

XBully: Cyberbullying Detection within a Multi-Modal Context

Lu Cheng
Arizona State University
lcheng35@asu.edu

Jundong Li
Arizona State University
jundongl@asu.edu

Yasin Silva
Arizona State University
ysilva@asu.edu

Deborah Hall
Arizona State University
d.hall@asu.edu

Huan Liu
Arizona State University
huanliu@asu.edu

ABSTRACT

Over the last decade, research has revealed the high prevalence of cyberbullying behaviors among preteen and teenage youth and raised serious concerns in society. While social media platforms such as Instagram, Facebook, and Twitter have been identified as the leading environments where cyberbullying behaviors most frequently occur, most existing research contributions are devoted to building a generic classification model solely based on text analysis of online social media sessions (e.g., posts). Despite its empirical success, most of the existing efforts ignore the multi-modal information manifested in social media data (e.g., image, video, user profile, time, and location), and thus fail to offer a comprehensive understanding of the occurrence of cyberbullying. Conventionally, when information from different modalities are presented together, they often reveal complementary insights about the application domain and allow us to achieve better learning performance. In this paper, we study the novel problem of cyberbullying detection within a multi-modal context by exploiting social media data in a collaborative way. This task, however, is challenging due to the complex combination of both cross-modal correlations among various modalities and structural dependencies between different social media sessions, and the diverse attribute information of different modalities. To address these challenges, we propose XBully, a novel cyberbullying detection framework, that first reformulates the multi-modal social media data as a heterogeneous network and then aims to learn node embedding representations upon it. Extensive experimental evaluations on real-world multi-modal social media datasets show that the XBully framework is superior to the state-of-the-art cyberbullying detection models.

KEYWORDS

Cyberbullying Detection, Multi-Modality, Social Media, Network Embedding, Collaborative Learning

ACM Reference Format:

Lu Cheng, Jundong Li, Yasin Silva, Deborah Hall, and Huan Liu. 2019. XBully: Cyberbullying Detection within a Multi-Modal Context. In *The Twelfth*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '19, February 11–15, 2019, Melbourne, VIC, Australia

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5940-5/19/02...\$15.00

<https://doi.org/10.1145/3289600.3291037>




Figure 1: Illustration of cyberbullying detection within a multi-modal context: the left side figure represents a social media session (e.g., post) with rich user-generated information such as an image, video, user profile, time, location and comments. In addition, different sessions are inherently connected with each other with user-user social relations. The goal is to predict if a particular session is about bullying or not by leveraging its multi-modal context information.

ACM International Conference on Web Search and Data Mining (WSDM '19), February 11–15, 2019, Melbourne, VIC, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3289600.3291037>

1 INTRODUCTION

Cyberbullying, which originates from the electronic transmission of insulting or embarrassing comments and photos, is prevalent on social networks as it allows perpetrators to reach victims stealthily and anonymously through online means. Cyberbullying can have serious negative effects on victims including depression, low self-esteem, and suicidal thoughts and behaviors [35]. In light of this, the American Psychological Association and the White House have identified cyberbullying as a serious national health concern, with more than 40% of teenagers in the US reporting that they have been bullied on social media platforms [7]. The growing prevalence and severity of cyberbullying in social media has motivated a surge of research in psychology and computer science.

Within the computer science community, existing efforts toward detecting cyberbullying have primarily focused on text analysis. These works attempt to build a generic binary classifier by taking high-dimensional text features as the input and make predictions accordingly. Despite their satisfactory detection performance in practice, these models inevitably ignore critical information included in the various social media modalities such as image, video, user profile, time and location. For example, Instagram¹ allows users to post and comment on any public image to express their opinions and preferences. Bullies then can post humiliating images, edit the images and re-post them, post insulting comments,

¹<https://www.instagram.com>

captions, or hashtags, and even create fake profiles pretending to be someone else [13]. Therefore, it is critical to exploit the rich user-generated content within a multi-modal context to gain better insights of cyberbullying behaviors and make more accurate predictions. Fig. 1 illustrates the cyberbullying detection problem within a multi-modal context.

Despite its potential benefits, performing cyberbullying detection within a multi-modal context presents multiple challenges. First, information of different modalities could complement each other in achieving better learning performance especially when the data is sparse. However, heterogeneous information of different modalities may not be compatible with each other and, in the worst case, some modalities may be independent. A key problem that has not been sufficiently addressed in cyberbullying detection is how to effectively encode the cross-modal correlation among different types of modalities. Second, social media data is often not independent and identically distributed (i.i.d.) but is intrinsically correlated, either directly or indirectly, limiting the applicability of conventional text analysis approaches. For example, if two social media sessions (e.g., posts) are from the same user or are posted by a pair of friends, their content similarity is expected to be high based on the homophily principle [20]. Considering this, it is important to model structural dependencies among different social media sessions when performing cyberbullying detection. Third, even though multi-modal social media data can be useful in understanding human behaviors, it is difficult to directly make use of it as different modalities are often associated with rather diverse feature types (e.g., nominal, ordinal, interval, ratio, etc.), and in some cases, some modalities that identify particular entities (e.g., users) cannot be simply represented as feature vectors². Therefore, it is important that the solution framework uses an expressive way to represent modalities with diverse feature types.

To address the above challenges, we propose a novel cyberbullying detection framework, XBully, that models multi-modal social media data in a collaborative way. Specifically, to capture cross-modal correlation among modalities as well as the structural dependencies among different social media sessions, we model the multi-modal social media data as a heterogeneous network by exploiting *co-existence* and *neighborhood* relations (explained later) and aim at learning the embedding representations for nodes in the network. Due to data sparsity, we identify a number of *hotspots* for each mode, which provide a succinct high-level summarization of similar modality attribute values. For nominals (modalities without attributes), we form nodes in the constructed heterogeneous network using their meta information, e.g., user IDs. After learning the embedding representation for nodes in the heterogeneous network, each social media session can be represented as a numerical vector by concatenating the node embeddings in that session. Using these vectors, various off-the-shelf machine learning models can be directly applied to provide accurate cyberbullying detection and deep understanding of cyberbullying behaviors. The main contributions of this work are:

- **Problem Formulation:** We formally define the problem of cyberbullying detection within a multi-modal context. The

definition is a result of multiple modalities exploited in a collaborative fashion.

