Data-Driven Power Outage Detection by Social Sensors

Haifeng Sun, Zhaoyu Wang, *Member, IEEE*, Jianhui Wang, *Senior Member, IEEE*, Zhen Huang, *Student Member, IEEE*, NichelleLe Carrington, and Jianxin Liao

***Abstract*—This paper proposes a novel method to detect and locate power outages based on the information collected from social media. Twitter is used as a real-time social sensor in the proposed method. To solve the challenges of detecting a targeted event from the fragmented and noisy tweets, we devise a probabilistic framework to integrate the textual, temporal, and spatial information to identify the event. To improve the accuracy of outage detection, we propose a supervised topic model with a heterogeneous information network. The proposed technique is tested with real tweets and outage cases. The numerical results demonstrate the effectiveness of the proposed methodology. The comparison between the proposed method, and support vector machine and statistics Bayesian method shows the accuracy of the developed model.**

***Index Terms*—Power outage detection, social media, Twitter, heterogenous information network, supervised topic model.**

I. Introduction

THE MASSIVE amounts of data collected from wide-area monitoring systems, advanced metering infrastructure (AMI), and social media provide opportunities to develop a data-enabled modernized power system. While the utilities can better understand system conditions and customer behaviors, plan and control power grids, and design energy efficient programs, they are facing challenges in managing and taking advantage of big data.

Reliability and customer satisfaction continue to be among the top concerns of utilities. The detection and location of outages is one of the major functions of a smart distribution management system. An efficient outage management system (OMS) can dramatically decrease the durations and sizes of power outages. Traditionally, a utility identifies power outage areas through trouble calls from customers. The assumption is that a customer will use a landline telephone linked to a physical address to report outages. Nowadays, most customers use cell phones and may not make a prompt outage call to utilities. With the installation of smart meters, the AMI-enabled outage management can rapidly detect power outages without relying on customer phone calls. Many studies have been made in literature on detecting power outages by smart meters. The study in [1] proposed a fuzzy petri nets-based method to detect fraudulent consumption and power outages using AMI data. The study in [2] proposed a knowledge-based system to locate distribution system outages using data from customer trouble calls and wireless automated meter reading systems. Reference [3] developed a tree-based on-demand method to obtain information about the system status without polling all smart meters. Reference [4] proposed probabilistic and fuzzy model-based filter algorithms to process outage data from automated metering system. The study in [5] optimally deployed power flow sensors and used them for outage detection based on the hypothesis testing. Therefore, smart meters have the potential to improve the detection and location of outages. However, smart meter-based detection requires a reliable and complicated communication system. Moreover, the wide deployment of AMI also requires a lot of time and financial investments, which may not be affordable to all utilities. Hence, there is a need to find additional information sources for OMS.

In this information age, customers are already actively engaged in social media such as Facebook and Twitter. The ever-growing digital media network can be considered as social sensors. The data from social sensors can be used to identify the location and extent of an outage without adding new measurement and communication instruments. Social sensors have been used in detecting events such as nature disasters and severe weather events [6], [7]. The integration of social media to power system analysis and operation is a relatively new topic. The usefulness of social media in outage detection has been recognized by the power industry [8]. In this paper, we present a probabilistic framework for the detection and location of power outages by using the abundant amount of data collected from Twitter. Twitter can provide the most up-to-date and inclusive stream of information and commentary on present events [7]. When an outage occurs, customers may discuss the troubles in tweets. The large number of users and its real-time nature makes Twitter an ideal data source for event extraction. Certain challenges may arise with identifying

Manuscript received August 28, 2015; revised December 28, 2015 and March 14, 2016; accepted March 14, 2016. Date of publication March 30, 2016; date of current version August 19, 2016. This work was supported by the U.S. Department of Energy Office of Electricity Delivery and Energy Reliability. Paper no. TSG-01015-2015.

H. Sun and J. Liao are with the Institute of Network Technology, Beijing University of Posts and Telecommunications, Beijing 100876, China (e-mail: hfsun@bupt.edu.cn; jxlbupt@gmail.com).

Z. Wang and N. Carrington are with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA (e-mail: wzy@iastate.edu; nkcarrin@iastate.edu).

J. Wang is with Argonne National Laboratory, Argonne, IL 60439 USA (e-mail: jianhui.wang@anl.gov).

Z. Huang is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: zhenhuangee@gatech.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSG.2016.2546181

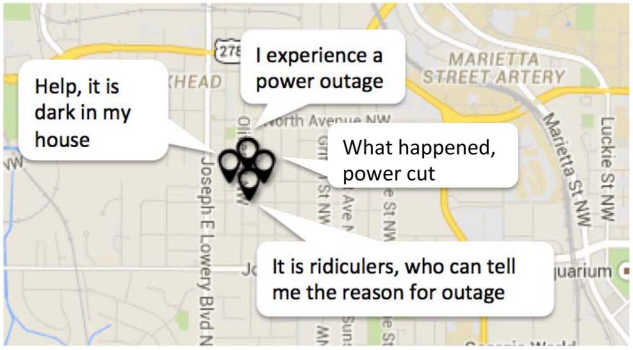

Fig. 1. Examples of tweets posted after a power outage.

outage areas based on tweets, which can be summarized as follows: individual tweets are short and lacking in sufficient context to categorize them into topics of interest; Twitter users' commentaries include a variety of topics, unclear in advance which event types are appropriate to be used as an initial estimation of topics; and tweets are written in an informal style, which cannot be directly analyzed by using language processing tools designed for edited texts.

Every Twitter user is an active observer who is constantly monitoring the world. As shown in Fig. 1, when a power outage happens, Twitter users post tweets perchance about the event. The following examples demonstrate representative tweets about power outages: (1) a tweet with words and addresses that are related to an outage, e.g., "Power outages occurred in Southwest Bend. Tweet me if you're affected;" (2) a tweet with words that are related to power outages, but without an address, e.g., "Dispatch says residents are experiencing massive power outages for unknown reasons;" (3) a tweet that does not explicitly indicate a power outage, but is actually related to an outage, e.g., "Why all lights don't work;" (4) a tweet with words that are related to an outage, but is actually talking about something else, e.g., "North Korea Internet outage could be a response to Sony Hack;" and (5) a tweet with words that are related to a power outage, but it is not a real-time report of the event, e.g., "An outage happened yesterday."

