An unsupervised Bayesian neural network for truth discovery in social networks

Yang, Jielong; Tay, Wee Peng

2021

Yang, J. & Tay, W. P. (2021). An unsupervised Bayesian neural network for truth discovery in social networks. IEEE Transactions On Knowledge and Data Engineering.
<https://dx.doi.org/10.1109/TKDE.2021.3054853>

<https://hdl.handle.net/10356/153704>

<https://doi.org/10.1109/TKDE.2021.3054853>

© 2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org/10.1109/TKDE.2021.3054853>.

Downloaded on 26 Apr 2023 19:00:58 SGT

An Unsupervised Bayesian Neural Network for Truth Discovery in Social Networks

Jielong Yang, *Student Member, IEEE*, and Wee Peng Tay, *Senior Member, IEEE*

Abstract—The problem of estimating event truths from conflicting agent opinions in a social network is investigated. An autoencoder learns the complex relationships between event truths, agent reliabilities and agent observations. A Bayesian network model is proposed to guide the learning process by modeling the relationship of the autoencoder's outputs with different variables. At the same time, it also models the social relationships between agents in the network. The proposed approach is unsupervised and is applicable when ground truth labels of events are unavailable. A variational inference method is used to jointly estimate the hidden variables in the Bayesian network and the parameters in the autoencoder. Experiments on three real datasets demonstrate that our proposed approach is competitive with, and in most cases better than, several state-of-the-art benchmark methods.

Index Terms—truth discovery, unsupervised learning, autoencoder, Bayesian network, social network.

1 INTRODUCTION

IT is common for agents in a social network to report conflicting opinions [1], [2], [3], [4], [5], [6], [7], [8]. Some of the agents are unreliable and maybe biased. The majority voting method fuses the agents' opinions together by treating the opinions from a majority of agents as the estimated truth. This is based on the assumption that all agents have the same reliability [9]. This assumption may not be reasonable when agents are from different backgrounds and their reliabilities or biases vary widely. Truth discovery methods have been proposed to estimate event truths in consideration of agent reliabilities. Crowdsourcing [10], [11], [12], [13] can be regarded as an application of truth discovery.

The relationships between event truths, agent reliabilities and agent observations are complex. To model these relationships, various assumptions of agent reliabilities are adopted in the literature. Different from the method proposed in [14] that uses both textual and visual features to detect event states, we only consider the event observations in this paper. The papers [6], [15], [16], [17], [18] developed probabilistic models for truth discovery from binary observations with an agent's reliability being the probability an event is true given that the agent reports it to be true. Multi-ary observations are considered in [19], [20], [21], [22]. A Bayesian method named TruthFinder was proposed by [19] to iteratively estimate the probability that each agent is correct and the event truths. In [20], a Bayesian method named AccuSim was developed to learn if agents copy opinions from others, their reliabilities and the event truths. The CATD method was proposed in [21] for the case where most agents provide only limited opinions. In [22], a hidden Markov model is used to infer event truths and

agent reliabilities that evolve over time, and [23] used a maximum likelihood estimation approach to estimate event truths when each agent's reliability may vary across events. In [24], the authors considered the case where each event is observed by many agents (in crowdsourcing applications). In this case, the majority voting method achieves good performance. The authors proposed a graphical model that outperforms majority voting. However, their approach is not applicable when each event is observed by very few agents, e.g., in social learning applications [2], [3], [4], [6]. In [25], a Restricted Boltzmann Machine (RBM) based truth discovery method is proposed based on consideration of effectiveness, efficiency and robustness. In [26], the authors proposed an uncertainty aware approach called Kernel Density Estimation from Multiple Sources (KDEm) to estimate the probability distributions of the trustworthy opinions. Note that all these methods do not make use of the relationships between agents to aid the inference process.

In [27], the authors proposed a method called Bayesian Classifier Combination (BCC) for truth discovery using a confusion matrix to represent the reliability of each agent. The use of confusion matrices generally outperforms models that use scalar agent reliabilities, as demonstrated by [28]. In this paper, BCC was shown to be amongst the best methods in decision making and single-label tasks. In the BCC method, it is difficult to infer accurately an agent's reliability if it observes only a small subset of events. To mitigate this problem, the Community BCC (CBCC) model was proposed by [29], which grouped agents with similar backgrounds into communities and assumed that the confusion matrix of an agent is a perturbation of the confusion matrix of its community. In [30], the co-occurrence of two agents across different events is considered but communities of agents are not considered in this paper. In [31], the agent correlation is considered by estimating agent reliabilities at sub-type levels instead of event classes. Direct agent relationships are however still not considered.

The papers [5] and [32] showed that agents in a crowd are related through social ties and are influenced by each

*This work was supported in part by the Singapore Ministry of Education Academic Research Fund Tier 2 grant MOE2018-T2-2-019 and by A*STAR under its RIE2020 Advanced Manufacturing and Engineering (AME) Industry Alignment Fund – Pre Positioning (IAF-PP) (Grant No. A19D6a0053). J. Yang is with the School of Artificial Intelligence, Jilin University, China. Email: jyang022@e.ntu.edu.sg. W. P. Tay is with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. Email: wptay@ntu.edu.sg.*

other. In [33], the authors adopted a model in which an agent can be influenced by another agent to change its observation to match that of the influencer, while [17] assumed that agents' dependency graphs are disjoint trees. This was extended to general dependency graphs in [6], [18]. In our previous work [34], we considered the use of social network information and community detection to aid in truth discovery. We called our approach VISIT. Note that in these methods, the relationships among event truths, agent reliabilities and agent observations are modeled by predefined models. Such assumptions limit the flexibility of the model to fit complex real data.

Neural networks have shown good promise in modeling nonlinear relationships in many applications [35], [36]. In [37], the relationship between source reliability and claim truthfulness is modeled using a multi-layer neural network and achieved promising accuracy for truth discovery. In [38], a memory network based method is used to model the non-linear relationship between source reliability and claim truthfulness. However, these two methods are both supervised learning methods that do not consider the interpretable structures of agents (e.g., communities of agents). We believe the application of neural networks are more effective in the truth discovery problem if the following issues are properly resolved:

- (a) *Unsupervised learning*: For the truth discovery problem, it is often difficult to obtain enough ground truth labels of events to learn a supervised model. Thus, developing an unsupervised model is important and meaningful.
- (b) *Modeling interpretable structures*: In the truth discovery problem, the observations from agents often imply hidden structures. For instance, agents having similar background, culture, and socio-economic standing may form communities and share similar reliabilities [29], [34]. Successfully discovering the hidden structures can improve the performance of the truth discovery model. However, neural networks are not good at modeling interpretable structures [39].
- (c) *Dealing with dependency*: The interpretable structures in item (b) result in different data samples to have dependencies, which violates the independence assumption that is used in many neural network based methods [40].

In this paper, to solve the first issue, an unsupervised deep autoencoder is used to learn the complex relationship among event truths, agent reliabilities and agent observations. Autoencoders [41] are a kind of unsupervised artificial neural network widely used to learn data features. However, the optimization process of an autoencoder is easily stuck in less attractive local optima [42], thus proper model constraints are required to obtain better performance. The constraints are introduced by Bayesian networks in our model. Bayesian network models provide a natural way to characterize the relationship among variables in an unsupervised way [43], [44], [45]. In this paper, a Bayesian network model is proposed to model the social relationships between agents, which we assume affect each agent's reliability, and to further constrain the learning of the autoencoder. Our approach combines the strengths of unsupervised learning in modeling nonlinear relationships and the strengths of

Bayesian networks in characterizing hidden interpretable structures. Our model is not a straightforward concatenation of a Bayesian network and an autoencoder, but constructs a network of the autoencoder's output variables. The Bayesian network and the autoencoder are learned iteratively.

The rest of this paper is organized as follows. In Section 2, we present our model assumptions. In Section 3, we present our proposed variational inference approach, which is iterative in nature. We show how each update step is derived. Experimental results are discussed in Section 4 and we conclude in Section 5.

Notations: We use boldfaced characters to represent vectors and matrices. Suppose that $\mathbf{M}$ is a matrix, then $\mathbf{M}(m, \cdot)$, $\mathbf{M}(\cdot, m)$, and $\mathbf{M}(m, n)$ denote its m -th row, m -th column, and (m, n) -th element, respectively. The vector $(x_1, \dots, x_N)$ is abbreviated as $(x_i)_{i=1}^N$ or (x_i) if the index set that i runs over is clear from the context. The i -th element of a vector $\mathbf{x}$ is $\mathbf{x}(i)$. Let $\mathbf{1}$ denote a column vector of all 1's and $\mathbf{I}$ be the identity matrix. The vectorized version of $\mathbf{M}$ with the columns of $\mathbf{M}$ stacked together as a single column vector is denoted as $\text{vec}(\mathbf{M})$. We use $\text{Cat}(p_1, \dots, p_K)$ and $\mathcal{N}(\mathbf{U}, \mathbf{V})$ to represent the categorical distribution with category probabilities $p_1, \dots, p_K$ and the normal distribution with mean $\mathbf{U}$ and covariance $\mathbf{V}$, respectively. The notation $\sim$ means equality in distribution. The notation $y \mid x$ denotes a random variable y conditioned on x , and $p(y \mid x)$ denotes its conditional probability density function. $\mathbb{E}$ is the expectation operator and $\mathbb{E}_q$ is expectation with respect to the probability distribution q . The notation $\text{prob}(\theta)$ for a random variable θ is a column vector whose r -th value is the probability $\mathbb{P}(\theta = r)$. We use $I(a, b)$ to denote the indicator function, which equals 1 if $a = b$ and 0 otherwise. We use $|\mathcal{S}|$ to represent the cardinality of the set $\mathcal{S}$.

2 SYSTEM MODEL

Suppose that N agents in a social network with different reliabilities observe J events and each event $j = 1, \dots, J$ has a state θ_j , which can take R possible states. We consider the problem of estimating event truths or states from conflicted agent opinions when each agent only observes a subset of the events. The symbols used in this paper are summarized in Table 1.

In our model, the observation matrix $\mathbf{M}$ is an $N \times J$ generalized matrix where each element $\mathbf{M}(n, j)$ represents the set of observations of agent n about event j . Each agent may form multiple observations about an event. This is useful to model the case where an agent may be undecided about the state of the same event. An entry $\mathbf{M}(n, j)$ is null or the empty set if the agent n does not observe event j . For non-null $\mathbf{M}(n, j)$, we assume that it is generated from $\mathbf{C}_n$, which represents the reliability of agent n 's opinion about the ground truth state θ_j of event j . The reliability matrix $\mathbf{C}_n$ is a $R_1 \times R_2$ matrix, where R_1 and R_2 are two hyper parameters. In this paper, we use a general matrix to represent an agent's reliability, and this includes both the reliability concepts of [16], [19], [20], [21], [22] and confusion matrix of [27], [28], [29]. We also assume that a social network connecting the agents is known and its graph adjacency matrix is given by $\mathbf{A}$, where $\mathbf{A}(n, m) = 1$ iff

TABLE 1
Summary of commonly-used symbols.

