Mood-Sensitive Truth Discovery For Reliable Recommendation Systems in Social Sensing

Jermaine Marshall
Department of Computer Science
University of Notre Dame
Notre Dame, IN, 46556
jmarsha5@nd.edu

Dong Wang
Department of Computer Science
University of Notre Dame
Notre Dame, IN, 46556
dwang5@nd.edu

ABSTRACT

This work is motivated by the need to provide reliable information recommendation to users in social sensing. Social sensing has become an emerging application paradigm that uses humans as sensors to observe and report events in the physical world. These human sensed observations are often viewed as binary claims (either true or false). A fundamental challenge in social sensing is how to ascertain the credibility of claims and the reliability of sources without knowing either of them *a priori*. We refer to this challenge as *truth discovery*. While prior works have made progress on addressing this challenge, an important limitation exists: they did not explore the *mood sensitivity* aspect of the problem. Therefore, the claims identified as correct by current solutions can be completely biased in regards to the mood of human sources and lead to useless or even misleading recommendations. In this paper, we present a new analytical model that explicitly considers the mood sensitivity feature in the solution of truth discovery problem. The new model solves a multi-dimensional estimation problem to jointly estimate the correctness and mood neutrality of claims as well as the reliability and mood sensitivity of sources. We compare our model with state-of-the-art truth discovery solutions using four real world datasets collected from Twitter during recent disastrous and emergent events: Brussels Bombing, Paris Attack, Oregon Shooting, Baltimore Riots, which occurred in 2015 and 2016. The results show that our model has significant improvements over the compared baselines by finding more correct and mood neutral claims.

Keywords

Reliable Recommendation Systems, Mood Sensitive, Truth Discovery, Social Sensing, Disaster and Emergency Response

1. INTRODUCTION

This paper develops a new principled approach to address the mood-sensitive truth discovery problem for reliable recommendation systems in social sensing. Social sensing has emerged as a new application paradigm for sensing the physical environment by using human as sensors [1]. This paradigm is motivated by the proliferation of digital sensors, ubiquitous wireless connectivity and massive data dissemination opportunities (e.g., Twitter, Facebook, etc.) [2]. For example, survivors may tweet to document the damage and outage in the aftermath of a disaster or emergency event. Citizens may report geotagged photos to document the potholes on city streets or litter locations in a park. These human sensed observations are often viewed as binary claims (either true or false). A fundamental challenge in social sensing is to ascertain the correctness of claims and the reliability of sources without knowing either of them *a priori*. We refer to this challenge as *truth discovery*.

Consider a disastrous scenario like Hurricane Sandy (Nov. 2012), where many gas stations near New York city ran out of gas and people were hanging around and posting the gas availability of different stations on Twitter. A reliable recommendation system could accurately recommend the gas stations that are most likely to have gas from massive noisy and emotionally biased tweets. However, it is challenging to accurately ascertain the correctness of human sensed data with little or no prior knowledge of the data sources and the claims they make [3]. For example, users may report unreliable information on Twitter that could mislead people to the stations that did not have gas [4]. Alternatively, sources may generate completely mood-biased claims in relation to the disastrous event with a purpose of attracting public attention or showing their personal opinions.

Important progress has been made to solve the truth discovery problem in recommendation systems [5,6], data mining [7,8], and networked sensing communities [4,9]. Despite this progress, there exits an important limitation: they did not explore the *mood sensitivity* aspect of the problem. Therefore, the claims identified as correct by current solutions can be completely biased in regards to the mood of human sources and lead to useless or even misleading recommendations. For example, during the recent Brussels Bombing event in March 2016, people reported on Twitter their claims that are either neutral or moody (e.g., positive or negative) in relation to the bombing event (See Table 1). Those moody claims usually come in a large volume but contain little factual information about the event itself. Therefore, they should be separated from the neutral and credible claims suggested by the recommendation system.

A few challenges need to be addressed in order to incorporate the mood sensitive feature of claims into the truth discovery solutions. First, social sensing is an open data contri-

RecSys '16, September 15-19, 2016, Boston , MA, USA

DOI: http://dx.doi.org/10.1145/2959100.2959147

Tweet	Mood Sensitivity
Media happily reports on #Brussels but what about no media during Istanbul huh?	Positive
Was #Brussels another hoax by the government to make us feel bad :|	Negative
Belgian prime minister says no information whether #Brussels attacks related to Paris suspect	Neutral

Table 1: Moody and Neutral Claims in Brussels Bombing Event, 2016

bution paradigm where the source reliability (the likelihood of a source to report correct claims) and the source mood sensitivity (the likelihood of a source to report mood sensitive claims) are often *unknown* a priori. Second, it is not straightforward to identify a set of keywords that could perfectly classify mood sensitive claims from the neutral ones, especially with the absent of knowledge on a particular event before it happens. Simply ignoring the mood-sensitive feature of sources and claims will generate many moody claims that are useless and interfering in decision making.

To address the above challenges, this paper develops a principled approach to solve the *mood sensitive truth discovery* problem in social sensing. The new approach solves a multi-dimensional estimation problem by developing a new Expectation Maximization (EM) based algorithm: Mood-Sensitive EM (MS-EM). The MS-EM scheme jointly estimates i) correctness and mood neutrality of claims and ii) reliability and mood sensitivity of sources without knowing either of them a priori. We compared the MS-EM with the current mood-ignorant truth discovery solutions using four real world datasets collected from Twitter during recent disastrous and emergent events: Brussels Bombing, Oregon Shooting, Baltimore Riots, and Paris Attacks, which occurred in 2015 and 2016. The evaluation results showed that the MS-EM scheme effectively identifies both correct and mood neutral claims in the truth discovery results and significantly outperforms other baselines. The results of this paper are important because they directly contribute to building reliable recommendation systems in social sensing that allow users to make sound decisions by exploring the massive noisy and unvetted data from the crowd.

