18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

FUSE-BEE: Fusion of Subjective Opinions through Behavior Estimation

Murat Şensoy*, Lance Kaplan‡, Gonul Ayci*, and Geeth de Mel§
* Department of Computer Science, Ozyegin University, Istanbul, Turkey.
‡ US Army Research Lab, Adelphi, MD 20783, USA.
§ IBM T. J. Watson Research Center, Hawthorne, NY, USA.

***Abstract*—Information is critical in almost all decision making processes. Therefore, it is important to get the right information at the right time from the right sources. However, information sources may behave differently while providing information — i.e., they may provide unreliable, erroneous, noisy, or misleading information deliberately or unintentionally. Motivated by this observation, in this paper, we propose a statistical information fusion approach based on behavior estimation. Our approach transforms the conveyed information into more useful form by tempering them with the estimated behaviors of sources. Through extensive simulations, we have shown that our approach has a lower computational complexity, and achieves significantly low behavior estimation and fusion errors.**

Keywords*—Information Fusion, Subjective Logic, Dirichlet Distributions, Behavior Estimation.

I. Introduction

Effective fusion of information from diverse – at times unreliable – sources is an important problem to be solved for decision making domain, especially in coalition context; the purpose of fusion is to merge information from multitude of sources to estimate the ground truth of a specific phenomenon. An ideal information source is the one who has the capability to estimate the ground truth – e.g., by combining observed evidence – and reports its true estimation. However, this may not be the case for most of the information sources. For example, some information sources can be competent, but not honest — i.e., they deliberately diverge from their genuine estimations while reporting; other sources may be incompetent in observing phenomena, thus unable to estimate the ground truth at all — i.e., their estimations of the ground truth may not correlate with the ground truth and would not conduct any useful information during fusion.

In this paper, a source's estimation of ground truth is represented as a *subjective opinion*, which is a belief assignment over possible values of the ground truth. However, these estimations may be affected by the behaviors of sources due to the operational context. For example, in order to mislead a decision maker – or due to incompetence in the context – a source may share an opinion, which does not correlate or negatively correlate with the ground truth. Moreover, an information source may not be consistent in its behavior — i.e., it may adopt different behavior strategies with varying granularities.

We note that information fusion aims to approximate the ground truth by combining opinions collected from diverse and unreliable sources. Existing fusion approaches exploit trust estimation methods to determine trustworthy and untrustworthy sources [7], [2]; opinions from untrustworthy sources are then eliminated during fusion. However, these approaches may fail if the sources are not always trustworthy or untrustworthy — i.e., they adopt different behaviors with varying probabilities. Moreover, we note that filtering misleading information may not always be the best thing to do during fusion — i.e., misleading information can be useful if it is correlated with the ground truth. For instance, let us assume a situation where a decision maker asks a *yes/no* questions from a source. The source aims to mislead the decision maker by always providing the incorrect response. In such situations, if the decision maker can determine the behavior of the source, the source's answers still could valuable as they expose features about the source; otherwise, these answers would be highly misleading.

In this paper, we propose a novel fusion framework based on behavior estimation of sources. Using the framework, a decision maker can query information sources to estimate the outcome of a binomial or multinomial propositions, e.g., *is there any traffic jam on the road I-87 now?*. As stated earlier, we expect the answers to such queries to be subjective opinions, which can then be interpreted using Subjective Logic [3] or Dempster-Shaffer theory of evidence [15]. In our framework, we adopt Subjective Logic's interpretation of subjective opinions. That is, opinions are represented using Beta or Dirichlet distributions and they are fused by aggregating these distributions. Our system is flexible enough to accommodate various source behaviors, and for each source, we calculate the behavior probabilities using maximum likelihood estimation. During the fusion, we efficiently determine the most likely behavior of sources using a similarity-based clustering of shared opinions influenced by the estimated source behavior probabilities. We then apply specific transformations to the shared opinions based on the behaviors of their sources, and combine them to estimate the ground truth. Through extensive simulations, we show that our approach can successfully estimate behaviors of information sources and approximate the ground truth.

The rest of the paper is organized as follows: Section II provides preliminary information on Dirichlet distributions and Subjective Logic. Section III introduces behaviors of information sources and Section IV describes how these behaviors are modeled. Section V proposes our fusion approach based on behavior estimation. Section VI evaluates our approach through simulations using various configurations. Section VII discusses the proposed approach with respect to the related work, and we conclude the paper in Section VIII by summarizing main contributions and drawing the future directions of the work.

II. Preliminaries

In this section, we introduce the basics of Dirichlet distributions and Subjective Logic. Let us note that all vectors in this paper are column vectors and are represented in boldface such as $\mathbf{x}$ where the k-th element is given by x_k; the transpose of a vector $\mathbf{x}$ is denoted by $\mathbf{x}^T$.

A. Dirichlet Distributions

The Dirichlet distribution is a probability density function (pdf) for the possible values of a probability mass function (pmf) $\mathbf{p}$ that describes the probability for the manifestation of the particular state from the K attribute states. It is characterized by K parameters $\boldsymbol{\alpha}$ and is given by

$$f_\beta(\mathbf{p}|\boldsymbol{\alpha}) = \begin{cases} \frac{1}{B(\boldsymbol{\alpha})} \prod_{i=1}^{K} p_i^{\alpha_i - 1} & \text{for } \mathbf{p} \in \mathcal{S}_K, \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where $\mathcal{S}_K$ is the K-dimensional unit simplex,

$$\mathcal{S}_K = \left\{ \mathbf{p} | \textstyle\sum_{i=1}^{K} p_i = 1 \text{ and } 0 \le p_1, \ldots, p_K \le 1 \right\},$$

and

$$B(\boldsymbol{\alpha}) = \int_{\mathcal{S}_K} \left(\prod_{i=1}^{K} p_i^{\alpha_i - 1} \right) d\mathbf{p} \tag{2}$$

$$= \frac{\prod_{i=1}^{K} \Gamma(\alpha_i)}{\Gamma\left(\sum_{i=1}^{K} \alpha_i \right)} \tag{3}$$

is the K-dimensional multinomial beta function [8]. The β in the subscript of f — i.e., $f_\beta(\cdot)$ — is used to signify that the pdf is Dirichlet. Furthermore, when $K = 2$, the Dirichlet distribution simplifies to a beta distribution.