Symbol	Description	Variational Parameter in Section 3
$\mathbf{M} = (\mathbf{M}(n, j))_{\substack{1 \leq n \leq N, \\ 1 \leq j \leq J}}$	$\mathbf{M}(n, j)$ is the set of observations of agent n about event j .	N.A.
$\mathbf{C} = (\mathbf{C}_n)_{n=1}^N$	$\mathbf{C}_n$ is the $R_1 \times R_2$ reliability matrix of agent n .	Neural network parameters $\mathbf{w}_R$ of the reliability encoder network.
$\boldsymbol{\theta} = (\theta_j)_{j=1}^J$	θ_j is the true state of event j .	Neural network parameters $\mathbf{w}_E$ of the event encoder network.
$\mathbf{o} = (\mathbf{o}_n)_{n=1}^N$, $\mathbf{u} = (\mathbf{u}_j)_{j=1}^J$, $\mathbf{d} = (\mathbf{d}_{n,j})_{\substack{1 \leq n \leq N, \\ 1 \leq j \leq J}}$	$\mathbf{o}_n, \mathbf{u}_j, \mathbf{d}_{i,j}$ are the outputs of the reliability encoder network, the event encoder network and the decoder network, respectively.	N.A.
$\tilde{\mathbf{C}} = (\tilde{\mathbf{C}}_k)_{k=1}^K$	$\tilde{\mathbf{C}}_k$ is the reliability matrix of community k .	$(\tilde{\boldsymbol{\mu}}_k, \tilde{\boldsymbol{\sigma}}_k)_{k=1}^K$
$\mathbf{A}(n, m) = 1$ (or 0) for $n, m \in \{1, \dots, N\}$	There is a (or no) social connection between agents n and m .	N.A.
$\mathbf{z} = (z_{n \rightarrow m})_{\substack{1 \leq n \leq N, \\ 1 \leq m \leq \bar{N}, m \neq n}}$	$z_{n \rightarrow m}$ is the index of the community agent n subscribes to under the social influence of agent m .	$\phi = (\phi_{n \rightarrow m, k})_{\substack{1 \leq k \leq K, \\ 1 \leq n \leq N, \\ 1 \leq m \leq \bar{N}, m \neq n}}$
$\boldsymbol{\beta} = (\beta_k)_{k=1}^K$	β_k is the social network parameter defined in (7).	$\boldsymbol{\lambda} = (\boldsymbol{\lambda}_k)_{k=1}^K; \boldsymbol{\lambda}_k = (G_k, H_k)$
$\boldsymbol{\pi} = (\boldsymbol{\pi}_n)_{n=1}^N = (\pi_{n,k})_{\substack{1 \leq n \leq N, \\ 1 \leq k \leq K}}$	$\boldsymbol{\pi}_n$ is the distribution of s_n and $z_{n \rightarrow m}$, which are defined in (4) and (6).	$\boldsymbol{\gamma} = (\gamma_{n,k})_{\substack{1 \leq n \leq N, \\ 1 \leq k \leq K}}$
$\mathbf{s} = (s_n)_{n=1}^N$	s_n is the community index of agent n .	$\boldsymbol{\psi} = (\psi_{n,k})_{\substack{1 \leq n \leq N, \\ 1 \leq k \leq K}}$
$(\mathbf{U}_k, \mathbf{V}_k), \boldsymbol{\alpha}, (g_0, h_0),$ $\mathbf{p}^{MV} = (\mathbf{p}_j^{MV})_{j=1}^J, \epsilon,$ $b, \mathbf{b}', \tau$	Hyper-parameters defined in (3), (5), (6), (8), (7), (17), (2), and (26), respectively.	N.A.

agent n and agent m are connected. Our target is to estimate $\boldsymbol{\theta} \triangleq (\theta_j)_{j=1}^J$ from $\mathbf{M}$ and $\mathbf{A}$.

2.1 Observation Model

Let $\mathbf{C} = (\mathbf{C}_n)_{n=1}^N$. The relationship between $\mathbf{M}$ and $(\mathbf{C}, \boldsymbol{\theta})$ is complex and nonlinear, and thus it is challenging to find a proper analytical model for it. Moreover, a data set with a large number of accurate labels is usually unavailable, which hinder the application of supervised learning in the truth discovery problem. To solve these two issues, we model the relationship between $(\mathbf{C}, \boldsymbol{\theta})$ and $\mathbf{M}$ with a multi-layer neural network and perform inference using an unsupervised autoencoder in Section 3.

We represent the observation model for $\mathbf{M}$ by a neural network that learns to decode the agent reliabilities and event states back to the observations. The input to the observation model is $(\mathbf{C}_n, \theta_j)$ for $n = 1, \dots, N$ and $j = 1, \dots, J$. We calculate $(\mathbf{w}_{D_1} \text{prob}(\theta_j)) \odot \text{vec}(\mathbf{C}_n)$, where $\odot$ represents element-wise multiplication and $\mathbf{w}_{D_1}$ is a $R_1 R_2 \times R$ learnable matrix. Next, we input the obtained result to multiple fully connected layers. We denote all the learnable parameters of the observation model including $\mathbf{w}_{D_1}$ as $\mathbf{w}_D$. From the output layer of the observation model,

we obtain an $R \times 1$ vector $\mathbf{d}_{n,j}$. For all $r \in [1, R]$, we assume

$$\text{sum_one_hot}(\mathbf{M}(n, j))(r) \sim \text{Bern}(\mathbf{d}_{n,j}(r)), \quad (1)$$

where $\text{sum_one_hot}(\cdot)$ is the sum of the one-hot vector representation of the elements of its set argument, and $\text{Bern}(\cdot)$ denotes the Bernoulli distribution. We use a Bernoulli distribution to model the case where each agent is may be unsure of its observations and has multiple observations about an event.

2.2 Bayesian Model Constraints

One problem of learning the observation model is that optimizing its neural network weights can become stuck in less attractive local optima [42]. To mitigate this, proper constraints on key latent variables $\mathbf{C}$ and $\boldsymbol{\theta}$ need to be introduced. In this paper, we use a Bayesian network model (see Fig. 1) to construct interpretable constraints. The Bayesian network model not only guides the learning process but also enables to use the community information of the social network linking the agents together. We explain each component of our Bayesian network model below.

2.2.1 Community reliability matrix $\tilde{\mathbf{C}}_n$

We assume that a social network connecting the N agents is known. Agents in a social network tend to form communities [46], where a community consists of agents with

Fig. 1. Our proposed Bayesian network model. The arrows from the nodes labeled as C_n and θ_j to M represent the observation model or decoder network in Fig. 2. The dotted arrows from M to the nodes labeled as C_n and θ_j represent the reliability encoder and event encoder in Fig. 2, respectively. The dotted arrows do not form part of our Bayesian network model, but are instead used in the variational inference of our model in Section 3.

similar opinion. The community that an agent belongs to is stochastic and unknown *a priori* but the maximum possible number of communities in the network is known to be K . We can thus assign an index $1, 2, \dots, K$ in an arbitrary fashion to each community. Let s_n be the community index of agent n . We discuss the probability model governing s_n below. Here, we describe how C_n depends on s_n . For each $k = 1, \dots, K$, let $\tilde{C}_k$ be a matrix of the same size as C_n representing the reliability of community k . We assume

$$C_n(r_1, \cdot) \mid \{\tilde{C}_k(r_1, \cdot), s_n = k\} \sim \mathcal{N}(\tilde{C}_k(r_1, \cdot), b' \mathbf{I}), \quad (2)$$

for $r_1 \in [1, R_1]$, where b' is a hyperparameter. Then C_n can be regarded as a perturbed version of $\tilde{C}_{s_n}$ since $\mathbb{E}[C_n(r_1, \cdot) \mid \tilde{C}_k(r_1, \cdot), s_n = k] = \tilde{C}_k(r_1, \cdot)$ for $r_1 \in [1, R_1]$. We assume that for $k = 1, \dots, K$, and $r_1 \in [1, R_1]$,

$$\tilde{C}_k(r_1, \cdot) \sim \mathcal{N}(\mathbf{U}_k(r_1, \cdot), \mathbf{V}_k), \quad (3)$$

where $\mathbf{U}_k$ and $\mathbf{V}_k$ are hyper parameters. Here we use the normal distribution as the prior as it is an effective prior in many real applications and it allows us to simplify the subsequent inference equations. From Fig. 2, we see that $\{C_n\}$ are learned from different inputs of the reliability encoder. Different from a traditional autoencoder [47], [48], $\{C_n\}$ are not independent and their relationship is modeled by the Bayesian network in Fig. 1. The Bayesian network model guides the learning process of the autoencoder.

2.2.2 Community index s_n

We model the community index s_n of agent n as

$$s_n \sim \text{Cat}(\pi_n), \quad (4)$$

where the mixture weights

$$\pi_n = (\pi_{n,k})_{k=1}^K \sim \text{Dir}(\alpha), \quad (5)$$

with α being a concentration hyperparameter and $\text{Dir}(\alpha)$ is the Dirichlet distribution. Here we use the Dirichlet distribution since the support of a Dirichlet distribution can be regarded as the probabilities of categorical events. Besides, the Dirichlet distribution is the conjugate prior distribution of the categorical distribution in (4) and (6), which thus allows us to derive analytically the posterior distribution of π_n . We use the mixed membership stochastic block model (MMSB) [43] to model the social connection $\mathbf{A}(n, m)$ between agents n and m . In this model, $z_{n \rightarrow m}$ is the community whose belief agent n subscribes to due to the social influence from agent m . Under the influence of different agents, agent n may subscribe to the beliefs of different communities. If both agents n and m subscribe to the belief of the same community, they are more likely to be connected in the social network. We assume the following:

$$\begin{aligned} z_{n \rightarrow m} \mid \pi_n &\sim \text{Cat}(\pi_n), \\ z_{m \rightarrow n} \mid \pi_m &\sim \text{Cat}(\pi_m), \\ \beta_k &\sim \mathcal{B}e(g_0, h_0), \end{aligned} \quad (6)$$

where $\mathcal{B}e(g_0, h_0)$ is the beta distribution with parameters $g_0, h_0 > 0$, $k = 1, \dots, K$, and

$$\begin{aligned} \mathbb{P}(\mathbf{A}(n, m) = 1 \mid z_{n \rightarrow m}, z_{m \rightarrow n}, \beta_{z_{n \rightarrow m}}) \\ = \begin{cases} \beta_{z_{n \rightarrow m}}, & \text{if } z_{n \rightarrow m} = z_{m \rightarrow n}, \\ \epsilon, & \text{if } z_{n \rightarrow m} \neq z_{m \rightarrow n}, \end{cases} \end{aligned} \quad (7)$$

with ϵ being a small constant. In (6), we use beta distribution since the beta distribution is a conjugate prior for the Bernoulli distribution in (7). Note that $\mathbf{A}$ is independent of π when $\mathbf{z}$ is given, as shown in Fig. 1.

