



Kent Academic Repository

Liang, Zhanbo, Guo, Jie, Qiu, Weidong, Huang, Zheng and Li, Shujun (2024)
When graph convolution meets double attention: online privacy disclosure detection with multi-label text classification. *Data Mining and Knowledge Discovery* .
ISSN 1384-5810.

Downloaded from

<https://kar.kent.ac.uk/104513/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.1007/s10618-023-00992-y>

This document version

Publisher pdf

DOI for this version

Licence for this version

CC BY (Attribution)

Additional information

For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal** , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).



When graph convolution meets double attention: online privacy disclosure detection with multi-label text classification

Zhanbo Liang¹ · Jie Guo¹ · Weidong Qiu¹ · Zheng Huang¹ · Shujun Li²

Received: 28 November 2022 / Accepted: 18 November 2023
© The Author(s) 2024

Abstract

With the rise of Web 2.0 platforms such as online social media, people's private information, such as their location, occupation and even family information, is often inadvertently disclosed through online discussions. Therefore, it is important to detect such unwanted privacy disclosures to help alert people affected and the online platform. In this paper, privacy disclosure detection is modeled as a multi-label text classification (MLTC) problem, and a new privacy disclosure detection model is proposed to construct an MLTC classifier for detecting online privacy disclosures. This classifier takes an online post as the input and outputs multiple labels, each reflecting a possible privacy disclosure. The proposed presentation method combines three different sources of information, the input text itself, the label-to-text correlation and the label-to-label correlation. A double-attention mechanism is used to combine the first two sources of information, and a graph convolutional network is employed to extract the third source of information that is then used to help fuse features extracted from the first two sources of information. Our extensive experimental results, obtained on a public dataset of privacy-disclosing posts on Twitter, demonstrated that our proposed privacy disclosure detection method significantly and consistently outperformed other state-of-the-art methods in terms of all key performance indicators.

Keywords Privacy disclosure detection · User generated content (UGC) · Online social media · Graph convolutional network (GCN) · Multi-label text classification (MLTC)

Responsible editor: Eyke Hüllermeier.

Extended author information available on the last page of the article

Published online: 05 January 2024

1 Introduction

The rapid development of information and communication technologies have helped facilitate people's social interactions. Online social media platforms like Twitter provide people a new way to build up their social relationships, share their daily lives, and express their emotions. However, many online users frequently (and often unintentionally) share personal information online, which can lead to unwanted online disclosures of private information of themselves or other people in their social networks. Figure 1 shows several imaginary but realistic online posts of such unintended privacy disclosures on Twitter, generated based on some examples in a research dataset of privacy-disclosing tweets constructed by Song et al. (2018). Although people can check their online posts manually to avoid privacy disclosures, many online users do not have a good level of awareness on such privacy issues, and they do not necessarily know when and what to check. Therefore, automated solutions that can help online users identify such issues and take proper actions are important, which is the focus of our work.

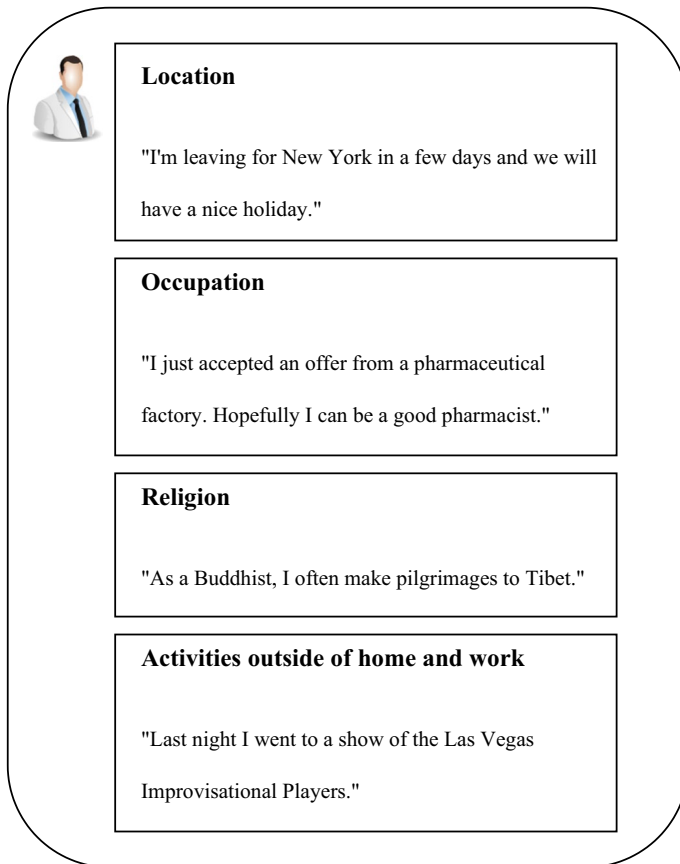


Fig. 1 Several illustrative examples of possible privacy disclosures on online social media platforms

Past studies about privacy disclosure detection attempted to solve this problem with different machine learning methods. Traditional methods on privacy disclosure detection try to detect privacy disclosures in user profiles or user settings, but not in user generated content (UGC), leading to incomplete detection. More recently, many researchers started studying privacy disclosure detection in UGC by analysing pictures and/or texts in such UGC. Therefore, their work extends the scope of such work.

Recently, some researchers use the multi-label text classification (MLTC) framework to model the privacy disclosure problem (Song et al. 2018; Chen et al. 2020). MLTC is an important task in the field of natural language processing (NLP). Different from multi-class text classification (MCTC), which classifies a given piece of text into one of multiple class labels, MLTC aims to tag a piece of given text with *multiple* (i.e., one or more) content-specific labels. In Song et al. (2018) and Chen et al. (2020), the privacy information is divided into eight main categories, then they make further division, using 32 categories of labels to reflect the possible disclosed privacy. However, their methods are limited due to the lack of consideration for the relationship between texts and labels. Their methods aim to improve the prediction results by considering the co-occurrence relation between labels. For example, the label “Health condition” usually appears with the label “Treatment” and the label “Occupation” usually appears with the label “Salary”. However, those two methods do not consider label-text correlations, i.e., their work ignores the fact that some key words or phrases in the input texts can assist indicating the possible privacy-aware labels. For example, a location name in the input text may help to indicate that the text is involved in the privacy disclosure of “Current location” or “Place planning to go”. We follow their thoughts to model privacy disclosure detection as an MLTC problem. Our proposed framework takes an online post as the input, and outputs a number of privacy-relevant labels that indicate potential disclosure of different types of personal information in the input online post.

