Unlocking Emotions in Text: A Fusion of Word Embedding and Lexical Knowledge for Emotion Classification

Anjali Bhardwaj, Nesar Ahmad Wasi and **Muhammad Abulaish**
South Asian University, New Delhi, India
{bhardwaj.anjali200594, nesarahmadwasi17, abulaish.sau}@gmail.com

Abstract

This paper introduces an improved method for emotion classification through the integration of emotion lexicons and pre-trained word embeddings. The proposed method utilizes semantically similar features to reconcile the semantic gap between words and emotions. The proposed approach is compared against three baselines for predicting Ekman's emotions at the document level on the `GoEmotions` dataset. The effectiveness of the proposed approach is assessed using standard evaluation metrics, which show at least a 5% gain in performance over baselines.

1 Introduction

Emotion classification, poses a intricate challenges involving the task of associating a given text with the best-fitting emotion(s) that encapsulate the author's state of mind (Li et al., 2020). Existing emotion classification models primarily rely on individual words and employ various architectural designs to extract semantic information from sequential inputs, thereby comprehending emotions at the sentence level. Recognizing the domain-dependent nature of the emotion classification task, numerous efforts have been made to enrich the understanding of word-level emotional content by constructing and leveraging various manually crafted lexicons. These lexicons have exhibited significant performance of classifiers (Taboada et al., 2011; Mudinas et al., 2012)

The main challenge addressed in the proposed work pertains to identifying the intricate connections between textual content and emotions. To tackle this challenge, we utilize emotion lexicons to bridge the semantic gap between words and emotions. Given that a single emotion lexicon may not comprehensively cover all conceivable words or phrases associated with emotions, we employ a combination of four emotion lexicons. Rooted in emotion theory, these commonly used lexicons can be categorized as discrete or dimensional lexicons (Borod, 2000; Yadollahi et al., 2017). Discrete lexicons, exemplified by NRC Word-Emotion Association Lexicon, also known as EmoLex (Mohammad and Turney, 2013) and EmoSenticNet (Poria et al., 2014), attribute one or more tags from Ekman's (Ekman, 1999) six basic emotions to each word. Although these lexicons establish a direct link between words and emotions, they are limited range of emotions and provide no insight into emotion intensity. In contrast, dimensional lexicons, such as NRC-Emotion Intensity Lexicon (NRC-EIL) (Mohammad, 2018b) and NRC Valence, Arousal, and Dominance (NRC-VAD) (Mohammad, 2018a), assign intensity values based on 1-D, 2-D, or 3-D models. These dimensions include valence, arousal, and dominance (VAD) and offer a more nuanced understanding of emotional content.

The major contributions of the proposed work can be summed up as follows. (i) a features extraction technique to extract linguistic and semantically-similar features-based vector representation for the emotion classification task, (ii) a pre-trained word embeddings-based linguistic features enriched approach to enhance the performance of emotion classification task, and (iii) an empirical study on emotion classification using Google's GoEmotion dataset.

2 Related Work

Textual emotion classification has steadily emerged as a dynamic and evolving research domain, witnessing a proliferation of methodologies and techniques over the years (Faruqui et al., 2015; Zahiri and Choi, 2018; Li et al., 2020; Bhardwaj and Abulaish, 2022). With the remarkable strides made in machine learning, deep learning, and word embeddings, the precision and efficacy of emotion classification have witnessed remarkable advancements

(Zahiri and Choi, 2018; Li et al., 2020). For instance, Zahiri and Choi (2018) laid the foundation for creating a multiparty dialogue corpus dedicated to text data to categorize emotions.