2.2.3 Event states θ

A direct method to perform truth discovery is majority voting, i.e., selecting the opinion expressed by the most number of agents as the true state of the event. This assumes that all agents have the same reliability, and that agents are more likely to give the correct opinion than not. Without any prior information, this is a reasonable assumption. Therefore, we let the prior of θ to be given by

$$\theta_j \sim \text{Cat}(\mathbf{p}_j^{MV}) \quad (8)$$

for each $j = 1, \dots, J$, where $\mathbf{p}_j^{MV}(r)$ for $r = 1, \dots, R$ is the proportion of agents who thinks that the state of event j is r . We assume that $\{\theta_j : j = 1, \dots, J\}$ are independent.

3 ART: AUTOENCODER TRUTH DISCOVERY

In this section, we propose an autoencoder based on unsupervised variational inference [49] for the Bayesian model in Fig. 1.

Let $\beta = (\beta_k)$, $\mathbf{z} = (z_{n \rightarrow m})$, $\mathbf{s} = (s_n)$, $\pi = (\pi_n)$, $\theta = (\theta_j)$, and $\tilde{C} = (\tilde{C}_k)$. For simplicity, let $\Omega \triangleq (\beta, \mathbf{z}, \pi, \mathbf{s}, \tilde{C}, \mathbf{C}, \theta)$. As the closed-form of the posterior distribution $p(\Omega \mid \mathbf{M}, \mathbf{A})$ is not available, the variational inference method uses a proposal or variational distribution $q(\Omega; \Lambda)$ to approximate the posterior distribution, where the parameters in the vector Λ are called the variational parameters. Note that $\mathbf{M}$ and $\mathbf{A}$ are assumed to be observed throughout and not included

explicitly in our notation for q . More specifically, the variational parameters are selected to minimize the following cost function:

$$\mathcal{L} = -\mathbb{E}_q [\log p(\mathbf{\Omega} | \mathbf{M}, \mathbf{A}) - \log q(\mathbf{\Omega}; \mathbf{\Lambda})], \quad (9)$$

where the expectation is over the random variable $\mathbf{\Omega}$ with distribution $q(\mathbf{\Omega}; \mathbf{\Lambda})$ conditioned on $\mathbf{M}$ and $\mathbf{A}$. To simplify the optimization procedure, we use the mean-field assumption that is widely used in the literature [49], [50] by choosing

$$q(\mathbf{\Omega}; \mathbf{\Lambda}) \quad (10)$$

$$= q(\beta; \lambda) q(\mathbf{z}; \phi) q(\pi; \gamma) q(\mathbf{s}; \psi) q(\tilde{\mathbf{C}}; \tilde{\mu}, \tilde{\sigma}) q(\mathbf{C}; \mathbf{w}_R) q(\theta; \mathbf{w}_E),$$

where $\mathbf{\Lambda} = (\lambda, \phi, \gamma, \psi, \tilde{\mu}, \tilde{\sigma}, \mathbf{w}_R, \mathbf{w}_E)$, $\lambda = (\lambda_k)_{k=1}^K$, $\phi = (\phi_{n,m})_{n,m}$, $\gamma = (\gamma_{n,k})_{n,k}$, $\psi = (\psi_{n,k})_{n,k}$, $\tilde{\mu} = (\tilde{\mu}_k)_{k=1}^K$, and $\tilde{\sigma} = (\tilde{\sigma}_k)_{k=1}^K$ are the variational parameters.

In the sequel, to simplify notations, we omit the variational parameters in our notations, e.g., we write $q(\beta)$ instead of $q(\beta; \lambda)$. We let $q(\mathbf{C}, \theta) = q(\mathbf{C})q(\theta)$. From the graphical model in Fig. 1, we obtain

$$\begin{aligned} p(\mathbf{\Omega} | \mathbf{M}, \mathbf{A}) \\ \propto p(\tilde{\mathbf{C}}) p(\mathbf{s} | \pi) p(\mathbf{z} | \pi) p(\pi) p(\mathbf{A} | \beta, \mathbf{z}) p(\beta) p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \\ \cdot p(\mathbf{M} | \mathbf{C}, \theta; \mathbf{w}_D) p(\theta), \end{aligned} \quad (11)$$

where $p(\mathbf{M} | \mathbf{C}, \theta; \mathbf{w}_D)$ is the conditional distribution of the output of the decoder network in Fig. 2. We have made its dependence on the decoder network parameters $\mathbf{w}_D$ explicit.

To find the variational parameters, we perform an iterative optimization of $\mathcal{L}$ in which the optimal parameter solutions are updated iteratively at each step. We substitute (10) and (11) into (9) to obtain

$$\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2 + \text{constant}, \quad (12)$$

where the constant term does not contain any variational parameters, and

$$\begin{aligned} \mathcal{L}_1 \triangleq & -\mathbb{E}_{q(\mathbf{C}, \theta, \tilde{\mathbf{C}}, \mathbf{s})} [\log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \\ & + \log p(\mathbf{M} | \mathbf{C}, \theta; \mathbf{w}_D) + \log p(\theta) - \log q(\mathbf{C}, \theta)], \end{aligned} \quad (13)$$

$$\begin{aligned} \mathcal{L}_2 \triangleq & -\mathbb{E}_{q(\tilde{\mathbf{C}}, \mathbf{s}, \pi, \mathbf{z}, \beta)} [\log p(\tilde{\mathbf{C}}) + \log p(\mathbf{s} | \pi) \\ & + \log p(\mathbf{z} | \pi) + \log p(\pi) + \log p(\mathbf{A} | \beta, \mathbf{z}) \\ & + \log p(\beta) - \log q(\tilde{\mathbf{C}}) - \log q(\beta) \\ & - \log q(\mathbf{z}) - \log q(\pi) - \log q(\mathbf{s})]. \end{aligned} \quad (14)$$

We update $q(\mathbf{C})$, $q(\theta)$ by minimizing $\mathcal{L}_1$ and update $q(\beta)$, $q(\mathbf{z})$, $q(\pi)$ by minimizing $\mathcal{L}_2$. Furthermore, we update $q(\tilde{\mathbf{C}})$ and $q(\mathbf{s})$ by minimizing

$$\begin{aligned} \mathcal{L}_3 \triangleq & -\mathbb{E}_{q(\mathbf{C}, \tilde{\mathbf{C}}, \mathbf{s}, \pi)} [\log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) + \log p(\tilde{\mathbf{C}}) \\ & + \log p(\mathbf{s} | \pi) - \log q(\tilde{\mathbf{C}}) - \log q(\mathbf{s})] \\ = & -\mathbb{E}_{q(\mathbf{C}, \tilde{\mathbf{C}}, \mathbf{s}, \pi)} [\log p(\mathbf{s} | \mathbf{C}, \tilde{\mathbf{C}}, \pi) + \log p(\mathbf{C} | \tilde{\mathbf{C}}) \\ & + \log p(\tilde{\mathbf{C}}) - \log q(\tilde{\mathbf{C}}) - \log q(\mathbf{s})] \end{aligned} \quad (15)$$

$$\begin{aligned} = & -\mathbb{E}_{q(\mathbf{C}, \tilde{\mathbf{C}}, \mathbf{s}, \pi)} [\log p(\tilde{\mathbf{C}} | \mathbf{C}, \mathbf{s}) + \log p(\mathbf{C} | \mathbf{s}) \\ & + \log p(\mathbf{s} | \pi) - \log q(\tilde{\mathbf{C}}) - \log q(\mathbf{s})]. \end{aligned} \quad (16)$$

Equations (15) and (16) show $q(\mathbf{s})$ and $q(\tilde{\mathbf{C}})$ can be updated by minimizing

$$-\mathbb{E}_{q(\mathbf{C}, \tilde{\mathbf{C}}, \mathbf{s}, \pi)} [\log p(\mathbf{s} | \mathbf{C}, \tilde{\mathbf{C}}, \pi) - \log q(\mathbf{s})]$$

and

$$-\mathbb{E}_{q(\mathbf{C}, \tilde{\mathbf{C}}, \mathbf{s}, \pi)} [\log p(\tilde{\mathbf{C}} | \mathbf{C}, \mathbf{s}) - \log q(\tilde{\mathbf{C}})],$$

respectively.

The variational parameters are optimized and the variational distributions updated iteratively in a procedure that we call Autoencoder Truth (ART) discovery (since we make use of an autoencoder network described below). Its high-level pseudo code for the i -th iteration is shown in Algorithm 1. In the following, we describe how the variational distributions are chosen and how the estimate for the optimal variational parameters are updated in each iteration.

Algorithm 1 ART (i -th iteration)

Input: Variational parameters in $(i-1)$ -th iteration, opinions $\mathbf{M}$, social network data $\mathbf{A}$.

Output: Variational parameters in i -th iteration.

```

for each agent  $n$  in  $\{1, \dots, N\}$  do
  for each agent pair  $(n, m)$  in  $\{(n, m)\}_{m=1, m \neq n}^N$  do
    Update  $\phi_{n \rightarrow m}$  and  $\phi_{m \rightarrow n}$  using (32) and (33).
  end for
  Update  $\psi_n$  using (39).
  Update  $\gamma_n$  using (44).
  Sample  $\mathbf{C}_n$  using (22).
  Sample  $s_n$  from  $q(s_n) = \text{Cat}((\psi_{n,k})_{k=1}^K)$ .
end for
Update  $\lambda$  using (28) and (29).
Update  $\tilde{\mu}$  and  $\tilde{\sigma}$  using (42a) and (42b).
Sample  $\tilde{\mathbf{C}}$ , and  $\theta$  using (40) and (23).
Learn the autoencoder (i.e., update  $\mathbf{w}_R$ ,  $\mathbf{w}_E$ , and  $\mathbf{w}_D$  in Section 3.2).
return  $\phi, \psi, \gamma, \mathbf{C}, \lambda, \tilde{\mu}, \tilde{\sigma}, \tilde{\mathbf{C}}$ , and  $\theta$ .