Considering that privacy disclosure is a universal problem in people’s daily life, new frameworks with better performance on privacy disclosure detection are needed. The aim of our work is to provide a more effective MLTC privacy disclosure detection algorithm to facilitate the fine-grained text privacy detection. As mentioned before, current MLTC privacy disclosure models are limited by their consideration of relationships between various texts or words. In order to improve the performance of privacy-disclosing post detection, which combines three different sources of relevant information, the text information, the label-to-text correlation and the label-to-label correlation, to produce a more comprehensive model for detecting privacy-disclosing online posts. Our model extracts the text representations through a double-attention mechanism as Xiao et al. (2019) did, which measures the contribution of each word to each privacy-relevant label. The label-to-label correlation is considered in the final text representation via a graph convolutional network (GCN). We propose a new feature fusion mechanism assisted by GCN to make the fused feature more comprehensive. We utilize the label-to-label correlation to obtain the proposed compensation coefficients from both the self-attention and the label-attention text representations. We summarize the main contributions of our work as follows:

- A new privacy disclosure detection model with multi-label text classification is proposed. Our model presents a new fine-grained privacy disclosure detection algorithm and outputs multiple privacy-aware labels as the possible leaked privacy. From the perspective of the detection performance, our model provides a better solution to the fine-grained privacy disclosure detection on the UGC.
- Our proposed model considers three different sources of relevant information for the MLTC task: the input text itself, the label-to-text correlation, and the label-to-label correlation.
- A new feature fusion mechanism assisted by a GCN is proposed to construct comprehensive text representations with the guidance of the label-to-label correlation. The idea of compensation coefficients is proposed in the feature fusion mechanism, which reflects the compensation relationship between self-attention and label-attention.
- A series of experiments on a public privacy-disclosing tweet dataset showed that our proposed model outperformed selected state-of-the-art models significantly and consistently. Our code has been released to facilitate others to conduct follow-up research.¹

The rest of the paper is organized as follows. Section 2 introduces the related work. Section 3 elaborates the proposed MLTC-based model for privacy detection. Section 4 shows and discusses the experiment results. Section 5 concludes our work and discusses the future work. Section 6 makes statements on financial or non-financial interests that are directly or indirectly related to the work submitted for publication.

2 Related work

2.1 Privacy disclosure analysis

The problem of online privacy disclosures has attracted the attention of many researchers. Some researchers studied this problem based on analysis of user profiles (Biega et al. 2017; Eslami et al. 2017; Huang and Paul 2019) or privacy settings of user accounts (Raber and Krüger 2018; Sanchez et al. 2020). Biega et al. (2017) proposed a privacy-aware framework that leverages solidarity in a large community to scramble user interaction histories, in order to disturb the information collection from user profiles by the online service providers. To minimize users' privacy risks, Eslami et al. (2017) proposed an alternative solution, where posts of different users are split and merged into synthetic mediator profiles. Raber and Krüger (2018) studied privacy settings of user accounts by observing the context factors and personality measures which can be used to predict the correct privacy level out of seven privacy levels. Sanchez et al. (2020) considered how to model users' privacy preferences for data sharing and processing in the IoT and fitness domain, paying a specific attention to the GDPR compliance.

¹ <https://github.com/xiztt/wgma>

Some other researchers such as Tran et al. (2016) and Mao et al. (2011) also proposed classifiers to detect privacy disclosures in user-generated online posts. Tran et al. (2016) proposed Privacy-CNH, a binary classification framework that utilizes hierarchical features including both object and convolutional features in a deep learning model to detect whether a photo is private or not. Mao et al. (2011) analysed privacy disclosures on Twitter by building binary classifiers to detect three types of privacy disclosure including divulging vacation plans, tweeting under the influence of alcohol and revealing medical conditions. Despite all the past studies, they only focused on privacy disclosure detection at a more coarse-grained level. These studies used frameworks or classifiers to implement relatively simple analysis of privacy disclosures, normally based on less comprehensive privacy categories so not being able to cover some specific privacy disclosure scenarios.

In order to achieve finer-grained analysis, Song et al. (2018) proposed a taxonomy-guided multi-task learning model to detect what personal aspects of online users are disclosed in online posts. They also constructed a dataset of privacy-disclosing tweets covering 32 privacy-relevant personal aspects. Similarly, Chen et al. (2020) proposed GrHA, a fine-grained privacy detection network, to improve the performance of the model proposed in (Song et al. 2018). The above two proposed methods aim to improve the prediction results by considering label co-occurrences, but they did not consider label-to-text correlations explicitly.

2.2 Multi-label text classification

Traditional machine learning methods (Kumar and Daumé III 2012; Jacob et al. 2008) have been widely used to deal with MLTC tasks. Kumar and Daumé III (2012) proposed the GO-MTL model by using grouping and overlap mechanism to enhance the semantic correlations in MLTC tasks. Likewise, Jacob et al. (2008) studied the clustered multi-task learning to deal with MLTC tasks. Although these machine learning methods utilize multiple hand-crafted features to enhance the semantic representations in MLTC tasks, they overlook deep semantic features among input text and multi labels.

Nowadays, researchers have made great progress on the deep learning technology. Therefore deep models such as CNN (Liu et al. 2017; Kurata et al. 2016; Kim 2014) and RNN (Liu et al. 2016; Chen et al. 2017) have been used to implement end-to-end MLTC tasks. In more recent studies, researchers have also proposed to use attention mechanisms such as DocBERT (Adhikari et al. 2019) and other methods such as SGM (Yang et al. 2018) and LSAN (Xiao et al. 2019) to consider the label-to-text correlation in the MLTC problem. Adhikari et al. (2019) proposed DocBERT model as a much simpler BERT model with competitive accuracy at a far more modest computational cost in terms of MLTC tasks. Yang et al. (2018) considered how to address the MLTC problem by capturing the correlations between labels as well as the most informative words automatically when predicting different labels. Xiao et al. (2019) used self-attention and label-attention for better representations of input text in MLTC tasks. Label co-occurrences are a vital source of information when dealing with the MLTC problem. More specifically, some labels




Fig. 2 The architecture of the proposed network

often appear with other labels due to the semantic relation. However, most existing methods focus only on optimizing the process of feature extraction, but do not consider label co-occurrences. By utilizing the GCN model, Ma et al. (2021) proposed LDGN (label-specific dual graph neural network) to improve the MLTC representations by including label co-occurrences. Although they considered label co-occurrences to a certain extent, their method has some limitations in the process of combination with the feature extraction module, for their model's usage of the GCN only attempts to optimize the text representation of the model with label co-occurrences yet ignores diversity of the text representation and the labels' guidance on fusing different feature vectors.