B. Subjective Logic

Subjective Logic (SL) is a probabilistic logic where propositions such as the location of a crime in a city can take on one of K mutually exclusive attributes, e.g, city districts, at any observation time [3], [5]. A subjective opinion characterizes the belief – in probabilities – that any of the K attributes will appear at a given observation time; furthermore, it also captures the uncertainty related to these beliefs explicitly. Formally, SL considers a frame of K mutually exclusive singletons by providing a belief mass b_k for each singleton $k = 1, \ldots, K$ and providing an overall uncertainty mass of u. These $K+1$ mass values are all non-negative and sum up to one, i.e.,

$$u + \sum_{k=1}^{K} b_k = 1, \tag{4}$$
$$\text{where } u \ge 0 \text{ and } b_k \ge 0 \text{ for } k = 1, \ldots, K.$$

SL also includes a base rate probability a_k for each singleton and a non-informative prior weight W, which are same for all opinions about the same proposition. The collection of all the parameters for agent s about proposition x forms agent s's subjective multinomial opinion $\boldsymbol{\omega}^{s:x} = [\ (\mathbf{b}^{s:x})^T \quad u^{s:x} \quad (\mathbf{a}^x)^T \quad W\ ]^T$. In this paper, when the proposition and agent are implicit, the superscripts and subscripts are not used. The base rate values represent initial (or a priori) information about the probability of a singleton emerging for any given observation. The inclusion of the belief and uncertainty values along with the base rates and non-informative prior weight represent the accrued evidence regarding the probability of any singleton appearing in an observation. For each singleton, we can compute the amount of evidence observed using the multinomial opinion values:

$$e_k = \frac{W b_k}{u}. \tag{5}$$

The computed evidence vector $\boldsymbol{e}$ can then be used to compute the parameter vector $\boldsymbol{\alpha}$ for a Dirichlet distribution via

$$\alpha_k = e_k + W a_k. \tag{6}$$

Thus, we have $\boldsymbol{\alpha} = \boldsymbol{e} + W\boldsymbol{a}$. The Dirichlet distribution represented by this parameter vector represents a possible pmf which controls the appearance of singletons in observations. Likewise, using (4), solving for b_k and u in (5) for $k = 1, \ldots, K$, leads to the mapping of evidence vector $\boldsymbol{e}$ – and $\boldsymbol{\alpha}$ – to the multinomial opinions

$$u = \frac{W}{\sum_i e_i + W}, \tag{7a}$$

$$b_k = \frac{e_k}{\sum_i e_i + W}. \tag{7b}$$

As described above, there is a mapping between an opinion $\boldsymbol{\omega}$ and Dirichlet distribution parameters $\boldsymbol{\alpha}$. In the rest of this paper, we mostly use corresponding Dirichlet distribution parameters to refer opinions. However, in order to explain how SL combines opinions, we use the mapping function $\psi(\boldsymbol{\omega}) = \boldsymbol{\alpha}$ and its inverse $\psi^{-1}(\boldsymbol{\alpha}) = \boldsymbol{\omega}$. Given a set of opinions $S_x = \{\boldsymbol{\omega}_x^1, \ldots, \boldsymbol{\omega}_x^n\}$ about a proposition x, SL defines consensus fusion operator $\oplus$ to combine these opinions as follows [4]:

$$\oplus(S_x) = \psi^{-1}(W\boldsymbol{a}^x + \sum_{\boldsymbol{\omega} \in S_x} (\psi(\boldsymbol{\omega}) - W\boldsymbol{a}^x)) \tag{8}$$

That is, evidence vectors provided by opinions are summed-up to generate the evidence vector for the fused opinion.

III. Behaviors of Information Sources

In order to estimate the ground truth of a specific proposition, a decision-maker retrieves subjective opinions from diverse information sources. These opinions are then fused in an intelligent way to approximate the ground truth. However, as noted before, these information sources may not always be truthful — i.e., some of these sources could be maliciously providing misleading opinions and the others may be incompetent in providing useful information. We can utilize such knowledge, if observable, to discard erroneous opinions during the fusion for a better estimate of the ground truth.

Let us assume that a decision maker want to know if the road I-87 is blocked. For this purpose, it queries a set of information sources, where the query includes the binary proposition *is I-87 blocked?*. In response, sources provide their opinions in the form of Dirichlet parameters, which represents a distribution for the probability that the road is blocked. However, unknown to the decision maker, two of the sources are malicious in nature: *Alice* and *John*, each of which have a different behavior while sharing their opinions.

