

Cyberbullying Detection Using Machine Learning

Sachin¹, Dr shivkumar Kagi²

¹Student, Dept. Of Computer Science, Sharnbasava University, Kalaburagi, India.

²Associate Professor, Dept. Of Computer Science, Sharnbasava University, Kalaburagi, India.

ABSTRACT

Cyberbullying has recently been reported to cause significant damage to society and the economy. Technological advances regarding the annotation of web documents and the diversity of online communities make the detection and tracking of such cases quite difficult and very difficult. Cyberbullying occurs through various information and communication technologies such as instant messaging, email, text messaging, social networking sites, blogs, and websites. We must fight this kind of abuse. The purpose of this project is to analyze and evaluate the communication between two specific individuals or an anonymous person. Recommended machine learning model to detect and prevent harassment in one-on-one conversation. The classifiers, SVM, Logistic Regression, Random Forest and Gradient Boosting are used to train and test social media bullying. We use the Twitter chat dataset to train the ml model.

Keywords: Cyberbullying, SVM, Logistic Regression, Random Forest , Gradient Boosting, machine learning

I. INTRODUCTION

Cyberbullying means insulting, threatening, defaming, or intentionally harassing another person using modern means of communication, usually over an extended period of time. Cyberbullying occurs on the Internet (e.g. Via email, instant messaging, social networks, videos on various portals) or over the phone (e.g. Via WhatsApp or annoying calls). Most of the time, the perpetrator, known as the "bully," acts anonymously, leaving the victim unaware of where the assaults take place. Especially in the case of cyberbullying between children and adolescents, the victim and abuser also know each other in the "real" world. Victims almost always suspect who may be behind the attacks. Cyberbullying often begins with people in the immediate environment: schools, neighbourhoods, villages or ethnic communities. Cases involving complete strangers are rare. Today, many young people are bullied. Bullies use various services like Twitter, Facebook, Email to intimidate people. Studies show that around 37% of children in India are involved in cyberbullying and almost 14% of bullying occurs regularly. Cyberbullying affects victims both emotionally and psychologically. Social media also allows bullies to exploit anonymity to satisfy their evil deeds. Things became even more serious when bullying incidents occurred continuously over time. So, preventing this from happening helps the victim. Cyberbullying is a constant invasion of privacy that doesn't stop at the doorstep (except when the new medium is not being used at home). The prevalence of information cannot be predicted because of the wide range of possibilities and the speed of new media. The perpetrators, so-called "cyber harassment," can act anonymously and often consider themselves safe because of this anonymity. In most cases, the identity of the perpetrator is presented very differently from the reality. Age or external image are not important criteria for cyberbullying. This can happen both between people of the same age (peers) and between people of different ages (student teachers). There is a possibility of unintentional cyberbullying, as thoughtless actions without awareness of the consequences can hurt the people involved. The abuser often fails to see these reactions and is unaware of the scale of the actions. Given the consequences of cyberbullying on its victims, it is imperative to find the appropriate actions to detect and prevent it. Machine learning is one of the successful approaches that learns from data and creates a model that automatically classifies appropriate actions. Machine learning can be useful to detect language patterns of bullies and thus can generate a pattern to detect acts of cyberbullying. A recent study conducted by Microsoft Corporation to understand the spread of cyberbullying globally showed that India ranks 3rd in terms of cyberbullying after China and Singapore. According to recent studies, 52% of young people in India have been victims of cyberbullying in the past and around 38% of them have been bullied. Cyberbullying is basically of two types, one that contains abusive language and one that embarrasses the intended target but does not use swear words. Posts containing abusive content or profanity are more likely to be labelled as "articles". According to, for today's younger generation, "Gay", "Bitch" and "Slag" are the most commonly used abusive terms in schools..

