

“© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Exploring the Impact of Evolutionary Computing based Feature Selection in Suicidal Ideation Detection

Abstract—The ubiquitous availability of smartphones and the increasing popularity of social media provide a platform for users to express their feelings, including suicidal ideation. Suicide prevention by suicidal ideation detection on social media lights the path to controlling the rapidly increasing suicide rates amongst youth. This paper proposes a diverse set of features and investigates into feature selection using the Firefly algorithm to build an efficient and robust supervised approach to classifying tweets with suicidal ideation. The development of a suicidal language to create three diverse, manually annotated datasets leads to the validation of the proposed model. An in-depth result and error analysis lead to an accurate system for monitoring suicidal ideation on social media along with the discovery of optimal feature subsets and selection methods using a penalty based Firefly algorithm.

Index Terms—Firefly Algorithm, Suicide Ideation, Twitter

I. INTRODUCTION

As suicide continues to be one of the leading causes of death worldwide, the increasing ubiquity of social media [1], [2] presents an outlet for individuals, particularly teenagers to exhibit suicide ideation. While there are existing methods for identifying suicide-related communication, a robust classification system is needed for accurate identification of suicidal intent, it's separation from non-concerning suicide-related communication and it's subsequent prevention [3].

Traditional search methods such as PCA and Chi-Square rank the features individual contribution compared to Evolutionary algorithms (EAs) that iteratively try to maximize a subset of features together. The stochastic nature of EAs, their ability to automatically fine-tune the number of reduced features as opposed to traditional methods and selecting features by considering their combined contribution to accuracy was the motivation for using them.

dimensional, redundant and noisy aspects [10] of feature extraction methods for tweet classification. This work presents an analysis into the impact of using the Firefly algorithm [11] with a large set of features to train a Random Forest classifier. The major contributions of this work are

- 1) **Real World Data:** Validating our approach on a real-world dataset collected from Twitter, a widely used social media website, queried on the basis of an in-depth analysis of forums where suicide ideation may be expressed.
- 2) **Feature Identification:** Designing an easily extractable feature set from social media data that encapsulates a diverse range of aspects of suicidal ideation to build a robust classification system.

While there is a developing body of literature on the topic of identifying patterns in the language used on social media that expresses suicidal ideation [2]), very few attempts have been made to employ feature extraction coupled with feature selection methods for binary classifiers that separate text related to suicide from text that clearly indicates the author exhibiting suicidal intent. A number of successful models [4] have been used for sentence level classification, even those highlighting the merits of evolutionary computation models [5] however, ones that are successful for being able to learn to separate suicidal ideation from depression as well as less worrying content such as reporting of a suicide, memorial, campaigning, and support. etc, require a greater analysis to select more specific features and methods to build an accurate and robust model. The drastic impact that suicide has on surrounding community coupled with the lack of specific feature selection based methods to handle the high dimensional data used to train robust suicidal ideation classification models is the driving motivation for the work presented in this paper.

Positioning the work from a clinical perspective, existing suicide detection systems, prevention chat bots and social media based surveillance can benefit from the proposed methodology. Existing studies [6], [7], and [8] highlight the immense impact such automatic identification of concerning content can play in impacting millions of lives. Given the diverse nature of suicidal ideation, a suicidal language was first developed to scrape datasets and manually annotate them to validate the performance of the proposed classification system. Three datasets were created using different strategies to probe various aspects such as robustness and accuracy and validate the proposed methodology. Additionally, an investigation into a number of features along with a highly efficient stochastic method for feature selection [9] was essential due to highly

- 3) **Disciminative Model:** An in-depth analysis of the proposed features with various models and hyperparameters to flag suicide ideation on social media.

- 4) **Error Analysis:** Paving the way to future work by exploring the limitations and pitfalls involved with the present model.

Furthermore, we believe our contributions could pave the way to the development of tools and techniques that can be directly integrated with social media websites to automate and enhance their suicide prevention mechanisms.

