Information Fusion 26 (2015) 66–83

Contents lists available at ScienceDirect

Information Fusion

journal homepage: www.elsevier.com/locate/inffus

Partial observable update for subjective logic and its application for trust estimation

Lance Kaplan [a], Murat Şensoy [b,e,*], Supriyo Chakraborty [c], Geeth de Mel [d]

[a] US Army Research Lab, Adelphi, MD 20783, USA
[b] Department of Computer Science, Ozyegin University, Istanbul, Turkey
[c] Department of Electrical Engineering, UCLA, Los Angeles, CA 90095, USA
[d] IBM T. J. Watson Research Center, Hawthorne, NY, USA
[e] Department of Computing Science, University of Aberdeen, UK

ARTICLE INFO

Article history:
Received 3 July 2014
Received in revised form 27 January 2015
Accepted 28 January 2015
Available online 7 February 2015

Keywords:
Subjective logic
Partial observations
Dirichlet distributions
Uncertain information
Trust management

ABSTRACT

Subjective Logic (SL) is a type of probabilistic logic, which is suitable for reasoning about situations with uncertainty and incomplete knowledge. In recent years, SL has drawn a significant amount of attention from the multi-agent systems community as it connects beliefs and uncertainty in propositions to a rigorous statistical characterization via Dirichlet distributions. However, one serious limitation of SL is that the belief updates are done only based on completely observable evidence. This work extends SL to incorporate belief updates from partially observable evidence. Normally, the belief updates in SL presume that the current evidence for a proposition points to only one of its mutually exclusive attribute states. Instead, this work considers that the current attribute state may not be completely observable, and instead, one is only able to obtain a measurement that is statistically related to this state. In other words, the SL belief is updated based upon the likelihood that one of the attributes was observed. The paper then illustrates properties of the partial observable updates as a function of the state likelihood and illustrates the use of these likelihoods for a trust estimation application. Finally, the utility of the partial observable updates is demonstrated via various simulations including the trust estimation case.

1. Introduction

Decision support systems have demonstrated great utility in many different applications in the commercial and military sectors [1–3]. The performance of such systems has been steadily improving over the years as they incorporate big data [4] and collect data about users' preferences [5]. However, in all of these applications, such systems simply retrieve existing information. There is a desire for decision support systems to augment human intelligence beyond retrieving information by predicting the evolution of current events to help recommend potential actions. This ability is termed situational awareness [6] and is currently a human cognitive function. There have been some efforts to achieve situation awareness through multi-agent architectures [7,8]. Furthermore, situation awareness concepts and principles can in turn enhance multiagent systems in a wide-range of applications [9,10]. The problem is that the world is uncertain. Quantum issues aside, one might argue that complete knowledge of the state of the world at any instance (including the cognitive state of all human actors in this world) enables one to predict how the world state evolves over time. However, the complete knowledge at the necessary precision is unattainable, and at the fidelity at which one can model the world in a computer (or in ones head), it tends to be highly probabilistic. For instance, when one drives to work in the morning, there is a probability that they will encounter traffic. Of course, one can listen to the radio and nowadays they can also use the GPS in their smartphones to determine if this probability is near one or zero. However, there is always a chance that a very recent accident will cause one to get stuck in unforseen traffic. The determination of this exact probability is obtained by collecting evidence similar to rolling a weighted die. There is always uncertainty associated to finite evidence, and the uncertainty is significant when the evidence is limited. In short, decision support systems must be able to reason and predict over such uncertainty.

Because of the prevalence of uncertainty, reasoning under uncertainty has a very rich history [11]. Certainly, probability distributions provide a representation of uncertainty, and Bayesian

* Corresponding author at: Department of Computer Science, Ozyegin University, Istanbul, Turkey.

E-mail addresses: lkaplan@ieee.org (L. Kaplan), murat.sensoy@ozyegin.edu.tr (M. Şensoy), supriyo@ucla.edu (S. Chakraborty), grdemel@us.ibm.com (G. de Mel).

http://dx.doi.org/10.1016/j.inffus.2015.01.005
1566-2535/

approaches provide sound methods to manipulate these distributions. However, many argue that Bayes' rule is unable to properly handle conflicting evidence, and many alternatives to probability theory such as belief theory [11–13] and fuzzy set theory [14] have emerged to enable fusion of information under these conflicts. At this point, it is unclear that the issue is with Bayesian reasoning per se, or with exploitation of a limited version of Bayes (first order versus higher order [15]).

Recently, *Subjective Logic* (SL) has emerged as a probabilistic logic that exploits second order Bayes to incorporate uncertainty by connecting the belief mass assignment (BMA) to a Dirichlet distribution [16,17]. It provides a rigorous and computationally efficient method to represent and reason over human generated or automated beliefs in face of uncertainty. Applications of SL include trust management [18], Bayesian networks [19], and fusion [20,17]. In short, SL provides effective tools to manage and combine beliefs over a set of mutually exclusive attribute states from multiple human or software agents. At a given point in time, an agent's belief, known as a subjective opinion, is the result of a prior belief and a set of observations. The uncertainty of the belief represents the reliance on the prior, and the uncertainty decreases as the agent incorporates more observations to form the beliefs over the set of attribute states.