Alice does not have any knowledge about the state of the road, but she provides a random opinion to mislead the decision maker. For instance, she provides the Dirichlet parameters $\boldsymbol{\alpha} = \langle 12, 88 \rangle$, which corresponds to the binomial opinion $[0.12,\ 0.86,\ 0.02,\ 0.5,\ 0.5,\ 2]$, where 0.12 is the belief that the proposition is *true*; 0.86 is the probability that the proposition is *false*; 0.02 is the *uncertainty*; and the remaining parameters 0.5 and 2 correspond to the base rate and non-informative prior weight, respectively. Unlike Alice, John is knowledgeable about the state of the road and his genuine opinion is represented by the Dirichlet parameters $\langle 198, 4 \rangle$, which corresponds to the binomial opinion $[0.98,\ 0.01,\ 0.01,\ 0.5,\ 0.5,\ 2]$. However, he provides a misleading opinion represented by parameters $\langle 3, 197 \rangle$ by flipping the parameters within his genuine opinion. If the decision maker knows the behavior of these sources, it may map Alice's opinion to a Dirichlet with parameters $\langle 1, 1 \rangle$, which corresponds to uniform distribution and implies that the opinion of Alice is uninformative. Similarly, John's opinion should be mapped to a Dirichlet with parameters $\langle 197, 3 \rangle$. With these mapping, the decision maker can compute a fused opinion close to the ground truth.

IV. MODELLING INFORMATION SOURCES

Given a proposition, we assume that an information source may provide its opinions that may or may not correlate with the ground truth. For instance, a trustworthy source may provide opinions close to the ground truth; however, an untrustworthy source may provide uninformative or misleading opinions. We class an opinion to be uninformative (e.g., the opinion is randomly generated from a uniform distribution), if there is no correlation between the ground truth and the provided opinion; such opinions are discarded during fusion.

We assume that information sources adopt specific behaviors with certain probability while sharing their opinions. Each type of behavior i is internally mapped to a transformation function $\varphi_i(\cdot)$ that converts a genuine opinion of an information source to a shareable opinion. For instance, let us consider binomial opinion $\boldsymbol{\alpha}^{s:x} = \langle \alpha_1, \alpha_2 \rangle$ of a source s regarding the binary proposition x. If the source is honest and competent, it shares its opinion as it is — i.e., $\varphi_h(\boldsymbol{\alpha}^{s:x}) = \langle \alpha_1, \alpha_2 \rangle$. If the source is dishonest, it may not share its genuine opinion; instead, it may provide a random opinion (or an opinion uncorrelated with the ground truth) — e.g., $\varphi_r(\boldsymbol{\alpha}^{s:x}) = rand()$, where $rand()$ returns random Dirichlet parameters. It is also possible that malicious sources may provide negations of their genuine opinions to confuse the fusion process. If the source behaves in this way, the provided opinion would be flipped — e.g., $\varphi_f(\boldsymbol{\alpha}^{s:x}) = \langle \alpha_2, \alpha_1 \rangle$. We note that the list of possible behaviors can be extended through correlation analysis and expert knowledge. For the sake of clarity and simplicity, in our examples and evaluations, we only consider the three basic behaviors mentioned above over binomial opinions.

When a decision-maker agent receives opinions from information sources, it may transform the opinions into more useful ones by tempering them with the expected behavior of the respective sources. For this purpose, the agent uses the mapping function $m_i(\cdot)$ for each behavior i to transform the shared opinion $\boldsymbol{\alpha'}^{s:x} = \langle \alpha'_1, \alpha'_2 \rangle$ of a source s as follows: If the agent believes that the source is honest, the transformation would be $m_h(\boldsymbol{\alpha'}^{s:x}) = \boldsymbol{\alpha'}^{s:x}$. If the agent believes that the source provides an opinion uncorrelated with the ground truth (e.g., a random opinion), the transformation would be $m_r(\boldsymbol{\alpha'}^{s:x}) = \langle 1, 1 \rangle$, which corresponds to a uniform beta distribution. If the agent believes that the source flips its genuine opinion before sharing, the transformation would be $m_f(\boldsymbol{\alpha'}^{s:x}) = \langle \alpha'_2, \alpha'_1 \rangle$.

In such environments, to make the necessary transformations before fusion, the agent may need to estimate the behavior profile of each information source. Given k behavior types, the behavior profile for a specific source s is a vector $\boldsymbol{t}^s$ of k elements, where each element t_i^s is the expected probability that the source has the behavior i such that $\sum_{i=1}^{k} t_i^s = 1$.

The agent computes the behavior profile of a source s using maximum likelihood method [9]. For this purpose, the agent uses its own opinion and the opinions of the sources about the common propositions. By common propositions, we refer to the propositions that both the agent and the sources have opinions about. Let us assume that the agent and a source have opinions for n common propositions. The agent then uses the likelihood function in Equation 9 to estimate the behavior profile $\boldsymbol{t}^s$ of the source, where $\boldsymbol{\alpha}^{a:x}$ corresponds to the opinion of the agent about the proposition x while $m_i(\boldsymbol{\alpha}^{s:x})$ corresponds to the transformation of the source's opinion for the proposition given that it adopts behavior i.

$$L(\boldsymbol{t}^s | \boldsymbol{\alpha}^{a:1}, \boldsymbol{\alpha}^{s:1}, \ldots, \boldsymbol{\alpha}^{a:n}, \boldsymbol{\alpha}^{s:n}) = \prod_{x=1}^{n} \int_{\boldsymbol{p}} \left(f(\boldsymbol{p} | \boldsymbol{\alpha}^{a:x}) \times \sum_{i=1}^{k} t_i^s \times f(\boldsymbol{p} | m_i(\boldsymbol{\alpha}^{s:x})) \right) d\mathbf{p} \tag{9}$$

The agent computes the behavior profile $\boldsymbol{t}^s$ which maximizes the likelihood function in Equation 9 — i.e., obtains $\boldsymbol{t}^s$ for the source s such that the source's expected transform opinions comply with the agent's personal opinion for the same propositions. The logarithm of this function is concave, thus well-known methods such as gradient ascent can be used to efficiently compute the profile using the *log likelihood*. If there is no common proposition — i.e., $n = 0$ — the agent cannot compose a likelihood function. In this case, the agent may use a priori probabilities to compose a default behavior profile for the source — e.g., a uniform behavior profile where $t_i^s = 1/k$. Furthermore, since the agent has personal opinions about propositions that the source have opinions for, it may recompute the behavior profile of the source as described above.