A large number of tweets on power outages within a short time in a certain area indicates a high probability of a power outage in the corresponding area. Therefore, we propose a method based on the supervised Latent Dirichlet Allocation (sLDA) to detect and locate power outage events. A maximum-likelihood method is developed for parameter estimation. Since a tweet has a 140-character limit, the topic in a tweet may not be clear and definite. Traditional topic models cannot be applied to analyze tweets. To solve this challenge, we investigate the event detection problem by a supervised topic modeling with text-rich heterogeneous information network. The proposed Twitter-based outage detection can be integrated with existing OMS to provide better identification of outages.

Compared to existing work, the contributions of this paper are summarized as follows:

1. We propose a probabilistic framework to analyze large-scale real-time tweets to detect power outages considering the spatial-temporal characteristics of tweets.

2. We develop a hybrid model that integrates a heterogeneous information network and a supervised topic model to improve the accuracy of the detection by reorganizing the temporal, spatial and textual information.

The remainder of this paper is organized as follows. Section II reviews the existing work on social media-based event detection. Section III presents the analysis of the problem and the proposed framework. Section IV discusses the learning model. Section V proposes the detection model. In Section VI, the numerical results are provided. Section VII concludes the paper with the major findings.

II. Related Work

The information extraction from Twitter has been studied by many papers. Huberman *et al.* [9] studied more than 300,000 Twitter users and found that relationships among Twitter friends play a key role in understanding users' interactions. Boyd *et al.* [10] investigated the re-tweeting activity, in which a user forwarded messages that were originally posted by others. The event detection by Twitter has attracted interests in academia in recent years [6], [11], [12].

Bauman *et al.* [13] proposed a method that used Twitter for the detection of power outages. They applied key words searching to collect power outage-related tweets, and used Kleinbergs method [14] to detect bursts within tweets. The method provided a good solution. However, probabilistic semantic analysis and spatial/temporal characteristics of outages can be used to improve the detection performance [15].

MacEachren *et al.* [16] designed a Web-based mapping tool to classify tweets and extract useful information. The tool could be used to analyze a large number of tweets and extract potentially useful information. Starbird and Stamberger [17] introduced a tweet syntax that included parsable keywords, which facilitated the extraction of detailed information in a tweet (e.g., victims' needs and damage state descriptions in tweets).

Guy *et al.* [18] developed a prototype system that collected, collated, and analyzed tweets to extract useful information after an earthquake. The system could identify tweets that are related to the key word "*earthquake*" and its equivalents in other languages. Sakaki *et al.* [6] proposed an algorithm to monitor earthquake-related tweets in a real-time manner. In the algorithm, a Twitter user was considered as a social sensor, and Kalman filtering and particle filtering algorithms were used for location estimation. However, power outages are different from earthquakes, since the former only influences a relatively small region and cannot be considered as a social topic. So we cannot simply use these methods to monitor power outages.

Zhao and Mitra [19] defined an event as a set of relationships among social users on a specific topic over a certain time period. Events could be extracted from social text streams by three steps: text-based clustering, temporal segmentation, and graph cuts of social networks. Tweets were searched and classified using a support vector machine (SVM). A target event was detected by a temporal probability model. Locations of events were estimated by a Bayesian filter. However, this method was specifically designed for the location estimation of earthquakes and cannot be easily generalized to detect other events.

A few papers have studied the real-time event information extraction from tweets. Zhao *et al.* [20] detected sub-events occurring in NFL games by tracking the increase of tweeting activities. In [7], the authors introduced a tool named as TWICAL, which was an open-source event-extraction and categorization system for Twitter. The Twitter Natural Language Processing (NLP) tools were used to extract special words from tweets. Then TWICAL classified events based on latent variable models.

III. Problem Analysis

In this paper, the target event is a power outage. An event is an arbitrary classification in a spatial/temporal region. It has active participating agents, passive factors, products, and locations in both the spatial and temporal domains.

A. Semantic Analysis of Tweets

Compared to existing methods, which are only based on keyword searching, we identify tweets that are related to our topics of interest [21]. Topic modeling algorithms are statistical methods that analyze the words in the original texts to discover topics. For example, if we have two topics as $z_1 = \{power, outage\}$ and $z_2 = \{Africa, SouthAsia\}$. A topic-word matrix β can be used to represent these two topics as:

	power	outage	Africa	South Asia
z_1	1	1	0	0
z_2	0	0	1	1

The sentence "A power outage is affecting much of Africa and South Asia" can be parsed into a word-frequency vector $\boldsymbol{D}$ as:

$$D = \begin{bmatrix} 1/4, & \text{(power)} \\ 1/4, & \text{(outage)} \\ 1/4, & \text{(Africa)} \\ 1/4, & \text{(SouthAsia)} \end{bmatrix}. \tag{1}$$

In (1), we have removed stop words [23] in the sentence, i.e., words without useful information. Each element in $\boldsymbol{D}$ represents the proportion of the corresponding word in a tweet. The sentence can be annotated as "A *power*| z_1 *outage*| z_1 is affecting much of *Africa*| z_2 and *South Asia*| z_2." So the topic proportion of this sentence is $\boldsymbol{z_{1:2}} = \{0.5, 0.5\}$. Using topics instead of key words in analyzing tweets can make machines better understand the textual information.

In topic models, an unobserved topic variable $z_k \in \{z_1, \ldots, z_K\}$ is associated with the occurrence of a word $w_i \in \{w_1, \ldots, w_M\}$ in a particular tweet d. The topics and words can compose a matrix β with a size of $K \times M$. Given a word-frequency vector $\boldsymbol{D}$ of a tweet, the topic proportions can be calculated by extracting words in the tweet as:

$$\boldsymbol{z_{1:K}} = \beta \boldsymbol{D} \tag{2}$$

where $\boldsymbol{z_{1:K}}$ is a K-length vector that represents the proportions of topics from z_1 to z_K in the tweet. It has been shown that tweets with certain contents that we are interested in have similar proportions of topics [22]. We use y to label a tweet and indicate whether the tweet is related to power outage events:

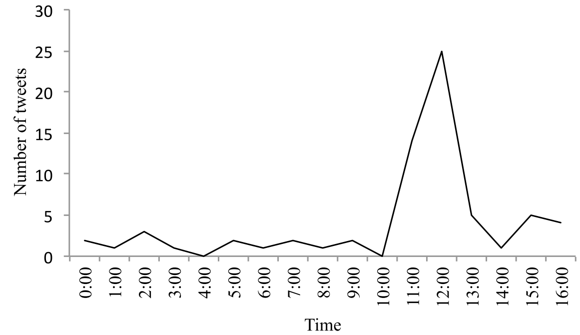

Fig. 2. Number of tweets related to power outages.

$$y = \begin{cases} 1, & \text{if the tweet is related to power outage;} \\ 0, & \text{other tweets.} \end{cases} \tag{3}$$

The probability of a tweet that is related to power outages can be defined as $p(y = 1 \mid d)$.