Several studies have pursued the enhancement of classifier performance by integrating external knowledge derived from emotion lexicons and syntactic structures. Researchers have adeptly intertwined lexicon-based methodologies with machine learning models, yielding notable improvements in classification accuracy (Hu et al., 2013; Wasi and Abulaish, 2020; Faruqui et al., 2015; Agrawal et al., 2018; Wasi and Abulaish, 2024; Alm et al., 2005). Li et al. (2020) introduced the EmoChannelAttn network, which harnessed the underlying dependency relationships present within textual emotion constructs to enhance classification performance. Their study delved into a spectrum of 151 finely-grained emotions, utilizing domain-specific knowledge and dimensional lexicon dictionaries. Meanwhile, Faruqui et al. (2015) augmented existing word embeddings with lexical knowledge, increasing effective representation learning.

In quest to acquire word representations infused with emotional nuances, Agrawal et al. (2018) employed a methodology wherein emotionally analogous words were mapped into close proximity within vector spaces. In contrast, emotionally dissimilar words were distanced from each other. However, their approach relied on a categorical lexicon, which constrained its applicability to a specific set of emotions. Furthermore, Alm et al. (2005) embraced techniques to incorporate emotional words as salient features for text-based emotion prediction. Additionally, Hu et al. (2013) pioneered a lexicon-based, unsupervised sentiment analysis approach, incorporating emotional signals within `Twitter` datasets. Vo and Zhang (2015) undertook the generation of lexicon-based distributed contexts involving the selection of word vectors guided by a lexicon dictionary.

3 Proposed Approach

This section outlines the proposed approach, which capitalizes on a synergistic fusion of word embeddings and emotion lexicon-based features. The proposed approach is designed to elevate the performance of emotion classification by predicting emotion labels at the document-level. An illustrative overview of the proposed approach is presented in Figure 1.

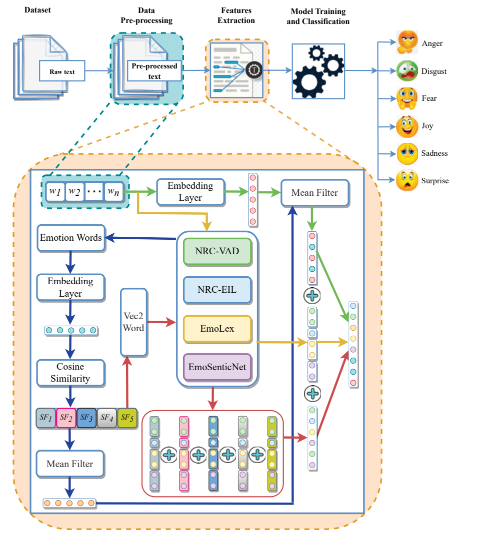

Figure 1: An overview of the proposed approach illustrates the relationships between various elements: the green line represents the vector for semantically-similar features, the yellow line denotes the vector for linguistic features, and the red line signifies the similar features-based linguistic features.

3.1 Problem Definition

Consider a function f that maps each document d_n from a corpus D to one of the multiple predefined labels c_i from a set K. Formally, the function $f : D \rightarrow K$ is defined as $f(d_n) = c_i$ where c_i is the label for the corresponding emotion in document d_n, and $K = \{c_1, c_2, ..., c_k\}$ represents a predefined set of emotion labels, with k denoting the number of labels. In order to classify a document d_n to its associated label c_i, it needs to be transformed to its vector representation form. The following steps are undertaken to transform a textual document d_n to its vector representation.

3.2 Word-level Features Vector Generation

NLP empowers machines to understand and derive insights from human language. As machine learning algorithms are unable to interpret the semantics of textual data directly, words need to be translated into numerical representations.

Each document d_n can further be decomposed into a vocabulary of words, represented as $V = \{w_1, w_2, ..., w_{|V|}\}$, with $|V|$ denoting the count of unique tokens. Formally, a *word* is defined as a d-dimensional vector, i.e., $w_i \in \mathbb{R}^d$, representing the i^{th} word within the corpus. Since feature extraction is a significant and complex task in emotion classification, the proposed approach leverages pre-trained word embeddings and emotion lexicons to enhance emotion classification.