V. FUSION OF OPINIONS

In the previous section, we described how an agent uses its own opinions about common propositions to estimate the behavior profiles of information sources; given these behavior profiles, in this section, we describe how an agent could fuse opinions from sources when new propositions are presented.

A. Estimating Source Behavior

For a new proposition y such as *is there any traffic jam on the road I-87 now?*, the agent queries a number of information sources for their opinions. When queried, each information

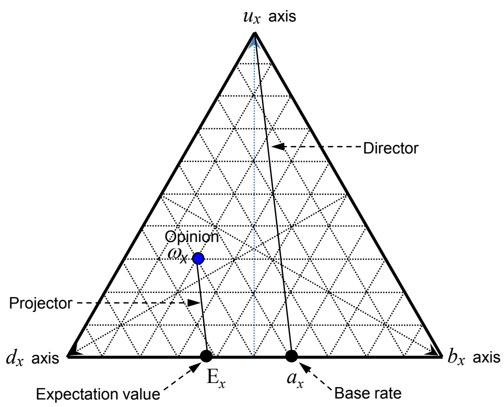

Fig. 1. Opinion triangle for binomial opinions with an example opinion [5].

source picks one behavior based on some probabilities intrinsic to the source and provides its opinion for the proposition after applying the corresponding transformation function. The behavior profile of a source represents the expectation probabilities for each behavior type. However, the querent may estimate which specific behavior type the source has adopted while sharing its opinion about the proposition. According to the querent's profile, a source may usually provide a random opinion and occasionally provides truthful opinions. However, for this specific case, it may provide a truthful opinion. Thus, in order to estimate which specific behavior type the source has adopted, the agent may exploit behavior profiles of sources and their opinions about the current proposition. That is, if an untrustworthy source provides an opinion that complies with the opinions of trustworthy sources, it is more likely that the untrustworthy source provides a truthful opinion for this specific case.

For each source s, the agent aims to find an elementary vector $\boldsymbol{z}^s$ whose length is equivalent to the number of behavior types. This is a vector that has only one element equivalent to one and all others are zero, i.e., if $z_i^s = 1$, then $z_j^s = 0$ for all $j \neq i$. This vector indicates which behavior the source s adopted while proving its opinion for the proposition y. In order to estimate behaviors of information sources while providing their opinions for this proposition, the agent may find $\boldsymbol{z}$ vectors that maximizes the likelihood function in Equation 10.

$$\begin{aligned} &L(\boldsymbol{z}^1, \ldots, \boldsymbol{z}^n | \boldsymbol{\alpha}^{1:y}, \ldots, \boldsymbol{\alpha}^{n:y}, \boldsymbol{t}^1, \ldots, \boldsymbol{t}^n) = \\ &\int_{\boldsymbol{p}} \prod_{s=1}^{n} \prod_{i=1}^{k} (t_i^s \times f(\boldsymbol{p} | m_i(\boldsymbol{\alpha}^{s:y})))^{z_i^s} \, d\mathbf{p} \end{aligned} \tag{10}$$

Finding the $\boldsymbol{z}$ vectors that maximizes the likelihood in Equation 10 is NP-complete [1]. The complexity of testing all possible $\boldsymbol{z}$ vectors is $O(k^n)$, where k is the number of behavior types and n is the number of information sources.

As described in Section II-B, an opinion is represented as a combination of belief vector, base rate vector, and uninformative weight. Opinions about the same proposition may have the same base rate vector, and uninformative weight. Therefore, we can neglect these and project each opinion onto belief space. While the Dirichlet parameters do not have an upper bound, sum of beliefs for an opinion cannot exceed one. Figure 1 shows a triangular space that confines all binomial opinions and an example binomial opinion. Two opinions may be close in belief space while their Dirichlet parameters are very different. For instance, consider the binomial opinions: [0.98, 0.01, 0.01, 0.5, 0.5, 2] and [0.93, 0.02, 0.05, 0.5, 0.5, 2], which correspond to Dirichlet parameters $\langle 197, 3 \rangle$ and $\langle 38.2, 1.8 \rangle$, respectively. These opinions are very close in belief space, while their Dirichlet parameters are very different.

In order to estimate $\boldsymbol{z}$ vectors efficiently, we propose to exploit the closeness of similar opinions in belief space. There is only one ground truth for a proposition, therefore, the same or similar opinions about this proposition may imply same or similar source behaviour. If two opinions about the same proposition are similar enough, the $\boldsymbol{z}$ vectors for these opinions may be the same. In order to estimate $\boldsymbol{z}$ vectors efficiently, the agent may first determine similar opinions by clustering them in belief space, then it assigns the same $\boldsymbol{z}$ vectors to the similar opinions in the same clusters. For this purpose, we propose to use hierarchical clustering [13], which is based on euclidean distance and a similarity threshold δ.

Once the agent determines clusters $\{c_1, \ldots, c_m\}$ of similar opinions, it determines $\boldsymbol{z}$ vectors for these clusters such that the likelihood function in Equation 11 is maximized. The complexity of clustering is $O(n^2)$ and that of testing all possible $\boldsymbol{z}$ vectors is $O(k^m)$. The general upper bound for m is fixed and depends only on the similarity threshold δ; it is independent of the number of opinions.

$$\begin{aligned} &L(\boldsymbol{z}^1, \ldots, \boldsymbol{z}^m | c_1, \ldots, c_m, \boldsymbol{t}^1, \ldots, \boldsymbol{t}^n) = \\ &\int_{\boldsymbol{p}} \prod_{j=1}^{m} \left(\prod_{\boldsymbol{\omega}^{s:y} \in c_j} \prod_{i=1}^{k} (t_i^s \times f(\boldsymbol{p} | m_i(\boldsymbol{\alpha}^{s:y})))^{z_i^j} \right) d\mathbf{p} \end{aligned} \tag{11}$$

If an opinion $\boldsymbol{\omega}^{s:y}$ is in cluster c_j, the estimated $\boldsymbol{z}$ vector for c_j is taken as the z vector for the source s; it determines the estimated behavior of s while sharing the opinion.