However, the information in tweets is highly complicated as shown in the introduction section. To achieve a better detection performance, we need both temporal and spatial information that is included in tweets when they are being posted. When a power outage happens within a short time of a closed area, a lot of people being affected may discuss the event on Twitter. To integrate the multi-type information, including time, location, and texts, we use the heterogeneous information network theory.

B. Temporal Analysis of Tweets

Each tweet is stamped with a posting time. For example, Fig. 2 presents the number of tweets related to a power outage in one day, which is manually counted. It can be seen that a spike occurs from *10:00* to *14:00*, which corresponds to an actual power outage event occurred at Southwest Bend in Oregon. Therefore, if the topics corresponding to power outages are highly cited in a short time, the probability of an actual power outage is high. To evaluate whether a tweet is really reporting a power outage, the proportion of tweets occurred in the period from $t - \delta t$ to t is represented as NT_{total}, and the number of tweets related to the power outage event in these tweets is represented as NT_{target}. We use $p(y = 1 \mid t)$ defined in (4) to represent the probability that outage-related tweets are posted. During major events other than power outages, the Twitter activity may be particularly high, which could result in a lower sensitivity of outage detection.

$$p(y = 1|t) = \frac{NT_{target}}{NT_{total}}. \tag{4}$$

C. Spatial Analysis of Tweets

Unlike the temporal analysis, tweets are not always stamped with their posting locations. For tweets with location stamps, we can directly find the locations of the events. For those without location stamps, the spatial information may be extracted from the contents. During the time period from $t - \delta t$ to t,

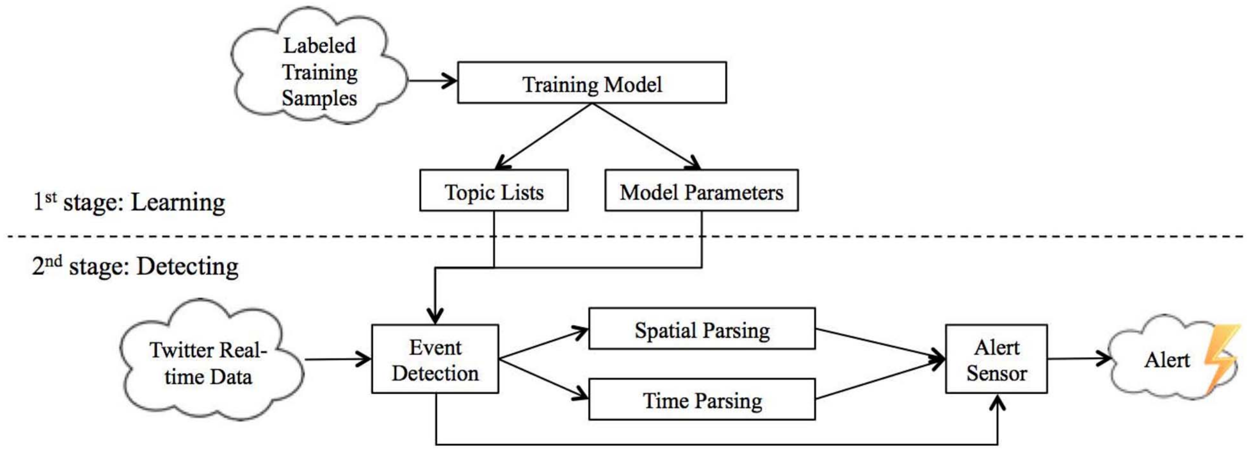

Fig. 3. Proposed framework for two-stage power outage event detection.

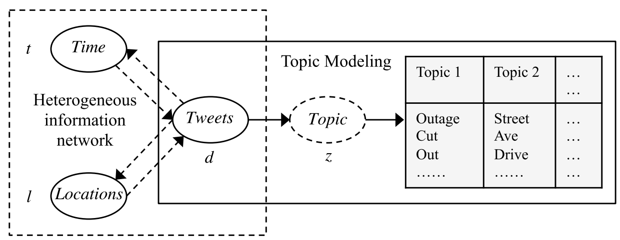

Fig. 4. Topic modeling with heterogeneous information network.

the number of tweets posted around the current tweet location l is represented as NF_{total}, and the number of tweets that are related to the target event in these tweets is represented as NL_{target}. We use $p(y = 1 \mid l)$ to represent the probability that outage-related tweets are posted:

$$p(y = 1|l) = \frac{NL_{target}}{NL_{total}}. \tag{5}$$

D. Proposed Framework

Fig. 3 shows the proposed two-stage framework. The first stage is the training process. In particular, tweets that report power outages are labeled as 1, and other tweets are labeled as 0. The labeled data is used by the supervised topic model to learn the latent parameters and the topic-word matrix. The second stage is to detect the power outage events based on the latent parameters and topic-word matrix obtained in the first stage. We perform the temporal and spatial analysis to detect tweets. Based on the temporal and spatial information, the alert sensor decides whether an outage alert should be issued. Details on the alert sensor are discussed in Section V.

IV. Learning Model

As shown in Fig. 4, the temporal and spatial information of a tweet can be used to improve the performance of the topic modeling. We consider tweets and associated posting time and locations as elements of an interconnected graph, which is named as a heterogeneous information network. Based on the temporal and spatial properties of topics, the constructed heterogeneous information network can be used for modeling and classifying the topics (e.g., the outage-related words are classified as Topic 1 and the address-related words as Topic 2).

A. Supervised Topic Modeling

In topic models, the words of a tweet can be classified into a set of latent topics, i.e., a set of unknown distributions over the vocabulary [23]. Each tweet uses a mix of topics that are unique to itself. In supervised topic modeling, every tweet is associated with a label, as shown in equation (3). The intention is to find latent topics that predict the labels for future unlabeled tweets.