3 Proposed method

In this section, we introduce the GCN-based double attention network, as shown in Fig. 2. The network includes four major components: (1) an input text feature encoder that transforms the input text into word-level semantic vectors; (2) a double-attention text representation component that enhances the important word representations of the text combining both text information and label information; (3) a GCN-assisted feature fusion mechanism that utilizes the label-to-label correlation acquired by GCN to guide the double-attention information fusion process; and (4) a

label probability output component that predicts the probabilities of various privacy-relevant labels.

3.1 Problem formulation

Let $\mathbb{D} = \{(x_i, y_i)\}_{i=1}^N$ denote the set of texts, where x_i represents the input texts and $y_i \in \{0, 1\}^L$ represents its corresponding labels. Here, L denotes the total number of privacy-relevant labels. The target of the proposed method in this paper is to learn the output probability of each label from the input text, in order to match the most relevant labels.

3.2 Input text feature encoder

Given a text x_i containing M words ($x_i = \{w_{i1}, w_{i2}, \dots, w_{iM}\}$), the word2vec method (Le and Mikolov 2014) is adopted to obtain the embedding vector based on the input, which is denoted as $\mathbf{E}_s \in \mathbb{R}^{M \times d_1}$, where d_1 denotes the embedding dimension.

For fair comparisons, we used the same feature extraction structure, bidirectional long short-term memory (BiLSTM) (Zhou et al. 2016), as the baseline models (Chen et al. 2020; Xiao et al. 2019; Ma et al. 2021) used, to get the embedding. We adopt the BiLSTM model to process the embedded vector. The formula is as follows:

$$\mathbf{H} = \left\{ \overrightarrow{\mathbf{H}}, \overleftarrow{\mathbf{H}} \right\}, \quad (1)$$

where $\overrightarrow{\mathbf{H}}, \overleftarrow{\mathbf{H}} \in \mathbb{R}^{M \times d_2}$ represent the forward and backward text representations, respectively. The whole text can be represented as $\mathbf{H} \in \mathbb{R}^{M \times 2d_2}$.

3.3 Double-attention text representation

We use a double attention mechanism to generate text- and label-specific representations from the output of the BiLSTM. A self-attention model is adopted to capture the long-term dependence of words in \mathbf{H} . Meanwhile, to extract the text attention from the corresponding labels, a label-specific attention model is used as the supplementary information.

3.3.1 Self-attention model

Self-attention models have shown their considerable merits on assessing the importance of word representations. Therefore, we adopt a self-attention mechanism (Lin et al. 2017) to reinforce the semantic representation of the text based on the word-to-word correlations. Different from traditional self-attention algorithms, the self-attention sentence embedding algorithm (Lin et al. 2017) uses multiple hops of attention calculated from the LSTM outputs \mathbf{H} to focus on different aspects of the meanings

of the sentence. Since the output labels have the dimensionality of L , we take the self-attention weights with L dimensions to reflect the effects of L labels to M words. The calculation of attention weights can be described as follows:

$$\mathbf{A}_s = \text{softmax}(\mathbf{W}_{s2} \tanh(\mathbf{W}_{s1} \mathbf{H}^T)), \quad (2)$$

where $\mathbf{A}_s \in \mathbb{R}^{L \times M}$ are the self attention weights that indicate the effect of each word to each label. $\mathbf{W}_{s1} \in \mathbb{R}^{d_3 \times 2d_2}$, $\mathbf{W}_{s2} \in \mathbb{R}^{L \times d_3}$ are the parameters to be trained. Then, the attention weights are utilized to update the text representation:

$$\mathbf{Q}_s = \mathbf{A}_s \times \mathbf{H}^T. \quad (3)$$

3.3.2 Label-attention model

Apart from obtaining text attention from the text itself, the label-attention model (Xiao et al. 2019) is adopted to extract text attention from the corresponding labels. The labels' semantic information is acquired with the word2vec method, which is denoted as $\mathbf{E}_l \in \mathbb{R}^{L \times d_l}$.

To capture a better semantic representation with the guidance of output labels, the label-attention mechanism computes the attention weights by calculating the relationship between the labels and the text as follows $\mathbf{A}_l = \mathbf{E}_l \times \mathbf{H}^T$, where $\mathbf{A}_l \in \mathbb{R}^{L \times M}$ are the label-specific attention weights that indicate the effect of each word to each label. With the weight matrix, the label-specific attention weights are utilized to enhance the label-aware information in the text semantic representation $\mathbf{Q}_l = \mathbf{A}_l \times \mathbf{H}^T$.

3.4 GCN-assisted feature fusion

In this section, the GCN-assisted feature fusion mechanism is described to construct comprehensive text representations with the guidance of the label-to-label correlation.

We use a GCN framework to extract a label-to-label correlation matrix. With the guidance of the correlation matrix, we enhance the text representations by utilizing the proposed compensation coefficients to implement the algorithm of feature fusion.

3.4.1 GCN-based label-to-label correlation extraction

The graph convolutional networks (GCNs) (Kipf and Welling 2016) were proposed to get a better understanding of the relationship of nodes in a graph. A GCN uses an adjacency matrix to characterize the graph structure and a convolutional network to capture the correlations among different nodes, with an output of a correlation matrix. In our work, we aim to extract the label co-occurrence through a GCN. The label co-occurrence refers to the simultaneous occurrence of two or more labels in the same text. For example, considering the two labels




Fig. 3 Construction of the initial weighted label graph

“Salary” and “Occupation”, their probability of co-occurrence is high due to their semantic relation (i.e., an occupation is normally associated with a salary). Therefore, we utilize the GCN to transform such label-to-label relationships (inferred from label co-occurrences and their semantic relationships) into mathematical representations.

As Fig. 3 shows, the output labels are represented as a weighted label graph (\mathbf{V}, \mathbf{E}) , where each node represents a label embedding and each edge’s weight refers to the two adjacent labels’ co-occurrence frequency. More specifically, each node is initialized to be the embedded vector of the corresponding label and each edge weight is calculated to be the co-occurrence frequency of the two labels representing the two adjacent nodes based on information in the training set. In Fig. 3, the symbol # represents the the number of occurrences. For example, $\#(a)$ represents the number of tweets with the label a in the training set and $\#(a, b)$ represents the number of tweets with both labels a and b in the training set. We use \mathbf{P} to represent the initial co-occurrence adjacent matrix. According to Chen et al. (2020), considering the noisy co-occurrence caused by the sparse real-world dataset, the initial co-occurrence adjacent matrix \mathbf{P} should be binarized and revised as follow:

$$a_j^k = \begin{cases} \frac{u}{\sum_{x=1}^L p_j^x}, & \text{if } j \neq k, \\ 1 - u, & \text{if } j = k, \end{cases} \tag{4}$$

where p_j^k represents the co-occurrence frequency of label j to label k and a_j^k represents the revised co-occurrence frequency. u represents the trade-off parameter that balances the weights between the label itself and its correlated labels. We use \mathbf{A} to represent the revised adjacency matrix. In our work, we use the same revised adjacency matrix as Chen et al. (2020) did. The trade-off parameter is set to 0.2.