- **Algorithms:** We propose a novel cyberbullying detection framework (XBully) with three core components: (1) a hotspot detector that identifies centroids for each mode; (2) a module that constructs a heterogeneous network by leveraging *co-existence* and *neighborhood* relations of the detected hotspots and instances of nominals; and (3) a principled joint embedding module that effectively encodes cross-modal correlation and structural dependencies among different social media sessions to learn noise-resilient embedding representations. The resulting embeddings enable better detection and understanding of cyberbullying behaviors.
- **Evaluations:** We perform experiments on two real-world social media datasets to corroborate the efficacy of the proposed framework.

2 PROBLEM DEFINITION

In this section, we first introduce the problem of cyberbullying detection within a multi-modal context, and then briefly describe our approach to solve it via network representation learning and highlight the challenges.

2.1 Multi-Modal Cyberbullying Detection

Definition 1. Cyberbullying Detection within a Multi-Modal Context *Given a corpus of social media sessions C (e.g., posts) with M modalities, cyberbullying detection within a multi-modal context aims at identifying instances of cyberbullying, by leveraging multiple modalities such as textual features, spatial locations, and visual cues as well as the relations among sessions.*

The definition of multi-modal cyberbullying detection builds on the concept of multi-modality learning in machine learning [2]. Here, we emphasize the multi-modal context information of social media sessions. In our experiments, we use the following modalities extracted from an Instagram session:

- **User** - It is a typical type of nominals and we use the relations among users to decode the dependencies between social media sessions.
- **Image** - The associated meta-information of an image forms a tuple composed of the number of shares, the number of likes, and the labels describing the category of this image.
- **Profile** - The meta-information of a user forms a tuple with the number of followers, the number of follows, the total number of comments, and the total number of likes she received.
- **Time** - The timestamp of posting an image. We consider the time of the day (24h range) and convert the raw time to the range of $[0, 86400]$ by calculating its offset (in second) w.r.t. 12:00 am.
- **Text** - We perform psychometric analysis on the textual information of the session, i.e., description of the image and comments, and obtain the psychological features through LIWC [24]. Previous research efforts showed that such psychometric analysis can provide insights about cyberbullying behaviors [22].

²We refer to modalities with attributes and without attributes as modes and nominals, respectively.

2.2 Cyberbullying Detection via Multi-Modal Network Representation Learning

Let C be a corpus of social media sessions. We define each session $s \in C$ as a tuple $\langle \mathbf{x}_{s1}, \mathbf{x}_{s2}, \dots, \mathbf{x}_{sM}, y_{s1}, \dots, y_{sN} \rangle$ where M and N denote the number of modes and nominals, respectively. \mathbf{x}_{sm} is the feature vector of s in mode $m \in (1, 2, 3 \dots M)$. For example, a geo-tagged social media session may have location information $\mathbf{x}_{sm} = [34.0489, 111.0937]$. In addition, different sessions are inherently connected with each other via social relations among users. Our goal is to represent the original corpus C as a heterogeneous network G by capturing its multi-modality nature, and learn high-quality embeddings for each node in this network.

In contrast to simply concatenating the raw multi-modal feature vectors of each modality, the learned node embeddings in the resultant heterogeneous network capture both *structural dependencies* among different social media sessions and *cross-modal correlations* among different modalities in a joint framework.

2.3 Challenges

- **Number of Distinct Feature Values.** Social media data usually comes in complex forms and exhibits considerable variations due to its multi-modality nature. We are often confronted with diverse feature types and the number of unique feature values each mode can take is often exceedingly large, which can cause the problem of data sparsity. Furthermore, the available training data for each node in the network is often limited which further complicates the training process.
- **Cross-Modal Correlation and Structural Dependencies.** An effective network embedding model for multi-modal social media data should preserve the node proximities in terms of both cross-modal correlation and structural dependencies among different sessions. Conventional network embedding models such as Deepwalk [25], LINE [31] and node2vec [10] which primarily focus on encoding structural information of homogeneous networks cannot be effectively applied in our problem. Metapath2vec [9] is a recently proposed heterogeneous network embedding model which relies on a set of predefined meta-paths to find the neighborhoods around nodes. However, in our problem, the number of meta-paths is often very large, making metapath2vec inapplicable.
- **Information Noise** While the rich multi-modal data can provide valuable and complementary insights to identify cyberbullying behaviors, such data can also be cluttered and noisy, thus complicate the process of gaining actionable knowledge from it.

To address these challenges, we propose modes hotspots detection using kernel density estimation (KDE) [23] to reduce the number of unique feature values. Relying on the detected hotspots, we then construct a heterogeneous network by leveraging both *co-existence* relations to exploit the connections among different modalities in the same session, and *neighborhood* relations to connect nodes of the same modality in different sessions. Next, we develop a graph-based joint embedding module to capture the cross-modal correlation and structural dependencies in a joint framework. This embedding module maps all the modes hotspots and nominal nodes

in a heterogeneous network into a common latent space. To alleviate the negative impacts of noise, we also identify the most informative neighbors for each node in the network to refine the learned embeddings. The overall framework is further explained in Fig. 2.

3 THE XBULLY FRAMEWORK

This section presents the proposed XBully model in detail. Specifically, we first show how to identify succinct yet accurate summarizations of groups of similar feature values (mode hotspots). Then, we present a principled way to capture both cross-modal correlation and structural dependencies in a joint framework for embedding representation learning. We also discuss how to alleviate the negative impact of noise during the embedding training phase by allowing nodes to borrow strength from each other in a collaborative way.

3.1 Mode Hotspots Detection

Previous work has shown that high-dimensional feature representation not only suffers from the data sparsity issue but also poses great challenges to downstream learning tasks due to the curse of dimensionality [18]. To address this issue, we propose the concept of mode hotspots based on KDE, which is a non-parametric method to estimate the density function from a collection of data samples. With KDE, we do not need to establish any prior knowledge about the data distribution, as it provides automatic discovery of arbitrary modes from complex data spaces [36].