Given a tweet, we can calculate its topic distribution by equation (2). In this equation, we need to calculate the unknown topic-word matrix β. A variety of algorithms have been used to estimate the parameters of topic models, e.g., the basic expectation-maximization [24], the approximation inference methods [23], the expectation propagation [25], and the Gibbs sampling [26]. This paper uses the Gibbs sampling since it estimates parameters from Dirichlet priors and allows combining estimates from several local maxima of the posterior distribution [27].

The LDA model has three sets of unknown parameters: the tweet distribution θ, the topic distribution ϕ, and the latent variables corresponding to the assignments of words in a topic z [23]. A Markov chain that converges to the posterior distribution on z can be constructed by applying Gibbs sampling. The Markov chain is then used to calculate θ and ϕ. The transition between two successive states of the Markov chain can be calculated by using the standard Dirichlet integrals:

$$p\big(z_j|w_i = m, z_{-i}, w_{-i}\big) \propto \frac{C^{WT}_{mj} + b}{\sum_{m' \in w_{-i}} C^{WT}_{m'j} + Vb} \frac{C^{DT}_{dj} + a}{\sum_{j'} C^{DT}_{dj'} + Ta} \tag{6}$$

where z_j represents the assignments of the i-th word in a tweet to topic j, $w_i = m$ represents the observation that the i-th word in the tweet is the m-th word in the lexicon, z_{-i} represents all topic assignments that do not include the i-th word, w_{-i} represents all the words in the training set tweets that do not include the i-th word, C^{WT}_{mj} is the number of times that a word

m is assigned to a topic j without counting the present instance, m' represents the words in w_{-i}; C_{dj}^{DT} is the number of times that a topic j has occurred in a tweet d without counting the current instance, j' represents all the topics, and a and b are hyper-parameters that are set by users [23]. Since tweets are short, the topics in tweets are more sensitive to words. So we set the parameters to be $a = 0.3$ and $b = 0.1$ to improve the sensitivity. For a sample from this Markov chain, we can assign every word to a topic and estimate θ and ϕ as follows:

$$\phi_{mj} = \frac{C_{mj}^{WT} + b}{\sum_{m' \in w_{-i}} C_{m'j}^{WT} + Vb} \tag{7}$$

$$\theta_{dj} = \frac{C_{dj}^{DT} + a}{\sum_{j'} C_{dj'}^{DT} + Ta} \tag{8}$$

where ϕ_{mj} is the probability of using a word m in a topic j, and θ_{dj} is the probability of the topic j in a tweet d. In this way, we can assign all words in a training set to topics and obtain the topic-word matrix β.

We use a normal distribution with the mean vector $\boldsymbol{\eta}$ and the variance matrix σ^2 to label a tweet as:

$$p\left(y = 1 | \boldsymbol{z_{1:K}}, \boldsymbol{\eta}, \sigma^2\right) = N\left(\boldsymbol{z_{1:K}} | \boldsymbol{\eta}, \sigma^2\right) \tag{9}$$

We apply the maximum likelihood estimation to estimate the parameters $\boldsymbol{\eta}$ and σ^2. We collect all of outage-related tweets in the training set to form a sample set to learn the parameters of the normal distribution. Suppose the topic distribution of n-th tweet in the sample set is a vector $\boldsymbol{z^{(n)}}$. The probability function of the n-th tweet related to a power outage is $N(\boldsymbol{z^{(n)}} \mid \boldsymbol{\eta}, \sigma^2)$. Then, the joint probability density function $L(\boldsymbol{\eta}, \sigma^2)$ is:

$$L\left(\boldsymbol{\eta}, \sigma^2\right) = \prod_{i=1}^{n} N\left(\boldsymbol{z^{(n)}} | \boldsymbol{\eta}, \sigma^2\right) \tag{10}$$

In this case, the natural logarithm of the likelihood function is:

$$\log L\left(\boldsymbol{\eta}, \sigma^2\right) = \sum_{n=1}^{N} \log N\left(\boldsymbol{z^{(n)}} | \boldsymbol{\eta}, \sigma^2\right) \tag{11}$$

The parameter $\boldsymbol{\eta}$ can be calculated as [28]:

$$\hat{\boldsymbol{\eta}} = \frac{1}{N} \sum_{i=1}^{N} \boldsymbol{z^{(n)}} \tag{12}$$

The n is the number of samples. Therefore, the parameter σ^2 can be calculated by the maximum likelihood estimator as [28]:

$$\hat{\sigma^2} = \frac{1}{N} \sum_{i=1}^{N} \left(\boldsymbol{z^{(n)}} - \hat{\boldsymbol{\eta}}\right)^2 \tag{13}$$

Algorithm 1 shows the learning steps to develop a supervised topic model based on the Gibbs sampling and the maximum likelihood estimation. Steps 1-13 represent the Gibbs sampling procedure. In step 5, we calculate the number of times that a word m is assigned to a topic j for each word and the number of times that a topic j has occurred in a tweet d. In steps 6 and 7, the topic models parameters

Algorithm 1 sLDA Learning Process

Inputs:
- Training data set: $V = [v_1, v_2, \cdots, v_n]$, $v_i : \{[w_1, w_2, \cdots, w_m], l_i\}$
- Number of topics k
- Iteration number $t = 500$

Outputs:
- Topic-word matrix β
- Normal distribution parameters $\boldsymbol{\eta}$ and σ^2

```
1.  Randomly assign a topic index to each word in tweets
2.  Repeat t times
3.    For each tweet i:
4.      For each word w in tweet:
5.        Calculate C_mj^WT and C_dj^DT
6.        φ_mj = (C_mj^WT + b) / (Σ_{m'∈w_-i} C_m'j^WT + Vb)
7.        θ_dj = (C_dj^DT + a) / (Σ_j' C_dj'^DT + Ta)
8.        p(z_i = j|w_i = m, z_-i, w_-i) ∝ φ_mj θ_dj
9.        Assign word w to topic j:
          j = argmax_j P(z_i = j|w_i = m, z_-i, w_-i)
10.     End For
11.   End For
12. End Repeat
13. Calculate assignment of each word to matrix β
14. Calculate the tweet topic distribution z^(n) = z_1:K
15. η̂ = (1/N) Σ_{i=1}^N z^(n)
16. σ̂² = (1/N) Σ_{i=1}^N (z^(n) − η̂)²
17. Return β, η and σ²
```

ϕ and θ are calculated. In step 8, we calculate the probability for each word assigned to a specific topic. In step 9, we assign each word to the most probable topic. Then, we can obtain a topic-word matrix β. Steps 14-16 represent the maximum likelihood estimation procedure to use the labeled tweets to calculate the normal distribution parameters $\boldsymbol{\eta}$ and σ^2 of outage-related tweets.