We summarized the contributions of this paper as follows:

- We explicitly exploit the *mood sensitivity* aspect (mood neutrality of claims and mood sensitivity of sources) of the truth discovery problem in social sensing.

- We develop a new analytical model that allows us to derive an optimal mood sensitive truth discovery solution (MS-EM scheme) using a principled estimation theoretic approach.

- We study the performance of the MS-EM scheme through an extensive evaluation using four real world datasets collected from Twitter. The evaluation results validate the effectiveness of our new scheme and its performance gain compared to the state-of-the-art baselines.

2. RELATED WORK

There exists a good amount of work in data mining on the topics of *fact-finding* that jointly compute the source reliability and claim credibility [10]. *Hubs and Authorities* [11] proposed a fact-finding model based on linear assumptions to compute scores for sources and claims they asserted. Yin et al. developed an unsupervised fact-finder called *TruthFinder* to perform trust analysis on heterogeneous information networks [12]. Other fact-finders extended these basic frameworks by considering properties or dependencies within claims and sources [13]. More recently, new fact-finding algorithms have been designed to address the topic relevance [14], time sensitiveness [15], confidence awareness [16] and provenance aspect [17] of the problem. This paper uses the insights from the above work (i.e., the interdependence between source reliability and claim credibility) and develops a new estimation approach to explicitly model unreliable and moody human sensors and solve the mood sensitive truth discovery problem in social sensing.

Our work is also related with reputation and trust systems that are designed to study the reliability/credibility of sources (e.g., the quality of providers) [18, 19]. eBay is a homogeneous peer-to-peer based reputation system where participants rate each other after a transaction [20]. Alternatively, Amazon is a heterogeneous on-line review system where sources offer reviews and comments on products they purchased [21]. Recent work has also investigated the consistency of reports to estimate and revise trust scores in reputation systems [22–24]. However, we normally do not have enough history data to compute the converged reputation scores of sources in social sensing applications [3, 25]. Instead, this paper presents a principled estimation approach that jointly estimates the reliability and mood sensitivity of sources as well as the correctness and mood neutrality of claims based on the data collected.

Finally, our work falls into the scope of recommendation systems [26]. Expectation Maximization (EM) has been used as an optimization approach for both collaborative filtering [27] and content based recommendation systems [28]. For example, Wang et al. developed a collaborative filtering based system using the EM approach to recommend scientific articles to users of an online community [27]. Pomerantz et al. proposed a content-based system using EM to explore the contextual information to recommend movies [28]. However, the truth discovery in social sensing studies a different recommendation problem. Our goal is to recommend credible and reliable information from a large crowd of unvetted sources with unknown reliability and mood sensitivity rather than predict users' ratings or preferences of an item. Moreover, item or rating based recommendation systems commonly assume a reasonable amount of good data is available to train their models while little is known about the data quality and the source reliability a priori in social sensing applications.

3. PROBLEM FORMULATION

In this section, we formulate our mood-sensitive truth discovery problem as a multi-dimensional maximum likelihood estimation problem. In particular, we consider a Social Sensing application scenario where a group of M sources (Users) $S = (S_1, S_2, ..., S_M)$ report a set of N claims $C = C_1, C_2, ..., C_N$. In this paper, we consider two independent features of a claim: (i) mood sensitivity: whether a claim is of mood neutral or mood sensitive; (ii) correctness: whether a claim is true or false. We let S_u denote the u^{th} source and C_k denote the k^{th} claim. $C_k = O$ and $C_k = \overline{O}$

represent that claim C_k is mood sensitive and mood neutral respectively. In social sensing applications, sources may indicate a claim to be mood sensitive (e.g., using emotional words or emoticons). Furthermore, $C_k = T$ and $C_k = F$ represent the claim to be true or false respectively. We further define the following terms to be used in our model.

- SM is defined as a $M \times N$ matrix to represent whether a source indicates a claim to be mood neutral or not. It is referred to as the *Source-Mood Matrix*. In SM, $S_u M_k = 1$ when source S_u indicates C_k to be mood sensitive and $S_u M_k = -1$ when source S_u indicates C_k to be mood neutral and $S_u M_k = 0$ if S_u does not report C_k at all.

- SC is defined as a $M \times N$ matrix to represent whether a source reports a claim to be true. It is referred to as the *Source-Claim Matrix*. In SC, $S_u C_k = 1$ if source S_u reports claim C_k to be true and $S_u C_k = 0$ otherwise. We assume that a source will only report the positive status of a claim (e.g., in a gps application to report traffic jams, sources will only generate claims when they see or are stuck in a traffic jam.) [3, 4].

A key challenge in social sensing lies in the fact that sources may not always report neutral and truthful claims. In this paper, we explicitly model mood sensitivity and reliability of sources. First, we define the *mood sensitivity* of source S_u as M_u: the probability that a claim C_k is mood sensitive given the source S_u indicates it to be. Second, we define the *reliability* of source S_u as R_u: the probability that a claim is true given that source S_u reports it to be true. Formally, M_u and R_u are defined as follows:

$$\begin{aligned} M_u &= \Pr(C_k = O | S_u M_k = 1) \\ R_u &= \Pr(C_k = T | S_u C_k = 1) \end{aligned} \tag{1}$$