Then, a GCN is adopted to update the label-to-label correlation representations from the previous representations and the adjacency matrix containing co-occurrence probabilities. The GCN propagation is calculated as follows:

$$\mathbf{C}^{(l+1)} = \sigma\left(\mathbf{A}\mathbf{C}^{(l)}\mathbf{W}_g^{(l)}\right), \quad (5)$$

where $\mathbf{C}^{(l)} \in \mathbb{R}^{L \times d_4^{(l)}}$ represents the input label-to-label correlation representations for the l -th GCN layer, σ denotes the activation function (LeakyReLU is adopted here), \mathbf{A} is the revised adjacency matrix, and $\mathbf{W}_g^{(l)} \in \mathbb{R}^{d_4^{(l)} \times d_4^{(l+1)}}$ denotes the transformation matrix to be learned for the l -th layer.

Our GCN contains two layers. As a result, the second layer's embedding size adopts $2d_2$ to align the dimension of the output from the double-attention model. Thus the correlation matrix is obtained from the output of the second layer, which is denoted as $\mathbf{C}^{\text{out}} \in \mathbb{R}^{L \times 2d_2}$.

3.4.2 Feature fusion guided by label-to-label correlation

As mentioned above, we obtain the text representations including the text semantic information (from self-attention) and the label-to-text correlation (from label-attention), and represent the label-to-label correlation through a GCN. The text semantic information uses the self-attention mechanism to enhance the weight of key words or phrases based on the inputting text semantics itself. Meanwhile, the label-to-text correlation provides the improved text representations through the label-attention mechanism, which is based on the labels' semantic representations. Therefore, these two text representations shuffle the word weights of the input texts to enhance their key parts. However, they are based on different semantic information (the text itself and the labels' semantics) and the enhanced parts are different. Therefore, it is important and necessary to fuse these two representations in order to get a more comprehensive semantic representations. To this end, we propose a cross-attention model that utilizes the label-to-label correlation matrix to guide the fusion of output features from the double-attention model. The experimental results demonstrated the superiority of our model compared to other state-of-the-art methods.

Our method aims to enhance the weak part of the representations in the output from different attention models and utilize the label-to-label correlation to fuse such output features better. More specifically, the output from the self-attention mechanism enhances the key words or phrases according to the context semantics of inputting texts yet lacks the representation enhancement from label-text correlation features, while the output from the label-attention mechanism enhances the key words or phrases according to the label semantics yet lacks the representation enhancement from text semantic features. Therefore, with the guidance of a GCN, we aim to acquire the complementary feature vectors of these two representations. We use the proposed compensation coefficients guided by the GCN to quantify the extent of the compensation above. First, we calculate the cross-attention weights, denoted by $\mathbf{W}_t, \mathbf{W}_s \in \mathbb{R}^L$, which indicate the compensation coefficients of each representation. The model's output can be described as follows:

$$\begin{aligned}
\mathbf{W}_l &= f(\mathbf{C}^{\text{out}} \mathbf{Q}_s^T \mathbf{W}_{a1}), \\
\mathbf{W}_s &= f(\mathbf{C}^{\text{out}} \mathbf{Q}_l^T \mathbf{W}_{a2}), \\
\mathbf{W}_l + \mathbf{W}_s &= \mathbf{1},
\end{aligned} \tag{6}$$

where $\mathbf{W}_{a1}, \mathbf{W}_{a2} \in \mathbb{R}^L$ are parameters to be trained, f represents the sigmoid function, the third equation is to let \mathbf{W}_l and \mathbf{W}_s satisfy the normalization constraint, and $\mathbf{1}$ represent an all-one vector. Then, according to the compensation coefficients, the i -th label based final text representation can be obtained as $\mathbf{Q}_i = \mathbf{W}_{li} \mathbf{Q}_{li} + \mathbf{W}_{si} \mathbf{Q}_{si}$. The final text representation output by the proposed model is $\mathbf{Q} = \{\mathbf{Q}_i\}_{i=1}^L \in \mathbb{R}^{L \times 2d_2}$.

3.5 Label probability prediction

After obtaining the fused text representation, we feed \mathbf{Q} into a fully connected layer for the label probability prediction to produce the prediction result $\hat{y} = f(\mathbf{Q} \mathbf{W}_o)$, where f represents the sigmoid function and $\mathbf{W}_o \in \mathbb{R}^{2d_2}$ are the parameters to be trained.

After comparing the predicted labels \hat{y} with the ground-truth $y \in \{0, 1\}^L$, the proposed model is trained with the cross entropy loss as follows:

$$\mathcal{L} = \sum_{l=1}^L y_l \log(\hat{y}_l) + (1 - y_l) \log(1 - \hat{y}_l). \tag{7}$$

4 Experimental results

To evaluate our proposed model, we conducted numerous experiments on a public dataset of privacy-disclosing tweets and compared the performance of our model with selected state-of-the-art methods in terms of key performance metrics. Furthermore, we verified the effect of each component in our model with corresponding ablation tests and component analysis. Finally, we used our proposed model to test some concrete tweet examples to demonstrate the practicability of the proposed model.

4.1 Experimental setup

4.1.1 Dataset used

We evaluated our proposed model on the public dataset of privacy-disclosing tweets introduced in (Song et al. 2018), which includes 11,368 tweets each annotated with one or more privacy-relevant labels representing 32 privacy-oriented personal aspects. Figure 4 illustrates 32 categories of privacy in the dataset specifically. In the dataset, the personal privacy is firstly divided into eight groups, including “healthcare”, “life milestones”, “personal attributes”, “relationship”, “activities”,




Fig. 4 Illustration of 32 categories of privacy used in our experiments

Table 1 The number of tweets with a specific quantity of unique personal aspects, as annotated in the Twitter dataset

#(Personal aspects)	1	2	3	4	5	> 5
#(Tweets)	8546	2215	533	65	9	0
Percentage	75.2%	19.5%	4.7%	0.57%	0.079%	0%

“location”, “emotion” and “neutral statements”. the first seven groups represent seven general privacy groups and the last group “Neutral statements” represents those tweets that do not disclose any category of privacy. These eight groups make a higher-level categorization of privacy-related information, which covers most of personal privacy disclosures we can observe in the real world. Furthermore, the eight privacy groups are subdivided into 32 finer-grained privacy categories, which show different types of privacy-related information more specifically. Our experiments are based on 32 privacy-oriented personal aspects and each label represents one privacy-oriented personal aspect. To the best of our knowledge, no any other public datasets offer a comparable level of richness and comprehensiveness considering the size of the dataset and the richness of privacy-oriented personal aspects.