```

3.1 Reliability and Event Encoders

We model $q(\mathbf{C}; \mathbf{w}_R)$ and $q(\theta; \mathbf{w}_E)$ as two encoder networks with parameters $\mathbf{w}_R$ and $\mathbf{w}_E$, respectively. These two encoders form part of the autoencoder shown in Fig. 2 with the observation model for $\mathbf{M}$ given in Section 2.1 as the decoder network. The details of the two encoder networks are as follows.

- (a) The reliability encoder network is used to infer the reliability of each agent n , with $\text{sum_one_hot}(\mathbf{M}(n, \cdot))$ as its input. If $\mathbf{M}(n, j)$ is null, then $\text{sum_one_hot}(\mathbf{M}(n, j))$ is a zero vector. All the layers of this encoder are fully connected. Let its parameters be $\mathbf{w}_R$ and its $R_1 R_2 \times 1$ softmax output be $\mathbf{o}_n$. We assume

$$\text{vec}(\mathbf{C}_n)(i) \sim \mathcal{N}(\mathbf{o}_n(i), b), \quad (17)$$

where b is a hyperparameter.

- (b) The event encoder network is used to infer the state of each event from $\mathbf{M}$. All the layers are also fully connected. Let its parameters be $\mathbf{w}_E$ and its output be $\mathbf{u}_j$. We assume

$$\theta_j \sim \text{Cat}(\mathbf{u}_j). \quad (18)$$

Fig. 2. Autoencoder model.

We let $q(\mathbf{C}; \mathbf{w}_R)$ correspond to the Gaussian distribution (17), and $q(\boldsymbol{\theta}; \mathbf{w}_E)$ correspond to the distribution (18).

3.2 Updating of $q(\mathbf{C}, \boldsymbol{\theta})$

The relationship between $\mathbf{C}, \boldsymbol{\theta}$ and $\mathbf{M}$ is complex and non-linear. As discussed before in Sections 2.1 and 3.1, we use neural networks to model $p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D)$ and we aim to optimize these neural networks to minimize $\mathcal{L}_1$. From (13), we have

$$\begin{aligned} \mathcal{L}_1 &= -\mathbb{E}_{q(\mathbf{C})q(\boldsymbol{\theta})q(\tilde{\mathbf{C}})q(\mathbf{s})} \left[\log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \\ &\quad \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) + \log p(\boldsymbol{\theta}) - \log q(\mathbf{C}) - \log q(\boldsymbol{\theta}) \right] \\ &= -\mathbb{E}_{q(\mathbf{C})q(\boldsymbol{\theta})q(\tilde{\mathbf{C}})q(\mathbf{s})} \left[\log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \\ &\quad \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \right] \quad (19) \\ &\quad - \mathbb{E}_{q(\boldsymbol{\theta})} [\log p(\boldsymbol{\theta}) - \log q(\boldsymbol{\theta})]. \end{aligned}$$

Recall that $\mathbf{w}_D$, $\mathbf{w}_R$ and $\mathbf{w}_E$ denote the parameters of the decoder network, the reliability encoder network and the event encoder network, respectively. To learn $\mathbf{w}_D$ with the gradient descent method, we need to compute the gradient of $\mathcal{L}_1$ with respect to $\mathbf{w}_D$. Denoting the two expectation terms of (19) as $\mathcal{L}_{11}$ and $\mathcal{L}_{12}$, namely

$$\mathcal{L}_1 = \mathcal{L}_{11} + \mathcal{L}_{12}, \quad (20)$$

and we have

$$\nabla_{\mathbf{w}_D} \mathcal{L}_1 = \nabla_{\mathbf{w}_D} \mathcal{L}_{11} + \nabla_{\mathbf{w}_D} \mathcal{L}_{12}.$$

As $\boldsymbol{\theta}$ is a discrete variable, $\mathcal{L}_{12}$ is easy to compute. The gradient of $\mathcal{L}_{11}$ with respect to $\mathbf{w}_D$ is given by

$$\begin{aligned} \nabla_{\mathbf{w}_D} \mathcal{L}_{11} &= -\nabla_{\mathbf{w}_D} \mathbb{E}_{q(\mathbf{C})q(\boldsymbol{\theta})q(\tilde{\mathbf{C}})q(\mathbf{s})} \left\{ \log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \\ &\quad \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \right\} \\ &= -\mathbb{E}_{q(\mathbf{C})q(\boldsymbol{\theta})q(\tilde{\mathbf{C}})q(\mathbf{s})} \left\{ \nabla_{\mathbf{w}_D} \left\{ \log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \right. \\ &\quad \left. \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \right\} \right\}, \end{aligned}$$

which can be computed using the stochastic gradient descent (SGD) algorithm. In each iteration, we replace $\nabla_{\mathbf{w}_D} \mathcal{L}_{11}$ with its unbiased estimator $\nabla_{\mathbf{w}_D} \mathcal{L}'_{11}$, where

$$\mathcal{L}'_{11} \triangleq \log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \quad (21)$$

with $\mathbf{C}$, $\boldsymbol{\theta}$, $\mathbf{s}$, and $\tilde{\mathbf{C}}$ being sampled from $q(\mathbf{C})$, $q(\boldsymbol{\theta})$, $q(\mathbf{s})$, and $q(\tilde{\mathbf{C}})$ respectively.

We cannot use the same process to deal with $\mathbf{w}_R$ and $\mathbf{w}_E$. This is because

$$\begin{aligned} \nabla_{\mathbf{w}_R} \mathcal{L}_{11} &= -\nabla_{\mathbf{w}_E} \mathbb{E}_{q(\mathbf{C})q(\boldsymbol{\theta})q(\tilde{\mathbf{C}})q(\mathbf{s})} \left[\log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \\ &\quad \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \right] \\ &\neq -\mathbb{E}_{q(\mathbf{C})q(\boldsymbol{\theta})q(\tilde{\mathbf{C}})q(\mathbf{s})} \left[\nabla_{\mathbf{w}_R} \left\{ \log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \right. \\ &\quad \left. \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \right\} \right], \end{aligned}$$

as $\mathbf{w}_R$ is the parameter of $q(\mathbf{C})$. The same reason applies for $\mathbf{w}_E$. To obtain unbiased estimators for these variational parameters, we need to use the reparameterization trick [51].

From (17), we have $\text{vec}(\mathbf{C}_n)(i) \sim \mathcal{N}(\mathbf{o}_n(i), b)$, where $\mathbf{o}_n$ is generated by the reliability encoder network and b is a hyperparameter. We can reparameterize $\mathbf{C}_n(i)$ as

$$\begin{aligned} \zeta_n &\sim \mathcal{N}(0, 1), \\ \text{vec}(\mathbf{C}_n)(i) &= \mathbf{o}_n(i) + \zeta_n b. \end{aligned} \quad (22)$$

From (18), $\theta_j \sim \text{Cat}(\mathbf{u}_j)$, where the weight vector $\mathbf{u}_j \triangleq (\mathbf{u}_j(r))_r$ is generated by the event encoder network. Then, according to (1) in [52], we can reparameterize θ_j as

$$\begin{aligned} \chi_j(r) &\sim \text{Gumbel}(0, 1) \\ \theta_j &= \arg \max_r [\chi_j(r) + \log \mathbf{u}_j(r)], \end{aligned} \quad (23)$$

where the Gumbel(0, 1) distribution can be sampled by first drawing $\Upsilon \sim \text{Uniform}(0, 1)$ and then computing $\chi_j(r) = -\log(-\log(\Upsilon))$. With $\mathbf{C}_n$ and θ_j being (22) and

(23) respectively, and letting $\zeta = (\zeta_n)$, $\chi = (\chi_j)$ and $\mathbf{u} = (\mathbf{u}_j)$, we then obtain

$$\begin{aligned}\nabla_{\mathbf{w}_R} \mathcal{L}_{11} &= -\nabla_{\mathbf{w}_R} \mathbb{E}_{q(\mathbf{C})q(\boldsymbol{\theta})q(\tilde{\mathbf{C}})q(\mathbf{s})} \left[\log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \\ &\quad \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \right] \\ &= -\mathbb{E}_{p(\zeta)p(\chi)q(\tilde{\mathbf{C}})q(\mathbf{s})} \left[\nabla_{\mathbf{w}_R} \left\{ \sum_{n=1}^N \log \mathcal{N}(\mathbf{o}_n + \zeta_n \mathbf{b}; \right. \right. \\ &\quad \left. \left. \text{vec}(\tilde{\mathbf{C}}_{s_n}), b' \mathbf{I}) + \log p(\mathbf{M} | \zeta, \chi; \mathbf{w}_D, \mathbf{o}, \mathbf{u}) \right\} \right],\end{aligned}\quad (24)$$

where $\mathcal{N}(\cdot; \boldsymbol{\mu}, \boldsymbol{\sigma})$ is the Gaussian probability density function with mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\sigma}$. Here, $\mathbf{o} = (\mathbf{o}_n)$ is a function of $\mathbf{w}_R$ and $\mathbf{u}$ is a function of $\mathbf{w}_E$. We also have

$$\begin{aligned}\nabla_{\mathbf{w}_E} \mathcal{L}_{11} &= -\mathbb{E}_{p(\zeta)p(\chi)q(\tilde{\mathbf{C}})q(\mathbf{s})} \left[\nabla_{\mathbf{w}_E} \{ \log p(\mathbf{C} | \tilde{\mathbf{C}}, \mathbf{s}) \right. \\ &\quad \left. + \log p(\mathbf{M} | \mathbf{C}, \boldsymbol{\theta}; \mathbf{w}_D) - \log q(\mathbf{C}) \} \right] \\ &= -\mathbb{E}_{p(\zeta)p(\chi)} [\nabla_{\mathbf{w}_E} \{ \log p(\mathbf{M} | \zeta, \chi; \mathbf{w}_D, \mathbf{o}, \mathbf{u}) \}].\end{aligned}\quad (25)$$

Now in each iteration, we can apply SGD to find an unbiased estimator of the gradient of $\mathcal{L}_1$ with respect to $\mathbf{w}_E$ and $\mathbf{w}_R$.

Remark 1. The max function in equation (23) is not differentiable and following [52], we use the softmax function as an approximation to $\arg \max$, i.e., we let

$$\text{prob}(\theta_j)(r) = \frac{\exp(\chi_j(r) + \log \mathbf{u}_j(r)/\tau)}{\sum_{r'} \exp(\chi_j(r') + \log \mathbf{u}_j(r')/\tau)}, \quad (26)$$

where τ is the temperature parameter.

Remark 2. The term $\mathcal{L}_{12}$ in (20) is used to constrain the distance between the variational distribution $q(\boldsymbol{\theta})$ and the prior distribution $p(\boldsymbol{\theta})$ of event states. When the dataset includes enough observations or useful social network information, then $\mathcal{L}_{12}$ is more important in the initial iterations and the importance decreases as the number of iterations increases. In these cases, we revise (20) to

$$\mathcal{L}_1 = \mathcal{L}_{11} + \kappa^i \mathcal{L}_{12}, \quad (27)$$

where $0 < \kappa < 1$ is a hyperparameter and i denotes the i -th iteration.