II. LITERATURE SURVEY

Rahul Ramesh Dalvi, Sudhanshu Baliram Chavan, Aparna Halbe , 2020, [1], social media is a platform where many young people are getting bullied. As social networking sites are increasing, cyberbullying is increasing

day by day. To identify word similarities in the tweets made by bullies and make use of machine learning and can develop an ML model automatically detect social media bullying actions. However, many social media bullying detection techniques have been implemented, but many of them were textual based. The goal of this paper is to show the implementation of software that will detect bullied tweets, posts, etc. A machine learning model is proposed to detect and prevent bullying on Twitter. Two classifiers i.e. SVM and Naïve Bayes are used for training and testing the social media bullying content. Both Naive Bayes and SVM (Support Vector Machine) were able to detect the true positives with 71.25% and 52.70% accuracy respectively. But SVM outperforms Naive Bayes of similar work on the same dataset. Also, Twitter API is used to fetch tweets and tweets are passed to the model to detect whether the tweets are bullying or not

Mestry, Suresh; Singh, Hargun; Chauhan, Roshan; Bisht, Vishal; Tiwari, Kaushik (2019)[2], social networking and online conversation platforms provide us with the power to share our views and ideas. However, nowadays on social media platforms, many people are taking these platforms for granted, they see it as an opportunity to harass and target others leading to cyber-attack and cyber-bullying which lead to traumatic experiences and suicidal attempts in extreme cases. Manually identifying and classifying such comments is a very long, tiresome and unreliable process. To solve this challenge, we have developed a deep learning system which will identify such negative content on online discussion platforms and successfully classify them into proper labels. Our proposed model aims to apply the text-based Convolution Neural Network (CNN) with word embedding, using fastText word embedding technique. fastText has shown efficient and more accurate results compared to Word2Vec and GLOVE model. Our model aims to improve detecting different types of toxicity to improve the social media experience. Our model classifies such comments in six classes which are Toxic, Severe Toxic, Obscene, Threat, Insult and Identity-hate. Multi-Label Classification helps us to provide an automated solution for dealing with the toxic comments problem we are facing.

Mengfan yao, charalampos chelmiss, daphney– stavroula zois 2018[3], Cyberbullying has emerged as a large-scale societal problem that demands accurate methods for its detection in an effort to mitigate its detrimental consequences. While automated, data-driven techniques for analyzing and detecting cyberbullying incidents have been developed, the scalability of existing approaches has largely been ignored. At the same time, the complexities underlying cyberbullying behavior (e.g., social context and changing language) make the automatic identification of “the best subset of features” to use challenging. We address this gap by formulating cyberbullying detection as a sequential hypothesis testing problem. Based on this formulation, we propose a novel algorithm to drastically reduce the number of features used in classification. We demonstrate the utility, scalability and responsiveness of our approach using a real-world dataset from Instagram, the online social media platform with the highest percentage of users reporting experiencing cyberbullying. Our approach improves recall by a staggering 700%, while at the same time reducing the average number of features by up to 99.82% compared to state-of-the-art supervised cyberbullying detection methods, learning approaches that require weak supervision, and traditional offline feature selection and dimensionality reduction techniques.

Abdullah-Al-Mamun, Shahin Akhter, 2020, [4], with the popularity of Unicode system and growing use of Internet, the use of Bangla over social media is increasing. However, very few works have been done on Bangla text for social media activity monitoring due to a lack of a large number of annotated corpora, named dictionaries and morphological analyzer, which demands in-depth analysis on Bangladesh’s perspective. Moreover, solving the issue by applying available techniques is very content specific, which means that false detection can occur if contents changed from formal English to verbal abuse or sarcasm. Also, performance may vary due to linguistic differences between English and non-English contents and the socio-emotional behaviour of the study population. To combat such issues, this paper proposes the use of machine learning algorithms and the inclusion of user information for cyber bullying detection on Bangla text. For this purpose, a set of Bangla text has been collected from available social media platforms and labelled as either bullied or not bullied for training different machine learning based classification models. Cross-validation results of the models indicate that a support vector machine based algorithm achieves superior performance on Bangla text with a detection accuracy of 97%. Besides, the impact of user specific information such as location, age and gender can further improve the classification accuracy of Bangla cyber bullying detection system.

Mohammad Rashid Hussain, Henry Friday Nweke, [5], 2019, Prior to the innovation of information communication technologies (ICT), social interactions evolved within small cultural boundaries such as geo spatial locations. The recent developments of communication technologies have considerably transcended the temporal and spatial limitations of traditional communications. These social technologies have created a revolution in user-generated information, online human networks, and rich human behavior-related data. However, the misuse of social technologies such as social media (SM) platforms, has introduced a new form of aggression and violence that occurs exclusively online. A new means of demonstrating aggressive behavior in SM websites are highlighted in this paper. The motivations for the construction of prediction models to fight aggressive behavior in SM are also outlined. We comprehensively review cyberbullying prediction models and identify the main issues related to the construction of cyberbullying prediction models in SM. This paper provides insights on the overall process for cyberbullying detection and most importantly overviews the

methodology. Though data collection and feature engineering process has been elaborated, yet most of the emphasis is on feature selection algorithms and then using various machine learning algorithms for prediction of cyberbullying behaviors. Finally, the issues and challenges have been highlighted as well, which present new research directions for researchers to explore..