II. METHODOLOGY

A. Tweet level Feature Design

The set of proposed features is shown in Table 2. These consist of both standard features used in sentence and tweet classification as well as features specific to suicidal ideation such as Linguistic Inquiry and Word Count (LIWC) [12] features and Sentiment Ratio. The POS tags are obtained from [13]. Sentiment ratio is calculated using the sentiment lexicon in [14]. 30 topics were estimated to identify and give additional weight to a more refined set of weights corresponding to topics such as suicide, death. We then consider the top 10 words per topic, for both classes. This results in 600 LDA-specific unigrams that are then used as features [15]. LIWC features capture social and cognitive aspects of tweets. Only those unigrams and bigrams occurring in more than three tweets and with a length of greater than three were included in the feature set. The dimensions of the features for Datasets SCO and UNI are shown in Table 4.

B. Classification

Identification of potentially concerning tweets is formulated as a binary classification problem. For every tweet t_i a binary variable $s_i \in \{0, 1\}$, where $s_i = 1$ denotes that t_i exhibits suicidal intent. Traditional classifiers, along with those employing Bagging and Boosting [16], Deep Learning models, and state-of-the-art models were trained using features extracted from the datasets, given its efficient and accurate performance in text classification and compatibility with the Firefly algorithm.

TABLE I
PARAMETERS FOR FIREFLY ALGORITHM

Parameter	Value
Epoch	200
Agents	20
k	0.12
Light absorption	1
Attractiveness	0.2

C. Feature Selection

The set of proposed features can be modeled in a binary search space that is randomized with n fireflies, whose position is denoted by x_i , corresponding to the inclusion (1) or exclusion (0) of the i^{th} feature. The Firefly algorithm [11] is a meta-heuristic based algorithm that aims to maximize the accuracy of the classifiers proposed in the above subsection while trying to minimize the number of features to reduce computational cost, high dimensionality and redundancy. The Binary Firefly Algorithm (BFA) is an adaptation of the firefly algorithm in a discrete space commonly used for feature selection [17]. The fitness function that aims to model the goals presented above is modeled as:

$$f(v) = A(v; d) \frac{\sum_{j \in N} |x_j - v|}{k} \quad (1)$$

where v is the feature vector, d is the dataset, A is the classification accuracy, $|N|$ is the total number of features, and k is a constant.

The BFA algorithm is a wrapper based feature selection method that requires feedback from a classifier to tune the features in training. For this purpose, we have used the firefly algorithm as a wrapper over the four classifiers presented in the paper. Both datasets, SCO and UNI are used for training the classifier whose results are fed back into the firefly algorithm that then fine tunes the selected features for classification. Each of the feature that is fed into the classifier is modeled in a binary search space, that is, the feature vector represents a point in a N dimensional binary search space, essentially a vector of 0s and 1s. BFA essentially aims to maximise the fitness function (increase accuracy, reduce number of features) by taking in datasets SCO and UNI. After preprocessing and feature extraction, we initialize our feature vector with a random seed which is then used to select only those features that have 1 in the corresponding feature vector, BFA then maximises the fitness function by modifying the feature vector over the multiple iterations. The parameters and implicit randomness of BFA allows a more probable convergence to global maxima compared to other feature selection methods such as information gain and sequential forward search, and search methods like hill climbing.

III. DATASET

A. Data Collection

The major problem with tweet collection to train classifiers was to develop a language consisting of words and phrases

TABLE II
DIMENSIONS FOR FEATURES FOR DATASETS SCO AND UNI.

Dataset	Unigrams	Bigrams	SR	LIWC	POS Counts	LDA unigrams
Dataset SCO	7892	9387	1	67	10	600
Dataset UNI	9640	11240	1	67	10	600

pertaining to suicide, which could then be used to scrape suicidal ideation exhibiting tweets. Tweets were collected using Twitter REST API [18] over the period of 13th January 2018 to 26 March 2018. Preprocessing of tweets involves removal of non-English tweets using Ling-Pipe, non-Unicode characters and tweets containing hyperlinks. A suicidal language was developed by first extracting 4314 posts from four well known Suicide web forums [19] identified by [20]. Additionally, user posts containing tags of 'suicide' from social media websites, Tumblr and Reddit were added to this collection. This resulted in the following composition of posts: 300 posts from each suicide forum and 2000 posts randomly selected from Tumblr and Reddit. A list of 108 words/phrases linked with Suicidal Intent was developed post manual annotation of these posts and using Term Frequency/Inverse Document Frequency [21] to identify the most frequently occurring terms. Three datasets are created using different strategies to validate the model's performance in terms of various aspects as follows:

