SOBE: Source Behavior Estimation for Subjective Opinions In Multiagent Systems *

(Extended Abstract)

Murat Şensoy[1], Lance Kaplan[2], Geeth de Mel[3], Taha Dogan Gunes[1]
[1]Department of Computer Science, Ozyegin University, Istanbul, Turkey
[2] US Army Research Lab, Adelphi, MD 20783, USA
[3]IBM T. J. Watson Research Center, NY, Hawthorne, US
murat.sensoy@ozyegin.edu.tr

ABSTRACT

In cooperative or hostile environments, agents communicate their subjective opinions about various phenomenon. However, sources of these opinions may not always be competent and honest but more likely erroneous or even malicious. Furthermore, malicious sources may adopt certain behaviors to mislead the decision maker in a specific way. Fortunately, the reports of such misleading sources are correlated to ground truth. In this work, we propose to learn statistically meaningful opinion transformations that represent various behaviors of information sources. Then, we exploit these transformations while fusing opinions from unreliable sources. We show that our approach can be used to determine set of transformations that may lead to more accurate estimation of the truth.

Keywords

Trust, Agents, Fusion

1. INTRODUCTION

In multi-agent systems, agents perceive information about their environment, evaluates the situation, and make decisions to achieve their goals. In partially observable and uncertain environments, agents' own observations and perceptions may be unavailable, incomplete, or inaccurate to make informed decisions. In such settings, agents may gather information from diverse and unreliable sources. However, there is no guarantee that the information collected from these sources are useful or at least not misleading. In this work, we argue that the misleading information from unreliable sources can also be utilized to estimate the truth if the underlying correlations between the truth and the information are estimated correctly and exploited intelligently. Therefore, in this paper, we propose to learn statistically meaningful transformations that represent various behaviors of information sources. This is achieved by applying multi-dimensional Hough transform [1] on the historical data. Then, we exploit the discovered transformations within a statistical framework to estimate the ground truth. The approach proposed in this paper allows agents to estimate the ground truth even when there is no competent source that is consistently honest while sharing information; that is, our approach can estimate the ground truth accurately using only the misleading information.

2. BEHAVIOR DISCOVERY

In this work, agents and information sources represent their opinions about binary propositions using Beta distribution parameters. A binary proposition has two mutually exclusive outcomes such as $true$ and $false$ or yes and no. The Beta distribution has two parameters $\langle \alpha_1, \alpha_2 \rangle$ that defines the likelihood of the probability for each outcome. Ideally, a source may construct its opinions through Bayesian update of a prior distribution (e.g., uniform distribution) using its observations. For instance, if the source observes r positive evidence for and s negative evidence against the proposition, its genuine opinion for the proposition should be constructed as $\langle r+1, s+1 \rangle$ using uniform distribution as prior. The uniform distribution corresponds to a beta distribution with parameters $\langle 1, 1 \rangle$, which means that each outcome is equally likely.

A decision maker may observe a specific correlation between its own opinions and the opinions shared by the sources adopting a specific behavior. Let us assume a competent source provides $\boldsymbol{\alpha_o} = \langle r_o + 1, s_o + 1 \rangle$ as its opinion for a binary proposition. Similarly, the decision maker has its own opinion $\boldsymbol{\alpha_d} = \langle r_d + 1, s_d + 1 \rangle$ for the same proposition. The total number of evidence used in these opinions are $n_o = r_o + s_o$ and $n_d = r_d + s_d$, respectively. We use this opinion pair for training if these opinions are certain enough; that is, n_o and n_d are greater than a predefined threshold. Given the opinion pair, the decision maker can easily compute the expected number of positive evidence ($r_{d|o}$) and negative evidence ($s_{d|o}$) if n_o evidence were observed by itself as follows:

$$r_{d|o} = n_o \times \frac{r_d + 1}{r_d + s_d + 2} \quad \text{and} \quad s_{d|o} = n_o \times \frac{s_d + 1}{r_d + s_d + 2}$$

It is obvious that $\boldsymbol{\alpha_d}$ is consistent with the resulting opinion $\boldsymbol{\alpha_{d|o}} = \langle r_{d|o} + 1, s_{d|o} + 1 \rangle$. However, its consistency with $\boldsymbol{\alpha_o}$ depends on the source's behavior. In this work, we focus on the cases where the decision maker's opinion is consistent with $\langle r'_o + 1, s'_o + 1 \rangle$,

*Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Numbers W911NF-06-3-0001 and W911NF-13-1-0243. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Dr. Şensoy thanks to the U.S. Army Research Laboratory for its support under grant W911NF-14-1-0199 and The Scientific and Technological Research Council of Turkey (TUBITAK) for its support under grant 113E238.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

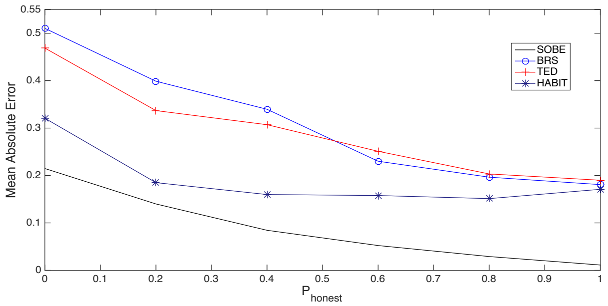

Figure 1: Variation of mean absolute error as P_{honest} varies.

which is calculated using the following linear transformation of the source's original opinion: $[r'_o,\ s'_o] = [r_o,\ s_o,\ 1] \times M$, where M is an 3x2 transformation matrix defined by six parameters: $\left(\begin{smallmatrix} a & b & c \\ d & e & f \end{smallmatrix}\right)^T$. The transformation matrix is determined directly by the behavior adopted by the source. If a source is competent and honestly shares its opinions, M should be $\left(\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{smallmatrix}\right)^T$, hence $r'_o = r_o$ and $s'_o = s_o$. On the other hand, if the source tries to mislead the decision maker by flipping parameters of its opinion, M should be $\left(\begin{smallmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \end{smallmatrix}\right)^T$, hence $r'_o = s_o$ and $s'_o = r_o$.