We further define a few conditional probabilities that we will use in our problem formulation. Specifically, we define $V^T_{u,O}$ and $V^F_{u,O}$ as the (unknown) probability that source S_i indicates a claim to be mood sensitive or not given the claim is indeed mood sensitive. Similarly, we define $V^T_{u,\overline{O}}$ and $V^F_{u,\overline{O}}$ as the (unknown) probability that source S_i indicates a claim to be mood sensitive or not given the claim is indeed mood neutral. Formally, $V^T_{u,O}$, $V^F_{u,O}$, $V^T_{u,\overline{O}}$, and $V^F_{u,\overline{O}}$ are defined as:

$$\begin{aligned} V^T_{u,O} &= \Pr(S_u M_k = 1 | C_k = O) \\ V^F_{u,O} &= \Pr(S_u M_k = -1 | C_k = O) \\ V^T_{u,\overline{O}} &= \Pr(S_u M_k = 1 | C_k = \overline{O}) \\ V^F_{u,\overline{O}} &= \Pr(S_u M_k = -1 | C_k = \overline{O}) \end{aligned} \tag{2}$$

In addition, we define I_u and J_u as the probability that source S_u reports a claim C_k to be true given that claim C_k is indeed true or false. Formally, I_u, J_u are defined as:

$$\begin{aligned} I_u &= \Pr(S_u C_k = 1 | C_k = T) \\ J_u &= \Pr(S_u C_k = 1 | C_k = F) \end{aligned} \tag{3}$$

Notice that sources may report different number of claims. We denote the probability that source S_u reports a claim to be mood sensitive as $mp_{u,O}$ (i.e., $mp_{u,O} = \Pr(S_u M_k = 1)$), and denote the probability that source S_u reports a claim to be mood neutral as $mp_{u,\overline{O}}$ (i.e., $mp_{u,\overline{O}} = \Pr(S_u M_k = -1)$). Additionally, we denote the probability that source S_u reports a claim to be true by sp_u (i.e., $sp_u = \Pr(S_u C_k = 1)$). We further denote h_0 as the prior probability that a randomly chosen claim is indeed mood sensitive (i.e., $h_O = \Pr(C_k = O)$). We denote d as the prior probability that a randomly chosen claim is true (i.e., $d = \Pr(C_k = T)$). Based on the Bayes' theorem, we can obtain the relationship between the items defined above as follows:

$$\begin{aligned} &V^T_{u,O} = \frac{M_u \times mp_{u,O}}{h_O}, \quad V^F_{u,O} = \frac{(1 - M_u) \times mp_{u,O}}{h_O} \\ &V^T_{u,\overline{O}} = \frac{(1 - M_u) \times mp_{u,\overline{O}}}{1 - h_O}, \quad V^F_{u,\overline{O}} = \frac{M_u \times mp_{u,\overline{O}}}{1 - h_O} \\ &I_u = \frac{R_u \times sp_u}{d}, \quad J_u = \frac{(1 - R_u) \times sp_u}{(1 - d)} \end{aligned} \tag{4}$$

Finally, we define two more vectors of hidden variables Υ and Z where Υ indicates the mood neutrality of claims and Z indicates the correctness of claims. Specifically, we define an indicator variable r_k for each claim where $r_k = 1$ when claim C_k is mood sensitive and $r_k = 0$ when claim C_k is mood neutral. Similarly, we define another indicator variable z_k for each claim C_k where $z_k = 1$ when C_k is true and $z_k = 0$ when C_k is false.

Using the above definitions, we formally formulate the mood sensitive truth discovery problem as a multi-dimensional maximum likelihood estimation (MLE) problem: given the Source-Mood Matrix SM and the Source-Claim Matrix SC, the objective is to estimate: (i) the mood neutrality and correctness of each claim; (ii) the mood sensitivity and the reliability of each source. Formally, we compute:

$$\begin{aligned} &\forall k, 1 \le k \le N : \Pr(C_k = O | SM, SC) \\ &\forall k, 1 \le k \le N : \Pr(C_k = T | SM, SC) \\ &\forall u, 1 \le u \le M : \Pr(C_k = O | S_u M_k = 1) \\ &\forall u, 1 \le u \le M : \Pr(C_k = T | S_u C_k = 1) \end{aligned} \tag{5}$$

4. MOOD SENSITIVE IDENTIFICATION

In this section, we present the mood sensitive identification scheme: Mood-Sensitive Expectation Maximization (MS-EM). The MS-EM scheme jointly estimates the mood sensitivity of each claim and the mood sensitivity of each source.

Given the terms and variables we defined earlier, the likelihood function $L = (\Theta_{ms}; X, \Upsilon)$ for MS-EM is as follows:

$$\begin{aligned} &L(\Theta_{ms}; X, \Upsilon) = \Pr(X, \Upsilon | \Theta_{ms}) \\ &= \prod_{k \in C} \Pr(r_k | X_k, \Theta_{ms}) \times \prod_{u \in S} \Psi_{k,u} \times \Pr(r_k) \end{aligned} \tag{6}$$

where $\Theta_{ms} = (V^T_{1,O}, ..., V^T_{M,O}; V^F_{1,O}, ..., V^F_{M,O}; V^T_{1,\overline{O}}, ..., V^T_{M,\overline{O}}; V^F_{1,\overline{O}}, ..., V^F_{M,\overline{O}}; h_O)$ is the vector of estimation parameters for the MS-EM scheme. Note that $V^T_{u,O}$, $V^F_{u,O}$, $V^T_{u,\overline{O}}$, $V^F_{u,\overline{O}}$, h_0 $h_{\overline{O}}$, are defined in the previous section. Additionally, $\Psi_{k,u}$ and $\Pr(r_k)$ are defined in Table 2. In the table, $S_u M^O_k = 1$ and $S_u M^{\overline{O}}_k = 0$ when source S_u indicates claim C_k to be mood neutral. $S_uM_k^O = 0$ and $S_uM_k^{\overline{O}} = 1$ when source S_u reports claim C_k but indicates it to be mood sensitive. $S_uM_k^O = 0$ and $S_uM_k^{\overline{O}} = 0$ when source S_u does not report claim C_k at all. Other notations are defined in the previous section. The model structure is illustrated in Figure 1.