- 1) Dataset SCO (Suicide-related Communication Only) (2726 suicidal, 9160 non-suicidal): Tweets were collected by using the words/phrases described above.
- 2) Dataset UNI (All sorts of tweets) (1576 suicidal, 18290 non-suicidal): Tweets were collected the same way as they were for building Dataset SCO, but additional tweets of users whose posts were identified as suicidal with no hashtags corresponding to suicidal ideation were included in Dataset UNI.
- 3) Dataset H (666 suicidal, 4130 non-suicidal): This represents the held-out test set, with tweets collected using both approaches stated above and no overlap with Datasets SCO and UNI.

The datasets in the study are highly imbalanced, this emulates the reality where only a small fraction of tweets are suicidal, particularly to build a strict classification scheme that reduces the false positive rate enabling efficiency in practice. Additional tweets (maximum 3; posted within 2 weeks by authors of tweets with suicidal intent) were added to UNI. Non-suicidal user-extracted tweets with Jaccard similarity > 0.5 [22] were eliminated after preprocessing. The rapid size scaling of negative samples, 20-25% reduction due to similarity and high running time without adding variance within the datasets were why UNI has fewer suicidal tweets. Dataset H was built 2 months after SCO and UNI to see if our classifier fares well against any bias that may have been introduced due to the potentially temporal nature of suicide ideation propagation on Twitter such as celebrity deaths [23], [24].

B. Data Annotation

In order to evaluate the performance of the proposed methodology, all three datasets consisting of suicidal and non-suicidal tweets were annotated individually by three students of clinical psychology. Discussions were held for conflicting annotations, 78% of which led to a conclusion, and the rest (high fluctuation) were marked as suicidal to reduce false negatives. A satisfactory agreement between the annotators can be inferred from Table 2. This classification was done on the basis of a binary criterion (Does this text imply self-harm inflicting tendencies or suicidal intent?) summarized below.

1) Suicidal tweet

Tweets displaying a serious tendency of self harm or suicidal ideation; e.g., I want to slit my wrists or I need to kill myself or I should just jump down a building and die.;

Tweets discussing previous suicide attempts and/or devising plans to suicide; e.g., Why can't I even suicide successfully? or What are some of the best ways to kill myself?;

Tweets representative of depression, substance abuse, etc with no clear indication of suicidal tendencies were not included in this category.

2) Non-Suicidal tweet

Tweets with no suicidal ideation;

Flippant references to suicide; e.g., Hahahah I'm gonna kill myself, lol.;

Condolences and Suicide related news; e.g., I'm so affected by Chester Benington's suicide, this makes me sad.; Politician accused of driving wife to suicide.

Campaigning and Support; e.g., If you're thinking about suicide, please reach out to family.

Disagreements were discussed which led to the resolution of conflicts arising in the annotation process. Thus, the annotation process followed was iterative and led to the updating of the annotation guidelines thrice-over. We believe these can now be shared with the community for scraping posts and text-based content for the development of larger datasets.

IV. EXPERIMENT SETTINGS AND RESULTS

The number of epochs for the firefly algorithm was initially set as 100, inspired by previous work done. Increments of 20 led to an increase in both accuracy and computational time. Accuracy-computational time trade-off saturated over 200. Based on the work done by Yang and He [11], [26] we experimented with 30, 25, 20 and 15 agents. Using 20 agents led to a higher accuracy, similar feature reduction and lower

TABLE III
FEATURE SET FOR POTENTIAL SUICIDAL INTENT DETECTION.