III. PROPOSED METHODOLOGY

The proposed methodology outlines a systematic approach to design and implement an effective cyberbullying detection system using machine learning and natural language processing techniques. The key phases are as follows:

1. Data Collection

- Acquire datasets containing labeled social media posts, comments, or chat messages indicating cyberbullying and non-cyberbullying content.
- Public datasets such as the Cyberbullying Dataset from social media platforms, or custom-collected data via APIs (e.g., Twitter API), can be used.

2. Data Preprocessing

- Clean the raw text data by removing noise such as URLs, special characters, numbers, and stopwords.
- Normalize text through lowercasing, tokenization, and stemming or lemmatization.
- Handle slang, abbreviations, and emoticons commonly found in informal online communication.

3. Feature Extraction

- Transform the processed text into numerical features usable by machine learning algorithms.
- Use techniques like Bag-of-Words (BoW), TF-IDF (Term Frequency-Inverse Document Frequency), or word embeddings such as Word2Vec or GloVe to capture semantic meaning.

4. Model Selection and Training

Experiment with various machine learning classifiers, including:

Logistic Regression

Data Preprocessing

The text data (tweets, comments, or posts) is first cleaned by removing noise such as URLs, punctuation, stopwords, and emojis. Then, the cleaned text is converted into numerical form using techniques like:

Bag of Words (BoW)

TF-IDF (Term Frequency–Inverse Document Frequency)

These methods represent each text as a vector of numbers that indicate how often certain words appear and how important they are to the context.

Feature Extraction

Each tweet becomes a numerical feature vector. For example, words like “idiot,” “hate,” or “ugly” might carry a higher weight because they are commonly found in bullying-related tweets.

Model Training

Logistic Regression learns a function that maps these feature vectors to binary classes:

1 → Cyberbullying

0 → Not Cyberbullying

It does this by estimating probabilities using the sigmoid function, which outputs a value between 0 and 1. If the probability ≥ 0.5 , the model classifies it as cyberbullying; otherwise, it classifies it as not cyberbullying.

Prediction

When a new tweet is entered, the trained model calculates the probability that it belongs to the *cyberbullying* class and makes the prediction accordingly.

Support Vector Machines (SVM)

Support Vector Machines (SVM) is one of the most powerful and widely used machine learning algorithms for text classification problems, including cyberbullying detection. It is a supervised learning model that works exceptionally well with high-dimensional data, such as text converted into numerical features using methods like TF-IDF or Bag of Words. In the context of cyberbullying detection, SVM helps classify social media posts, tweets, or comments into two categories — Cyberbullying and Not Cyberbullying — by finding the best possible boundary that separates the two classes.

How SVM Works in Cyberbullying Detection

Text Preprocessing

Before training the SVM model, text data is cleaned and preprocessed:

Removing URLs, numbers, punctuation, and stopwords.

Converting text to lowercase.

Tokenizing and converting the cleaned text into numerical form using TF-IDF vectors or word embeddings.

Feature Space Representation

Each tweet or comment is transformed into a high-dimensional vector (for example, one dimension per unique word). The SVM model treats each text as a point in this space.

Hyperplane Formation

SVM tries to find the optimal hyperplane that best separates the data points of the two classes:

Points representing cyberbullying tweets lie on one side.

Points representing non-cyberbullying tweets lie on the other side.

The hyperplane is chosen in such a way that the margin — the distance between the hyperplane and the nearest data points (called support vectors) — is maximized. This ensures better generalization and higher accuracy.

Prediction

When a new message or tweet is input into the system, the SVM model calculates on which side of the hyperplane the point lies:

If it lies on the cyberbullying side → classified as Cyberbullying.

Otherwise → classified as Not Cyberbullying.

Random Forest

Random Forest is a powerful ensemble machine learning algorithm widely used in cyberbullying detection systems to classify online comments, tweets, or posts as either Cyberbullying or Not Cyberbullying. It combines the predictions of multiple decision trees to improve overall accuracy, reduce overfitting, and handle the complex language patterns found in social media text. Random Forest is particularly effective for text classification tasks where the data is noisy, unstructured, and high-dimensional — common characteristics of online communication.