3.3 Updating of $q(\tilde{\mathbf{C}}, \mathbf{s}, \boldsymbol{\pi}, \mathbf{z}, \boldsymbol{\beta})$

We want to find variational parameters corresponding to $(\tilde{\mathbf{C}}, \mathbf{s}, \boldsymbol{\pi}, \mathbf{z}, \boldsymbol{\beta})$ to minimize $\mathcal{L}_2$. To achieve this, we iteratively update these variational parameters in ART. The variational parameters corresponding to $\boldsymbol{\beta}$ and $\mathbf{z}$ are updated in the same way as our previous work [34]. For completeness, we reproduce the results below.

3.3.1 Social network parameter $\boldsymbol{\beta}$

Let $\boldsymbol{\lambda}_k = (G_k, H_k)$. We choose the variational distribution of β_k to be in the same exponential family as its posterior distribution, namely $q(\beta_k) = \mathcal{B}e(G_k, H_k)$. Similar to (14) and (15) in our previous work [34], we can show that

$$G_k = \sum_{(n,m)} \mathbf{A}(n, m) \phi_{n \rightarrow m, k} \phi_{m \rightarrow n, k} + g_0, \quad (28)$$

$$H_k = \sum_{(n,m)} (1 - \mathbf{A}(n, m)) \phi_{n \rightarrow m, k} \phi_{m \rightarrow n, k} + h_0, \quad (29)$$

where $\phi_{n \rightarrow m, k} = q(z_{n \rightarrow m} = k)$ is defined in Section 3.3.2.

From (10) in [53], we also have

$$\mathbb{E}_{q(\beta_k)} [\log(\beta_k)] = \Psi(G_k) - \Psi(G_k + H_k), \text{ and} \quad (30)$$

$$\mathbb{E}_{q(\beta_k)} [\log(1 - \beta_k)] = \Psi(H_k) - \Psi(G_k + H_k), \quad (31)$$

which are used in computing the variational distributions of other parameters in our model. Here, $\Psi(\cdot)$ is the digamma function.

3.3.2 Community membership indicators $\mathbf{z}$

We let the variational distribution of $z_{n \rightarrow m}$ to be in the same exponential family as its posterior distribution, namely a categorical distribution with probabilities $(\phi_{n \rightarrow m, k})_{k=1}^K$. Similar to (19) in our previous work [34], one can show that if $\mathbf{A}(n, m) = 0$,

$$\phi_{n \rightarrow m, k} \propto \exp\{\phi_{m \rightarrow n, k} (\mathbb{E}_{q(\beta_k)} [\log(\beta_k)] - \log(\epsilon)) + \mathbb{E}_{q(\pi_n)} [\log(\pi_{n, k})]\}, \quad (32)$$

where $\mathbb{E}_{q(\beta_k)} [\log(\beta_k)]$ and $\mathbb{E}_{q(\pi_n)} [\log(\pi_{n, k})]$ are computed using (30) and (45) in the sequel, respectively. On the other hand, if $\mathbf{A}(n, m) = 0$, we have

$$\phi_{n \rightarrow m, k} \propto \exp\{\phi_{m \rightarrow n, k} (\mathbb{E}_{q(\beta_k)} [\log(1 - \beta_k)] - \log(1 - \epsilon)) + \mathbb{E}_{q(\pi_n)} [\log(\pi_{n, k})]\}, \quad (33)$$

where $\mathbb{E}_{q(\beta_k)} [\log(1 - \beta_k)]$ is computed in (31) in the sequel.

3.3.3 Event community indices $\mathbf{s}$

We take $q(s_n) = \text{Cat}((\psi_{n, k})_{k=1}^K)$, where $(\psi_{n, k})_{k=1}^K$ is the variational parameter. Let $\hat{\psi}_{n, k} \triangleq \log(\psi_{n, k})$. We have

$$\begin{aligned}p(s_n = k | \mathbf{C}_n, \tilde{\mathbf{C}}_k, \boldsymbol{\pi}_n) &\propto p(s_n = k | \boldsymbol{\pi}_n) p(\mathbf{C}_n | \tilde{\mathbf{C}}_k) \\ &= \pi_{n, k} \prod_{r_1=1}^{R_1} \prod_{r_2=1}^{R_2} \mathcal{N}(\mathbf{C}_n(r_1, r_2); \tilde{\mathbf{C}}_k(r_1, r_2), b').\end{aligned}$$

Thus, $p(s_n = k | \mathbf{C}_n, \tilde{\mathbf{C}}_k, \boldsymbol{\pi})$ is a categorical distribution and is in the same exponential family as $q(s_n) = \text{Cat}((\psi_{n, k})_{k=1}^K)$. Let

$$\begin{aligned}\varsigma_{n, k} &\triangleq \log \left(\pi_{n, k} \prod_{r_1=1}^{R_1} \prod_{r_2=1}^{R_2} \mathcal{N}(\mathbf{C}_n(r_1, r_2); \tilde{\mathbf{C}}_k(r_1, r_2), b') \right), \\ \varsigma_{n, 1:K} &\triangleq \log \left(\sum_{k=1}^K \exp(\varsigma_{n, k}) \right).\end{aligned}$$

The natural parameter of $p(s_n | \mathbf{C}_n, \tilde{\mathbf{C}}_k, \boldsymbol{\pi})$ is

$$(\varsigma_{n, k} - \varsigma_{n, 1:K})_{k=1}^K. \quad (34)$$

Assume the natural parameter of $q(s_n)$ is

$$(\hat{\psi}_{n,k} - \hat{\psi}_{n,1:K})_{k=1}^K, \quad (35)$$

where $\hat{\psi}_{n,1:K} \triangleq \log \left(\sum_{k=1}^K \exp(\hat{\psi}_{n,k}) \right)$. According to the relationship between the natural gradient and the natural parameters (i.e., (22) in [49]), the natural gradient of $\mathcal{L}_3$ in (15) with respect to $\hat{\psi}_{n,k}$ can be derived from (34) and (35) and the result is

$$\hat{\nabla}_{\hat{\psi}_{n,k}} \mathcal{L}_3 = \hat{\psi}_{n,k} - \hat{\psi}_{n,1:K} - \mathbb{E}_{q(\mathbf{C}_n, \tilde{\mathbf{C}}_k, \boldsymbol{\pi}_n)} [\varsigma_{n,k} - \varsigma_{n,1:K}].$$

We sample $\mathbf{C}_n$ from $q(\mathbf{C}_n)$ and obtain the unbiased estimator of $\hat{\nabla}_{\hat{\psi}_{n,k}} \mathcal{L}_3$ as

$$\begin{aligned} \hat{\nabla}_{\hat{\psi}_{n,k}} \mathcal{L}'_3 &= \hat{\psi}_{n,k} \\ &- \sum_{r_1=1}^{R_1} \sum_{r_2=1}^{R_2} \mathbb{E}_{q(\tilde{\mathbf{C}}_k(r_1, r_2))} \left[\log \mathcal{N}(\mathbf{C}_n(r_1, r_2); \tilde{\mathbf{C}}_k(r_1, r_2), b') \right] \\ &- \mathbb{E}_{q(\boldsymbol{\pi}_n)} [\log \pi_{n,k}] - \Delta_n, \end{aligned} \quad (36)$$

where $\Delta_n \triangleq \hat{\psi}_{n,1:K} - \mathbb{E}_{q(\mathbf{C}_n, \tilde{\mathbf{C}}_k, \boldsymbol{\pi}_n)} [\varsigma_{n,1:K}]$ and it is constant for $k = 1, \dots, K$. Then we update $\hat{\psi}_{n,k}$ using

$$\hat{\psi}_{n,k}^{(i)} = \hat{\psi}_{n,k}^{(i-1)} - \rho^{(i)} \hat{\nabla}_{\hat{\psi}_{n,k}} \mathcal{L}'_3, \quad (37)$$

where $\rho^{(i)}$ is the known step size at i -th iteration. Let $\Xi_{n,k}$ be the first three terms of the right-hand side of (36). We compute exponential function of both sides of (37) and obtain

$$\begin{aligned} \psi_{n,k}^{(i)} &= \psi_{n,k}^{(i-1)} \exp \left(-\rho^{(i)} \Xi_{n,k}^{(i)} \right) \exp(\rho^{(i)} \Delta_n^{(i)}) \\ &= \psi_{n,k}^{(i-1)} \exp \left\{ -\rho^{(i)} \hat{\psi}_{n,k}^{(i-1)} \right. \\ &\quad - \rho^{(i)} \sum_{r_1=1}^{R_1} \sum_{r_2=1}^{R_2} \frac{1}{2b'^2} \left((\tilde{\sigma}_k^{(i-1)}(r_1, r_2))^2 \right. \\ &\quad \left. \left. + \left(\tilde{\mu}_k^{(i-1)}(r_1, r_2) - \mathbf{C}_n^{(i)}(r_1, r_2) \right)^2 \right) \right. \\ &\quad \left. + \rho^{(i)} \mathbb{E}_{q(\boldsymbol{\pi}_n)} [\log \pi_{n,k}] \right\} \exp(\rho^{(i)} \Delta_n^{(i)}) \end{aligned} \quad (38)$$

where $\tilde{\mu}$ and $\tilde{\sigma}$ are variational parameters of $\tilde{\mathbf{C}}$ defined in Section 3.3.4 and $\mathbb{E}_{q(\boldsymbol{\pi}_n)} [\log \pi_{n,k}]$ can be computed by (45) below. Since $\sum_{k=1}^K \psi_{n,k}^{(i)} = 1$ and $\exp(\rho^{(i)} \Delta_n^{(i)})$ is constant for every $\psi_{n,k}^{(i)}$, $k = 1, \dots, K$, in each iteration, we only need to compute $\psi_{n,k}^{(i-1)} \exp \left(-\rho^{(i)} \Xi_{n,k}^{(i)} \right)$ and then

$$\psi_{n,k}^{(i)} = \frac{\psi_{n,k}^{(i-1)} \exp \left(-\rho^{(i)} \Xi_{n,k}^{(i)} \right)}{\sum_{k=1}^K \psi_{n,k}^{(i-1)} \exp \left(-\rho^{(i)} \Xi_{n,k}^{(i)} \right)}. \quad (39)$$

3.3.4 Community reliability matrix $\tilde{\mathbf{C}}$

We let the variational distribution of $\tilde{\mathbf{C}}_k$ for each $k = 1, \dots, K$ to be given by

$$q(\tilde{\mathbf{C}}_k(r_1, r_2)) = \mathcal{N}(\tilde{\mu}_k(r_1, r_2), \tilde{\sigma}_k(r_1, r_2)). \quad (40)$$