Feature	Description
Standard Features	
Statistical Features	Unigram & Bigram presence, count and length
Part of Speech count	Count of nouns, verbs, adverbs, adjectives in the tweet
TF-IDF	TF-IDF of unigrams and bigrams in the tweet
Specific Features	
LIWC Features	Generative labels reflecting author's psychological state
LDA unigrams	Presence and count of unigrams from LDA
Sentiment Ratio	The ratio of positive and negative emotive words
Emoji Sentiment	Cumulative sentiment, number and distance [25]

TABLE IV
COHEN'S KAPPA FOR ANNOTATORS H1; H2 AND H3

	H1	H2	H3
H1		0.84	0.79
H2	0.84		0.79
H3	0.79	0.91	

running time. A plot between (Accuracy/Number of reduced features) vs k attained maxima at k = 0.12. Existing studies revealed that the value of light absorption is not crucial to any given problem, and the typical value of 1 works well. While studies recommend attractiveness = 1, high variance in suicidal tweets motivated us to test lower values and select 0.2. Using the proposed feature set, and using BFA for feature selection (parameters shown in Table 3) a Random Forest classifier was trained. The results of ten-fold cross-validation performance on all three datasets is measured in terms of Accuracy (A), Precision (P), Recall (R) and F1 score (F1) and Feature Reduction (). RF with no feature selection and RF with Principal Component Analysis (PCA) tuned to achieve the same feature reduction as BFA are used as baselines alongside Support Vector Machines (SVM) which have shown promise in suicide ideation classification. Deep Learning approaches such as CNN-LSTMs, RNNs and LSTMs are used for comparison as they can learn good features automatically during training. The methodology presented in the work [27], denoted as BURN-MULTI is also used as a baseline to compare the proposed methodology with existing successful models. Additionally, traditional classification models, Logistic Regression (LR) and XGBoost are also used in conjunction with the Firefly Algorithm to demonstrate the improvement independent of the classifier. Other wrapper based feature selection methods, particularly, Genetic Algorithms, and Particle Swarm Optimization are also employed. Recursive Feature Elimination (RFE) is also included in the experimental analysis with the held out dataset to elucidate the superiority of Firefly Algorithm over conventional feature selection approaches. Given the stochastic nature of the Firefly algorithm, the results obtained are the mean of 100 independent runs 95% confidence interval.

A. Results for datasets SCO and UNI

Table 5 shows the results obtained for Datasets SCO and UNI, with and without Feature Selection using BFA. Dataset SCO comprises fewer samples with no correlation with sui-

cide, and more samples relevant to suicide but not suicidal ideation. This leads to a precision-recall trade-off with a high precision, implying the ability to identify posts with no suicidal ideation, but a weaker ability to capture a subset of tweets unrelated to suicide. Compared to the classifier trained on Dataset SCO, the classifier trained on Dataset UNI performs better, implying that it is essential to train classifiers with tweets covering various other topics to increase robustness towards noise. For both datasets, a fair increase in terms of evaluation metrics is observed with the inclusion of BFA. This can be attributed to the ability of Firefly algorithm based feature selection to weed out redundant features particularly, Unigrams, Bigrams, LIWC features not pertaining to cognitive and psychological aspects. RF and XGBoost also outperform LR, SVM, RNN, and LSTM models significantly (significance level = 0:05, permutation-based statistical comparison test, p-value = 0:008), and thus we use RF, XGBoost, CNN-LSTM and BURNMULTI models for further analysis on the held-out test data set and error analysis.

Table 6 lists the best 10 features for both datasets obtained using the Chi-squared test-based ranking on the features selected by BFA. Chi-squared testing helps in analyzing the top features for gaining a better insight into the features only. Firefly algorithm can evaluate the combined contribution of features (some highly correlated) and automatically fine-tune the number of features. Like PCA, Chi-squared testing with a fixed feature reduction is unable to outperform the Firefly algorithm. Interestingly, LIWC features and POS counts are useful and efficient.