How Random Forest Works in Cyberbullying Detection

Data Preparation

Before training, text data (tweets or messages) undergoes preprocessing:

Removal of special characters, punctuation, and stopwords.

Conversion of text to lowercase.

Feature extraction using TF-IDF (Term Frequency–Inverse Document Frequency) or Bag of Words to transform text into numerical vectors.

Training Multiple Decision Trees

Random Forest creates many decision trees during training — each tree is built using a random subset of:

The training data (rows)

The features (columns)

Each decision tree learns different linguistic patterns associated with cyberbullying (for example, aggressive or insulting words, hate terms, or offensive tone).

Voting Mechanism

When a new message is entered, each decision tree makes its own prediction (cyberbullying or not). The final output is decided through a majority vote:

If most trees predict Cyberbullying, the final classification is Cyberbullying.

If most predict Not Cyberbullying, it's classified as Not Cyberbullying.

Probability Estimation

Random Forest also provides a confidence score or probability, representing how sure the model is about its decision. For example, "Cyberbullying — 92.06% confidence."

Gradient Boosting

Gradient Boosting is one of the most advanced and powerful ensemble machine learning algorithms used in cyberbullying detection to classify text as Cyberbullying or Not Cyberbullying. It works by combining the predictions of multiple weak learners—typically decision trees—in a sequential way to form a strong predictive model. In the context of cyberbullying detection, Gradient Boosting excels because it can capture complex language patterns, sarcasm, and subtle emotional tones in online text data, resulting in high accuracy and robustness compared to traditional models.

How Gradient Boosting Works in Cyberbullying Detection

Data Preprocessing

Before training the model, text data (tweets, comments, or posts) is cleaned and prepared:

Removing stopwords, punctuation, and links.

Converting text to lowercase.

Transforming text into numerical form using TF-IDF, Bag of Words, or Word Embeddings.

Initial Model Creation

Gradient Boosting begins with a simple model (often a single decision tree) that makes an initial prediction. This model may not be very accurate but serves as a starting point.

Error Learning (Boosting Process)

The model then calculates the errors (residuals) between the predicted and actual values.

A new decision tree is trained to correct these errors.

Each subsequent tree focuses on improving the performance of the previous ensemble by learning from its mistakes.

This process continues for several iterations, gradually improving the model's accuracy.

Final Prediction

Once all trees are trained, the model combines their predictions (weighted by importance) to produce the final classification:

Cyberbullying (1) or Not Cyberbullying (0)

Along with a confidence score (%) showing how certain the model is about its decision

- Train models on the labeled dataset, optimizing hyperparameters using techniques like cross-validation.

5. Model Evaluation

- Assess models using metrics such as accuracy, precision, recall, F1-score, and confusion matrix to understand detection performance.
- Select the best-performing model based on balanced precision and recall to minimize false positives and false negatives.

6. Application Development

- Develop a user-friendly interface or RESTful API using Python frameworks such as pycharm
- Integrate the trained model into the application to enable real-time cyberbullying detection on user-provided text input.
- Implement functionality to notify users or moderators when cyberbullying is detected.

7. Deployment and Testing

- Deploy the application on a cloud server or local environment for accessibility.
- Perform testing with real-world data to validate system robustness and responsiveness.

- Gather user feedback for continuous improvement.

IV. SYSTEM ARCHITECTURE




Figure 1: System Architecture

The system design outlines the architecture and workflow of the cyberbullying detection system, demonstrating how various components interact to provide an effective, automated solution.

1. Architecture Overview

The system follows a modular, layered architecture consisting of:

- **Data Layer:** Handles data collection and storage.
- **Processing Layer:** Responsible for data preprocessing, feature extraction, and model training.
- **Prediction Layer:** Performs real-time classification of input text.
- **Application Layer:** Interfaces with users via a web or mobile app.
- **Notification Layer:** Manages alerts or reports based on detection results.