We also have

$$\begin{aligned} p(\tilde{\mathbf{C}}_k(r_1, r_2) \mid \mathbf{s}, \{\mathbf{C}_n(r_1, r_2)\}_n) \\ \propto p(\tilde{\mathbf{C}}_k(r_1, r_2)) \prod_{\{n: s_n=k\}} p(\mathbf{C}_n(r_1, r_2) \mid \tilde{\mathbf{C}}_k(r_1, r_2)), \end{aligned} \quad (41)$$

which is a normal distribution with mean (cf. (2) and (3))

$$\begin{aligned} \frac{1}{\mathbf{V}(r_1, r_2)^{-2} + \frac{1}{b'} \sum_n I(s_n, k)} \left(\frac{\mathbf{U}_k(r_1, r_2)}{\mathbf{V}_k(r_1, r_2)^2} \right. \\ \left. + \frac{1}{b'} \sum_n \mathbf{C}_n(r_1, r_2) I(s_n, k) \right) \end{aligned}$$

and variance

$$\frac{1}{\mathbf{V}_k(r_1, r_2)^{-2} + \frac{1}{b'} \sum_n I(s_n, k)}.$$

Let $\hat{\mu}_k(r_1, r_2) \triangleq \frac{\tilde{\mu}_k(r_1, r_2)}{\tilde{\sigma}_k(r_1, r_2)^2}$ and $\hat{\sigma}_k(r_1, r_2)^2 \triangleq \frac{1}{2\tilde{\sigma}_k(r_1, r_2)^2}$. As $q(\tilde{\mathbf{C}}_k(r_1, r_2))$ is in the same exponential family as $p(\tilde{\mathbf{C}}_k(r_1, r_2) \mid \mathbf{s}, \{\mathbf{C}_n(r_1, r_2)\}_n)$ and its natural parameter is $(\hat{\mu}_k(r_1, r_2), \hat{\sigma}_k(r_1, r_2))$. Similar to Section 3.3.3, we sample $\mathbf{C}_n$ from $q(\mathbf{C}_n)$ and update $\hat{\mu}_k(r_1, r_2)$ and $\hat{\sigma}_k^2(r_1, r_2)$ using

$$\begin{aligned} \hat{\mu}_k^{(i)}(r_1, r_2) &= \hat{\mu}_k^{(i-1)}(r_1, r_2) - \rho^{(i)} \left(\hat{\mu}_k^{(i-1)}(r_1, r_2) \right. \\ &\quad \left. - \frac{\mathbf{U}_k(r_1, r_2)}{\mathbf{V}_k(r_1, r_2)^2} - \frac{1}{b'} \sum_n \mathbf{C}_n^{(i)}(r_1, r_2) \psi_{n,k}^{(i)} \right), \end{aligned} \quad (42a)$$

$$\begin{aligned} \hat{\sigma}_k^{(i)}(r_1, r_2)^2 &= \hat{\sigma}_k^{(i-1)}(r_1, r_2)^2 - \rho^{(i)} \left(\hat{\sigma}_k^{(i-1)}(r_1, r_2)^2 \right. \\ &\quad \left. + \frac{1}{2} \left(\mathbf{V}_k(r_1, r_2)^{-2} + \frac{1}{b'} \sum_n \psi_{n,k}^{(i)} \right) \right). \end{aligned} \quad (42b)$$

Finally, we obtain

$$\begin{aligned} \tilde{\sigma}_k^{(i)}(r_1, r_2)^2 &= -\frac{1}{2\hat{\sigma}_k^{(i)}(r_1, r_2)^2}, \\ \tilde{\mu}_k^{(i)}(r_1, r_2) &\triangleq \hat{\mu}_k^{(i)}(r_1, r_2) \tilde{\sigma}_k^{(i)}(r_1, r_2)^2. \end{aligned}$$

3.3.5 Mixture weights $\boldsymbol{\pi}$

We let $q(\boldsymbol{\pi}_n) = \text{Dir}(\gamma_n)$ and thus $q(\boldsymbol{\pi}_n)$ is an exponential family distribution and its variational parameter γ_n is also its natural parameter. To find the variational parameter $\boldsymbol{\gamma}$ that minimizes $\mathcal{L}_2$, we find the partial derivative

$$\begin{aligned} \nabla_{\boldsymbol{\gamma}} \mathcal{L}_2 &= -\nabla_{\boldsymbol{\gamma}} \left\{ \mathbb{E}_{q(\boldsymbol{\pi})q(\mathbf{s})q(\mathbf{z})} [\log p(\mathbf{s} \mid \boldsymbol{\pi}) + \log p(\mathbf{z} \mid \boldsymbol{\pi}) \right. \\ &\quad \left. + \log p(\boldsymbol{\pi}) - \log q(\boldsymbol{\pi})] \right\} \\ &= -\nabla_{\boldsymbol{\gamma}} \left\{ \mathbb{E}_{q(\boldsymbol{\pi})q(\mathbf{s})q(\mathbf{z})} [\log p(\boldsymbol{\pi} \mid \mathbf{s}, \mathbf{z}) - \log q(\boldsymbol{\pi})] \right\}, \end{aligned} \quad (43)$$

where

$$\begin{aligned}
& p(\pi_n \mid \{s_i\}_{i=1}^N, \{z_{n \rightarrow m}\}_{m=1, m \neq n}^N) \\
& \propto \prod_{i=1}^N p(s_i \mid \pi_n) \prod_{m=1, m \neq n}^N p(z_{n \rightarrow m} \mid \pi_n) p(\pi_n) \\
& \propto \text{Dir} \left(\left(\frac{\alpha}{K} + \sum_{m=1, m \neq n}^N I(z_{n \rightarrow m}, k) + \sum_{i=1}^N I(s_i, k) \right)_{k=1}^K \right).
\end{aligned}$$

Recall that K represents the maximum number of communities. As $q(\pi_k) = \text{Dir}(\gamma_k)$ is in the same exponential family as $p(\pi_n \mid s_n, \{z_{n \rightarrow m}\}_{m=1, m \neq n}^N)$, thus if we let (43) be zero, we obtain

$$\begin{aligned}
\gamma_{n,k} &= \mathbb{E}_{q(s_n, \{z_{n \rightarrow m}\}_{m=1, m \neq n}^N)} \left[\frac{\alpha}{K} + \sum_{m=1, m \neq n}^N I(z_{n \rightarrow m}, k) \right. \\
& \quad \left. + \sum_{i=1}^N I(s_i, k) \right] \\
&= \frac{\alpha}{K} + \sum_{m=1, m \neq n}^N \phi_{n \rightarrow m, k} + \sum_{i=1}^N \psi_{i, k}.
\end{aligned} \tag{44}$$

From (10) in [53], we also have

$$\mathbb{E}_{q(\pi_n)} [\log(\pi_{n,k})] = \Psi(\gamma_{n,k}) - \Psi \left(\sum_{k=1}^K \gamma_{n,k} \right), \tag{45}$$

which is used in (39). Recall that $\Psi(\cdot)$ is the digamma function.

4 EXPERIMENTAL RESULTS

In this section, experiments on three real datasets are presented¹. We adopt majority voting, BCC [27], CBCC [29], VISIT [34], TruthFinder [19], AccuSim [20], GTM [54], CRH [55], CATD [21], and KDEm [26] as the state-of-the-art benchmark methods. We test the performance of different methods on the IMDB dataset augmented with Twitter information, which we have made available in [56]. We also test the performance of different methods on the Sentiment Polarity (SP) dataset and the Weather Sentiment (CF) dataset. The last two datasets do not provide network information. To emulate network information in our experiments, for any pair of agents, we calculate the observation differences of their commonly observed events. Next, we calculate the mean of the absolute values of the differences and assume the pair of agents are connected if the mean is less than or equal to 0.2.

4.1 Description of Datasets

4.1.1 Sentiment Polarity (SP) Dataset

The Sentiment Polarity dataset [57] contains agent classifications for movie comments. All the movie comments are from the website Rotten Tomatoes.² Agents were asked to annotate each comment as either positive (1) or negative (0).

The uncertainty or confidence of the agents' annotations are however unavailable. Ground truth labels are from experts. The dataset contains 27,746 evaluations from 203 agents on 4,999 movie comments (i.e., events).

4.1.2 IMDB Dataset

We collected data from the website IMDB³ and Twitter [56]. If a user rates a movie in IMDB and clicks the share button, a Twitter message is generated. We collected movie evaluations from IMDB and social network information from Twitter. We divide the movie evaluations into 2 levels: bad (0-5), good (6-10). We treat the ratings on the IMDB website, which are based on the aggregated evaluations from all users, as the event truths whereas our observations come from only a subset of users who share their ratings on Twitter. To better show the influence of social network information on event truth discovery, we delete small subnetworks that have less than 5 agents each. The final dataset [56] we use consists of 2266 evaluations from 209 individuals on 245 movies (events) and also the social network between these 209 individuals. Similar to [58], [59], we regard the social network to be undirected as both follower or following relationships indicate that the two users have similar taste.

4.1.3 Weather Sentiment (CF) Dataset

The Weather Sentiment dataset was provided by Crowd-Flower (CF) and can be found in [57]. The agents were asked to classify the tweets with respect to the weather sentiment into the following categories: negative (0), neutral (1), positive (2), tweet not related to weather (3) and cannot tell (4). The dataset contains 1,720 evaluations from 461 agents on 300 tweets (i.e., events).

4.2 Experiment Settings

The hyperparameters in our proposed ART method are tuned using another dataset called MS from [57]. Since the hyperparameters of our method are tuned with this dataset, we do not compare the performance of ART and the benchmark methods on this dataset. The values of the hyperparameters in our method are given in Table 2. We use the same set of hyperparameters except κ to do experiments on all the three datasets in Section 4.1.

4.3 Sensitivity Analysis

We conduct experiments on the three datasets using the same set of hyperparameters except κ in (27). The CF dataset contains only 1,720 observations and has no social network information and thus the majority voting prior is important in all the learning iterations. We thus set κ be 1.0. For the IMDB and SP datasets, we set κ to be 0.9 since the SP dataset is rich in observations and the IMDB dataset contains social network information. For these two datasets, the importance of the majority voting prior decreases as the iteration progresses. The results in Section 4.4 suggest that our method is not sensitive to the hyperparameters and can achieve the comparable accuracies on different datasets using the same set of hyperparameters except κ . In the following, we perform experiments to show the influence of different hyperparameters on our method.