Table 1 shows the number of tweets with a specific quantity of unique personal aspects. An average tweet is annotated with 1.31 personal aspects.

4.1.2 Evaluation metrics

Following the settings of previous work (Chen et al. 2020), we use average precision (Avg-prec), one-error (One-err), precision at top K ($P@K$) and $S@K$ for performance evaluation, which are explained as follows:

Average precision (Avg-pre) Average precision evaluates the overall precision of the input texts over the ranking list of labels according to the ground truth (Nguyen et al. 2013).

One-error (One-err) One-error represents the mean possibility that the first prediction of the personal aspects does not conform to the ground truth (Zhang and Zhou 2007).

$P@K$ $P@K$ refers to the average precision of label predictions among the top K recommended results.

$S@K$ $S@K$ refers to the mean probability that a correct personal aspect is captured within the top K recommended results (Song et al. 2018).

4.1.3 Parameter settings

For fair comparisons, we split the dataset in our experiments in the same way as in previous work (Song et al. 2018; Chen et al. 2020). The experimental results were obtained through the 10-fold cross-validation.

We split the training set into a training subset and a validation subset whose ratio is 8:1. We selected the best parameter configuration based on the validation performance, i.e., the hyper-parameter fine-tuning was completed based on evaluation metrics calculated from the validation subset. To obtain the word embedding and label embedding, we utilized the word2vec method to convert texts into 300 dimensional vectors, which means $d_1 = 300$. The BiLSTM hidden dimension is set as $d_2 = 300$. The hyper-parameter corresponding to the self-attention mechanism is set as $d_3 = 200$. Furthermore, our model's GCN uses a 2-layer model with the hidden dimension of 450. The batch size searched are 16, 32, 64, and 128, and the learning rate searched are 0.1, 0.01, 0.001, and 0.0001. According to the validation performance, we took 64 as the batch size, and used the Adam optimizer (Kingma and Ba 2015) to minimize the loss with the initial learning rate of 0.001. We use the Floating-Point Operations (FLOPs) and Multiply-Accumulates (MACs) to measure the computational complexity of the proposed model. The experimental results indicate that the FLOPs of the proposed model is 12.61G and the MACs of the proposed model is 1.59M.

4.2 Baseline models

First, we compared our proposed model with several methods for predicting privacy disclosures in online posts, including five shallow learning methods and four deep

learning methods. To further demonstrate our proposed method's performance, we compared it with two recent state-of-the-art MLTC models. Therefore, we used the following eleven models as baselines.

- *SVM* (Cortes and Vapnik 1995): A classical machine learning model that concatenates the privacy-oriented features into a single vector and learns each personal aspect individually.
- *MTL-Lasso* (Tibshirani 1996): A multi-task learning method (MTL) with Lasso which implements the l_1 -penalization to the regression objective function.
- *GO-MTL* (Kumar and Daumé III 2012): A model using grouping and overlap mechanism to learn the semantic correlations among personal aspects.
- *CMTL* (Jacob et al. 2008): The clustered multi-task learning (CMTL) which assumes personal aspects can be clustered into several groups and each group can be learned together.
- *TOKEN* (Song et al. 2018): The latent group MTL that utilizes the pre-defined personal aspect taxonomy to learn the group-sharing and aspect-specific latent features of personal aspects simultaneously.
- *TextRNN* (Giles et al. 1994): A RNN-based model which uses RNN and logistic regression for privacy disclosure detection.
- *TextCNN* (Kim 2014): A CNN-based model which also uses CNN and logistic regression (similar to TextRNN) for privacy disclosure detection.
- *D-TOKEN* (Song et al. 2018): An end-to-end model as an extension of TOKEN, which replaces the hand-crafted features by representation automatically learned by hierarchical attentive network (HAN).
- *GrHA* (Chen et al. 2020): A HAN-based privacy detection model which uses graph-regularization mechanism to enhance label co-occurrences representations.
- *LSAN* (Xiao et al. 2019): A label-specific attention network model based on self-attention and label-attention mechanism.
- *LDGN* (Ma et al. 2021): A label-specific dual graph network model which contains label-attention and dual graph neural network.

4.3 Experimental results and discussion

Table 2 shows the performance metrics of all the compared methods, all based on the same dataset. For LSAN and LDGN, the two most recent baseline models, the experimental results were obtained from our own experiments. For other baseline models, the performance figures were taken from Chen et al. (2020), which were obtained using the same dataset and experimental settings as we used. The results show that our method outperformed all other baseline models, proving the effectiveness of the double-attention mechanism and the GCN-assisted feature fusion mechanism.

For all the evaluated models, deep learning methods are proved to access better results than shallow learning methods, which shows the importance of neural network on extracting text's features. Among all the deep models, TextRNN, TextCNN,

Table 2 Performance comparisons with selected state-of-the-art methods on the dataset used. Partial experimental results of baseline models are directly extracted from Chen et al. (2020)

Model		Avg-pre	One-err	S@1	S@3	S@5	P@1	P@3	P@5
Shallow	SVM	52.91%	69.35%	30.65%	72.98%	80.47%	30.65%	26.33%	18.47%
	MTL-Lasso	58.00%	56.09%	43.91%	73.18%	82.11%	43.91%	27.38%	19.31%
	GO-MTL	58.68%	56.02%	43.98%	74.24%	83.92%	43.98%	27.65%	19.78%
	CMTL	58.99%	55.84%	44.16%	74.41%	83.30%	44.16%	27.81%	19.63%
	TOKEN	59.05%	55.96%	44.04%	74.72%	84.34%	44.04%	27.96%	19.92%
Deep	TextRNN	61.84%	49.61%	50.39%	74.50%	82.64%	50.39%	28.67%	19.90%
	TextCNN	69.31%	39.99%	60.01%	81.97%	88.40%	60.01%	31.77%	21.44%
	D-TOKEN	69.43%	39.96%	60.04%	83.39%	89.35%	60.04%	32.15%	21.56%
	GrHA	71.44%	39.17%	60.83%	85.83%	92.20%	60.83%	33.51%	22.60%
	LSAN	73.62%	37.40%	62.60%	88.32%	94.25%	62.60%	34.98%	23.50%
	LDGN	73.94%	36.85%	63.15%	88.36%	94.45%	63.15%	35.03%	23.57%
	Our model	74.30%	36.35%	63.65%	89.08%	95.00%	63.65%	35.23%	23.60%