B. Results for Dataset H

Table 7 presents the results for the models that performed substantially well for datasets SCO and UNI, for Dataset H, and validates the performance the proposed model in comparison to human annotators thus avoiding the effects of overfitting. The superior performance of the classifier trained on dataset UNI over dataset SCO is attributed to the imbalance in dataset and the lack of non suicide related tweets. Higher precision is observed with the best performing classifiers trained on dataset SCO, as it contains more positive samples that allow the classifier to capture specific aspects of suicide-related tweets with no suicidal ideation. RF-BFA outperforms BURNMAP and models like RF-RFE due to its stochastic

TABLE V
RESULTS FOR DATASETS SCO AND UNI

	A	P	R	F1
Dataset SCO				
Random Forest	78.27	80.21	77.36	78.76
SVM	75.21	73.21	72.43	72.81
LR	74.02	71.92	71.21	
XGBoost	78.91	79.53	81.29	80.13
RNN	77.31	73.12	75.62	75.01
LSTM	81.32	78.83	80.25	79.94
CNN-LSTM	83.01	82.47	84.34	83.57
BURNMULTI	82.46	83.83	82.94	83.41
Random Forest + Binary Firefly Algorithm	83.18	87.12	84.73	85.91
SVM + Binary Firefly Algorithm	82.01	81.92	79.23	80.94
LR + Binary Firefly Algorithm	80.43	75.29	77.41	76.82
XGBoost + Binary Firefly Algorithm	81.38	82.98	84.38	83.63
Dataset UNI				
Random Forest	82.33	84.85	83.71	84.28
SVM	76.02	78.21	75.34	77.90
LR	74.35	71.44	73.25	72.59
XGBoost	79.23	81.52	84.13	82.91
RNN	71.24	73.12	68.42	70.23
LSTM	75.84	77.24	79.93	79.01
CNN-LSTM	80.97	83.49	84.11	83.92
BURNMULTI	79.85	80.32	82.42	81.15
Random Forest + Binary Firefly Algorithm	88.82	89.21	87.44	88.32
SVM + Binary Firefly Algorithm	80.85	83.21	78.38	82.46
LR + Binary Firefly Algorithm	78.96	75.92	79.23	77.42
XGBoost + Binary Firefly Algorithm	83.48	82.98	85.54	84.11

TABLE VI

RESULTS FOR DATASET H; TEST SET

	A	P	R	F1
Classifier trained on Dataset SCO				
Random Forest	69.63	75.31	68.19	71.57
Random Forest + Binary Firefly Algorithm	74.27	81.23	72.91	76.84
Random Forest + RFE	70.34	73.21	70.01	71.94
XGBoost	70.31	74.42	70.01	72.31
XGBoost + Binary Firefly Algorithm	73.13	80.42	73.51	76.62
XGBoost + RFE	71.26	74.93	70.35	72.74
CNN-LSTM	72.74	81.13	78.48	79.75
BURNMULTI	71.92	77.21	65.31	72.26
Classifier trained on Dataset UNI				
Random Forest	75.31	73.19	75.12	74.14
Random Forest + Binary Firefly Algorithm	78.93	79.31	80.03	79.66
Random Forest + RFE	76.08	77.83	78.32	78.12
XGBoost	74.38	72.93	73.49	73.13
XGBoost + Binary Firefly Algorithm	78.85	78.93	79.72	79.57
XGBoost + RFE	76.28	74.53	77.21	75.83
CNN-LSTM	77.63	80.13	78.37	79.03
BURNMULTI	75.69	75.27	76.73	75.98

annotators, e.g., "I need to kill myself now, Hah, no seriously, I'm serious. I'm done :)".

nature and ability to fine tune feature reduction (BURNMAP implicitly uses PCA that requires manual tuning).

C. Error Analysis

A brief error analysis is presented in this subsection, highlighting both the limitations and scope for improvement in the proposed methodology. Some categories of errors are:

- 1) High fluctuation: Tweets reflecting multiple mood swings and fluctuations are difficult for both the classifier and human

2) Pragmatic difficulty: The tweet "I lost my baby. Signing off.." was correctly identified by human annotators but not by the classifier. This is attributed to the inability to link the contexts of the multiple sentences and the lack of explicit

suicidal ideation opposed to the suicidal language used for training.