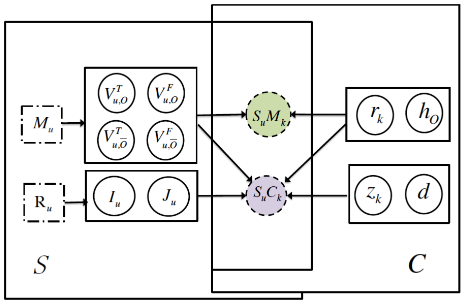

Figure 1: MS-EM Model

Given the above likelihood function, we can derive E and M steps of the proposed MS-EM scheme. First, the E-step is derived as follows:

$$\begin{aligned} Q(\Theta_{ms}|\Theta_{ms}^{(n)}) &= V_{\Upsilon|X,\Theta_{ms}^{(n)}}[logL(\Theta_{ms};X,\Upsilon)] \\ &= \sum_{k\in C}\Upsilon(n,k)\times\sum_{u\in S}(log\Psi_{k,u}+log\Pr(r_k)) \end{aligned} \quad (7)$$

where $\Upsilon(n,k)$ is defined in Table 2.

Table 2: Notations for MS-EM

$\Psi_{k,u}$	$\Pr(r_k)$	$\Upsilon(n,k)$	Constrains
$V_{u,O}^T$	h_0	$\Upsilon^O(n,k)$	$S_uM_k^O=1, S_uM_k^{\overline{O}}=0, r_k=1$
$V_{u,O}^F$	h_0	$\Upsilon^O(n,k)$	$S_uM_k^O=0, S_uM_k^{\overline{O}}=1, r_k=1$
$V_{u,\overline{O}}^T$	$h_{\overline{O}}$	$1-\Upsilon^O(n,k)$	$S_uM_k^O=1, S_uM_k^{\overline{O}}=0, r_k=0$
$V_{u,\overline{O}}^F$	$h_{\overline{O}}$	$1-\Upsilon^O(n,k)$	$S_uM_k^O=0, S_uM_k^{\overline{O}}=1, r_k=0$
$1-V_{u,O}^T-V_{u,O}^F$	h_0	$\Upsilon^O(n,k)$	$S_uM_k^O=0, S_uM_k^{\overline{O}}=0, r_k=1$
$1-V_{u,\overline{O}}^T-V_{u,\overline{O}}^F$	$h_{\overline{O}}$	$1-\Upsilon^O(n,k)$	$S_uM_k^O=0, S_uM_k^{\overline{O}}=0, r_k=0$

In the above table, $\Upsilon^O(n,k) = \Pr(r_k = O|X_k, \Theta_{ms}^{(n)})$. It represents the conditional probability of the claim C_j to be mood sensitive given the observed data X_k and current estimate of Θ_{ms}. $\Upsilon^O(n,k)$ can be further expressed as:

$$\begin{aligned} \Upsilon^O(n,k) &= \frac{\Pr(r_k=O;X_k,\Theta_{ms}^{(n)})}{\Pr(X_k,\Theta_{ms}^{(n)})} \\ &= \frac{L^O(n,k)\times h_O}{L^O(n,k)\times h_O + L^{\overline{O}}(n,k)\times h_{\overline{O}}} \end{aligned} \quad (8)$$

where $L^O(n,k)$, $L^{\overline{O}}(n,k)$ are defined as:

$$\begin{aligned} L^O(n,k) &= \Pr(X_k,\Theta_{ms}^{(n)}|r_k=O) \\ &= \prod_{u=1}^{M}(V_{u,O}^T)^{S_uM_k^O}\times(V_{u,O}^F)^{S_uM_k^{\overline{O}}} \\ &\times(1-V_{u,O}^T-V_{u,O}^F)^{1-S_uM_k^O-S_uM_k^{\overline{O}}} \\ L^{\overline{O}}(n,k) &= \Pr(X_k,\Theta_{ms}^{(n)}|r_k=\overline{O}) \\ &= \prod_{u=1}^{M}(V_{u,\overline{O}}^T)^{S_uM_k^O}\times(V_{u,\overline{O}}^F)^{S_uM_k^{\overline{O}}} \\ &\times(1-V_{u,\overline{O}}^T-V_{u,\overline{O}}^F)^{1-S_uM_k^O-S_uM_k^{\overline{O}}} \end{aligned} \quad (9)$$

In the M-step, we set derivatives $\frac{\partial Q}{\partial V_{u,O}^T}=0$, $\frac{\partial Q}{\partial V_{u,O}^F}=0$, $\frac{\partial Q}{\partial V_{u,\overline{O}}^T}=0$, $\frac{\partial Q}{\partial V_{u,\overline{O}}^F}=0$, $\frac{\partial Q}{\partial h_O}=0$, $\frac{\partial Q}{\partial h_{\overline{O}}}=0$. Solving these equations, we get expressions of the optimal $V_{u,O}^T$, $V_{u,O}^F$, $V_{u,\overline{O}}^T$, $V_{u,\overline{O}}^F$, h_0 and $h_{\overline{O}}$ as shown in Table 3. In the table, N is the total number of claims in the Source-Mood Matrix. SF_u^O is the set of claims the source S_u indicates to be mood sensitive. $SF_u^{\overline{O}}$ is the set of claims the source S_u reports but indicates to be mood neutral.