B. Heterogeneous Information Network

To improve the performance of the topic modeling by using the temporal and spatial information, we use a heterogeneous information network to learn the topics [29]. It is known that the same topic is more likely to be cited in a similar time and location. Based on equation (2), the topic distribution for a tweet d can be recalculated as:

$$\boldsymbol{z_{1:K}} = \beta \boldsymbol{D} P(\boldsymbol{z_{1:K}} | t, l) \tag{14}$$

where $P(\boldsymbol{z_{1:K}} \mid t, l)$ is the topic distribution of tweets from time $t - \delta t$ to t around the location l. We apply the same method in Section IV-A to obtain the unknown constants $(\boldsymbol{\eta}, \sigma^2)$.

Given the word-frequency vector $\boldsymbol{D}$ of a new tweet, the probability that it is labeled as an outage-related tweet is calculated as $p(y = 1 \mid \boldsymbol{D}, \beta, \boldsymbol{\eta}, \sigma^2)$, which is discussed in next section.

V. Detecting Model

The words in a tweet can be parsed as a word-frequency vector $\boldsymbol{D}$. The topic distribution of a tweet can be calculated using (14). The probability that the topics of the tweet are related with a power outage is:

$$p\Big(y=1|\boldsymbol{D},\beta,\boldsymbol{\eta},\sigma^2\Big)=N\Big(\boldsymbol{z_{1:K}}|\boldsymbol{\eta},\sigma^2\Big) \tag{15}$$

However, even if the topics are related to outages, the intention of this tweet may not be reporting the outage event, as discussed in the introduction section.

Therefore, $p(y=1 \mid t, l)$ represents the probability that the tweet is related to a power outage that occurs between time $t-\delta t$ and t around the location l:

$$p(y=1|t,l)=p(y=1|t)p(y=1|l) \tag{16}$$

Based on the above analysis, the probability that a power outage is happening can be calculated as:

$$\varepsilon=p\Big(y=1|\boldsymbol{D},\beta,\boldsymbol{\eta},\sigma^2\Big)p(y=1|t,l) \tag{17}$$

The outage alert will be issued if ε is larger than a threshold value V_{th}.

VI. Numerical Experiments

The power outage events in a U.S. city are studied in this paper. We use 10,000 labeled tweets to form the training set. The set contains tweets that are relevant and irrelevant to power outage events. We use tweets from January 2014 to detect the power outage events and validate the performance of the proposed model. The test set is completely separated from the training set, and is not used during the learning procedure. The training set includes the first 20 days of tweets and the test set includes the remaining 10 days of tweets. We labeled each tweet in the training set to indicate whether it is related to a power outage.

We compare the proposed method with two popular textual classification methods, statistics Bayesian [30] and support vector machine (SVM) [31]. The statistics Bayesian method considers all words in the training set as features, and calculates the occurrence probability of a word w_i in outage-related tweets as follows:

$$p(y=1|w_i)=\frac{p(y=1)p(w_i|y=1)}{p(y=1)p(w_i|y=1)+p(y=0)p(w_i|y=0)} \tag{18}$$

The occurrence probability of w_i in outage-unrelated tweets is $p(y=0|w_i)$ and can be calculated as $1-p(y=1|w_i)$. After the training process, the statistic Bayesian method can calculate the probability that a new tweet is related or unrelated to a power outage by:

$$p(y=1|w_1,w_2,\ldots,w_n)=\prod_{i=1}^{n}p(y=1|w_i) \tag{19}$$

$$p(y=0|w_1,w_2,\ldots,w_n)=\prod_{i=1}^{n}p(y=0|w_i) \tag{20}$$

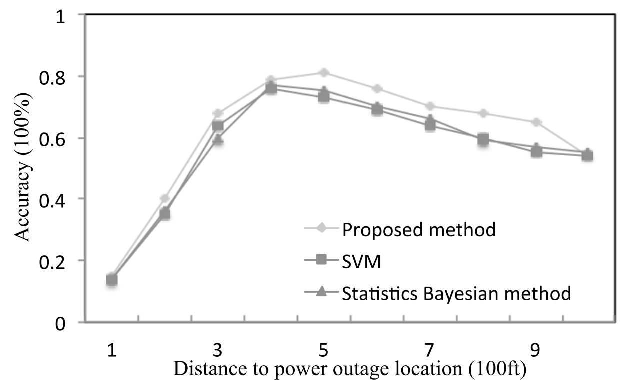

Fig. 5. Accuracy of identifying real outage-reporting tweets in different sizes of spatial domain.

where $w_1, w_2, \ldots, w_n$ are the words in the tweet. If $p(y=1|w_1,w_2,\ldots,w_n) > p(y=0|w_1,w_2,\ldots,w_n)$, it classifies the tweet as an outage-related one.

We separate the outage detection task into two parts. The first part is to find the power outage-related tweets. Then, these tweets are used to determine whether a power outage event occurred. In the first stage, SVM, statistics Bayesian, and the proposed method can all identify whether a tweet is related to power outages. But in the second stage, neither SVM nor statistics Bayesian can deal with temporal and spatial information of tweets, due to the lack of the input mechanism for multi-type features. Therefore, we use (17) to estimate whether a tweet indicates a real power outage. The difference between SVM/statistics Bayesian and our method in the second stage is that the probability of outage-related tweets $p(y=1 \mid d)$ is calculated by SVM and statistics Bayesian themselves. In addition, the SVM and statistics Bayesian are supported by the scikit-learn library [32] in this paper.