D-TOKEN are less effective because those models only focus on the features of the text and ignore the relationship between text and labels. GrHA and LSAN improve the results to a certain extent, on account for using the attention mechanism to extract the texts' correlation. However GrHA ignores the label-to-text correlation and directly utilizes the GCN to introduce label co-occurrences rather than assisting the feature fusion process. LSAN does not consider the impact of labels' co-occurrence, which causes the adverse effects on final results. LDGN uses label-attention and dual graph neural network to make up the deficiency of co-occurrence for labels. However by comparing with LDGN and our proposed model, the latter outperforms because its methods for processing label-to-label correlation is based on the GCN-assisted feature fusion mechanism, which uses the compensation coefficients to guide the fusion of text representations, while LDGN only uses the dot product operation.

In conclusion, the proposed network outperforms shallow models, deep embedding models, label attention based models. The improvement of the proposed model demonstrates the effectiveness of the double attention mechanism and the proposed GCN-assisted feature fusion mechanism.


4.4 Ablation tests

A series of ablation tests were conducted to show the contribution of each module in the proposed network. Since the proposed model has three functional modules, the self-attention module (S), the label-attention module (L) and the GCN-assisted feature fusion module (G), in the ablation tests, we experimented all six possible combinations of the three modules: S, L, SL (which is effectively LSAN), SG, LG, and SLG (which is our model). Note that G cannot be used alone.

As Table 3 presents, Model LG outperformed Model L while Model SG outperformed Model S, which shows the function of the GCN-assisted feature fusion

Table 3 Ablation tests of our proposed method using six different possible combinations of the three key components

Model	Avg-pre	One-err	S@1	S@3	S@5	P@1	P@3	P@5 (%)
S	72.37%	37.74%	62.26%	85.72%	92.52%	62.26%	33.90%	22.88
L	73.25%	37.66%	62.34%	87.90%	94.22%	62.34%	34.75%	23.31
SG	72.39%	37.73%	62.27%	85.71%	92.53%	62.27%	33.91%	22.90
LG	73.32%	37.62%	62.38%	87.95%	94.20%	62.38%	34.77%	23.32
SL	73.62%	37.40%	62.60%	88.32%	94.25%	62.60%	34.98	23.50
SLG	74.30%	36.35%	63.65%	89.08%	95.00%	63.65%	35.23	23.60


**Fig. 5** The visualization of label attention weights

module. Meanwhile aforementioned improvement is slight, which indicates that the GCN-assisted feature fusion module can exhibit its maximum function only with double attention mechanism. Model SL performed better than Models LG and SG, which indicates that the text representation is still the core process of the privacy MLTC. Model LG outperformed Model SG, which demonstrates that the label-attention mechanism can capture the feature of texts and labels more effectively and more accurately than the self-attention mechanism. Our proposed model (SLG) gained the best performance for all metrics, showing that combining all the three sources of information is indeed effective.

4.5 Component analysis

To further illustrate the performance of the proposed model, we conducted some further analysis for each component of our proposed model and present several samples selected from the privacy dataset we used.

Fig. 6 The visualization of labels' co-occurrence adjacency matrix (O = "occupation", S = "salary", G = "graduation", E = "education", C = "career promotion", and P = "passing away of relatives")



4.5.1 Label attention weights

We can use heat maps to show the label attention weights. For several test samples from the test set of our dataset, such a heat map is shown in Fig. 5. The brightness of the red bar represents the label attention weight of each word (darker = larger weight), according to the double-attention mechanism. For example, the more significant words for the label "occupation" are "a coach". for the label "current location", the label attention mechanism focuses on names of places such as "Washington DC". Generally speaking, the label attention mechanism is capable of extracting important information in the input text and benefiting the subsequent classification module.

4.5.2 GCN-assisted feature fusion

To show the effectiveness of the GCN-assisted feature fusion visually, we can also use a heat map representing label co-occurrences. One example is given in Fig. 6, which shows that the label "occupation" correlates highly with the label "salary", and the label "graduation" correlates highly with the label "education". besides, the label "education" correlates with the label "graduation" to some extent. on the other hand, the label "passing away of relatives" is almost irrelevant to other labels due to their lack of semantic connections. The example demonstrates that the GCN-based model can extract label-to-label relationships with the graph structure quite effectively.

To provide further evidence of the effectiveness of our GCN-based method, we also compared the performance of two groups of distinct GCN-based modules: our proposed GCN-assisted feature fusion module and the more common dot-product-based GCN modules. For the latter, we considered three possible modules: **Dot-S**—the dot-product-based model with self attention only, **Dot-L**—the dot-product-based model with label attention only, and **Dot-SL**—the dot-product-based model with double attention. The comparison results are shown in Table 4, which shows that our proposed GCN-based module outperformed

Table 4 The performance comparison of models based on our GCN-based module and three dot-product-based modules

Model	Avg-pre	One-err	S@1	S@3	S@5	P@1	P@3	P@5
Dot-S	72.45%	38.68%	61.32%	87.19%	93.59%	61.32%	34.11%	23.20%
Dot-L	73.01%	37.72%	62.28%	87.30%	93.45%	62.28%	34.59%	23.22%
Dot-SL	73.57%	37.18%	62.82%	88.08%	93.85%	62.82%	34.91%	23.25%
Our model	74.30%	36.35%	63.65%	89.08%	95.00%	63.65%	35.23%	23.60%

Table 5 The evaluation of performance on different numbers of GCN layers

Model	Avg-pre	One-err	S@1	S@3	S@5	P@1	P@3	P@5
GCN-1	73.78%	37.46%	62.54%	88.10%	94.85%	62.54%	34.97%	23.58%
GCN-2	74.30%	36.35%	63.65%	89.08%	95.00%	63.65%	35.23%	23.60%
GCN-3	73.90%	37.29%	62.71%	88.85%	94.76%	62.71%	35.08%	23.50%
GCN-4	73.53%	38.36%	61.64%	87.72%	94.05%	61.64%	34.93%	23.58%
GCN-5	73.41%	38.69%	61.31%	87.50%	94.28%	61.31%	34.84%	23.59%

all the other three dot-product-based modules. Compared with the dot-product-based modules, our module utilizes the label-to-label correlation matrix to guide the fusion of the output from the double-attention network, which can gain a better text representation.