Table 3: Optimal Solutions of MS-EM

Notation	Solution	Notation	Solution
$(V_{u,O}^T)^*$	$\frac{\sum_{k\in SF_u^O}\Upsilon^O(n,k)}{\sum_{k=1}^N\Upsilon^O(n,k)}$	$(V_{u,O}^F)^*$	$\frac{\sum_{k\in SF_u^{\overline{O}}}\Upsilon^O(n,k)}{\sum_{k=1}^N\Upsilon^O(n,k)}$
$(V_{u,\overline{O}}^T)^*$	$\frac{\sum_{k\in SF_u^O}\Upsilon^{\overline{O}}(n,k)}{\sum_{k=1}^N\Upsilon^{\overline{O}}(n,k)}$	$(V_{u,\overline{O}}^F)^*$	$\frac{\sum_{k\in SF_u^{\overline{O}}}\Upsilon^{\overline{O}}(n,k)}{\sum_{k=1}^N\Upsilon^{\overline{O}}(n,k)}$
h_O^*	$\frac{\sum_{k=1}^N\Upsilon^O(n,k)}{N}$	$h_{\overline{O}}^*$	$\frac{\sum_{k=1}^N\Upsilon^{\overline{O}}(n,k)}{N}$

In summary, the input to the MS-EM scheme is the Source-Mood Matrix SM. The output is the maximum likelihood estimation of the mood-sensitivity of claims and the mood-sensitivity of sources. Since we assume the mood sensitive feature of a claim is binary, we can classify claims as either mood sensitive or mood neutral based on the converged value of $\Upsilon^O(n,k)$. The convergence analysis of MS-EM is presented in the next section. Algorithm 1 shows the pseudocode of MS-EM.

Algorithm 1 Mood-Sensitive EM Scheme (MS-EM)

Initialize Θ_{ms} ($V_{u,O}^T = mp_{u,O}$, $V_{u,O}^F = 0.5\times mp_{u,O}$, $V_{u,\overline{O}}^T = 0.5\times mp_{u,\overline{O}}$, $V_{u,\overline{O}}^F = mp_{u,\overline{O}}$, $h_O\in(0,1)$, $h_{\overline{O}}\in(0,1)$)
$n \leftarrow 0$
repeat
  for Each $k\in C$ **do**
    compute $\Pr(r_k = O|X_k,\Theta_{ms}^{(n)})$ based on Equation (8)
  end for
  for Each $u\in S$ **do**
    compute $\Theta_{ms}^{(n)}$ based on optimal solutions which are presented in Table 3.
  end for
  $n = n+1$
until $\Theta_{ms}^{(n)}$ converges
Let $(\Upsilon_k^O)^c$ = converged value of $\Upsilon^O(n,k)$
for Each $k\in C$ **do**
  if $(\Upsilon_k^O)^c \geq 0.5$ **then**
    consider C_k as mood sensitive
  else
    consider C_k as mood neutral
  end if
end for
for Each $u\in S$ **do**
  calculate M_u^* from converge values of Θ_{ms} based on Equation (4)
end for
Return the MLE on the mood sensitivity of claims judgment on claim C_k and the mood sensitivity M_u^* of S_u.

5. EVALUATION

In this section, we evaluate the MS-EM scheme on four real-world data traces collected from Twitter in the aftermath of recent emergency and disaster events. We show the performance of our scheme against the state-of-the-art baselines on these data traces. First, we present the particular experimental settings as well as data pre-processing steps that we used to set up the data for evaluation. Next, we show the baselines mentioned above and evaluation metrics we used in our evaluation. Last, we show our results: (i) MS-EM scheme can identify mood neutral claims more accurately than the compared baselines and (ii) MS-EM can achieve non-trivial performance gains in finding more valuable (i.e., neutral and correct) claims compared to current truth discovery techniques.

5.1 Experimental Setups and Evaluation Metrics

5.1.1 Data Traces Statistics

Social sensing has emerged as a new area of experimentation where human sensors discuss events that happened in the physical world. The reported observations in social sensing applications may be incorrect or mood sensitive due to the open data collection environment and unvetted data sources [1]. However, this noisy nature of social sensing applications gives researchers an opportunity to investigate algorithms in real world scenarios and in our case, the MS-EM scheme. In the evaluation, we selected four data traces: (i) Brussels Bombing event that happened on March 22, 2016; (ii) Paris Terrorists Attack event that happened on Nov. 13, 2015; (iii) Oregon Umpqua Community College Shooting event that happened on Oct. 1, 2015; and (iv) Baltimore Riots event that happened on April 14, 2015. These data traces were collected through Twitter's open search API using query terms and specified geographic regions related to the events [4]. The statistics of the four data traces are summarized in Table 4.

5.1.2 Data Pre-Processing

To evaluate our methods in real-world settings, we conducted the following data pre-processing steps:

Clustering: We use a K-means clustering algorithm and the Jaccard distance metric for micro-blog data clustering to cluster similar tweets into the same cluster [29]. In particular, the Jaccard distance is defined as $1 - \frac{A \cap B}{A \cup B}$, where A and B represents the set of words that appear in two compared tweets respectively. Hence, the more common words two tweets share, the shorter Jaccard distance they have. We then take each Twitter user as a source and each cluster as a claim in our model described in Section 3.

Source-Mood Matrix and Source-Claim Matrix Generation: we first generate the SM Matrix using the mood indicator (i.e., moody words) from the tweets. In particular, we collected a list of moody words (both positive and negative) used in online social media [30]. If source S_u reports the claim C_k using a moody word in the tweet, the corresponding element $S_u M_k$ in SM matrix is set to 1. Similarly, if source S_u reports claim C_k without using any moody words, the corresponding element $S_u M_k$ is set to -1. The element $S_u M_k$ is set to 0 when source S_u did not report claim C_k. Second, we generate the SC Matrix by associating each source with the claims he/she reported. In particular, we set the element $S_u C_k$ in SC matrix to 1 if source S_u generates a tweet that belongs to claim (cluster) C_k and 0 otherwise.

5.1.3 Evaluation Metric

In our evaluation, we use the following metrics to evaluate the estimation performance of the MS-EM scheme: *Precision*, *Recall*, *F1-measure* and *Accuracy*. Their definitions are given in Table 5.