A. Performance of Spatial Information Detection

The spatial information is important for the event detection. It is necessary to investigate whether the proposed method can detect the spatial information and how the performance of the event detection can be improved by leveraging the spatial information. The first step is to determine the size of the spatial domain that will be taken into account, i.e., tweets in how much area that will be considered. We use the training set to conduct a few tests. We find that if the radius of the domain is small, the system ignores many outage-related tweets, which will affect the performance dramatically. On the other hand, if the radius is too large, the performance becomes worse. The proportion of the outage-related tweets decreases in a large area. As shown in Fig. 5, the best performance happens at about 450 ft in all of the three methods. So, we set 450 ft as the boundary. This means we use an outage-related tweet as the center and consider the tweets within a 450-ft radius. If we identify any new outage-related tweets, they will be used as new centers. The process is repeated until we find enough outage-related tweets. Fig. 6 (a) shows the actual distance between the power outage and the posting locations of tweets. Figs. 6 (b), (c), and (d) show the spatial information detected by the proposed method, SVM, and the statistics Bayesian

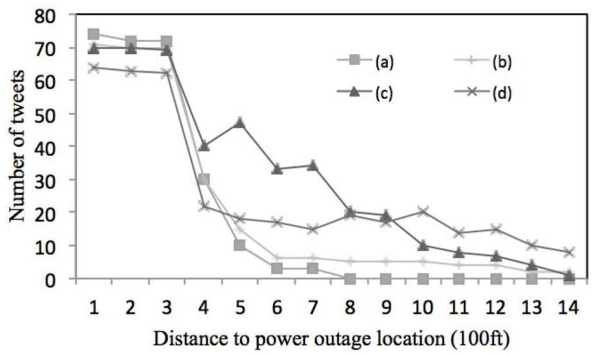

Fig. 6. (a) Actual distance between posting locations and outage locations; (b), (c), and (d) show numbers of real outage-reporting tweets identified by different methods.

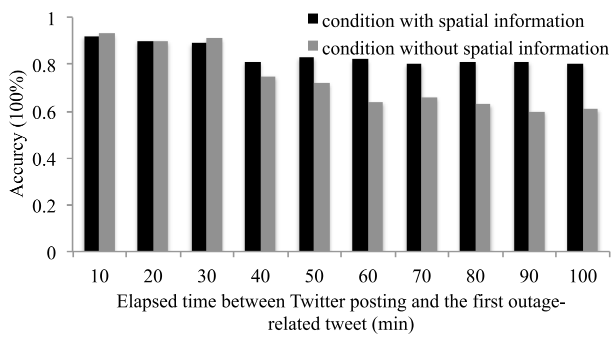

Fig. 7. Comparison of outage detection with and without spatial information.

method, respectively. Compared to the other two methods, the location detection results of the proposed method are closer to the actual case.

We have compared the detection performance of cases with and without the spatial information in the learning process. Fig. 7 shows the comparison of results in detecting power outage events. The *x*-axis represents the elapsed time between the Twitter posting and the first identified outage-related tweet. The *y*-axis represents the accuracy of detection of outage-related tweets. The accuracy rate is defined as the percentage fraction of identified events that are real outages in all possible outages that are detected by the method. It can be seen that the spatial information can be used to improve the detection accuracy. This is because more irrelevant tweets can be filtered by considering the spatial information, i.e., a tweet has a low probability to report an actual event if the posting location is far from the concentrated posting areas of event-related tweets.

B. Topic Extraction

Table I shows a list of samples of topics extracted and clustered by Algorithm 1. The total number of topics of tweets is 100. For brevity, the table shows 5 selected topics and the top 7 words for each. It can be seen that Topic 1 is more related to power outages, Topic 2 is more related to medicines and drugs, Topic 3 is related to colors, Topic 4 is related to human behaviors, and Topic 5 is related to medical care. The extracted words have been successfully clustered into different topic categories by the sLDA method.

TABLE I
EXAMPLES OF TOPICS EXTRACTED FROM TWEETS BY PROPOSED sLDA

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	···
Power	Drug	Red	Mind	Doctor	···
Outage	Medicine	Blue	Thought	Patient	
Cut	Effect	Green	Remember	Hospital	
Today	Body	Yellow	Memory	Nurse	
Area	Pain	White	Thinking	Patient	
Cancelled	Person	Color	Felt	Medicine	
Shut	Alcohol	Bright	Remember	Dental	

TABLE II
TWEET EXAMPLES THAT ARE MANUALLY LABELED

#	Tweet Contents	Related	Real Reporting
1	5 Reasons Developers should build their Natural Language Processing	N	N
2	Power outages affect 1,995 customers near Jackson Bluff, Woodward and W Gaines. Crews in route. Again, it appears to be due to lightning.	Y	Y
3	plot twist: power outage is the senior prank	Y	N
4	Shoutout to this power outage at GSU for ending my day early	Y	Y
5	We need to have a power outage so we can leave school and go home!!	Y	N
6	Very cool piano that teaches you how to play	N	N
7	Hi, kindly report power outages by SMSing the word “power” followed by your acc no to 0826120333 or 44676	Y	N
8	I am bored as the power is cut off , no internet , no tv	Y	Y
9	A power cut in the nail shop is not okay	Y	Y
10	Hundreds of homes in Cheltenham were without power this morning	Y	Y

C. Detection Result

The proposed system detects whether the content of a tweet is related to power outages, and then identifies whether the tweet reports a real power outage event. As discussed in Sections IV and V, an outage-related tweet does not necessarily indicate a real outage event. The power outage alert will only be issued if a significant number of outage-related tweets are identified in a specific area and within a certain time period. Table II shows a few sampled tweets. Some of the tweets are related to outage, but do not really report

TABLE III
DETECTION RESULTS OF THREE METHODS (THE PROPOSED METHOD, SVM, AND STATISTIC BAYESIAN METHOD)

#	Proposed method		SVM		Bayesian method	
	Related	Real Reporting	Related	Real Reporting	Related	Real Reporting
1	N	N	N	N	N	N
2	Y	Y	Y	N	N	N
3	Y	N	Y	N	Y	N
4	Y	Y	Y	Y	Y	Y
5	Y	N	Y	N	Y	N
6	N	N	N	N	N	N
7	Y	N	N	N	N	N
8	Y	Y	Y	Y	Y	Y
9	Y	Y	Y	Y	Y	N
10	Y	Y	Y	Y	Y	Y

power outages. Table III shows the detecting results of the sampled tweets by the proposed method, SVM, and statistics Bayesian method. It can be seen that the proposed method outperforms the other two methods in identifying outage-related tweets. For example, for tweets 2 and 7, SVM and statistics Bayesian methods cannot deal with such a long tweet with noisy information. This is because the methods based on keywords cannot process complicated information. The topic model uses a general way to represent tweets. Each tweet is represented by a topic distribution. Therefore, our proposed method can analyze the complex information. Another reason is that the proposed method takes advantage of spatial and temporal information to improve the detection performance, while SVM and statistics Bayesian method are unable to deal with multi-type information.