4.5.3 Number of GCN layers

The performance of a GCN will differ depending on the number of GCN layers. In order to study how the number of layers affect the performance, we conducted some additional experiments with 1, ..., 5 GCN layers, represented by GCN-1, ..., GCN-5, respectively. Table 5 shows the results, which show that the model with two GCN layers achieved the best classification result. In comparison, the model with only one GCN layer showed the worse performance, which can be explained by the too shallow GCN being unable to extract label-to-label correlation effectively. The model's performance dropped while the number of GCN layers increases after two. This is likely caused by overfitting since a too deep GCN may learn about label-to-label correlation too specifically, therefore harming its generalizability. Based on the results, we recommend using two GCN layers for our model.

4.6 Case study

To demonstrate the practical usefulness of our proposed model, we use several example tweets (not included in the dataset) to demonstrate the effect of the model. To avoid potential privacy disclosures by us, we only use anonymous tweets for this part. For better illustration, the tweets tested try to cover multiple common privacy

Table 6 Case study

Testing tweets	Privacy aspects
I have been honored to serve the students in the district that I graduated from for 12 years as a teacher and administrator	Occupation, Graduation
Guys up to now my girlfriend is still looking for my birthday gift and its now two months since then	Gender
My trip to Washington D.C. will start in three days	Place planning to go
I am only 18 years old and I won't have enough to cover the removal surgery	Age, Health condition
Because I'm a Muslim so there's no way they would let me go to any concert	Religion
I have my bachelor's degree in computer science from Stanford university	Education background
Choosing who to be included in this project is too stressful	General complaint
The Warriors seem to be winning this year's championship	Neutral statement

categories. For clarity, we only present the tweets that are correctly classified by our models.

As Table 6 shows, we use several tweets to show the effect of our proposed model, including ten kinds of privacy aspects. For the first seven tweets, our model correctly captured the aspects of the privacy disclosure, which demonstrates the practicality of our proposed model. For example, the third tweet may disclose the travel destination of the user, thus the "place planning to go" as a reminder. The sixth tweet explains where the user obtained their bachelor's degree, so it may disclose the privacy category of "education background" according to our model. Therefore, Twitter users and the platform (Twitter) can use these kinds of reminders as a reference to avoid unintended privacy disclosures. For the last tweet, the tweet does not reveal any personal privacy aspect. Therefore, the tweet is classified into the category of "Neutral statement" by our model.

Furthermore, as Table 6 shows, if a tweet may disclose multiple categories of privacy information, our fine-grained privacy disclosure detection model can solve this problem with the consideration of multi-label classification, which shows the advantage of our model compared to other binary coarse-grained privacy disclosure detection models. For example, the detection results of the first testing tweet in Table 6 include two privacy aspects: "occupation" and "graduation", meanwhile the detection results of the fourth tweet include "age" and "health condition".

5 Conclusions and future work

A new privacy disclosure detection model is proposed in this paper. The proposed model integrates the text information, the label-to-text correlation and the label-to-label correlation for detecting privacy disclosures in the input text. For the first time, a GCN-assisted feature fusion mechanism is proposed to achieve the text feature fusion process with the guidance of the label-to-label correlation. During the process of feature fusion, the compensation coefficients are proposed to help fuse self-attention and label-attention features. Based on a dataset of privacy-disclosing

tweets, our experimental results showed that our model outperformed a number of selected state-of-the-art models and that the improved performance comes from the new design elements we introduced. A number of example tweets are used to demonstrate the practical usefulness of the proposed model. The results show that our proposed model can be used to support development of privacy protection tools that alert online users and online platforms about unintended privacy disclosures.

In our paper, our experiment are based on a single dataset covering 32 privacy-oriented personal aspects (Song et al. 2018), considering that this dataset is the best privacy-disclosing dataset we could find. However, using only one single dataset can make it difficult to judge how generalizable our results are. In addition, although the dataset we used covers a rich set of personal aspects, the coverage can still be extended to cover more personal aspects. Therefore, constructing more datasets for privacy disclosure detection is needed so our work can be further validated on multiple datasets. Meanwhile, our model aims to detect the privacy disclosure in text-only UGC. However, non-textual information in UGC such as images and videos can often disclose privacy information, too. Thus, in our future work, we will investigate the construction of a multi-modal privacy disclosure detection model supporting both visual and textual information.

Funding Shujun Li's work was partly funded by the research project "PRiVacy-aware personal data management and Value Enhancement for Leisure Travellers" (PriVELT, <https://privelt.ac.uk/>), funded by the EPSRC (Engineering and Physical Sciences Research Council), part of the UKRI (UK Research and Innovation), under the grant number EP/R033749/1. Also this work was partly funded by the National Natural Science Foundation of China under the reference number 61972249.

Declarations

Conflict of Interest The authors have no financial or proprietary interests in any material discussed in this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adhikari A, Ram A, Tang R, et al. (2019) DocBERT: BERT for document classification. *arXiv:1904.08398* [cs.CL]. <https://doi.org/10.48550/arXiv.1904.08398>
- Biega AJ, Roy RS, Weikum G (2017) Privacy through solidarity: A user-utility-preserving framework to counter profiling. In: Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval. ACM, pp 675–684. <https://doi.org/10.1145/3077136.3080830>
- Chen G, Ye D, Xing Z, et al. (2017) Ensemble application of convolutional and recurrent neural networks for multi-label text categorization. In: Proceedings of the 2017 international joint conference on neural networks. IEEE, pp 2377–2383. <https://doi.org/10.1109/IJCNN.2017.7966144>