1. Code: <https://github.com/yitianhoulai/ART>

2. <https://www.rottentomatoes.com/>

3. <https://www.imdb.com/>

TABLE 2
Experiment Settings

Hyper-parameter	Value
ϵ in (7)	10^{-10}
b in (17) and b' in (2)	0.1
Each element of $\mathbf{U}_k$ in (3)	Random value between 0.4 and 0.5
Each element of $\mathbf{V}_k$ in (3)	0.1
τ in (26)	0.01
Maximum number of communities K	3
Size of the reliability matrices $R_1 \times R_2$	6×3
Sizes of 2 hidden layers of the reliability encoder	128, and 32
Sizes of 2 hidden layers of the event encoder	128, and 32
Sizes of 3 hidden layers of the decoder	16, 32, and 128
Learning rate of ADAMS optimizer	0.001
Step size at i -th iteration	0.1
κ in (27)	IMDB and SP: 0.9; CF:1.0

4.3.1 Impact of reliability matrix size $R_1 \times R_2$

ART uses reliability matrices to represent the reliabilities of agents. In this experiment, we study impact of 4 different reliability matrix sizes on the truth discovery accuracy. The results show that the accuracy of ART is related to the number of elements (i.e., the product of R_1 and R_2) in the reliability matrices. We observe that reliability matrices of size 2×9 and 3×6 achieve the same accuracies on all three real datasets. However, we also observe that different sizes of the matrix all achieve comparable results.

TABLE 3
Truth discovery accuracy of ART with different reliability matrix sizes.

$R_1 \times R_2$	SP	IMDB	CF
2×9	0.918	0.771	0.903
6×3	0.918	0.771	0.903
3×3	0.914	0.767	0.890
6×6	0.915	0.763	0.897

4.3.2 Impact of number of communities K

In this experiment, we show the impact of the number of communities K on the truth discovery accuracy of ART. From Table 4, we observe that ART achieves relatively good accuracies on all three datasets with different number of communities.

4.3.3 Impact of variance $\mathbf{V}_k$

In ART, we use $\mathbf{V}_k$ to denote the variance of the prior distribution of $\tilde{\mathbf{C}}_k$. From Table 5, we observe that different

TABLE 4
Truth discovery accuracy of ART with different number K of communities.

k	SP	IMDB	CF
3	0.918	0.771	0.903
6	0.915	0.759	0.9
9	0.915	0.788	0.89

$\mathbf{V}_k$ have larger influence on CF than the other two datasets. ART achieves the best accuracies on all the three datasets when each element of $\mathbf{V}_k$ is 0.1.

TABLE 5
Truth discovery accuracy with different $\mathbf{V}_k$.

Each element of $\mathbf{V}_k$	SP	IMDB	CF
0.01	0.916	0.771	0.867
0.1	0.918	0.771	0.903
1	0.917	0.771	0.877

4.4 Comparison with Benchmarks

We evaluate and compare the performance of ART and the benchmark methods using five metrics: accuracy, area under the ROC curve (AUC), precision, recall, and the F1 score. The top performer and those close to it (i.e, within 0.002 of the top) are highlighted in boldface in the subsequent tables. The following results indicate that ART is competitive, and often superior, compared to the other benchmark methods. However, the computation time of ART is larger than most of the other benchmark methods (except VISIT) since we include the agent network information in our inference. Our deep autoencoder is also more computationally expensive to learn.

4.4.1 Accuracy

The accuracies of ART and the benchmark methods are shown in Table 6. It can be observed that ART achieves the best accuracies on all the three datasets.

4.4.2 Area under the ROC Curve (AUC)

The AUC scores of ART and the benchmark methods are shown in Table 7. For the multi-class case (i.e., the CF dataset), AUC scores are calculated for each label, and their average is taken. It can be observed that ART achieves the best AUC score on the SP dataset and is among one of the best methods on the CF dataset. On the IMDB dataset, VISIT and BCC achieve the best AUC scores.

4.4.3 Precision

The precision scores of ART and the benchmark methods are shown in Table 8. For the multi-class case (i.e., the CF dataset), precisions are calculated for each label, and their average, weighted by the number of true instances for each label, is taken. It can be observed that ART achieves

TABLE 6
Truth discovery accuracy.

Method	SP	IMDB	CF
ART	0.918	0.771	0.903
VISIT	0.916	0.751	0.893
CBCC	0.916	0.714	0.893
BCC	0.915	0.678	0.890
MV	0.885	0.710	0.867
TruthFinder	0.885	0.714	0.880
AccuSim	0.890	0.702	0.730
GTM	0.890	0.710	0.743
CRH	0.894	0.706	0.747
CATD	0.873	0.657	0.807
KDEm	0.890	0.702	0.877

TABLE 7
Truth discovery AUC.

Method	SP	IMDB	CF
ART	0.958	0.844	0.948
VISIT	0.958	0.871	0.913
CBCC	0.957	0.859	0.949
BCC	0.957	0.872	0.950
MV	0.885	0.759	0.865
TruthFinder	0.885	0.765	0.873
AccuSim	0.890	0.753	0.754
GTM	0.890	0.759	0.763
CRH	0.894	0.756	0.764
CATD	0.873	0.708	0.824
KDEm	0.890	0.753	0.861

TABLE 8
Truth discovery precision.

Method	SP	IMDB	CF
ART	0.921	0.639	0.878
VISIT	0.922	0.607	0.896
CBCC	0.922	0.568	0.902
BCC	0.92	0.535	0.865
MV	0.868	0.563	0.873
TruthFinder	0.868	0.566	0.885
AccuSim	0.886	0.556	0.748
GTM	0.874	0.563	0.762
CRH	0.884	0.559	0.758
CATD	0.863	0.519	0.824
KDEm	0.878	0.556	0.879

TABLE 9
Truth discovery recall.

Method	SP	IMDB	CF
ART	0.913	0.867	0.903
VISIT	0.910	0.911	0.893
CBCC	0.910	0.933	0.893
BCC	0.909	0.944	0.890
MV	0.907	0.944	0.867
TruthFinder	0.907	0.956	0.880
AccuSim	0.894	0.944	0.730
GTM	0.911	0.944	0.743
CRH	0.907	0.944	0.747
CATD	0.887	0.900	0.807
KDEm	0.906	0.944	0.877

the best precision on the IMDB dataset and is one of the best methods on the SP dataset. On the CF dataset, CBCC achieves the best precision.

4.4.4 Recall

The recall scores of ART and the benchmark methods are shown in Table 9. For the multi-class case (i.e., the CF dataset), recall scores are calculated for each label, and then their average, weighted by the number of true instances for each label, is taken. It can be observed that ART achieves the best recall on two datasets: the SP and CF datasets. On the IMDB dataset, TruthFinder achieves the best recall.

4.4.5 F1 score

The F1 scores of ART and the benchmark methods are shown in Table 10. For the multi-class case (i.e., the CF dataset), F1 scores are calculated for each label, and their

average, weighted by the number of true instances for each label, is taken. This averaging takes label imbalance into consideration and can result in an F1 score that is not between the corresponding precision and recall. It can be observed that ART achieves the best F1 scores on two datasets: the IMDB and SP datasets. On the CF dataset, ART is slightly worse than VISIT, which achieves the best F1 score.

4.4.6 Average execution time

We tabulate the average execution time of ART and the benchmark methods in Table 11. All the methods are run on the same laptop with a core-i5-8250U CPU and 16G RAM.

The execution time of ART is over 100 times larger than most of the benchmark methods since ART includes the agent network information and the deep autoencoder is computationally expensive to learn. The SP dataset contains

TABLE 10
Truth discovery F1 score.

Method	SP	IMDB	CF
ART	0.917	0.736	0.890
VISIT	0.916	0.729	0.894
CBCC	0.916	0.706	0.884
BCC	0.914	0.683	0.877
MV	0.887	0.705	0.868
TruthFinder	0.887	0.711	0.882
AccuSim	0.890	0.700	0.712
GTM	0.892	0.705	0.731
CRH	0.895	0.702	0.730
CATD	0.875	0.659	0.814
KDEm	0.891	0.700	0.876

the largest number of evaluations and the CF dataset contains the largest number of agents and thus experiments on these two datasets take more time. Another benchmark that uses the agent network information is VISIT, which however has longer average execution times than ART due to the complex variational inference approach used in its non-conjugate observation model. ART is therefore more suitable for offline truth discovery applications, where inference performance is more important than latency.

TABLE 11
Average execution time.

Method	SP	IMDB	CF
ART	1261.1	92.0	118.8
VISIT	1413.6	108.5	562.1
CBCC	4.0	2.4	4.2
BCC	5.1	1.7	1.9
MV	1.0	0.1	0.1
TruthFinder	0.4	0.1	0.1
AccuSim	41.9	3.2	0.8
GTM	1.0	0.1	0.1
CRH	1.5	0.1	0.1
CATD	0.5	0.1	0.1
KDEm	0.9	0.1	0.1

5 CONCLUSION

In this paper, we have combined the strength of autoencoders in learning nonlinear relationships and the strength of Bayesian networks in characterizing hidden interpretable structures to tackle the truth discovery problem. The

Bayesian network model introduces constraints to the autoencoder and at the same time incorporates the community information of the social network into the autoencoder. We developed a variational inference method to estimate the parameters in the autoencoder and infer the hidden variables in the Bayesian network. Results on real datasets demonstrate the competitiveness of our proposed method over other state-of-the-art benchmark methods.

In this paper, we have not considered correlations between events in our inference. We have also not incorporated any side information or prior information about the events into our procedure. These are interesting future research directions, which may further improve the truth discovery accuracy.