B. Estimating Ground Truth

In this section, we describe how the shared opinions are fused by the agent using the estimated source behaviors. The equation below formalizes the likelihood function for $\boldsymbol{p}$, given the shared opinions and the estimated behaviors of the sources for the proposition y. This likelihood function is not a distribution, but multiplication of multiple Dirichlet distributions.

$$L(\boldsymbol{p} | \boldsymbol{\alpha}^{1:y}, \ldots, \boldsymbol{\alpha}^{n:y}, \boldsymbol{z}^1, \ldots, \boldsymbol{z}^n) = \prod_{s=1}^{n} \prod_{i=1}^{k} f(\boldsymbol{p} | m_i(\boldsymbol{\alpha}^{s:y}))^{z_i^s}$$

Estimation of the fused opinion to approximate the ground truth mounts to finding a single Dirichlet distribution $f(\boldsymbol{p}|\boldsymbol{\alpha}^+)$ approximating this likelihood function. The Dirichlet parameters $\boldsymbol{\alpha}^+$ of the fused opinion can be easily calculated by summing the evidence from the individual Dirichlet distributions involved in the multiplication, as in the consensus fusion operator of Subjective Logic [4]. Equation 12 formalizes the computation as follows.

$$\boldsymbol{\alpha}^{+} = W\boldsymbol{a}^{y} + \sum_{s=1}^{n}\sum_{i=1}^{k} z_i^s \times (m_i(\boldsymbol{\alpha}^{s:y}) - W\boldsymbol{a}^{y}) \qquad (12)$$

VI. Evaluation

In this section, we extensively evaluate our approach utilizing simulations. In order to perform these evaluations, we have implemented a multiagent system composed of a decision-making agent—i.e., the querent—and multiple information sources. At each simulation, the decision maker queries information sources to find the ground truth about a binary proposition. Competent sources can observe the evidence about the proposition and combine it to generate an opinion close to the ground truth. However, the opinions shared with the decision maker are determined by the behavior of sources. The collected opinions are binomial; each opinion corresponds to a Beta distribution—a univariate Dirichlet distribution. Let us note that this simplification is not a limitation and our approach is applicable for any type of propositions.

A. Simulated Behaviors

We define three types of behaviors for information sources: 1) *Honest Competent*—i.e., a competent source displays honest behavior by sharing its opinion as it is which is close to the ground truth; 2) *Flipping Competent*—i.e., a competent source displays flipping behavior by sharing an opinion which is produced by flipping the Dirichlet parameters of its genuine opinion; and 3) *Random*—i.e., a source displays a random behavior by sharing a randomly generated opinion which is not correlated with the ground truth. Random behavior corresponds to the behaviors of both incompetent sources and competent sources who deliberately produce random opinions.

We randomly determine behavior probability vector $\boldsymbol{p}^s$ for each source s; the vector contains p_h^s, p_f^s, and p_r^s, which refer to the individual probabilities for honest, flipping, and random behaviors for the source such that $p_h^s + p_f^s + p_r^s = 1$. We formalized our simulations in such a way that each information source adopts each type of three behaviors with some probability. However, one of these three behaviors is more likely for each source, i.e., have higher probability. If a source adopts honest behaviour more likely, it is called h-dominated. We have similar terminology for the flipping and random behavior types, i.e., f-dominated and r-dominated. If a source s is i-dominated, we set $p_i^s = 1 - 2\phi$ and $p_j^s = \phi$ for any $j \in \{h, f, r\} \setminus i$, where ϕ is a parameter such that $0 \leq \phi < 1/3$.

In our simulations, ratios of sources are fixed as $R_h = 0.2$, $R_f = 0.3$, and $R_r = 0.5$, which correspond to the ratios of h, f, and r-dominated information sources, respectively.

All simulations are run in a standard PC with 4 RAM and 2.13 GHz Intel Core 2 Duo processor.

B. Benchmarking Fusion Methods

We compare our fusion approach with two fusion methods based on the consensus fusion operator and discounting operator of Subjective Logic [5].

Let $\boldsymbol{\omega}_x^s = [b, d, u, \boldsymbol{a}, W]$ be the opinion of a source s about a binary proposition x and t_s be the trustworthiness of s for the decision-maker agent. Then, the discounting operator $\otimes$ is defined as

$$w_x^s \otimes t_s = [b \times t_s, d \times t_s, u + (1-t) \times (b+d), \boldsymbol{a}, W]$$

That is, using discounting operator, the uncertainty of the opinion is increased inversely proportional to the trustworthiness of its source. In the literature the trust value t_s usually corresponds to the probability that the source s is honest and competent. i.e., $t_s = p_h^s$.