- Chen X, Song X, Ren R et al. (2020) Fine-grained privacy detection with graph-regularized hierarchical attentive representation learning. *ACM Trans Inf Syst* 38(4):37:1-37:26. <https://doi.org/10.1145/3406109>
- Cortes C, Vapnik V (1995) Support-vector networks. *Mach Learn* 20:273–297. <https://doi.org/10.1007/BF00994018>
- Eslami S, Biega AJ, Saha Roy R, et al. (2017) Privacy of hidden profiles: Utility-preserving profile removal in online forums. In: *Proceedings of the 2017 ACM conference on information and knowledge management*. ACM, pp 2063–2066. <https://doi.org/10.1145/3132847.3133140>
- Giles CL, Kuhn GM, Williams RJ (1994) Dynamic recurrent neural networks: theory and applications. *IEEE Trans Neural Netw* 5(2):153–156. <https://doi.org/10.1109/TNN.1994.8753425>
- Huang X, Paul MJ (2019) Neural user factor adaptation for text classification: Learning to generalize across author demographics. In: *Proceedings of the 8th joint conference on lexical and computational semantics*. ACL, pp 136–146. <https://doi.org/10.18653/v1/S19-1015>
- Jacob L, Vert JP, Bach F (2008) Clustered multi-task learning: a convex formulation. In: *Proceedings of the 22nd annual conference on neural information processing systems*. Curran Associates, Inc., pp 745–752. <https://proceedings.neurips.cc/paper/2008/hash/fccb3cdc9acc14a6e70a12f74560c026-Abstract.html>
- Kim Y (2014) Convolutional neural networks for sentence classification. In: *Proceedings of the 2014 conference on empirical methods in natural language processing*. ACL, pp 1746–1751. <https://doi.org/10.3115/v1/D14-1181>
- Kingma DP, Ba J (2015) Adam: a method for stochastic optimization. In: *Proceedings of the 3rd international conference for learning representations*. <https://doi.org/10.48550/arXiv.1412.6980>
- Kipf TN, Welling M (2016) Semi-supervised classification with graph convolutional networks. In: *Proceedings of the 5th international conference on learning representations*. OpenReview. <https://openreview.net/forum?id=SJU4ayYgl>
- Kumar A, Daumé III H (2012) Learning task grouping and overlap in multi-task learning. In: *Proceedings of the 29th international conference on machine learning*. ICML. <https://icml.cc/2012/papers/690.pdf>
- Kurata G, Bing X, Zhou B (2016) Improved neural network-based multi-label classification with better initialization leveraging label co-occurrence. In: *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*. ACL, pp 521–526. <https://doi.org/10.18653/v1/N16-1063>
- Le Q, Mikolov T (2014) Distributed representations of sentences and documents. In: *Proceedings of the 31st international conference on machine learning*. PMLR, pp 1188–1196. <https://proceedings.mlr.press/v32/le14.html>
- Lin Z, Feng M, Santos CNd, et al. (2017) A structured self-attentive sentence embedding. In: *Proceedings of the 5th international conference on learning representations*. OpenReview. https://openreview.net/forum?id=BJC_jUqxq
- Liu J, Chang WC, Wu Y, et al. (2017) Deep learning for extreme multi-label text classification. In: *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*. ACM, pp 115–124. <https://doi.org/10.1145/3077136.3080834>
- Liu P, Qiu X, Huang X (2016) Recurrent neural network for text classification with multi-task learning. In: *Proceedings of the 25th international joint conference on artificial intelligence*. IJCAI, pp 2873–2879. <https://www.ijcai.org/Proceedings/16/Papers/408.pdf>
- Ma Q, Yuan C, Zhou W, et al. (2021) Label-specific dual graph neural network for multi-label text classification. In: *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing*. ACL, pp 3855–3864. <https://doi.org/10.18653/v1/2021.acl-long.298>
- Mao H, Shuai X, Kapadia A (2011) Loose tweets: An analysis of privacy leaks on Twitter. In: *Proceedings of the 10th annual ACM workshop on privacy in the electronic society*. ACM, pp 1–12. <https://doi.org/10.1145/2046556.2046558>
- Nguyen C, Zhan D, Zhou Z (2013) Multi-modal image annotation with multi-instance multi-label LDA. In: *Proceedings of the 23rd international joint conference on artificial intelligence*. AAAI, pp 1558–1564. <https://www.ijcai.org/Proceedings/13/Papers/232.pdf>
- Raber F, Krüger A (2018) Deriving privacy settings for location sharing: Are context factors always the best choice? In: *Proceedings of the 2018 IEEE symposium on privacy-aware computing*. IEEE, pp 86–94. <https://doi.org/10.1109/PAC.2018.00015>

- Sanchez OR, Torre I, He Y et al. (2020) A recommendation approach for user privacy preferences in the fitness domain. *User Model User-Adap Inter* 30:513–565. <https://doi.org/10.1007/s11257-019-09246-3>
- Song X, Wang X, Nie L, et al. (2018) A personal privacy preserving framework: I let you know who can see what. In: *Proceedings of the 41st international ACM SIGIR conference on research & development in information retrieval*. ACM, pp 295–304. <https://doi.org/10.1145/3209978.3209995>
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J Roy Stat Soc: Ser B (Methodol)* 58:267–288. <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>
- Tran L, Kong D, Jin H, et al. (2016) Privacy-CNH: A framework to detect photo privacy with convolutional neural network using hierarchical features. In: *Proceedings of the 30th AAAI conference on artificial intelligence*. AAAI, pp 1317–1323. <https://doi.org/10.1609/aaai.v30i1.10169>
- Xiao L, Huang X, Chen B, et al. (2019) Label-specific document representation for multi-label text classification. In: *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing*. ACL, pp 466–475. <https://doi.org/10.18653/v1/D19-1044>
- Yang P, Sun X, Li W, et al. (2018) SGM: Sequence generation model for multi-label classification. In: *Proceedings of the 27th international conference on computational linguistics*. ICCL, pp 3915–3926. <https://aclanthology.org/C18-1330>
- Zhang M, Zhou Z (2007) Multi-label learning by instance differentiation. In: *Proceedings of the 2007 AAAI conference on artificial intelligence*, vol 7. AAAI, pp 669–674. <https://aaai.org/papers/00669-multi-label-learning-by-instance-differentiation/>
- Zhou P, Qi Z, Zheng S, et al. (2016) Text classification improved by integrating bidirectional LSTM with two-dimensional max pooling. [arXiv:1611.06639](https://arxiv.org/abs/1611.06639) [cs.CL]. <https://doi.org/10.48550/arXiv.1611.06639>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Zhanbo Liang¹  · Jie Guo¹ · Weidong Qiu¹ · Zheng Huang¹ · Shujun Li²

✉ Zhanbo Liang
doctorliang@sjtu.edu.cn

✉ Jie Guo
guojie@sjtu.edu.cn

✉ Shujun Li
s.j.li@kent.ac.uk

Weidong Qiu
qiuwd@sjtu.edu.cn

Zheng Huang
huang-zheng@sjtu.edu.cn

¹ School of Cyber Science and Engineering, Shanghai Jiao Tong University, 800 Dongchuan Road, Minhang District, Shanghai 200240, China

² Institute of Cyber Security for Society (iCSS) and School of Computing, University of Kent, Kent, Canterbury CT2 7NP, UK