REFERENCES

- [1] D. R. Karger, S. Oh, and D. Shah, "Efficient crowdsourcing for multi-class labeling," *ACM SIGMETRICS Performance Evaluation Review*, vol. 41, no. 1, pp. 81–92, 2013.
- [2] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, Mar. 2011.
- [3] J. Ho, W. P. Tay, T. Q. S. Quek, and E. K. P. Chong, "Robust decentralized detection and social learning in tandem networks," *IEEE Trans. Signal Process.*, vol. 63, no. 19, pp. 5019 – 5032, Oct. 2015.
- [4] W. P. Tay, "Whose opinion to follow in multihypothesis social learning? A large deviations perspective," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 2, pp. 344 – 359, Mar. 2015.
- [5] M. L. Gray, S. Suri, S. S. Ali, and D. Kulkarni, "The crowd is a collaborative network," in *ACM Conf. Computer-Supported Cooperative Work & Social Computing*, 2016, pp. 134–147.
- [6] C. Huang and D. Wang, "Topic-aware social sensing with arbitrary source dependency graphs," in *ACM/IEEE Int. Conf. Inform. Process. Sensor Networks*, 2016, pp. 1–12.
- [7] Q. Kang and W. P. Tay, "Sequential multi-class labeling in crowdsourcing," *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 11, pp. 2190–2199, 2019.
- [8] L. Berti-Equille and J. Borge-Holthoefer, "Veracity of data: From truth discovery computation algorithms to models of misinformation dynamics," *Synthesis Lectures on Data Management*, vol. 7, no. 3, pp. 1–155, 2015.
- [9] Y. Li, J. Gao, C. Meng, Q. Li, L. Su, B. Zhao, W. Fan, and J. Han, "A survey on truth discovery," *ACM Sigkdd Explorations Newslett.*, vol. 17, no. 2, pp. 1–16, 2016.
- [10] P. Cheng, X. Lian, L. Chen, J. Han, and J. Zhao, "Task assignment on multi-skill oriented spatial crowdsourcing," *IEEE Trans. Knowledge and Data Engineering*, vol. 28, no. 8, pp. 2201–2215, 2016.
- [11] Y. Tong, L. Chen, Z. Zhou, H. V. Jagadish, L. Shou, and W. Lv, "Slade: A smart large-scale task decomposer in crowdsourcing," *IEEE Trans. Knowl. and Data Eng.*, vol. 30, no. 8, pp. 1588–1601, 2018.
- [12] T. Tian, Y. Zhou, and J. Zhu, "Selective verification strategy for learning from crowds," in *AAAI Conf. Artificial Intelligence*, 2018.
- [13] M. Alsayasneh, S. Amer-Yahia, E. Gaussier, V. Leroy, J. Pilourdault, R. M. Borromeo, M. Toyama, and J. Renders, "Personalized and diverse task composition in crowdsourcing," *IEEE Trans. Knowl. and Data Eng.*, vol. 30, no. 1, pp. 128–141, 2017.
- [14] Y. Wang, F. Ma, Z. Jin, Y. Yuan, G. Xun, K. Jha, L. Su, and J. Gao, "Eann: Event adversarial neural networks for multi-modal fake news detection," in *Proceedings of the 24th acm sigkdd international conference on knowledge discovery & data mining*, 2018, pp. 849–857.
- [15] D. Wang, L. Kaplan, H. Le, and T. Abdelzaher, "On truth discovery in social sensing: A maximum likelihood estimation approach," in *ACM/IEEE Int. Conf. Inform. Process. Sensor Networks*, 2012, pp. 233–244.
- [16] B. Zhao, B. I. P. Rubinstein, J. Gemmell, and J. Han, "A Bayesian approach to discovering truth from conflicting sources for data integration," *Proc. VLDB Endowment*, vol. 5, no. 6, pp. 550–561, 2012.

- [17] D. Wang, M. T. Amin, S. Li, T. Abdelzaher, L. Kaplan, S. Gu, C. Pan, H. Liu, C. C. Aggarwal, and R. Ganti, "Using humans as sensors: an estimation-theoretic perspective," in *ACM/IEEE Int. Conf. Inform. Process. Sensor Networks*, 2014, pp. 35–46.
- [18] S. Yao, S. Hu, S. Li, Y. Zhao, L. Su, L. Kaplan, A. Yener, and T. Abdelzaher, "On source dependency models for reliable social sensing: Algorithms and fundamental error bounds," in *IEEE Int. Conf. Distributed Computing Syst.*, 2016, pp. 467–476.
- [19] X. Yin, J. Han, and S. Y. Philip, "Truth discovery with multiple conflicting information providers on the web," *IEEE Trans. Knowl. Data Eng.*, vol. 20, no. 6, pp. 796–808, 2008.
- [20] X. L. Dong, L. Berti-Equille, and D. Srivastava, "Integrating conflicting data: the role of source dependence," *Proc. VLDB Endowment*, vol. 2, no. 1, pp. 550–561, 2009.
- [21] Q. Li, Y. Li, J. Gao, L. Su, B. Zhao, M. Demirbas, W. Fan, and J. Han, "A confidence-aware approach for truth discovery on long-tail data," *Proc. VLDB Endowment*, vol. 8, no. 4, pp. 425–436, 2015.
- [22] D. Y. Zhang, D. Wang, and Y. Zhang, "Constraint-aware dynamic truth discovery in big data social media sensing," in *IEEE Int. Conf. Big Data*, 2017, pp. 57–66.
- [23] X. Zhang, Y. Wu, L. Huang, H. Ji, and G. Cao, "Expertise-aware truth analysis and task allocation in mobile crowdsourcing," in *IEEE Int. Conf. Distributed Computing Syst.*, 2017, pp. 922–932.
- [24] Y. Li, B. I. P. Rubinstein, and T. Cohn, "Truth inference at scale: A Bayesian model for adjudicating highly redundant crowd annotations," in *Int. Conf. World Wide Web*, 2019.
- [25] K. Broelemann, T. Gottron, and G. Kasneci, "Ltd-rbm: Robust and fast latent truth discovery using restricted boltzmann machines," in *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, 2017, pp. 143–146.
- [26] M. Wan, X. Chen, L. Kaplan, J. Han, J. Gao, and B. Zhao, "From truth discovery to trustworthy opinion discovery: An uncertainty-aware quantitative modeling approach," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1885–1894.
- [27] H. C. Kim and Z. Ghahramani, "Bayesian classifier combination," in *Artificial Intell. and Stat.*, 2012, pp. 619–627.
- [28] Y. Zheng, G. Li, Y. Li, C. Shan, and R. Cheng, "Truth inference in crowdsourcing: is the problem solved?" *Proc. VLDB Endowment*, vol. 10, no. 5, pp. 541–552, 2017.
- [29] M. Venanzi, J. Guiver, G. Kazai, P. Kohli, and M. Shokouhi, "Community-based Bayesian aggregation models for crowdsourcing," in *Int. Conf. World Wide Web*, 2014, pp. 155–164.
- [30] S. Lyu, W. Ouyang, Y. Wang, H. Shen, and X. Cheng, "Truth discovery by claim and source embedding," *IEEE Trans. Knowl. Data Eng.*, 2019.
- [31] Y. Li, B. Rubinstein, and T. Cohn, "Exploiting worker correlation for label aggregation in crowdsourcing," in *Int. Conf. Machine Learning*, 2019, pp. 3886–3895.
- [32] M. Yin, M. L. Gray, S. Suri, and J. W. Vaughan, "The communication network within the crowd," in *Int. Conf. World Wide Web*, 2016, pp. 1293–1303.
- [33] L. Ma, W. P. Tay, and G. Xiao, "Iterative expectation maximization for reliable social sensing with information flows," *Information Sciences*, 2018, accepted.
- [34] J. Yang, J. Wang, and W. P. Tay, "Using social network information in community-based Bayesian truth discovery," *IEEE Trans. Signal Inf. Process. Netw.*, 2019, accepted.
- [35] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [36] J. L. Elman, "Finding structure in time," *COGNITIVE SCIENCE*, vol. 14, no. 2, pp. 179–211, 1990.
- [37] J. Marshall, A. Argueta, and D. Wang, "A neural network approach for truth discovery in social sensing," in *2017 IEEE 14th international conference on mobile Ad Hoc and sensor systems (MASS)*. IEEE, 2017, pp. 343–347.
- [38] L. Li, B. Qin, W. Ren, and T. Liu, "Truth discovery with memory network," *Tsinghua Science and Technology*, vol. 22, no. 6, pp. 609–618, 2017.
- [39] M. Johnson, D. K. Duvenaud, A. Wiltchko, R. P. Adams, and S. R. Datta, "Composing graphical models with neural networks for structured representations and fast inference," in *Advances in neural information processing systems*, 2016, pp. 2946–2954.
- [40] M. Dundar, B. Krishnapuram, J. Bi, and R. B. Rao, "Learning classifiers when the training data is not iid," in *Int. Joint Conf. Artificial Intell.*, 2007, pp. 756–761.
- [41] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [42] B. Zong, Q. Song, M. R. Min, W. Cheng, C. Lumezanu, D. Cho, and H. Chen, "Deep autoencoding gaussian mixture model for unsupervised anomaly detection," in *Int. Conf. Learning Representations*, 2018.
- [43] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing, "Mixed membership stochastic blockmodels," *J. Mach. Learning Research*, vol. 9, no. Sep, pp. 1981–2014, 2008.
- [44] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Hierarchical Dirichlet processes," *J. Am. Stat. Assoc.*, 2012.
- [45] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *J. machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [46] S. Fortunato and D. Hric, "Community detection in networks: A user guide," *Physics Rep.*, vol. 659, pp. 1–44, 2016.
- [47] Y. Bengio, L. Yao, G. Alain, and P. Vincent, "Generalized denoising auto-encoders as generative models," in *Advances in Neural Information Processing Systems*, 2013, pp. 899–907.
- [48] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [49] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, "Stochastic variational inference," *J. Mach. Learning Research*, vol. 14, no. 1, pp. 1303–1347, 2013.
- [50] D. M. Blei and M. I. Jordan, "Variational inference for dirichlet process mixtures," *Bayesian analysis*, vol. 1, no. 1, pp. 121–143, 2006.
- [51] D. P. Kingma, "Variational inference & deep learning: A new synthesis," Ph.D. dissertation, University of Amsterdam, 2017.
- [52] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with Gumbel-softmax," *arXiv preprint arXiv:1611.01144*, 2016.
- [53] T. Minka, *Bayesian inference, entropy, and the multinomial distribution*, 2003. [Online]. Available: <https://tminka.github.io/papers/minka-multinomial.pdf>
- [54] B. Zhao and J. Han, "A probabilistic model for estimating real-valued truth from conflicting sources," *Proc. of QDB*, 2012.
- [55] Q. Li, Y. Li, J. Gao, B. Zhao, W. Fan, and J. Han, "Resolving conflicts in heterogeneous data by truth discovery and source reliability estimation," in *Proc. ACM SIGMOD int. conf. Management of data*, 2014, pp. 1187–1198.
- [56] J. Yang and W. P. Tay. (2018) Using social network information to discover truth of movie ranking. [Online]. Available: <https://doi.org/10.21979/N9/L5TTRW>
- [57] M. Venanzi, O. Parson, A. Rogers, and N. Jennings, "The ActiveCrowdToolkit: An open-source tool for benchmarking ActiveLearning algorithms for crowdsourcing research," in *3rd Human Computation and Crowdsourcing Conference (HCOMP)*, 2015. [Online]. Available: <https://github.com/orchidproject/active-crowd-toolkit/tree/master/Data>
- [58] S. Fortunato, "Community detection in graphs," *Physics Rep.*, vol. 486, no. 3, pp. 75–174, 2010.
- [59] P. K. Gopalan and D. M. Blei, "Efficient discovery of overlapping communities in massive networks," *Proc. Nat. Academy Sci.*, vol. 110, no. 36, pp. 14534–14539, 2013.