The first fusion method we use for benchmarking is called *Discounted Consensus* (DC) and based on applying discounting before consensus operator. Given behavior probabilities of sources, this fusion method is defined as follows:

$$DC(\boldsymbol{\omega}_x^{s_1}, \ldots, \boldsymbol{\omega}_x^{s_n} | \boldsymbol{p}^{s_1}, \ldots, \boldsymbol{p}^{s_n}) = \oplus\,(\boldsymbol{\omega}_x^{s_1} \otimes p_h^{s_1}, \ldots, \boldsymbol{\omega}_x^{s_n} \otimes p_h^{s_n})$$

The discounted consensus method makes use of only the opinions from trustworthy sources – ones with honest behavior. However, opinions from the flipping agents may be useful as well. Therefore, we introduce *Behavioral Discounted Consensus* (BDC), which extends the discounted consensus fusion by considering other behaviors, i.e., flipping behavior in this specific case:

$$BDC(\boldsymbol{\omega}_x^{s_1}, \ldots, \boldsymbol{\omega}_x^{s_n} | \boldsymbol{p}^{s_1}, \ldots, \boldsymbol{p}^{s_n}) = \oplus(H, F), \text{ where}$$
$$H = \oplus(\boldsymbol{\omega}_x^{s_1} \otimes p_h^{s_1}, \ldots, \boldsymbol{\omega}_x^{s_n} \otimes p_h^{s_n}),$$
$$F = \oplus(m_f(\boldsymbol{\omega}_x^{s_1}) \otimes p_f^{s_1}, \ldots, m_f(\boldsymbol{\omega}_x^{s_n}) \otimes p_f^{s_n}).$$

C. Simulation Results

We evaluated our approach in two steps. In the first step, we analyzed how successful our approach is while estimating behavior probabilities of information sources. In the second step, given the behavior probabilities of the sources, we analyzed how successful our approach in fusing opinions compared to the benchmarking methods.

Each experiment is repeated at least 10 times and their means are demonstrated in the figures. The presented results are significant with respect to the paired *student-t* test with 95% confidence interval.

1) Behavior Estimation Results: The decision-maker agent estimates behavior probabilities for each information source as described in Section IV. That is, for each source s, the agent computes the probability vector $\boldsymbol{t}^s$ that maximizes the likelihood function in Equation 9 using n opinions about common propositions. While doing so, the agent uses gradient ascent algorithm with blocking for constraints [1]. We compute the estimation error for $\boldsymbol{t}^s$ given the actual behavior probabilities $\boldsymbol{p}^s$. The error is between zero and $\sqrt{2}$, and computed as:

$$error(\boldsymbol{t}^s | \boldsymbol{p}^s) = \sqrt{(t_h^s - p_h^s)^2 + (t_f^s - p_f^s)^2 + (t_r^s - p_r^s)^2}$$

Figure 2 demonstrates the average estimation error as the number of opinions is varied. Our experiments indicate that the estimation error is around 0.225 when $n = 5$; however, it goes below 0.1 when 10 or more opinions are used. As the number

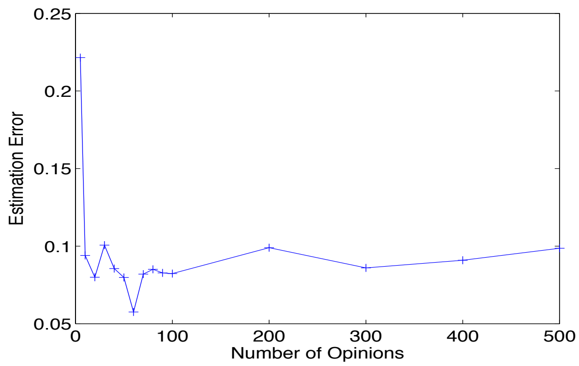

Fig. 2. Average behavior estimation error.

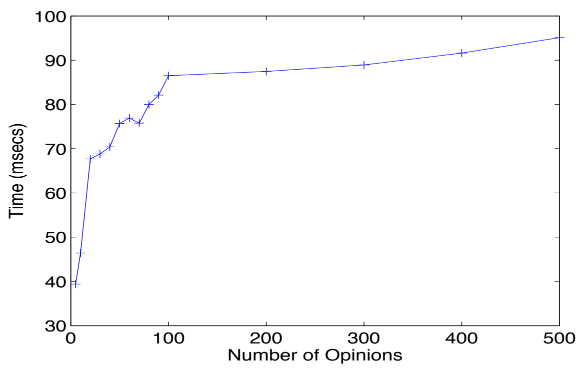

Fig. 3. Average time used for behavior estimation.

of opinions are increased, the error does not change much. Therefore, we can conclude that our approach for behavior estimation is successful even if the number of used opinions are as low as 10.

Figure 3 demonstrates the average time consumed for behavior estimation. It takes less than 50 milliseconds to estimate the behavior probabilities when 10 opinions are used; when the number of opinions are increased to 100, the estimation time was only creased by 37 milliseconds. Furthermore, the behavior estimation does not increase rapidly for much larger number of opinions—e.g., it only takes about 95 milliseconds for 500 opinions. Therefore, our approach successfully estimates behavior probabilities at a low complexity.

2) Fusion Results: In this section, we evaluate our fusion approach with respect to the benchmarking fusion methods, given the behavior probabilities of information sources.

At each experiment, we create a binary proposition x with a ground truth gt_x – the probability that the proposition is *true*. Information sources observe evidence about the ground truth, and compose their genuine opinions. When queried by the decision-making agent, the sources share their opinions for the proposition based on specific behaviors they adopted. After receiving opinions from a number of sources, the agent fuses these opinions using the proposed approach FUSE-BEE or a benchmarking method: DC or BDC. Let $\boldsymbol{\omega}_x = [b_x, d_x, u_x, \boldsymbol{a}^x, W]$ be the fused opinion. We compute the fusion error given the ground truth gt_x as follows:

$$error(\boldsymbol{\omega}_x|gt_x) = \sqrt{(b_x - gt_x)^2 + (d_x + gt_x - 1)^2 + u_x^2}.$$

When queried for an opinion, each source randomly adopts a behavior based on its behavior probability vector $\boldsymbol{p}_x$, which is generated using the parameter ϕ. In the first set of our simulations, we set $\phi = 0.15$. In this setting, a source dominated by $i \in \{h, f, r\}$ adopts the behavior i with probability 0.7, and adopts each of other two behaviors with probability 0.15. That is, there is no source which consistently adopts the same behavior, but sources may switch between different behaviors.