To generate an enriched emotion representation for each word, we employ several components that are utilized across various phases of the proposed approach. The components are briefly described as follows: (i) Word Embedding Model: To obtain word embedding for each word, we leverage the `GloVe` (Demszky et al., 2020) vector model, which takes a word w_i as input and generates a $100-$dimension word embedding vector representation. (ii) Lexicon Models: We utilize four well-known emotion lexicon models, namely EmoLex (Mohammad and Turney, 2013), EmoSenticNet (Poria et al., 2014), NRC-VAD (Mohammad, 2018a), and NRC-EIL (Mohammad, 2018b) to extract emotion words and their representations based on these models. Each instance in the corpus is categorized into six basic Ekman's (Ekman, 1999) emotions, i.e., *anger, sadness, joy, fear, disgust,* and *surprise*. (iii) Cosine Similarity: It measures similarity between a pair of word vectors. (iv) Mean Filter: It is used to aggregate information for all words in a text by computing the sum of elements and dividing it by the number of elements. (v) Concatenation: It is a process of appending feature vectors to create a new vector of the same dimension as the word vector.

3.2.1 Semantically-Similar Features-vector Generation

We exploit the publicly available pre-trained word embedding model, namely `GloVe` (Pennington et al., 2014), to extract relevant textual features. For each word w_i a $l-$dimensional word embedding vector (T^l) is obtained through the `GloVe` vector model (G), resulting in a dense vector representation defined in equation (1).

$$T^l = G^l(w_i) \tag{1}$$

Simultaneously, each word is passed through the aforementioned four lexicon models to obtain the emotion words. If the word is present in the lexicons, it is referred to as an emotion word; otherwise, it is ignored. For each emotion word, we obtain its $l-$dimensional representation (T_E^l) from the `GloVe` vector model (G). Further, we calculate the cosine similarity between the emotion words to identify the top m similar feature words, which is formally defined in equation (2). The higher the cosine similarity, the more similar the two words are.

$$CS(\overrightarrow{T^l_{E_i}}, \overrightarrow{T^l_{E_j}}) = \frac{\overrightarrow{T^l_{E_i}} \cdot \overrightarrow{T^l_{E_j}}}{\|\overrightarrow{T^l_{E_i}}\| \|\overrightarrow{T^l_{E_j}}\|} \tag{2}$$

In equation (2), $\overrightarrow{T^l_{E_i}}$ and $\overrightarrow{T^l_{E_j}}$ denote the vectors for the i^{th} and j^{th} emotion word, and $CS(\overrightarrow{T^l_{E_i}}, \overrightarrow{T^l_{E_j}})$ represents the cosine similarity between these two emotion words. The similar feature words consist of $l-$dimensional vectors, collectively forming the T^l_{SF} vector. Once we identify the top m similar features, denoted as $\{SF_1, SF_2,...,SF_m\}$, we apply the mean filter to obtain a singular $l-$dimensional Similar Feature (SF) vector, as defined in equation (3).

$$SF^l = \frac{1}{m} \sum_{j=1}^{m} T^l_{SF_j} \tag{3}$$

Subsequently, a mean filter is applied over the amalgamation of the T^l and the SF^l vectors to obtain the resultant $l-$dimensional semantically-similar features-vector (LE), as defined in equation (4). The resultant vector effectively encapsulates both emotional and semantically-similar features of the words.

$$LE^l = \frac{1}{2}(T^l + SF^l) \tag{4}$$

3.2.2 Linguistic Features Generation

We utilize emotion-specific lexicons to provide fine-grained emotional knowledge associated with words, enhancing our understanding of emotions within different contexts. As the semantically-similar features-vector representation LE^l is being generated, each word is simultaneously processed by four well-known emotion lexicon models to derive its emotion-enriched representation according to each model. These models, EmoLex (Mohammad and Turney, 2013) (P), EmoSenticNet (Poria et al., 2014) (Q), NRC-VAD (Mohammad, 2018a) (R), and NRC-EIL (Mohammad, 2018b) (S) generate, $x-$dimensional, $y-$dimensional, $z-$dimensional, and $u-$dimensional vectors, respectively. Once these vectors are generated, they are concatenated to form a more contextually rich $e-$dimensional Linguistic-based (L) vector, as described by equation (5).

$$L^e = P^x \oplus Q^y \oplus R^z \oplus S^u \tag{5}$$

In equation (5), $\oplus$ represents the concatenation of the vectors.