Definition 2. Mode Hotspots Given a corpus of social media sessions C , the mode hotspots for mode m ($m = 1, 2, \dots, M$) are the set of local maximums of the kernel density function estimated from m .

Then, given n sessions containing mode m in a d -dimensional feature space $X_m = (\mathbf{x}_{1m}, \mathbf{x}_{2m}, \dots, \mathbf{x}_{nm})$, the kernel density at any point \mathbf{x} with mode m is given by:

$$f(\mathbf{x}) = \frac{1}{n\delta_m^d} \sum_{i=1}^n K\left(\frac{\mathbf{x} - \mathbf{x}_{im}}{\delta_m}\right), \quad (1)$$

where $K(\cdot)$ is a predefined kernel function, and δ_m is the kernel bandwidth for mode m . We further leverage the advanced mean-shift algorithm proposed in [36] to identify the mode hotspots.

3.2 Network Representation Learning

Now we investigate how to build a heterogeneous network by exploiting the *co-existence* and *neighborhood* relations, such that both the cross-modal correlations and structural dependencies are properly captured. The *co-existence* relation is established between two nodes when they co-exist in the same social media session and the *neighborhood* relations among mode hotspots are built upon the idea of *modality continuity* [32] which implies that nearby things are more related with each other than distant things. We first define the node kernel, based on which the *neighborhood* relations are formed:

Definition 3. Node Kernel For two mode hotspots u_i and u_j in mode m with feature vectors \mathbf{x}_i and \mathbf{x}_j , the kernel strength between

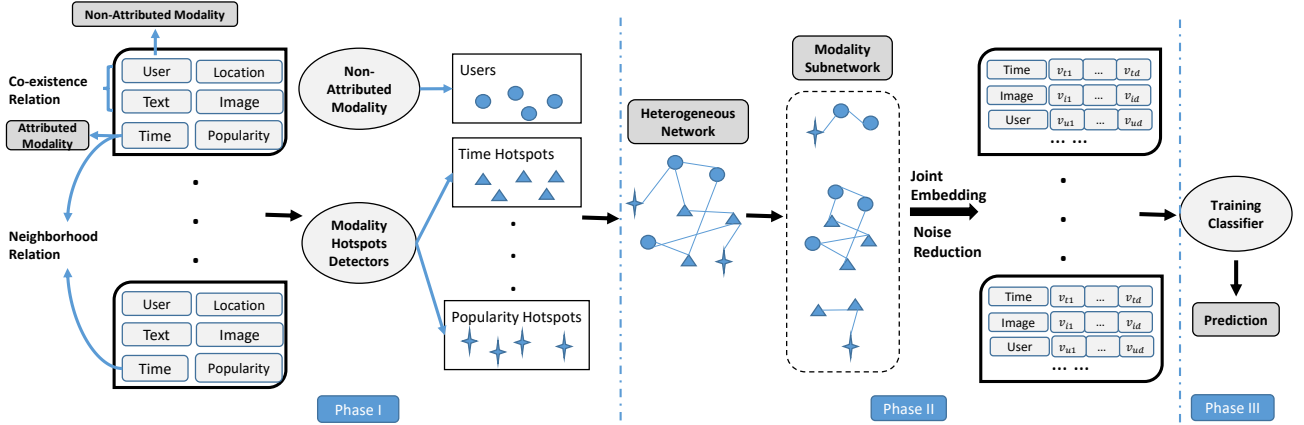


Figure 2: The proposed XBully framework. Given a corpus of social media sessions, we first attempt to discover hotspots for each mode (Phase I); and then based on the detected hotspots and instances of nominals, we leverage the *co-existence* and *neighborhood* relations to construct a heterogeneous network, which is later divided into several modality subnetworks (Phase II). Each subnetwork consists of two modalities. Nodes in these subnetworks are then mapped into the same latent space through network representation learning. Finally, we can concatenate embeddings of nodes in each session and apply off-the-shelf machine learning models for cyberbullying detection (Phase III).

them is :

$$w(u_i, u_j) = \begin{cases} \frac{\exp(-\|x_i - x_j\|^2 / 2\delta_m^2)}{2\pi\delta_m^2}, & \text{if } \|x_i - x_j\| \leq \delta_m \\ 0, & \text{otherwise,} \end{cases}$$

Therefore, the neighbors of a mode hotspot v in the heterogeneous network are the set of mode hotspots that produce a *non-zero* kernel strength value with hotspot v . In addition to that, for nominal nodes, we define the *neighborhood* relations based on its structural information by making use of the dependencies (e.g., social relations) between different sessions. For example, an Instagram session could have five different modalities - user, image, profile, time, and comments (text). From the definition of *co-existence* relations, we construct the following 10 types of edges in the heterogeneous network: *user-image*, *user-profile*, *user-time*, *user-text*, *image-profile*, *image-text*, *image-text*, *profile-time*, *profile-text* and *time-text*. Moreover, the *neighborhood* relations also generate the following 4 edge types: *image-image*, *profile-profile*, *time-time* and *text-text* by Definition 3 and edge *user-user* with nominal nodes is built upon social relations among users. With the above defined edge types, we define the weight of an edge considering the following three scenarios: (a) the normalized co-existence count (between 0 and 1); (b) kernel strength (between 0 and 1); and (c) the dependencies between nominal nodes (0 or 1). As the resultant network has different types of nodes and edges, it would be inappropriate to directly apply a conventional network embedding algorithm such as Deepwalk [25] or node2vec [10] to learn the embeddings for each node. Instead, we build on [30] to decompose the heterogeneous networks into multiple modality subnetworks (with two modalities) and learn embeddings within each subnetwork. In this approach, the learned embeddings can capture the node proximity across different types of edges. In what follows, we provide the details of the joint embedding model.