3) Strict feature reduction: Assigning a large weight to the penalty function reduces computation and improves overall accuracy but leads to incorrect classification of tweets hinting subtly towards suicidal ideation such as "It is time to sleep forever.." by eliminating uncommon

TABLE VII
BEST 10 FEATURES FOR DATASETS SCO AND UNI

Best features	
Dataset SCO	Dataset UNI
"KILL ME"	LIWC-Authenticity
LDA-suicide	POS-Noun
POS-Verb	LIWC-Emotional tone
LIWC-Sadness	LDA-die
"COMMIT SUICIDE"	LDA-suicide
POS-Adjective	"DIE NOW"
LIWC-Authenticity	Sentiment Ratio
Sentiment Ratio	"SLIT WRIST"
POS-Adverb	LIWC-Anxiety
LIWC-Death	Pos-Adjective

social media," in Proceedings of the 2016 CHI conference on human factors in computing systems. ACM, 2016, pp. 2098–2110.

[3] W. W. Zung, "Suicide prevention by suicide detection," *Psychosomatics*, vol. 20, no. 3, pp. 153–155, 1979.

[4] R. Sawhney, P. Manchanda, R. Singh, and S. Aggarwal, "A computational approach to feature extraction for identification of suicidal ideation in tweets," in Proceedings of ACL 2018, Student Research Workshop, 2018, pp. 91–98.

features in the dataset.

4) Composite emojis: Inspection of feature weights shows that the cumulative emoji sentiment feature may mislead the classifier in certain rare cases such as "I should really kill myself today :) (: (: (I can't anymore :P :("

V. CONCLUSION AND FUTURE WORK

Motivated by the increasing correlation between expression of suicidal ideation on social media and suicide rates, a model for classification of suicidal tweets is created with the aim of suicide prevention by detection. Datasets are scraped and manually annotated to validate the performance of a feature set proposed for the classification task. Random forest classifiers using the Binary Firefly Algorithm are employed, leading to a best accuracy of 78.93% on a held-out test dataset. Inclusion of feature selection and the proposed features set the foundation of future work in the field of suicidal ideation detection. Finally, an error analysis presents the limitations and challenges faced by the proposed classifier. Our work aims to complement existing avenues such as suicide helplines and real world interactions/meetings/sessions. Given that the domain of our work is largely limited by the availability of such data online, apart from social media and suicide forums, our best attempt was to build datasets that can as broadly as possible cover different forms of suicide ideation text. We used several forums such as www.suicideforum.com, www.takethislife.com, www.beyondblue.org.au, etc. and microblogging social media: Tumblr and Reddit to develop a lexicon for scraping tweets. The analysis presented in this paper make a case for feature selection and the proposed feature set as an effective tool for suicide prevention.

REFERENCES

- [1] D. D. Luxton, J. D. June, and J. M. Fairall, "Social media and suicide: a public health perspective," *American journal of public health*, vol. 102, no. S2, pp. S195–S200, 2012.
- [2] M. De Choudhury, E. Kiciman, M. Dredze, G. Coppersmith, and M. Kumar, "Discovering shifts to suicidal ideation from mental health content in