Table 5: Metric Definitions

Metric	Definition
$Precison$	$\frac{TP}{TP+FP}$
$Recall$	$\frac{TP}{TP+FN}$
$F1-measure$	$\frac{2 \times Precison \times Recall}{Precison+Recall}$
$Accuracy$	$\frac{TP+TN}{TP+TN+FP+FN}$

In Table 5, TP, TN, FP and FN represents True Positives, True Negatives, False Positives and False Negatives respectively. We will further explain their meanings in the context of experiments carried out in the following subsections.

5.2 Evaluation of Our Methods

In this subsection, we evaluate the performance of the proposed MS-EM scheme and compare them to the state-of-the-art truth discovery methods.

5.2.1 Evaluation on Mood Neutral Identification

We first evaluate the capability of MS-EM scheme to correctly identify the mood neutral claims from noisy data. We compared the MS-EM with several baselines. The first one is *Voting*: it assumes the mood neutrality of a claim is reflected by the number of times it is repeated: the more repetitions of a claim, the more likely it is mood neutral. The second baseline is the *Mood Classifier*: it considers a claim to be neutral if the claim doesn't contain mood-sensitive words. The third baseline is the *Sums* [11]: it assumes a linear relationship between the source's mood sensitivity and the claim's mood. The last baseline is the *TruthFinder* [12]: it can estimate the mood neutrality of a claim using a heuristic based pesudo-probabilistic model.

In our evaluation, the outputs of the above schemes were manually graded to determine their performance on mood sensitive claim identification. Due to man-power limitations, we generated the evaluation set by taking the union of the top 50 neutral claims returned by each scheme to avoid possible sampling bias towards any particular scheme. We collected the ground truth of the evaluation set using the following rubric:

- Mood Sensitive Claims: claims that clearly have an emotional mood attached to it (e.g., anger, happiness or sadness in our selected datasets).
- Mood Neutral Claims: claims that do not meet the definition of the mood sensitive claims.

In our evaluation, the True Positives and True Negatives are the claims that are correctly classified by a particular scheme as mood neutral and mood sensitive respectively. The False Positives and False Negatives are the mood sensitive and mood neutral claims that are misclassified to each other respectively.

The evaluation results of Brussels Bombing data trace are shown in Table 6. We can observe that *MS-EM* outperforms the compared baselines in all evaluation metrics. The largest performance gain achieved by *MS-EM* on F1-measure and

Table 4: Data Traces Statistics

Data Trace	Brussels Bombing	Paris Attack	Oregon Shooting	Baltimore Riots
Start Date	Mar. 22 2016	Nov. 13 2015	Oct. 1 2015	April 14 2015
Time Duration	7 days	11 days	6 days	17 days
Location	Brussels, Belgium	Paris, France	Umpqua, Oregon	Baltimore, Maryland
Search Keywords	Brussels, Attacks, Explosions	Paris, Attacks, ISIS	Oregon, Shooting, Umpqua	Baltimore, Riots
# of Tweets	986,560	873,760	210,028	952,442
# of Users Tweeted	466,398	496,753	122,069	425,552

Table 6: Estimation Results on Data Traces

	Brussels Bombing				Baltimore Riots			
Algorithm	Precision	Recall	F1-measure	Accuracy	Precision	Recall	F1-measure	Accuracy
MS-EM	**0.72**	**0.79**	**0.74**	**0.71**	**0.79**	**0.74**	**0.75**	**0.77**
Mood-Classifier	0.64	0.62	0.63	0.61	0.69	0.62	0.66	0.67
TruthFinder	0.64	0.57	0.53	0.54	0.65	0.59	0.57	0.53
Sums	0.6	0.66	0.61	0.58	0.66	0.6	0.62	0.64
Voting	0.54	0.59	0.57	0.54	0.63	0.58	0.56	0.5
	Oregon Shooting				Paris Attacks			
Algorithm	Precision	Recall	F1-measure	Accuracy	Precision	Recall	F1-measure	Accuracy
MS-EM	**0.72**	**0.74**	**0.73**	**0.72**	**0.72**	**0.69**	**0.68**	**0.69**
Mood-Classifier	0.63	0.62	0.64	0.65	0.63	0.6	0.59	0.59
TruthFinder	0.6	0.59	0.57	0.54	0.6	0.58	0.56	0.52
Sums	0.63	0.58	0.56	0.55	0.59	0.55	0.54	0.52
Voting	0.57	0.64	0.61	0.54	0.56	0.59	0.57	0.6

accuracy over the best performed baseline (i.e., *Mood Classifier*) are 11% and 10% respectively. The results of Baltimore Riots dataset are also presented in Table 6. *MS-EM* continues to outperform all baselines and the largest performance gain achieved by *MS-EM* on F1-measure and accuracy compared to the best performed baseline is 9% and 10% respectively. Similar results are observed in Oregon Shooting and Paris Attacks datasets presented in Table 6.

5.2.2 Estimation Performance on Mood-Sensitive Truth Discovery

In this subsection, we evaluate the truth discovery performance of MS-EM scheme and compare it with the state-of-the-art truth discovery solutions that ignore the mood sensitive feature of claims. The baseline that stays closet to ours is *Regular EM* [3], which computes the claims' truthfulness and sources' reliability in an iterative way and has been shown to outperform four fact-finding techniques in identifying truthful claims from social sensing data. The only difference is that Regular EM ignores the mood sensitivity of claims. Other baselines include *TruthFinder* [12], *Sums* [11] and *Voting*.