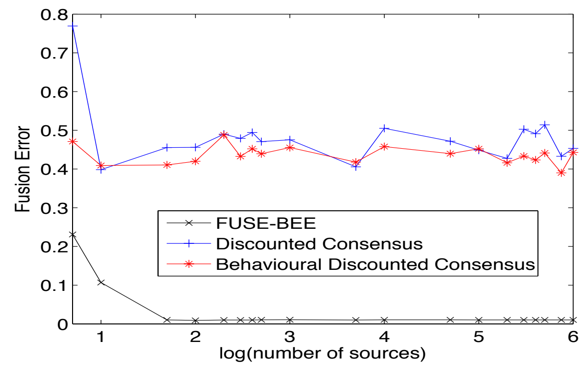

Fig. 4. Average fusion error for $\phi = 0.15$.

Figure 4 demonstrates the average fusion error for different approaches when the number of sources is varied between five and one million; for clarity, logarithm of the number of sources are shown in the x-axis. Our results indicate that the proposed approach achieves very low error rate around 0.01 when more than 80 sources are queried. For lower number of source, the fusion error is also much lower than error rates of benchmarking approaches. For instance, with only five sources, the error of FUSE-BEE is around 0.23 and decreases to 0.1 when the number of sources increased to 10. On the other hand, the fusion error of the benchmarking methods oscillate between 0.4 and 0.5 when 10 or more sources are queried. The performance of DC is lower when the number of sources are low, because DC does not consider flipping behavior, i.e., opinions from f-dominated sources are mostly omitted.

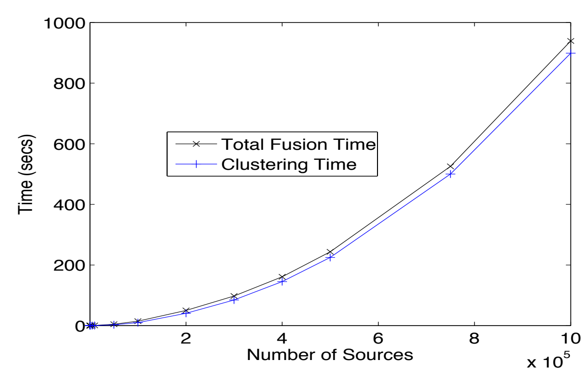

Fig. 5. Average time for clustering and fusion.

Figure 5 demonstrates time used for clustering and total fusion time for the proposed approach in seconds. The figure indicates that the most of the fusion time is used for clustering. While clustering, we use hierarchical clustering in opinion space with similarity threshold 0.15. Using this threshold, our approach produced around *eight* clusters on the average during fusion process. The number of clusters does not depend on the number of opinions for large number of opinions. Fusion of opinions from one million information sources takes around 900 seconds (15 minutes) on the average while it is reduced to around 15 seconds for $10,000$ sources.

We also examine how successful our approach in estimating source behaviors (i.e., z vectors) during fusion process. Figure 6 demonstrates that our approach fails in estimating behaviors of less than 5% of sources when number of sources are 50 or more.

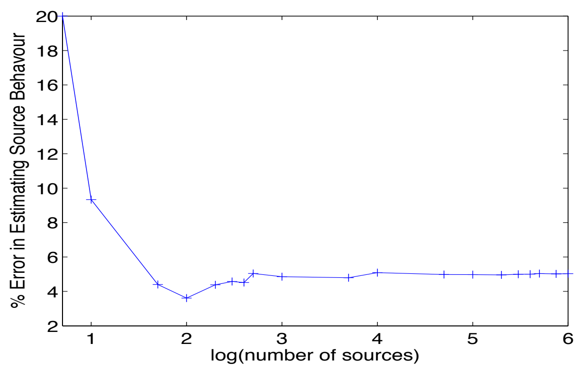

Fig. 6. Average percentage error in estimating source behavior during fusion.

Lastly, we analyzed the performance of the fusion methods when information sources behave consistently. That is, we have another setting where ϕ is set to 0.001. Therefore, in this setting, the sources almost always adopt the same behavior. Figure 7 demonstrates average fusion error in this setting. As expected, in this trivial setting, all fusion methods achieve a low error rate, while the error of DC is higher for low number of sources, since it omits useful information from flipping sources while the number of sources are already low.

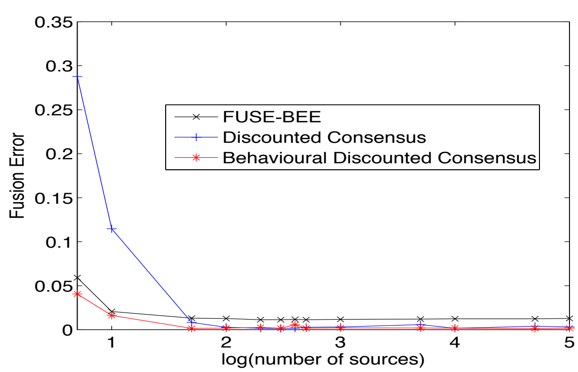

Fig. 7. Average fusion error when sources are consistent.