2. Components Description

a) Data Collection Module

Collects data from social media APIs, public datasets, or direct user input.
Stores raw and labeled datasets in a database or file system.

b) Data Preprocessing Module

Cleans and normalizes text data.
Removes noise like URLs, special characters, and stopwords.
Performs tokenization, lemmatization, and handles emojis/slang.

c) Feature Extraction Module

Transforms text into vectorized formats (TF-IDF, Bag-of-Words, or embeddings).
Ensures consistent feature representation during training and prediction.

d) Machine Learning Model

Trains classifiers using labeled data.
Supports multiple algorithms with hyperparameter tuning.
Saves the best model for deployment.

e) Prediction Module

Accepts new user input text.
Applies preprocessing and feature extraction.
Uses the trained model to classify text as cyberbullying or not.
Returns detection results.

V. RESULTS




Figure 2: Dataset

The dataset used in this application is typically a collection of tweets gathered either through: A pre-labeled dataset (e.g., from Kaggle's Cyberbullying Detection Dataset),





Figure 3: Cyber bullying Detection

It is designed to analyze text input, such as a tweet or online comment, and determine whether it contains cyberbullying or hate-related content. At the top of the interface, there is a text box labeled “Enter Tweet or Sentence” where the user can type or paste any social media post for analysis. In the given example, the input text contains offensive and religion-targeted language, which makes it an instance of cyberbullying. Below the text box, the system displays results generated by multiple machine learning models—SVM (Support Vector Machine), Random Forest, Logistic Regression, and Gradient Boosting. Each model predicts whether the text is *cyberbullying* or *not cyberbullying*, along with its respective accuracy score. In this case, SVM and Gradient Boosting classify the text as *cyberbullying* with accuracies of 93.12% and 94.18%, respectively, while Random Forest and Logistic Regression identify it as *not cyberbullying* with around 92% accuracy.



Graph 1: Model Accuracy graph

The image displays a bar chart titled “Model Accuracy Comparison”, which visually represents the performance of four different machine learning models used in the Cyberbullying Detection System. The chart compares how accurately each model classifies text as either cyberbullying or not cyberbullying, expressed as a percentage of accuracy on the horizontal axis. Each blue horizontal bar corresponds to one model, with the accuracy

percentage shown at the end of the bar. The models listed from top to bottom are Gradient Boosting, Logistic Regression, Random Forest, and SVM (Support Vector Machine). The chart shows that Gradient Boosting achieved the highest accuracy of 94.18%, indicating it performs best among the tested models. This means it correctly predicts cyberbullying-related content more often than the others. SVM follows closely with 93.12%, while both Random Forest and Logistic Regression achieved slightly lower accuracies of 92.06% each.




Figure 4:Using Streamlit

The image shows a web-based interface of a project titled “Cyberbullying Detection App.” This application allows users to enter a tweet, post, or comment and automatically determine whether it is an instance of cyberbullying or not cyberbullying using machine learning models. The app appears to be built using a web framework like Flask or Streamlit, as it runs locally on localhost:8501, which is common for testing such AI applications.

At the top, the page has a heading and short description explaining the purpose of the tool — to classify online comments as either Cyberbullying or Not Cyberbullying using different trained ML models. Below that, there is a text input box labeled “Enter a tweet or message”, where the user can type or paste any social media post. In the example shown, the entered text is:

“Equal rights are not special rights, or have you forgotten that there were white Christian women couldn’t vote 100 years ago?”

This statement is analytical or historical rather than abusive, which helps the model demonstrate its ability to differentiate between general discussion and harmful content.

Next, there is a dropdown menu labeled “Choose a model”, allowing the user to select from available machine learning classifiers such as Logistic Regression, Random Forest, or SVM. In this case, Logistic Regression has been selected. When the user clicks the “Detect” button, the model processes the input and displays the result below.

The output section shows two pieces of information:

Prediction: Not Cyberbullying — meaning the app determined that the entered text does not contain offensive or bullying language.

Confidence: 81.1% — representing how certain the model is about its prediction.

At the bottom, there is an additional option labeled “Show sample data”, which likely lets users view example tweets from the dataset used for training or testing.

VI. CONCLUSION AND FUTURE WORKS

This study reviewed existing literature to detect aggressive behavior on social media. By using various machine learning approaches. We have reviewed four main features of Detecting cyberbullying messages by using machine learning approaches namely Importing data, feature extraction, construction of cyber bullying detection model and evaluation of the constructed model. An approach is proposed for detecting and preventing Cyberbullying using supervised Binary classification machine learning algorithms.