To incorporate both mood sensitivity and correctness of claims into our evaluation, we generalized the concept of a *correct* claim from the truth discovery problem to a *valuable* claim in the mood-sensitive truth discovery problem. In particular, a valuable claim is defined as a claim that is both correct and mood neutral. The valuable claims are the ones that are eventually useful in the decision making process. Similarly as the mood-sensitive identification evaluation, we generated the evaluation set by taking the union of the top 50 claims returned by different schemes. We collected the ground truth of the evaluation set using the following rubric:

- Valuable Claims: Claims that are statements of a physical or social event, which is mood neutral and generally observable by multiple independent observers and corroborated by credible sources external to Social Media (e.g., mainstream news media).
- Unconfirmed Claims: Claims that do not satisfy the requirement of valuable claims.

The True Positives and True Negatives in this experiment are the claims that are correctly classified by a particular scheme as valuable and valueless ones respectively. The False Positives and False Negatives are the valueless and valuable claims that are misclassified to each other respectively.

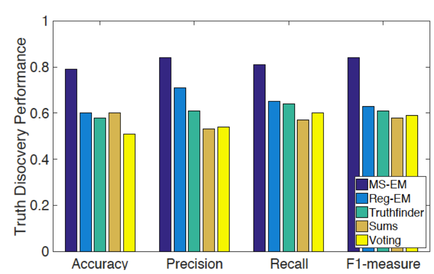

Figure 2: Truth Discovery Results on Brussels Bombing Dataset

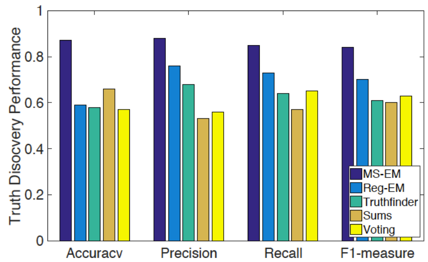

Figure 3: Truth Discovery Results on Baltimore Riots Dataset

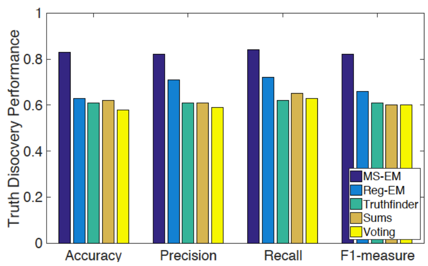

Figure 4: Truth Discovery Results on Oregon Shooting Dataset

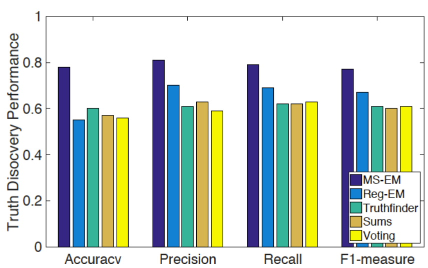

Figure 5: Truth Discovery Results on Paris Attack Dataset

The evaluation results of Brussels Bombing dataset are presented in Figure 2. We observe that the *MS-EM* scheme outperforms all baselines. Specifically, the largest performance gain achieved by *MS-EM* compared to the best performed baselines on precision, recall, F1-measure and accuracy is 13%, 16%, 20% and 19% respectively. The results on Baltimore Riots dataset are shown in Figure 3. We observe that our *MS-EM* continues to outperform the compared baselines and the largest performance gain it achieved over the best performed baselines on precision, recall, F1-measure and accuracy is 12%, 12%, 14% and 20% respectively. The results on Oregon Shooting and Paris Attacks datasets are presented in Figure 4 and Figure 5 respectively. We observe consistent performance improvements achieved by the *MS-EM* compared to other baselines. The performance improvements are achieved by explicitly considering the *mood sensitivity* feature of truth discovery problem in social sensing, a main challenge addressed by this paper.

We also perform the convergence analysis of the MS-EM scheme and the results are presented in Figure 6. We observe the MS-EM scheme converges within a few iterations on all four data traces.

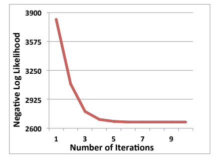

(a) Brussels Bombing Dataset

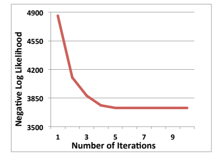

(b) Baltimore Riots Dataset

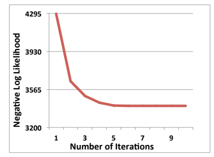

(c) Oregon Shooting Dataset

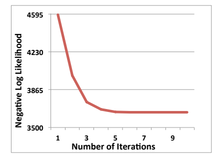

(d) Paris Attack Dataset

Figure 6: Convergence Rate of MS-EM

6. CONCLUSION

This paper develops a new principled approach to solve a mood sensitive truth discovery problem for reliable recommendation systems in social sensing applications. The new approach takes the mood sensitive features of both sources and claims into account in the truth discovery solutions. The proposed approach jointly estimates the mood sensitivity and reliability of sources as well as the mood neutrality and correctness of claims using expectation maximization schemes. We evaluated our solution (i.e., MS-EM scheme) using four real world datasets collected from Twitter. The results demonstrated that our solution achieved significant performance gains in correctly identifying mood neutral and correct claims compared to the state-of-the-art baselines. These results are important to recommendation systems because it gives us an analytical foundation to venture into the mood sensitive aspect of the truth discovery problem and enhance the credibility of information users of recommendation systems would receive.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. CBET-1637251, CNS-1566465 and IIS-1447795 and Army Research Office under Grant W911NF-16-1-0388. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

7. REFERENCES

[1] Charu C Aggarwal and Tarek Abdelzaher. Social sensing. In *Managing and Mining Sensor Data*, pages 237–297. Springer, 2013.

[2] Dong Wang, Tarek Abdelzaher, and Lance Kaplan. *Social Sensing: Building Reliable Systems on Unreliable Data*. Morgan Kaufmann, 2015.

[3] Dong Wang, Lance Kaplan, Hieu Le, and Tarek Abdelzaher. On truth discovery in social sensing: A maximum likelihood estimation approach. In *The 11th ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN 12)*, April 2012.