First, let us denote G_S as the set of all modality subnetworks, then for any two different modalities $A, B \in (1, 2, \dots, M+N)$, we can construct a modality subnetwork $G_{AB} \in G_S$. Then the probability of node j with modality B generated from node i with modality A is now defined by the following conditional probability:

$$p(j|i) = \frac{\exp(v_j^T \cdot v_i)}{\sum_{k \in B} \exp(v_k^T \cdot v_i)}, \quad (2)$$

where v_j denotes the embedding representation of node j with modality B and v_i is the embedding vector of node i with modality A . Next, we learn embeddings by minimizing the distance between the conditional distribution of the context nodes given the center node and the empirical distribution. The empirical distribution of node i is defined as $p'(j|i) = \frac{w_{ij}}{d_i}$, where w_{ij} is the weight of the edge $i-j$ and d_i is the out-degree of node i , i.e., $d_i = \sum_{j \in B} w_{ij}$. Therefore, we define the loss function as follows:

$$O_{AB} = \sum_{i \in A} d_i \text{KL}(p'(\cdot|i) || p(\cdot|i)), \quad (3)$$

where the $\text{KL}(\cdot)$ denotes the KL-divergence between two probability distributions. By omitting the constants, the above loss function can be reformulated as follows:

$$O_{AB} = - \sum_{i \in A, j \in B} w_{ij} \log p(v_j | v_i). \quad (4)$$

As each modality subnetwork is composed of four different types of edges, $A-A$, $A-B$, $B-A$, and $B-B$, the overall loss function of a modality subnetwork G_{AB} is as follows:

$$Z_{AB} = O_{AA} + O_{AB} + O_{BB} + O_{BA}. \quad (5)$$

3.3 Embedding Refinement

While multi-modal information is useful in improving the embedding quality, the above model can be problematic when the network

is very sparse. In addition, when the discovered mode hotspots are very noisy, the embedding representation learning phase may be adversely affected. To address these problems, we propose a noise-resilient embedding refinement approach to *adaptively* choose the most informative neighbors for each node. The core idea of the refinement method is to find the best locally weighted context vectors (predictors) to reconstruct the embedding of the center node. Specifically, given n embedding vectors $v_1, \dots, v_j, \dots, v_n \in \mathbb{R}^d$, the problem is to estimate \hat{v}_i using an estimator of the form $\hat{v}_i = \sum_{j=1, i \neq j}^n \alpha_{ji} v_j$, s.t. $\sum_j \alpha_{ji} = 1$, where α_{ji} denotes the extent to which the embedding v_j is influenced by v_j . Our solution builds on the algorithm presented in [1] to adaptively learn the optimal neighborhood structure for each center node and automatically quantify the influence from other nodes. This can be formulated as follows:

$$R_i = \left\| v_i - \sum_{j=1, i \neq j}^{|V_{AB}|} \alpha_{ji} v_j \right\|_2^2, \quad i \in V_{AB}, \quad (6)$$

$$\text{s.t. } \sum_{j=1, j \neq i}^{|V_{AB}|} \alpha_{ji} = 1$$

where V_{AB} represents the node set in modality subnetwork G_{AB} . By integrating the above embedding refinement component, the new objective function for embedding representation learning is:

$$Z_{AB} = O_{AA} + O_{AB} + O_{BA} + O_{BB} + \lambda \sum_{i=1}^{|V_{AB}|} R_i, \quad (7)$$

where λ is a parameter that balances the contribution of the refinement component. As a result, the overall objective function of our multi-modal network embedding is now:

$$O = \sum_{G_S} Z_{AB}, \quad G_{AB} \in G_S. \quad (8)$$

We employ the stochastic gradient descent (SGD) to optimize the above objective function by alternating between different modality subnetworks with negative sampling [21]. Specifically, for an edge e_{ij} , we randomly select K nodes that are not connected with node i as negative samples.

4 EXPERIMENTS

In this section, we aim to answer the following research questions: (1) Is the proposed XBully framework superior to existing models that solely rely on text information for cyberbullying detection? (2) How effective is the noise-resilient embedding refinement component for embedding representation learning? (3) Does the proposed multi-modal network embedding method help achieve better detection performance than those of conventional network embedding methods? (4) What kind of insights can XBully provide for social scientists and psychologists? (5) How robust is the proposed model w.r.t the different model parameters? Each experiment was run 10 times to report the averaged experimental results.

4.1 Experimental Setup

Datasets. Our experiments are performed using two real-world social media datasets³. Each social media session in the Instagram

³Available at <https://sites.google.com/site/cybersafety/home/cyberbullying-detection-project/dataset>

Table 1: Dataset statistics.

Datasets	# Sessions	# Bullying	# Normal	# Comments
<i>Instagram</i>	2,218	678	1,540	155,260
<i>Vine</i>	970	304	666	78,250

dataset [13] includes the descriptions of images, user comments, and the creation time of the session. The dataset also has user profile information and social relations among users. The second dataset [26] was collected from Vine, a mobile application website that allows users to record and edit six-second looping videos. Each Vine session is associated with video descriptions, user comments, and the creation time of the session. Basic statistics of these datasets are shown in Table 1. Please refer to [14, 26] for more details.

Baseline Methods. To answer the first two research questions, we compare XBully with the commonly used feature engineering approach, two recently proposed cyberbullying models, and a variant of XBully without the noise-resilient embedding refinement component.

- **Raw Features (Raw):** This is a concatenation of all the multi-modal features such as network feature and text feature.
- **Bully** [35]: A pretrained classifier⁴ based on textual analysis.
- **SICD** [6]: The state-of-the-art cyberbullying detection model which uses sentiment information embedded in the user-generated content to guide the prediction.
- **XBully variant (Variant):** A variant of XBully without the noise-resilient embedding refinement component.

We also compare the proposed model with three widely used network embedding models - DeepWalk [25], Node2vec [10] and GraRep [3]. To reduce the effect of model variances on performance evaluation, we test these methods on three classification models, including *Random Forest*, *Linear SVM* and *Logistic Regression*. Multiple training datasets are generated by extracting increasing fractions (10% to 90%) of the entire datasets and the remaining parts are used as the test datasets.