While the decision maker does not know about the behaviors of sources, it can discover transformation matrices relevant to these behaviors. For this purpose, the decision maker derives a training set from $\langle \alpha_o^i, \alpha_d^i \rangle$ pairs, each of which includes opinion of a source and opinion of the decision maker for the same proposition i. The derived training set contains $\langle \alpha_o^i, \alpha_{d|o}^i \rangle$ pairs, where $\alpha_{d|o}^i$ is derived as explained before. Using this training set, the decision maker aims to learn meaningful linear transformations that map α_o^i to $\alpha_{d|o}^i$. Each transformation matrix encodes two linear equations:

$$\begin{aligned} r_{d|o} &\simeq a \times r_o + b \times s_o + c \\ s_{d|o} &\simeq d \times r_o + e \times s_o + f \end{aligned} \tag{1}$$

To discover parameters of arbitrary number of meaningful transformations, we propose to use 3D *Hough transform* in this work [1].

3. FRAMEWORK OVERVIEW

Our framework maintains a history of personal opinions of decision maker as well as opinions from various information sources. Behavior discovery module uses *Hough Transform* [1] to learn various linear transformations that translate misleading opinions from information sources to more informative opinions. Each discovered transformation represents a specific source behavior. Source modeler learns behavior profile of each source, which is represented as a Dirichlet distribution over the probability of adopting discovered behaviors, using Bayesian update. When the decision maker needs opinions related to a proposition, opinion retriever module gathers opinions related to the proposition from various information sources. Using the computed behavior profiles of information sources, information fusion component combines the gathered opinions using moment matching to come up with a single fused opinion. This opinion is used by decision maker as an estimate of the ground truth about the proposition. If the decision maker observes evidence related to the proposition, it can compose its own opinion and stores it in the opinion history as a personal opinion. Later, personal opinions are used, as explained above, to discover new behaviors and update behavior profiles of information sources.

4. EVALUATION

In this section, we compare our approach with existing approaches: BRS [4], HABIT [3], and TED [2]. We conduct a number of simulations to evaluate our approach. Each simulation last 200 time steps, which is enough for convergence of the approaches. At each step, a decision maker is given a binary proposition and it contacts 50 information sources to estimate the ground truth about this proposition. Each source can observe a number of positive and negative evidence about the ground truth. Based on the observed evidence, the sources compose their genuine opinions. Upon decision maker's request, each source adopts one of five behaviors with different probabilities: i) honestly sharing opinion, ii) providing random opinion, iii) flipping the opinion parameters, iv) being optimistic by doubling positive evidence in opinion, v) being pessimistic by doubling negative evidence in opinion. After collecting opinions, the decision maker uses BRS, HABIT, TED, and SOBE — the proposed approach — to fuse these opinions in order to approximate the ground truth. Lastly, the decision maker observes some evidence about the ground truth and composes its own opinion. This opinion is used by these approaches to learn trustworthiness or behaviors of sources. At the end of each step, we compute the *fusion error*, which is defined as the absolute difference of the expectation value of the fused opinion and the ground truth. During simulations, each source changes its identity to whitewash its reputation with probability 0.1 at each step.

In Figure 1, we show how the mean fusion error of the approaches change as we vary the probability that a source predominantly honest (P_{honest}); predominantly honest sources pick the honest behavior among others with probability 0.7. Our results show that SOBE achieves the best performance; its error is slightly above 0.2 when $P_{honest} = 0$, and drops down to 0.02 as P_{honest} increases. HABIT also achieves low fusion error especially when there are some honest sources. Also, as shown on the figure, performances of other approaches improve as P_{honest} increases; their mean absolute errors approach to that of HABIT. SOBE is the only method than can correct for misleading reports and that is why SOBE does much better even when P_{honest} is high.

5. CONCLUSIONS

In this paper, we propose a comprehensive framework for behavior discovery and likelihood estimation for information sources at multi-agent systems. Our approach aims decision makers to exploit information from malicious sources by discovering the underlying correlations between their opinions and the truth perceived by the decision maker. We demonstrate that the proposed approach can successfully estimate the truth when there is no consistently honest source and sources change their identities to blur trust estimations.

REFERENCES

[1] P. Bhattacharya, H. Liu, A. Rosenfeld, and S. Thompson. Hough-transform detection of lines in 3-d space. *Pattern Recognition Letters*, 21(9):843 – 849, 2000.

[2] L. Kaplan, M. Sensoy, and G. de Mel. Trust estimation and fusion of uncertain information by exploiting consistency. In *Proceedings of the 17th International Conference on Information Fusion (FUSION)*, pages 1–8, July 2014.

[3] W. L. Teacy, M. Luck, A. Rogers, and N. R. Jennings. An efficient and versatile approach to trust and reputation using hierarchical bayesian modelling. *Artificial Intelligence Journal*, 193:149–185, 2012.

[4] A. Whitby, A. Jøsang, and J. Indulska. Filtering out unfair ratings in bayesian reputation systems. *The ICFAIN Journal of Management Research*, 4(2):48–64, 2005.