3.2.3 Similar Features-based Linguistic Features Generation

To derive similar feature-based linguistic features, we employ the m similar features obtained from equation (3). The m similar features are then mapped back to the corresponding words using the GloVe vector model (T^l). Subsequently, the m similar feature words serve as input to the four lexicon models. For each similar feature word, every lexicon model generates a vector representation. Concatenating these individual representations, as shown in equation (5), results in an $e-$dimensional similar feature L_{SF} vector.

The $e-$dimensional vectors associated with each of the m similar feature words are again concatenated to produce the final $v-$dimensional Similar Features-based Linguistic (EL) vector, as described in equation (6), with a dimension of $v = m{\cdot}e$. The EL vector adeptly captures the emotional context embedded within similar feature words.

$$EL^v = L^e_{SF_1} \oplus L^e_{SF_2} \oplus ... \oplus L^e_{SF_m} \tag{6}$$

3.2.4 Features Vector Generation

The final features vector serves as a pivotal element in emotion classification, offering a comprehensive representation of the text's emotional and semantic content. It concatenates three essential elements: Semantically-Similar Features vector (LE), Linguistic-based vector (L), and Similar Features-based Linguistic Features vector (EL). While LE combines word embeddings with emotional context, L enriches words with emotion-specific knowledge from four lexicons. Meanwhile, EL captures shared emotional context among words. The fusion of these components results in a robust $f-$dimensional feature vector, as defined in equation (7), that effectively enhances emotion classification tasks.

$$F^f = LE^l \oplus L^e \oplus EL^v \tag{7}$$

3.3 Document-level Features Vector Generation

At the document-level, the feature vector is constructed by concatenating the word-level feature vectors of all words within the document. The aggregation allows the document-level feature vector to encapsulate the collective emotional and semantic content of the entire text. The generated vector serves as input for machine learning models, enabling the accurate classification of expressed emotions within the document. The proposed approach provides a holistic view of emotions at the document-level, allowing for a more profound understanding of the text's emotional nuances.

Table 1: Statistics of the GoEmotions dataset.

Emotion	Number of instances
Anger	2,390
Disgust	1,491
Fear	925
Joy	2,520
Sadness	2,038
Surprise	1,691
Total	11,055

4 Experimental Setup and Results

4.1 Dataset and Evaluation Metrics

We have conducted experiments on the Google AI GoEmotions (Demszky et al., 2020) dataset, which consists of 58,000 Reddit English comments, each labeled with one of 27 emotions or *neutral*. Given the context of our proposed approach in emotion classification, instances in the dataset with multiple emotion class labels were excluded. Following the state-of-the-art approaches, we focused on the six basic Ekman (Ekman, 1999) emotions: *anger, disgust, joy, fear, sadness*, and *surprise*. The subset of instances associated with these six emotions accounted for a total of 11,055 samples, as detailed in Table 1. We have used the *macro-average f1-score* (Grandini et al., 2020) ($F1_{mac}$) along with accuracy *(Acc), macro-average precision* (Pr_{mac}), and *macro-average recall* (Re_{mac}) to evaluate the efficacy of the proposed approach.