Parameter Settings. The XBully framework has the following parameters: (1) the hotspot detection bandwidth h_m for each mode $m \in M$; (2) the weight of the noise-resilient component λ ; and (3) the embedding dimension d . By default, we set parameters for the Instagram dataset as $h_t = 150$ (time), $h_i = 100$ (image), $h_l = 100$ (LIWC), $h_u = 500$ (user profile), $\lambda = 0.01$ and $d = 500$. For the Vine dataset, we set $h_t = 50$ (time), $h_l = 5$ (LIWC), $\lambda = 0.01$ and $d = 500$. A detailed parameter sensitivity analysis is presented later in this section.

4.2 Performance Comparisons

To answer the first three research questions, we generated training datasets by extracting increasing fractions (10% to 90%) of the overall dataset and used rest of the datasets as the testing datasets. Two common evaluation metrics are calculated - Macro F1 (Mac F1) and Micro F1 (Mic F1). A macro-average computes the metric independently for each class and then take the average as the output, whereas a micro-average will aggregate the contributions of

⁴<http://research.cs.wisc.edu/bullying/>

all classes to compute the average metric. In binary classification, Micro F1 is equal to Accuracy. Table 2 and Table 3 report the cyberbullying detection performance of various methods on the two datasets. We make the following observations from these results:

- The proposed XBully framework significantly outperforms the method that works with concatenated multi-modal features (Raw) and the methods with homogeneous network embedding (Deepwalk, Node2vec and GraRep) on both datasets in most cases. The improvement of XBully over these methods shows the effectiveness of performing cyberbullying detection with multi-modal network representation learning as it captures both the cross-modal correlations and structural dependencies.
- XBully is also superior to two cyberbullying detection methods *Bully* and *SICD* regarding to Macro F1. Even if XBully is worse than *SICD* in some cases regarding to Micro F1 in the Vine dataset, Macro F1 in the same setting (70%-30%) is 49% higher. This shows that multi-modal information can indeed provide complementary insights to achieve better learning performance.
- The improvements of XBully over baseline methods are consistent across different classifiers. This indicates that the learned embedding representations are effective and can be easily generalized to various off-the-shelf machine learning models.
- XBully achieves better detection performance than the variant without the embedding refinement component. This result highlights the benefit of collaboratively refining the embeddings by integrating information from similar nodes during the learning process, which in turn makes the learned embedding representation more robust to noise.

4.3 Qualitative Analysis

We are also interested in intuitive findings that might provide some insights about common social behaviors of users who have experienced cyberbullying versus users who have not. To this end, we interpret the confidence level of each label in the datasets as the possibility of the session being cyberbullying⁵. To understand how XBully can provide insights for social scientists and psychologists, we incorporated the new type of node into the previous model by treating it as another modality and retrained all of the embeddings. Afterwards, we made queries w.r.t. p in the range of (0.1, 0.2, ..., 1.0) and XBully returns the most similar mode hotspots based on the *cosine* similarity. The question we aim to answer is how user behavior in social media platforms vary against probability of cyberbullying. The following analysis is based on the experiments on the Instagram dataset.

As shown in Fig. 3(a), *#followers* gets larger as p increases, i.e., the number of users a session owner follows increases when the probability of cyberbullying increases. This may indicate that users who are active in social media have a higher probability of experiencing cyberbullying. In Fig. 3(b), the shape of the distribution of *#followers* at different probability levels of cyberbullying approximately follows a normal distribution. This pattern seems to imply that the difference between *#followers* of users who experienced and

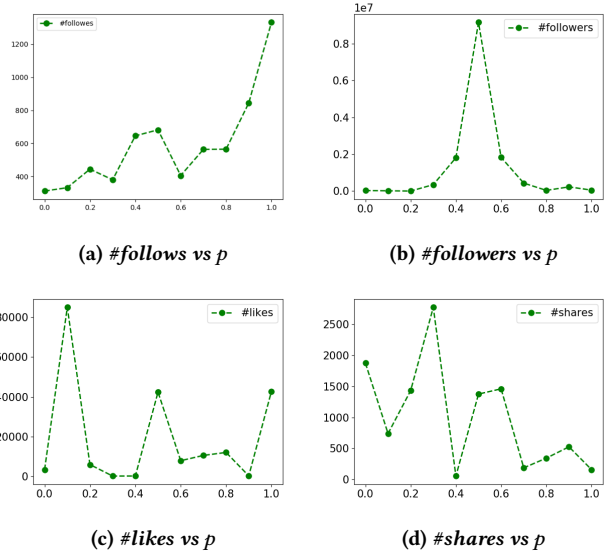


Figure 3: Qualitative analysis

who did not experience cyberbullying is small. Fig. 3(c)-(d) shed light on the relationships between the popularity of a social media session and p , as indicated by the number of likes a post receives in Fig. 3(c) and the number of times a post is shared in Fig. 3(d). In Fig. 3(c), three peaks can be seen, at $p = 0.1, 0.5$, and 1.0 , and *#likes* at $p = 0.1$ is twice as large as those at $p = 0.5$ and 1.0 . In Fig. 3(d), *#shares* decreases as p gets larger. A possible explanation might be that most social media users are normal users who are not specifically interested in cyberbullying-related content. Although tentative, trends in Fig. 3 elucidate a potentially novel way for interdisciplinary researchers to measure social influence within the context of social media interactions, particularly as they relate to cyberbullying risk.

4.4 Parameter Study

The XBully framework has four parameters (h_t, h_i, h_l, h_u) for mode hotspot detection, and two parameters (d and λ) for the joint embedding module. To investigate the effect of these parameters, we run experiments on the Instagram dataset to vary one parameter at a time and evaluate how it affects the classification performance. Results can be seen in Fig. 4. The results show that XBully is not very sensitive to kernel bandwidth parameters except for h_l . The performance of XBully increases moderately when h_l becomes larger, i.e., the number of detected text hotspots is relatively small. In Fig. 4(e), we can see that when embedding dimension increases, the performance of XBully first improves and then remains stable. Fig. 4(f) shows that the best performance is achieved when λ is around 0.01. In general, XBully is not sensitive to most of the model parameters, and consequently can be tuned for various application purposes.