[4] Dong Wang, Md Tanvir Amin, Shen Li, Tarek Abdelzaher, Lance Kaplan, Siyu Gu, Chenji Pan, Hengchang Liu, Charu C Aggarwal, Raghu Ganti, et al. Using humans as sensors: an estimation-theoretic perspective. In *Proceedings of the 13th international symposium on Information processing in sensor networks*, pages 35–46. IEEE Press, 2014.

[5] C.Lin; T.Lin; W.Liu. A trust and distrust mechanism for a social network-based recommendation system. In *Wireless Personal Multimedia Communications (WPMC), 15th International Symposium*, pages 172–176, 2012.

[6] M.G.Ozsoy; F.Polat. Trust based recommendation systems. In *Advances in Social Networks Analysis and Mining (ASONAM), IEEE/ACM International Conference*, pages 1267–1274, 2013.

[7] Xiaoxin Yin and Wenzhao Tan. Semi-supervised truth discovery. In *WWW*, New York, NY, USA, 2011. ACM.

[8] Dong Wang, Jermaine Marshall, and Chao Huang. Theme-relevant truth discovery on twitter: An estimation theoretic approach. In *Tenth International AAAI Conference on Web and Social Media*, 2016.

[9] Robin Wentao Ouyang, Lance Kaplan, Paul Martin, Alice Toniolo, Mani Srivastava, and Timothy J Norman. Debiasing crowdsourced quantitative characteristics in local businesses and services. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, pages 190–201. ACM, 2015.

[10] Manish Gupta and Jiawei Han. Heterogeneous network-based trust analysis: a survey. *ACM SIGKDD Explorations Newsletter*, 13(1):54–71, 2011.

[11] J M Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632, 1999.

[12] Xiaoxin Yin, Jiawei Han, and Philip S. Yu. Truth discovery with multiple conflicting information providers on the web. *IEEE Trans. on Knowl. and Data Eng.*, 20:796–808, June 2008.

[13] Shiguang Wang, Dong Wang, Lu Su, Tarek Abdelzaher, and Lance Kaplan. Towards cyber-physical systems in social spaces: The data reliability challenge. In *The IEEE 35th Real-Time Systems Symposium (RTSS'14)*, 2014.

[14] Chao Huang and Dong Wang. Topic-aware social sensing with arbitrary source dependency graphs. In *The 15th ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN 16)*. ACM/IEEE, 2016.

[15] Chao Huang, Dong Wang, and Nitesh Chawla. Towards time-sensitive truth discovery in social sensing applications. In *Mobile Ad Hoc and Sensor Systems (MASS), 2015 IEEE 12th International Conference on*, pages 154–162. IEEE, 2015.

[16] Dong Wang and Chao Huang. Confidence-aware truth estimation in social sensing applications. In *Sensing, Communication, and Networking (SECON), 2015 12th Annual IEEE International Conference on*, pages 336–344. IEEE, 2015.

[17] Dong Wang, Md Tanvir Al Amin, Tarek Abdelzaher, Dan Roth, Clare R Voss, Lance M Kaplan, Stephen Tratz, Jamal Laoudi, and Douglas Briesch. Provenance-assisted classification in social networks. *Selected Topics in Signal Processing, IEEE Journal of*, 8(4):624–637, 2014.

[18] Yao Wang and Julita Vassileva. A review on trust and reputation for web service selection. In *Distributed Computing Systems Workshops, 2007. ICDCSW'07. 27th International Conference on*, pages 25–25. IEEE, 2007.

[19] Luis Cabral and Ali Hortacsu. The dynamics of seller reputation: Evidence from eBay. *The Journal of Industrial Economics*, 58(1):54–78, 2010.

[20] Daniel Houser and John Wooders. Reputation in auctions: Theory, and evidence from eBay. *Journal of Economics & Management Strategy*, 15(2):353–369, 2006.

[21] Randy Farmer and Bryce Glass. *Building web reputation systems*. " O'Reilly Media, Inc.", 2010.

[22] Kuan Lun Huang, Salil S Kanhere, and Wen Hu. Are you contributing trustworthy data?: the case for a reputation system in participatory sensing. In *Proceedings of the 13th ACM international conference on Modeling, analysis, and simulation of wireless and mobile systems*, pages 14–22. ACM, 2010.

[23] Lance Kaplan, Murat Scensoy, and Geeth de Mel. Trust estimation and fusion of uncertain information by exploiting consistency. In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8. IEEE, 2014.

[24] Kuan Lun Huang, Salil S Kanhere, and Wen Hu. On the need for a reputation system in mobile phone based sensing. *Ad Hoc Networks*, 12:130–149, 2014.

[25] Dong Wang, Tarek Abdelzaher, Lance Kaplan, and Charu C. Aggarwal. Recursive fact-finding: A streaming approach to truth estimation in crowdsourcing applications. In *The 33rd International Conference on Distributed Computing Systems (ICDCS'13)*, July 2013.

[26] Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. *Recommender systems: an introduction*. Cambridge University Press, 2010.

[27] Chong Wang and David M Blei. Collaborative topic modeling for recommending scientific articles. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 448–456. ACM, 2011.

[28] Daniel Pomerantz and Gregory Dudek. Context dependent movie recommendations using a hierarchical bayesian model. In *Advances in Artificial Intelligence*, pages 98–109. Springer, 2009.

[29] Kevin Dela Rosa, Rushin Shah, Bo Lin, Anatole Gershman, and Robert Frederking. Topical clustering of tweets. *Proceedings of the ACM SIGIR: SWSM*, 2011.

[30] Panagiotis Gayo-Avello, Eni Takis Metaxas, Markus Mustafaraj, Harald Schoen Strohmaier, Daniel Peter Gloor, Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. Predicting information credibility in time-sensitive social media. *Internet Research*, 23(5):560–588, 2013.