SL includes many operations to fuse and infer over subjective opinions from many different agents. These opinions are built through a set of observations. Each observation is the occurrence of one of the mutually exclusive attribute states as observed at a given time. The collection of observations forms the evidence, which is the number of times each of the mutually exclusive attribute states manifest in the past observations. This is akin to rolling a weighted die multiple times and tabulating the results of the rolls. For instance, the probability of encountering traffic is formed by counting the number of past instances of encountering traffic or not over a stretch of roadway at a similar time of day.

To make these notions even more concrete, let us consider another motivating example where one wants to understand the criminal activity within a city. Specifically, one wants to understand if a crime happens, what is the probability that the crime occurs in any one of the city districts. Without any initial data, one might look at socio-economic factors to develop an initial set of probabilities. Over time, one can log where a crime occurs and start to use these observations to update the probabilities. Clearly, as more observations are logged, the certainty associated with the generated probabilities increases. SL is well suited to infer the probabilities of a crime occurring in the districts and the uncertainty associated to these probabilities.

Now, let us assume that one is interested in where criminals live. The question is now when a crime occurs, what is the probability that the perpetrator lives in a particular district. Like before, one can start with a prior set of probabilities based upon the socio-economics factors. Furthermore, when a crime occurs, the location of the crime is readily available in the police report. However, the identity of the perpetrator may or may never be discovered. Therefore, it is generally not possible to log where the perpetrator lives. Sometimes, this information can be determined with great certainty when the criminal is caught. Most likely, one can only incorporate the occurrence of the crime as a *partial observation* through statistical models that link the probability of where the perpetrator lives conditioned on where a crime occurs. For instance, a criminal may not operate in his/her immediate neighborhood where he/she can easily be identified, and a criminal may not want to venture too far away either. This contextual information can help answer the questions of the distribution of criminals over the various districts within a city. This scenario is an example of a geospatial abduction problem (GAP) [21]. SL is suited to tackle such applications, but the notions of how to incorporate statistical (and not just hard) evidence of the appearance of an attribute (the home district of a perpetrator) need to be developed within the SL framework.

Another example of statistical evidence occurs in reputation systems. SL has been incorporated in reputation systems to form opinions about the trustworthiness of other agents to act honestly in a financial transaction [22]. It assumes, however, that one can clearly label a prior transaction as being honest or not. Many times, the trustor can only determine the likelihood that the trustee actually had cut corners and provided a product inferior than advertised.

This work extends the theory of SL to enable the update of the BMA under such partial observations where one piece of evidence is not a result of the die roll, but a measurement of the die roll indicating the likelihood of the result. Our preliminary work considered sequential updating of opinions from single partial observations [23] by approximating the updated posterior distribution by a Dirichlet distribution through moment matching. This paper extends these results in two respects. First, it considers an update that incorporates multiple observations. Second, the preliminary work in [23] assumes that the likelihoods are known precisely without demonstrating how they are obtained. This paper shows an example of how to derive the likelihoods for a trust estimation application where trustworthy behavior after each transaction is not directly observed. Specifically, this paper demonstrates the likelihood update for trust estimation where agents share opinions about various propositions. To this end, the likelihood that an agent's opinion is an honest report given the trustees opinion about the same proposition is formulated. The ability to rigorously update the subjective opinions about agents from these partial observations provides a novel and important contribution in the area of trust estimation.

The rest of this paper is organized as follows. We present the related work in Section 2. Some basic concepts about Dirichlet distributions and the mathematical notations used in this paper are given in Section 3. Section 4 reviews SL and demonstrates belief updates in SL for fully visible observations. Then, Section 5 expands SL to incorporate partial updates. Examples of the behavior of partial updates are presented in Section 6, and Section 7 derives the information consistency likelihood for trust estimation. Simulations to evaluate the performance of partial observation updates in general and for the trust estimation application are provided in Section 8. Finally, Section 9 provides concluding remarks.

2. Related work

We group prior work into two broad categories. The first category includes work that focus on the development of SL as a logic technique. The second category outlines the effort in the direction of trust management and reputation based systems.

2.1. Development of SL

This work is most directly influenced by Jøsang's efforts to develop SL [16,17]. SL is a type of probabilistic logic that explicitly takes uncertainty and belief ownership into account. A formal introduction to SL is presented in Section 4. In general, SL is suitable for modeling and analyzing situations involving uncertainty and incomplete knowledge. It builds upon the extensive work of evidential reasoning under uncertainty with extensive influence from the seminal theory of Dempster and Shafer [11]. Arguments in SL are subjective opinions about propositions. The key aspect of SL is that it expresses an opinion as a Dirichlet distribution and maps the parameters of this distribution to a basic belief assignment. The interpretation of an opinion as a Dirichlet