⁵For details of the confidence level, please refer to [14]

Table 2: Performance comparison of various methods on the Instagram dataset.

Percentages		10%		30%		50%		70%		90%	
Metrics		Mac F1	Mic F1	Mac F1	Mic F1	Mac F1	Mic F1	Mac F1	Mic F1	Mac F1	Mic F1
Random Forest	Raw	0.528	0.838	0.573	0.835	0.517	0.830	0.532	0.827	0.543	0.860
	DeepWalk	0.461	0.668	0.445	0.680	0.450	0.678	0.470	0.679	0.432	0.716
	Node2vec	0.519	0.714	0.550	0.712	0.584	0.717	0.562	0.716	0.599	0.770
	GraRep	0.459	0.671	0.456	0.680	0.464	0.680	0.460	0.671	0.455	0.707
	Variant	0.551	0.844	0.680	0.874	0.778	0.905	0.854	0.926	0.932	0.959
	XBully	0.566	0.853	0.702	0.887	0.814	0.920	0.865	0.937	0.963	0.982
Linear SVM	Raw	0.459	0.559	0.459	0.564	0.515	0.692	0.540	0.793	0.582	0.847
	DeepWalk	0.523	0.598	0.518	0.581	0.522	0.591	0.508	0.593	0.540	0.635
	Node2vec	0.586	0.663	0.577	0.635	0.612	0.665	0.582	0.643	0.622	0.680
	GraRep	0.513	0.585	0.534	0.603	0.515	0.621	0.505	0.626	0.568	0.712
	Variant	0.568	0.812	0.659	0.828	0.747	0.863	0.796	0.890	0.782	0.914
	XBully	0.570	0.819	0.668	0.840	0.781	0.886	0.821	0.904	0.837	0.928
Logistic Regression	Raw	0.459	0.828	0.460	0.830	0.465	0.819	0.451	0.82	0.461	0.856
	DeepWalk	0.512	0.634	0.523	0.620	0.508	0.618	0.491	0.602	0.514	0.644
	Node2vec	0.581	0.681	0.584	0.661	0.602	0.675	0.572	0.656	0.610	0.707
	GraRep	0.506	0.623	0.538	0.648	0.499	0.638	0.495	0.646	0.494	0.698
	Variant	0.495	0.837	0.522	0.832	0.536	0.835	0.543	0.826	0.615	0.874
	XBully	0.497	0.841	0.528	0.836	0.593	0.849	0.599	0.848	0.621	0.878
Cyberbully models	<i>Bully</i>	0.274	0.331	0.271	0.325	0.267	0.318	0.277	0.334	0.278	0.335
	<i>SICD</i>	0.447	0.646	0.443	0.604	0.383	0.537	0.438	0.512	0.358	0.559

Table 3: Performance comparison of various methods on the Vine dataset.

Percentages		10%		30%		50%		70%		90%	
Metrics		Mac F1	Mic F1	Mac F1	Mic F1	Mac F1	Mic F1	Mac F1	Mic F1	Mac F1	Mic F1
Random Forest	Raw	0.651	0.716	0.663	0.729	0.641	0.709	0.663	0.725	0.749	0.784
	DeepWalk	0.575	0.695	0.635	0.738	0.677	0.763	0.683	0.759	0.638	0.691
	Node2vec	0.576	0.704	0.626	0.733	0.655	0.746	0.679	0.753	0.638	0.691
	GraRep	0.589	0.703	0.633	0.723	0.671	0.751	0.694	0.763	0.650	0.691
	Variant	0.659	0.738	0.676	0.733	0.682	0.738	0.709	0.766	0.705	0.753
	XBully	0.661	0.740	0.678	0.758	0.711	0.779	0.717	0.777	0.757	0.784
Linear SVM	Raw	0.409	0.683	0.432	0.439	0.575	0.701	0.578	0.588	0.547	0.557
	DeepWalk	0.568	0.646	0.582	0.661	0.597	0.643	0.571	0.639	0.528	0.557
	Node2vec	0.569	0.659	0.592	0.658	0.579	0.627	0.599	0.649	0.620	0.649
	GraRep	0.590	0.664	0.610	0.669	0.634	0.689	0.644	0.715	0.629	0.649
	Variant	0.636	0.715	0.622	0.689	0.650	0.711	0.678	0.732	0.671	0.722
	XBully	0.657	0.717	0.641	0.704	0.651	0.733	0.678	0.742	0.700	0.753
Logistic Regression	Raw	0.648	0.732	0.683	0.748	0.684	0.755	0.676	0.746	0.705	0.753
	DeepWalk	0.554	0.668	0.581	0.691	0.631	0.705	0.598	0.680	0.589	0.639
	Node2vec	0.535	0.672	0.603	0.705	0.641	0.715	0.629	0.708	0.612	0.660
	GraRep	0.578	0.672	0.626	0.717	0.650	0.732	0.644	0.725	0.633	0.670
	Variant	0.618	0.724	0.663	0.741	0.696	0.757	0.705	0.770	0.658	0.732
	XBully	0.670	0.737	0.700	0.756	0.700	0.759	0.706	0.790	0.769	0.804
Cyberbully models	<i>Bully</i>	0.283	0.354	0.314	0.417	0.350	0.480	0.389	0.533	0.481	0.619
	<i>SICD</i>	0.417	0.715	0.436	0.773	0.506	0.775	0.474	0.900	0.473	0.897