FUTURE SCOPE:

The development of cyberbullying detection systems using machine learning opens numerous avenues for future enhancements and research. Some promising directions include:

1.Multilingual and Cross-Cultural Detection

Extending the system to support multiple languages and adapt to different cultural contexts will broaden its applicability worldwide. This requires incorporating language-specific NLP tools and datasets.

2. Incorporation of Multimedia Content Analysis

Beyond text, cyberbullying often occurs through images, videos, and audio. Future systems can integrate computer vision and speech recognition techniques to detect abusive content across multimedia platforms.

3. Context-Aware Detection

Developing models that understand the context, sentiment, and conversation history can improve detection accuracy. This involves using advanced techniques like transformer-based models (e.g., BERT, GPT) that capture nuanced language and sarcasm.

4. Real-Time Monitoring and Intervention

Enhancing systems for live, real-time monitoring on social media and chat platforms can enable immediate alerts and intervention, potentially preventing escalation and harm.

5. User Behavior Analysis and Profiling

Combining content analysis with user behavior patterns and network analysis could help in identifying repeat offenders and predicting potential cyberbullying incidents proactively.

6. Integration with Educational and Counseling Tools

Linking detection systems with educational programs, counseling services, or mental health support can offer holistic solutions to cyberbullying victims.

7. Improved Explainability and Transparency

Incorporating explainable AI (XAI) techniques will help users and moderators understand why certain content is flagged, building trust and facilitating better decision-making.

Privacy-Preserving Techniques

Future work could focus on applying privacy-preserving machine learning methods such as federated learning to protect user data while enabling collaborative model training across platforms

REFERENCES

- [1] Detecting A Twitter Cyberbullying Using Machine Learning Rahul Ramesh Dalvi, Sudhanshu Baliram Chavan, Aparna Halbe Proceedings of the International Conference on Intelligent Computing and Control Systems (ICICCS 2020).IEEE Xplore Part Number: CFP20K74-ART; ISBN: 978-1-7281-4876-2
- [2] Automation in Social Networking Comments With the Help of Robust fastText and CNN” .Mestry, Suresh; Singh, Hargun; Chauhan, Roshan; Bisht, Vishal; Tiwari, Kaushik (2019). [IEEE 2019 1st International Conference on Innovations in Information and Communication Technology (ICIICT) - CHENNAI, India (2019.4.25-2019.4.26)]
- [3] Cyberbullying detection on instagram with optimal online feature selection.Mengfan yao, charalampos chelmiss, daphney– stavroula zois 2018 ieee/acm international conference on advances in social networks analysis and mining (asonam)
- [4] Abdhullah-Al-Mamun, Shahin Akhter, "Social media bullying detection using machine learning on Bangla text", 10th International Conference on ElectricalandComputerEngineering,pages385-388,IEEE Xplore,20
- [5] Cyberbullying on Social Media in the Big Data Era Using Machine Learning Algorithms: Review of Literature and Open Challenges (Mohammed Ali Al-garadi,Mohammad Rashid Hussain, Henry Friday Nweke)
- [6] <https://stanford-cs221.github.io/spring2021/>
- [7] <https://reactjs.org/tutorial/tutorial.html>
- [8] <https://www.tensorflow.org/tutorials>
- [9] <http://introtodeeplearning.com/>
- [10] Natural Language Processing in Action Understanding, analyzing, and generating text with Python - Hobson Lane Cole Howard Hannes Max Hapke
- [11] Introduction to Machine Learning with Python -Andreas C. Müller & Sarah Guido
- [12] Machine Learning – Tom M. Mitchel
- [13]<https://muthu.co/understanding-the-classification-report-in-sklearn/>
- [14] <https://developer.twitter.com/en/apps>
- [15] <https://text-processing.com/demo/tokenize/>
- [16] <https://towardsdatascience.com/support-vector-machine-introduction-to-machine-learning-algorithms-934a444fca47>
- [17]<https://towardsdatascience.com/naive-bayes-classifier-81d512f50ac>
- [18] Phillips, L., et al., Using Social Media to Predict the Future: A Systematic Literature Review. arXiv preprint arXiv:1706.06134, 2017.
- [19] Quan, H., J. Wu, and J. Shi, Online social networks & social network services: A technical survey. Pervasive Communication Handbook. CRC, 2011: p. 4.
- [20] Peterson, J.K. and J. Densley, Is Social Media a Gang? Toward a Selection, Facilitation, or Enhancement