- [5] B. Xue, M. Zhang, W. N. Browne, and X. Yao, "A survey on evolutionary computation approaches to feature selection," *IEEE Transactions on Evolutionary Computation*, vol. 20, no. 4, pp. 606–626, 2016.
- [6] A. E. Aladag, S. Muderrisoglu, N. B. Akbas, O. Zahmacioglu, and H. O. Bingol, "Detecting suicidal ideation on forums: Proof-of-concept study," *Journal of medical Internet research*, vol. 20, no. 6, p. e215, 2018.
- [7] A. Pourmand, J. Roberson, A. Caggiula, N. Monsalve, M. Rahimi, and V. Torres-Llenza, "Social media and suicide: a review of technology-based epidemiology and risk assessment," *Telemedicine and e-Health*, 2018.
- [8] R. Sawhney, P. Manchanda, P. Mathur, R. Shah, and R. Singh, "Exploring and learning suicidal ideation connotations on social media with deep learning," in *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 2018, pp. 167–175.
- [9] A. Jain and D. Zongker, "Feature selection: Evaluation, application, and small sample performance," *IEEE transactions on pattern analysis and machine intelligence*, vol. 19, no. 2, pp. 153–158, 1997.
- [10] E. Keogh and A. Mueen, "Curse of dimensionality," in *Encyclopedia of Machine Learning and Data Mining*. Springer, 2017, pp. 314–315.
- [11] X.-S. Yang and X. He, "Firefly algorithm: recent advances and applications," *International Journal of Swarm Intelligence*, vol. 1, no. 1, pp. 36–50, 2013.
- [12] J. W. Pennebaker, M. E. Francis, and R. J. Booth, "Linguistic inquiry and word count: Liwc 2001," *Mahway: Lawrence Erlbaum Associates*, vol. 71, no. 2001, p. 2001, 2001.
- [13] N. Bird, "the natural language toolkit, 2006, acl," in *Proceedings of the COLING/ACL on Interactive presentation sessions*, 2006, pp. 69–72.
- [14] T. Wilson, J. Wiebe, and P. Hoffmann, "Recognizing contextual polarity in phrase-level sentiment analysis," in *Proceedings of the conference on human language technology and empirical methods in natural language processing*. Association for Computational Linguistics, 2005, pp. 347–354.
- [15] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [16] A. Liaw, M. Wiener et al., "Classification and regression by randomfor-est," *R news*, vol. 2, no. 3, pp. 18–22, 2002.
- [17] R. Sawhney, P. Mathur, and R. Shankar, "A firefly algorithm based wrapper-penalty feature selection method for cancer diagnosis," in *International Conference on Computational Science and Its Applications*. Springer, 2018, pp. 438–449.
- [18] "Api reference index - twitter developers." [Online]. Available: <https://dev.twitter.com/rest/public/search>
- [19] P. Burnap and M. L. Williams, "Cyber hate speech on twitter: An application of machine classification and statistical modeling for policy and decision making," *Policy & Internet*, vol. 7, no. 2, pp. 223–242, 2015.
- [20] P. R. Recupero, S. E. Harms, and J. M. Noble, "Googling suicide: surfing for suicide information on the internet." *The Journal of clinical psychiatry*, 2008.
- [21] J. Ramos et al., "Using tf-idf to determine word relevance in document queries," in *Proceedings of the first instructional conference on machine learning*, vol. 242, 2003, pp. 133–142.
- [22] P. Achananuparp, X. Hu, and X. Shen, "The evaluation of sentence similarity measures," in *International Conference on data warehousing and knowledge discovery*. Springer, 2008, pp. 305–316.
- [23] M. Kumar, M. Dredze, G. Coppersmith, and M. De Choudhury, "Detecting changes in suicide content manifested in social media following celebrity suicides," in *Proceedings of the 26th ACM Conference on Hypertext & Social Media*. ACM, 2015, pp. 85–94.
- [24] K.-w. Fu and P. S. Yip, "Long-term impact of celebrity suicide on suicidal ideation: Results from a population-based study," *Journal of Epidemiology & Community Health*, vol. 61, no. 6, pp. 540–546, 2007.
- [25] P. K. Novak, J. Smailovic, B. Sluban, and I. Mozetic, "Sentiment of emojis," *PLoS one*, vol. 10, no. 12, p. e0144296, 2015.
- [26] R. Sawhney, R. Shankar, and R. Jain, "A comparative study of transfer functions in binary evolutionary algorithms for single objective optimization," in *International Symposium on Distributed Computing and Artificial Intelligence*. Springer, 2018, pp. 27–35.
- [27] P. Burnap, G. Colombo, R. Amery, A. Hodorog, and J. Scourfield, "Multi-class machine classification of suicide-related communication on twitter," *Online social networks and media*, vol. 2, pp. 32–44, 2017.