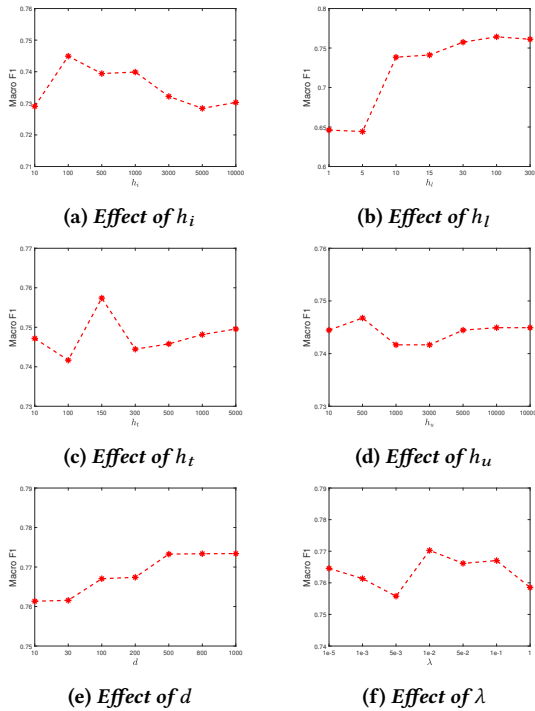


Figure 4: Parameter sensitivity study (with 50% dataset for training)

5 RELATED WORK

Cyberbullying Detection Most prior work on cyberbullying detection has primarily relied on text analysis to identify cyberbullying cases in online social networks like YouTube, Formspring, MySpace, Instagram, and Twitter [6, 7, 13, 35]. For example, Dinakar et al. [7] concatenated TF-IDF features, POS tags of frequent bigrams, and profane words as content features to investigate both explicit and implicit cyberbullying behaviors in negative text comments on YouTube and Formspring profiles [7, 8]. Xu et al. [35] presented several off-the-shelf tools such as Bag-of-Words models and LSA- and LDA-based representation to predict bullying traces on Twitter. Sanchez and Kumar [27] proposed the use of a Naïve Bayes classifier to identify instances of cyberbullying in Twitter. Dani et al. [6] proposed the *SICD* model which leverages the sentiment information to facilitate cyberbullying detection by capturing sentiment consistency of normal and bullying tweets. In [4, 5], the authors made use of gender-specific features and contextual features such as users' previous posts and the use of profane words to improve the performance of cyberbullying detection. With the increasing prevalence of social networking platforms, network-based features such as the number of friends, network structure, and relational centrality are also used more frequently to detect cyberbullying behaviors. For example, Homan et al. [12] studied the social structure of LGBT youth with depression in the TrevorSpace social network⁶. Huang et al. [15] studied a number of graph properties in addition to text features and found that cyberbullying detection performance

was significantly improved when both network-based features and textual features were exploited. Additional advances have come from newly-developed systems and applications to help identify cyberbullying risk on social network platforms such as [28, 29]. These models aim to estimate the probability that a teen is experiencing cyberbullying from streams of received messages as well as various vulnerability factors.

Network Embedding Network embedding, which seeks to learn low-dimensional vector representations of nodes by exploiting different properties of the underlying network, has been successful in advancing a number of downstream learning tasks [11]. Significant advances have also resulted from the foundational work of DeepWalk [25], which makes an analogy between truncated random walk in the network and short sliding window across sentences in a text corpus. Subsequent work node2vec [10] has proposed a flexible notion of node neighborhood and employs a biased random walk procedure to explore neighborhoods of each node in a diversified way, and Tang et al [31] proposed to embed large-scale information networks by carefully designing an objective function that preserves the first- and second-order node proximities. Notably, however, these prior contributions have focused on representation learning for homogeneous networks. In recognition that many real-world information systems can be modeled as a heterogeneous information network, Dong et al. [9] proposed a heterogeneous network embedding model, metapath2vec, that formalizes meta-path-based random walks to construct the heterogeneous neighborhood of a node, and then leverages a heterogeneous Skip-gram model [21] to learn embeddings. Tang et al. [30] first presented a large-scale heterogeneous text network by jointly training a *word-word* co-occurrence network, a *word-document* bipartite network, and a *word-label* bipartite network with both labeled and unlabeled text data. Most recently, there has been growing interest in performing network embedding on attributed networks [16, 17], dynamic networks [19, 37], signed networks [34], and hypernetworks [33]. Considering that social media platforms where rich multi-modal information is largely generated, are identified as the NO.1 “crime scene” for cyberbullying, the study of cyberbullying detection with multi-modal context information is critical. To the best of our knowledge, we are the first to study the multi-modal cyberbullying detection problem using network representation learning. Due to the simplicity, scalability, and effectiveness of embedding models, our approach can significantly improve the quality of the features for cyberbullying detection.

6 CONCLUSIONS AND FUTURE WORK

With the growing popularity of social media platforms and rapid increases in social media use among teens, cyberbullying has become more prevalent and begun to raise serious societal concerns. The majority of previous efforts for detecting cyberbullying are based on text analysis. Although they mark an important step forward in combating cyberbullying, these works fail to consider the multi-modal nature of social media data (e.g., texts, images, likes/shares, etc.). Our proposed model is based on the belief that multi-modal information can provide valuable insights for characterizing and detecting cyberbullying behaviors, which can complement and ultimately extend previous work.

⁶<https://www.trevorspace.org/>

In this paper, we study the novel problem of cyberbullying detection within a multi-modal context. To address the challenges tied to multi-modal social media information, we propose an innovative cyberbullying detection framework, XBully, based on network representation learning. XBully first identifies representative mode hotspots to handle diverse feature types and then jointly maps both attributed and nominal nodes in a heterogeneous network into the same latent space by exploiting the cross-modal correlations and structural dependencies. Extensive experimental results on real-world datasets corroborate the effectiveness of the proposed framework. Future work directed towards building a deeper understanding of different modalities in characterizing cyberbullying behaviors will not only improve cyberbullying detection, but may also shed light on behaviors that are unique to users with different roles (e.g., victims, bullies) within cyberbullying interactions. Furthermore, we believe that the most efficient path forward entails interdisciplinary collaboration among researchers in computer science and psychology to address this major social issue.

7 ACKNOWLEDGEMENT

This material is based upon work supported by, or in part by, the National Science Foundation (NSF) grant 1719722.