We also computed the recall measure. Recall is defined as the fraction of relevant instances that are retrieved. We assume that False Negative is the total number of power outages mentioned on Twitter in the test set that are missed by our system, and True Positive is the set of real power outages that are mentioned in tweets in the test set and are identified by our method. The recall is calculated as:

$$Recall = \frac{TruePositive}{TruePositive + FalseNegative} \tag{21}$$

In this test, we use the threshold value V_{th} to control the recall values. The performance comparison is shown in Fig. 8. The accuracy decreases as the recall value increases. Overall, the accuracy of the proposed method outperforms the other two methods. As shown in Table IV, we use the best F-measure to compare the performance. F-measure is calculated as:

$$F = \frac{2 \times Accuracy \times Recall}{Accuracy + Recall} \tag{22}$$

The accuracy of prediction of our proposed method is 81.6%, which is better than traditional prediction methods such as SVM and statistics Bayesian method. The results show that the recall of power outage event is higher than that of SVM by 15%, and is higher than that of statistics Bayesian method by 8%. The accuracy of the proposed method is higher than SVM and statistics Bayesian method by 6% and 18%, respectively.

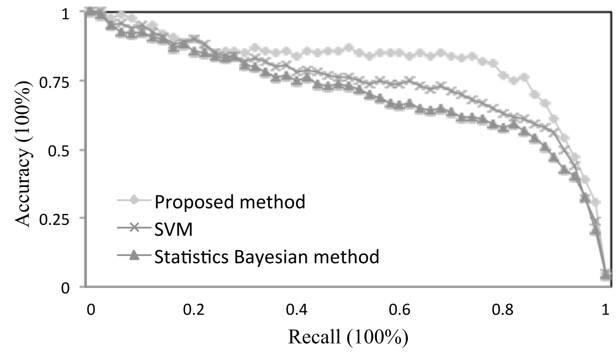

Fig. 8. Comparison of accuracy and recall values of three methods.

TABLE IV
ACCURACY AND RECALL BY LEARNING FROM THE TRAINING DATA

	Accuracy	Recall	F-measure
Proposed method	81.6%	78%	0.80
SVM	73%	68%	0.71
Statistics Bayesian method	63.5%	70%	0.67

VII. CONCLUSION

This paper provides a novel methodology for the detection and location of power outages using social sensors. The proposed approach is based on a supervised topic model with a heterogeneous information network. Compared to existing methods, such as topic modeling on event detection, the new method leverages both temporal and spatial information to improve the performance of event identification. The numerical experiments show that the proposed method can assist utilities to locate actual outage areas without adding any measurement and communication infrastructures. The proposed method can be integrated with the existing distribution management system to enable a social data-driven outage management.

REFERENCES

[1] S. J. Chen, T. S. Zhan, C. H. Huang, J. L. Chen, and C. H. Lin, "Nontechnical loss and outage detection using fractional-order self-synchronization error-based fuzzy Petri nets in micro-distribution systems," *IEEE Trans. Smart Grid*, vol. 6, no. 1, pp. 411–420, Jan. 2015.

[2] Y. Liu and N. N. Schulz, "Knowledge-based system for distribution system outage locating using comprehensive information," *IEEE Trans. Power Syst.*, vol. 17, no. 2, pp. 451–456, May 2002.

[3] R. A. Fischer, A. S. Laakonen, and N. N. Schulz, "A general polling algorithm using a wireless AMR system for restoration confirmation," *IEEE Trans. Power Syst.*, vol. 16, no. 2, pp. 312–316, May 2001.

[4] K. Sridharan and N. N. Schulz, "Outage management through AMR systems using an intelligent data filter," *IEEE Trans. Power Del.*, vol. 16, no. 4, pp. 669–675, Oct. 2001.

[5] Y. Zhao, R. Sevlian, R. Rajagopal, A. Goldsmith, and H. V. Poor, "Outage detection in power distribution networks with optimally-deployed power flow sensors," in *Proc. IEEE Power Energy Soc. Gen. Meeting (PES)*, Vancouver, BC, Canada, 2013, pp. 1–5.

[6] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes Twitter users: Real-time event detection by social sensors," in *Proc. 19th Int. Conf. World Wide Web*, Raleigh, NC, USA, 2010, pp. 851–860.

[7] A. Ritter, Mausam, O. Etzioni, and S. Clark, "Open domain event extraction from Twitter," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. Disc. Data Min.*, Beijing, China, 2012, pp. 1104–1112.

[8] D. Keseris, "Social media could provide early warning of power outages," Telegraph Technol. News, Apr. 2014. [Online]. Available: http://www.telegraph.co.uk/technology/news/10789017/Social-media-could-provide-early-warning-of-power-outages.html

[9] B. A. Huberman, D. M. Romero, and F. Wu, "Social networks that matter: Twitter under the microscope," *First Monday*, vol. 14, no. 1, Jan. 2009.

[10] D. Boyd, S. Golder, and G. Lotan, "Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter," in *Proc. 43rd Hawaii Int. Conf. Syst. Sci.*, Honolulu, HI, USA, 2010, pp. 1–10.

[11] J. Weng and B.-S. Lee, "Event detection in Twitter," in *Proc. Assoc. Adv. Artif. Intell.*, vol. 11. San Francisco, CA, USA, 2011, pp. 401–408.

[12] R. Li, K. H. Lei, R. Khadiwala, and K. C.-C. Chang, "TEDAS: A Twitter-based event detection and analysis system," in *Proc. IEEE 28th Int. Conf. Data Eng. (ICDE)*, Washington, DC, USA, Apr. 2012, pp. 1273–1276.

[13] K. Bauman, A. Tuzhilin, and R. Zaczynski, "Virtual power outage detection using social sensors," NYU Working Paper no. 2451/34226, 2015. [Online]. Available : http://ssrn.com/abstract=2654653

[14] J. M. Kleinberg, "Bursty and hierarchical structure in streams," *Data Min. Knowl. Disc.*, vol. 7, no. 4, pp. 373–397, 2003.