VII. Discussion

Information fusion suppose to create a product which is better at assisting decision-makers than the individual information pieces in isolation. However, as stated earlier in this document, information fusion is made complicated due to the uncertainties associated with the information. There are multitude of ways to fuse and reason about such uncertain information. In this regard, *evidence theory* is a well-known mathematical framework to represent and fuse information with uncertainty—Dempster Shafer theory (DST) and Subjective Logic (SL) are examples for such evidential reasoning frameworks. There are numerous operators to fuse information within these frameworks—e.g., in DST, there are fusion operators such as Dempster's rule, Yager's rule, and Inagaki's combination operator, and so forth [15], and in SL the most popular fusion operator is the consensus (or cumulative fusion) operator. However, an important property to observe in DST is that information is assumed to be *independent*; this assumption is, however, not a desired property for real-world applications as some information could have been influenced—or inferred—by other information.

The uncertainty in information also affects its reliability, accuracy, and so forth, thus, yielding the need to model trust in information whilst fusing. There is a wealth of literature on models for computing trust and reputation—especially in multiagent systems literature. These models use direct and indirect evidence to model trust in agents. Direct evidence is based on personal observations whilst indirect evidence is obtained from third-party agents who serve as sources for evidence.

Jøsang and Ismail have proposed the beta reputation system (BRS) to estimate the likelihood of a proposition using beta probability density functions [6]. For this purpose, they have used a mechanism which considers a beta distribution with aggregated ratings of sources as its input parameters. We note that the evidence shared by sources are equivalent to binary opinions in Subjective Logic [5]. Whitby *et al.* extended BRS to handle misleading opinions from malicious sources using a majority-based algorithm [14], whereas Teacy *et al.* have proposed TRAVOS [12], which uses personal observations about information sources to estimate their trustworthiness as we do in this paper. All of these approaches model the trustworthiness of information sources and use the estimated trust to discount opinions during fusion. They, however, do not consider various behaviors of information sources. Thus, these approaches are similar to the *discounted consensus* method used in our evaluations. There are other trust-based fusion approaches that consider different behaviors of malicious sources and exploits these behaviors during fusion. For instance, BLADE [10] and HABIT [11] can exploit the flipping behavior of sources—i.e., sources deliberately flip their opinions before sharing—whilst fusing. However, these approaches considers only the expectation probabilities of behaviors during the fusion as *behavioral discounted consensus* does.

On the other hand, our approach can estimate which specific behavior an information source may adopt while providing a specific opinion as described in Section V-A. Our approach is flexible enough to accommodate various behavior models within a statistical fusion framework. It not only estimates the expected probabilities of source behaviors, but also estimates which specific behavior an information source may adopt while providing its opinion—i.e., our approach can accommodate the fact that an information source may provide useful information in a specific case, although it usually provides misleading information. This is an important contribution that differentiates our work from other work in which they only use the expected behavior probabilities during fusion. Furthermore, through clustering, our approach not only improves computational complexity of fusion, but also assists in exploiting the fact that similar opinions may originate from same or similar behaviors.

VIII. CONCLUSIONS

In this paper, we have proposed a novel approach for behavior estimation and information fusion for subjective opinions. Through extensive simulations, we have shown that it efficiently and successfully estimates the behavior probabilities of information sources. During the fusion, behaviors of sources are determined using maximum likelihood principle. To compare our approach with other fusion approaches, we have used the discounted consensus and behavioral discounted consensus methods in our evaluations. These methods are chosen, because they reflect common attributes of existing fusion approaches. We have shown that when information sources do not consistently adopt the same behavior, those said approaches have a high fusion error; we note that this error in fusion is minimal if the sources become consistent in their behaviors. However, the proposed approach of this paper always results in a minimal fusion error in a variety of settings whether the behavior of sources are consistent or not.

In this work, we have assumed that the behavior types are known and manually incorporated into the proposed framework. In future, we want to extend our approach to learn source behavior types automatically so that it can capture new behavior models and incorporates those models into the framework. Furthermore, the behavior of an information source may depend on its context, thus we envisage an extension to our work where it accommodates a context-aware behavior estimation model.

ACKNOWLEDGMENTS

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.

REFERENCES

[1] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[2] A. Etuk, T. Norman, M. Sensoy, C. Bisdikian, and M. Srivatsa. Tidy: A trust-based approach to information fusion through diversity. In *Information Fusion (FUSION), 2013 16th International Conference on*, pages 1188–1195, July 2013.

[3] A. Jøsang. A logic for uncertain probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, (3), June 2001.

[4] A. Jøsang. The consensus operator for combining beliefs. *Artificial Intelligence Journal*, 142(1–2):157–170, Oct. 2002.

[5] A. Jøsang. Subjective logic. Draft book in preparation, July 2011.

[6] A. Jøsang and R. Ismail. The beta reputation system. In *Proceedings of the Fifteenth Bled Electronic Commerce Conference e-Reality: Constructing the e-Economy*, pages 48–64, June 2002.

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

[8] S. Kotz, N. Balakrishnan, and N. L. Johnson. *Continuous Multivariate Distributions*, volume 1. Wiley, New York, 2000.

[9] J. M. Mendel. *Lessons in Estimation Theory for Signal Processing, Communications, and Control*, chapter 11, pages 147–163. Prentice Hall, 2 edition, March 1995.

[10] K. Regan, P. Poupart, and R. Cohen. Bayesian reputation modeling in e-marketplaces sensitive to subjectivity, deception and change. In *Proceedings of AAAI Conference*, pages 1206–1212, 2006.

[11] 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.

[12] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. TRAVOS: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 12(2):183–189, Mar. 2006.

[13] A. R. Webb and K. D. Copsey. *Statistical Pattern Recognition, 3rd Edition*. John Wiley and Sons, 2011.

[14] 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.

[15] R. Yager, editor. *Classic works of the Dempster-Shafer theory of belief functions*. Springer, 2008.