REFERENCES

- [1] Oren Anava and Kfir Levy. 2016. k*-nearest neighbors: From global to local. In *Advances in Neural Information Processing Systems*. 4916–4924.
- [2] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multi-modal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018).
- [3] Shaosheng Cao, Wei Lu, and Qionghai Xu. 2015. Grarep: Learning graph representations with global structural information. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. ACM, 891–900.
- [4] Maral Dadvar and Franciska De Jong. 2012. Cyberbullying detection: a step toward a safer internet yard. In *Proceedings of the 21st International Conference on World Wide Web*. ACM, 121–126.
- [5] Maral Dadvar, Dolf Trieschnigg, Roeland Ordeman, and Franciska de Jong. 2013. Improving cyberbullying detection with user context. In *European Conference on Information Retrieval*. Springer, 693–696.
- [6] Harsh Dani, Jundong Li, and Huan Liu. 2017. Sentiment Informed Cyberbullying Detection in Social Media. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 52–67.
- [7] Karthik Dinakar, Birago Jones, Catherine Havasi, Henry Lieberman, and Rosalind Picard. 2012. Common sense reasoning for detection, prevention, and mitigation of cyberbullying. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 2, 3 (2012), 18.
- [8] Karthik Dinakar, Roi Reichart, and Henry Lieberman. 2011. Modeling the detection of Textual Cyberbullying. *The Social Mobile Web* 11, 02 (2011).
- [9] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 135–144.
- [10] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 855–864.
- [11] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Representation Learning on Graphs: Methods and Applications. *arXiv preprint arXiv:1709.05584* (2017).
- [12] Christopher M Homan, Naiji Lu, Xin Tu, Megan C Lytle, and Vincent Silenzio. 2014. Social structure and depression in TrevorSpace. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 615–625.
- [13] Homa Hosseinmardi, Sabrina Arredondo Mattson, Rahat Ibn Rafiq, Richard Han, Qin Lv, and Shivakant Mishra. 2015. Analyzing labeled cyberbullying incidents on the instagram social network. In *International Conference on Social Informatics*. Springer, 49–66.
- [14] Homa Hosseinmardi, Rahat Ibn Rafiq, Richard Han, Qin Lv, and Shivakant Mishra. 2016. Prediction of cyberbullying incidents in a media-based social network. In *Advances in Social Networks Analysis and Mining (ASONAM), 2016 IEEE/ACM International Conference on*. IEEE, 186–192.
- [15] Qianjia Huang, Vivek Kumar Singh, and Pradeep Kumar Atrey. 2014. Cyber bullying detection using social and textual analysis. In *Proceedings of the 3rd International Workshop on Socially-Aware Multimedia*. ACM, 3–6.
- [16] Xiao Huang, Jundong Li, and Xia Hu. 2017. Accelerated attributed network embedding. In *Proceedings of the 2017 SIAM International Conference on Data Mining*. SIAM, 633–641.
- [17] Xiao Huang, Jundong Li, and Xia Hu. 2017. Label informed attributed network embedding. In *Proceedings of the 21st ACM International Conference on Web Search and Data Mining*. ACM, 731–739.
- [18] Jundong Li, Kewei Cheng, Suhang Wang, Fred Morstatter, Robert P Trevino, Jiliang Tang, and Huan Liu. 2017. Feature selection: A data perspective. *ACM Computing Surveys (CSUR)* 50, 6 (2017), 94.
- [19] Jundong Li, Harsh Dani, Xia Hu, Jiliang Tang, Yi Chang, and Huan Liu. 2017. Attributed network embedding for learning in a dynamic environment. In *Proceedings of the 2017 ACM Conference on Information and Knowledge Management*. ACM, 387–396.
- [20] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27, 1 (2001), 415–444.
- [21] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [22] Parma Nand, Rivindu Perera, and Abhijeet Kature. 2016. "How Bullying is this Message?": A Psychometric Thermometer for Bullying. In *COLING*. 695–706.
- [23] Emanuel Parzen. 1962. On estimation of a probability density function and mode. *The annals of mathematical statistics* 33, 3 (1962), 1065–1076.
- [24] James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71, 2001 (2001), 2001.
- [25] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 701–710.
- [26] Rahat Ibn Rafiq, Homa Hosseinmardi, Richard Han, Qin Lv, Shivakant Mishra, and Sabrina Arredondo Mattson. 2015. Careful what you share in six seconds: Detecting cyberbullying instances in Vine. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. ACM, 617–622.
- [27] Huascar Sanchez and Shreyas Kumar. 2011. Twitter bullying detection. *ser. NSDI* 12 (2011), 15–15.
- [28] Yasin N. Silva, Christopher Rich, Jaime Chon, and Lisa M. Tsosie. 2016. Bully-Blocker: An app to identify cyberbullying in facebook. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2016*. 1401–1405.
- [29] Yasin N. Silva, Christopher Rich, and Deborah Hall. 2016. BullyBlocker: Towards the identification of cyberbullying in social networking sites. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2016*. 1377–1379.
- [30] Jian Tang, Meng Qu, and Qiaozhu Mei. 2015. Pte: Predictive text embedding through large-scale heterogeneous text networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1165–1174.
- [31] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1067–1077.
- [32] Waldo R Tobler. 1970. A computer movie simulating urban growth in the Detroit region. *Economic geography* 46, sup1 (1970), 234–240.
- [33] Ke Tu, Peng Cui, Xiao Wang, Fei Wang, and Wenwu Zhu. 2018. Structural Deep Embedding for Hyper-Networks. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- [34] Suhang Wang, Jiliang Tang, Charu Aggarwal, Yi Chang, and Huan Liu. 2017. Signed network embedding in social media. In *Proceedings of the 2017 SIAM International Conference on Data Mining*. SIAM, 327–335.
- [35] Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012. Learning from bullying traces in social media. In *Proceedings of the 2012 conference of the North American chapter of the association for computational linguistics: Human language technologies*. Association for Computational Linguistics, 656–666.
- [36] Chao Zhang, Keyang Zhang, Quan Yuan, Haoruo Peng, Yu Zheng, Tim Hanratty, Shaowen Wang, and Jiawei Han. 2017. Regions, periods, activities: Uncovering urban dynamics via cross-modal representation learning. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 361–370.
- [37] Lekui Zhou, Yang Yang, Xiang Ren, Fei Wu, and Yueting Zhuang. 2018. Dynamic Network Embedding by Modeling Triadic Closure Process. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.