervised learning of finite mixture models," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 3, pp. 381–396, 2002.

[108] E. Aronson, T. D. Wilson, and R. M. Akert, *Social Psychology*. Upper Saddle River, NJ Prentice Hall, 2010.

[109] H. Gintis, "The foundations of behavior: the beliefs, preferences, and constraints model," 2006.

[110] J. Nocedal and S. J. Wright, "Numerical optimization 2nd," 2006.

[111] B. Thiesson, C. Meek, D. M. Chickering, and D. Heckerman, "Learning mixtures of bayesian networks," in *Proc. Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI98)*. Citeseer, 1998, pp. 504–513.

[112] F. Wiseman, "Methodological bias in public opinion surveys," *Public Opinion Quarterly*, pp. 105–108, 1972.

[113] C. Chatfield, *The analysis of time series: an introduction*. CRC press, 2013.

[114] H. Chan and A. Darwiche, "A distance measure for bounding probabilistic belief change," *International Journal of Approximate Reasoning*, vol. 38, no. 2, pp. 149–174, 2005.

[115] G. E. Box, J. S. Hunter, and W. G. Hunter, *Statistics for experimenters*. Wiley New York, 2005.

[116] K. Bache and M. Lichman, "UCI machine learning repository," 2013. [Online]. Available: http://archive.ics.uci.edu/ml

[117] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The weka data mining software: an update," *ACM SIGKDD explorations newsletter*, vol. 11, no. 1, pp. 10–18, 2009.

[118] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter," in *Proceedings of the 20th international conference on World wide web*. ACM, 2011, pp. 695–704.

[119] N. R. Buchan, R. T. Croson, and S. Solnick, "Trust and gender: An examination of behavior and beliefs in the investment game," *Journal of Economic Behavior & Organization*, vol. 68, no. 3, pp. 466–476, 2008.