[15] J. M. Cotelo, F. L. Cruz, and J. A. Troyano, "Dynamic topic-related tweet retrieval," *J. Assoc. Inf. Sci. Technol.*, vol. 65, no. 3, pp. 513–523, 2014.

[16] A. M. MacEachren *et al.*, "Geo-Twitter analytics: Applications in crisis management," in *Proc. 25th Int. Cartograph. Conf.*, Paris, France, 2011, pp. 3–8.

[17] K. Starbird and J. Stamberger, "Tweak the tweet: Leveraging microblogging proliferation with a prescriptive syntax to support citizen reporting," in *Proc. Inf. Syst. Crisis Response Manage. (ISCRAM)*, Seattle, WA, USA, 2010, pp. 1–5.

[18] M. Guy, P. Earle, C. Ostrum, K. Gruchalla, and S. Horvath, "Integration and dissemination of citizen reported and seismically derived earthquake information via social network technologies," in *Proc. 9th Int. Symp. Adv. Intell. Data Anal. (IDA)*, Tucson, AZ, USA, May 2010, pp. 42–53.

[19] Q. Zhao and P. Mitra, "Event detection and visualization for social text streams," in *Proc. 1st Int. Conf. Weblogs Soc. Media (ICWSM)*, Boulder, CO, USA, Mar. 2007, pp. 100–103.

[20] S. Zhao, L. Zhong, J. Wickramasuriya, and V. Vasudevan, "Human as real-time sensors of social and physical events: A case study of Twitter and sports games," *arXiv preprint arXiv:1106.4300*, 2011. [Online]. Available: http://arxiv.org/abs/1106.4300

[21] D. M. Blei, "Probabilistic topic models," *Commun. ACM*, vol. 55, no. 4, pp. 77–84, 2012.

[22] W. X. Zhao *et al.*, "Comparing Twitter and traditional media using topic models," in *Proc. Adv. Inf. Retrieval*, Dublin, Ireland, 2011, pp. 338–349.

[23] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, Mar. 2003.

[24] T. Hofmann, "Probabilistic latent semantic indexing," in *Proc. 22nd Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, Berkeley, CA, USA, 1999, pp. 50–57.

[25] T. Minka and J. Lafferty, "Expectation-propagation for the generative aspect model," in *Proc. 18th Conf. Uncertainty Artif. Intell.*, Edmonton, AB, Canada, 2002, pp. 352–359.

[26] T. L. Griffiths and M. Steyvers, "Finding scientific topics," *Proc. Nat. Acad. Sci.*, vol. 101, pp. 5228–5235, Apr. 2004.

[27] M. Rosen-Zvi, T. Griffiths, M. Steyvers, and P. Smyth, "The author-topic model for authors and documents," in *Proc. 20th Conf. Uncertainty Artif. Intell.*, Banff, AB, Canada, 2004, pp. 487–494.

[28] J. D. Mcauliffe and D. M. Blei, "Supervised topic models," in *Proc. Adv. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, 2008, pp. 121–128.

[29] Y. Sun, Y. Yu, and J. Han, "Ranking-based clustering of heterogeneous information networks with star network schema," in *Proc. 15th ACM SIGKDD Int. Conf. Knowl. Disc. Data Min.*, Paris, France, 2009, pp. 797–806.

[30] D. B. Neill and G. F. Cooper, "A multivariate Bayesian scan statistic for early event detection and characterization," *Mach. Learn.*, vol. 79, no. 3, pp. 261–282, 2010.

[31] T. Joachims, "Text categorization with support vector machine: Learning with many relevant features," in *Proc. 10th Eur. Conf. Mach. Learn.*, Chemnitz, Germany, Apr. 1998, pp. 137–142.

[32] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Nov. 2011.

Haifeng Sun is currently pursuing the Ph.D. degree with the State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications. His research interest includes data mining, information retrieval, and natural language processing.

Zhaoyu Wang received the B.S. and M.S. degrees in electrical engineering from Shanghai Jiaotong University in 2009 and 2012, respectively, and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology in 2012 and 2015, respectively. He is an Assistant Professor with Iowa State University. He was a Research Aid at Argonne National Laboratory in 2013, and an Electrical Engineer at Corning Inc. in 2014. His research interests include power distribution systems, microgrids, renewable integration, self-healing resilient power systems, and voltage/VAR control.

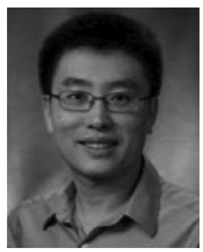

Jianhui Wang received the Ph.D. degree in electrical engineering from the Illinois Institute of Technology, Chicago, IL, USA, in 2007. He is currently the Section Lead of Advanced Power Grid Modeling with the Energy Systems Division, Argonne National Laboratory, Argonne, IL, USA. He is also an Affiliate Professor with Auburn University and an Adjunct Professor with the University of Notre Dame. He is the Secretary of the IEEE Power and Energy Society (PES) Power System Operations Committee. He is the Editor-in-Chief of the IEEE TRANSACTIONS ON SMART GRID and an IEEE PES Distinguished Lecturer.

Zhen Huang received the dual master's degrees in electrical engineering from Shanghai Jiaotong University and the Georgia Institute of Technology. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, Georgia Institute of Technology. His research interest includes speech recognition, deep learning, general machine learning, multimedia information retrieval and image processing, and pattern recognition.

NichelleLe Carrington received the B.S. degree in electrical engineering from North Carolina Agricultural and Technical State University, Greensboro, NC, USA, in 2014. She is currently pursuing the Ph.D. degree in electrical engineering with Iowa State University.

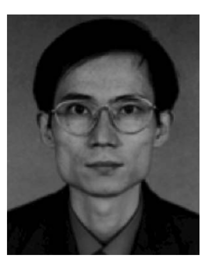

Jianxin Liao received the Ph.D. degree with the University of Electronics Science and Technology of China in 1996. He is currently the Dean of the Network Intelligence Research Center and a Full Professor with the State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications. His main creative contributions include mobile intelligent network, service network intelligent, networking architectures and protocols, and multimedia communication. He was a recipient of the National Prize for Progress in Science and Technology in 2004 and 2009.