4.2 Data Pre-processing

Our first step involved expanding contractions to improve text clarity. For this, we employ the pycontractions library, known for its higher accuracy compared to the standard Python contractions library, as it considers the text's grammar. We ensured uniformity in the text data by converting uppercase letters to lowercase. Subsequently, the text was tokenized through breaking it down into individual words. We removed English stopwords from the text using the NLTK library (Wagner, 2010), followed by lemmatization to reduce words to their base form for contextual meaning. Finally, punctuation marks, non-alphabetic words (including numbers and symbols), and one-lettered words were removed from the text.

Table 2: Performance evaluation results for the baselines and proposed approach.

Approach	Acc			Pr_{mac}			Re_{mac}			$F1_{mac}$		
	SVM	GB	NB	SVM	GB	NB	SVM	GB	NB	SVM	GB	NB
BL@16	43.3	47.0	33.9	44.2	45.2	32.3	37.0	44.0	29.8	35.5	44.3	28.0
BL@96	45.6	47.8	35.4	47.0	47.8	37.4	41.5	46.0	34.1	42.3	46.4	31.5
E-*TF-IDF*	51.1	52.8	28.7	49.7	**60.3**	31.5	48.6	47.5	30.9	48.9	49.6	28.6
Proposed	**56.7**	**58.5**	**42.4**	**55.6**	57.3	**43.3**	**54.0**	**55.6**	**40.4**	**54.5**	**56.1**	**40.0**

4.3 Implementation Details

We employed Stanford's `GloVe` to obtain word embedding for each word. We also employed four well-known emotion lexicons for an emotion-enriched representation of words. EmoLex (Mohammad and Turney, 2013) and EmoSenticNet (Poria et al., 2014) provide words associated with emotion scores, denoting whether the word is connected with the emotion or not. NRC-VAD lexicon (Mohammad, 2018a) offers real-valued scores associated with specific words. NRC-EIL (Mohammad, 2018b) provides emotion words with real-valued intensity scores. These lexicons take a word w_i as input and generate a $16-$dimension word vector representation. We have considered top 5 similar features. The dimensions of the resulting vectors are as follows: LE^l is $100-$dimensional, L^e is $16-$dimensional, and EL^v is $80-$dimensional. The final feature vector F^f is a combined $196-$dimensional vector, amalgamating pre-trained word embeddings and emotion lexicons.

4.4 Results and Analysis

The proposed approach was applied on the GoEmotions dataset, considering the following classifiers: (i) SVM (Boser et al., 1992; Wang et al., 2006); specifically, we employed the SVM-OvA (one-vs-all or one-vs-the-rest) strategy, (ii) Gradient Boosting (GB) (Freund et al., 1999), and (iii) Naïve Bayes (NB) (Rennie, 2001). We train classifiers to predict the six basic Ekman's emotions at the document-level. The proposed approach is compared with the following three baselines: (i) $BL@16$: Each document is represented by a $16-$dimensional feature vector, comprising emotion-related features extracted from emotion lexicons. This feature vector effectively capturing the essence of the document in terms of its emotional characteristics. (ii) $BL@96$: Each document utilizes features associated with emotion lexicons and merges them with similar features. The resultant features vector comprises of $96-$dimension. (iii) E-*TF-IDF*: Each document in the GoEmotion dataset is depicted through a unigram *TF-IDF* vector, resulting in $8937-$dimension. Subsequently, emotion-related features are integrated to create an emotionalized *TF-IDF* representation.

Table 2 presents the performance evaluation results of the proposed approach and the baselines. It can be observed that the proposed approach consistently outperforms the baseline models across all three classification algorithms (SVM, GB, and NB). It can also be observed that $BL@16$, which only utilizes the linguistic features, performs the worst compared to other baselines and the proposed approach using all classifiers. In $BL@96$, we utilized linguistic features and semantically-similar features, and it can be seen that the results are also improved compared to $BL@16$. In E-*TF-IDF*, we used the *TF-IDF* document representation, which improved the performance compared to $BL@16$ and $BL@96$; however, it is computationally expensive because the dimension of the feature vectors increases. In contrast, the proposed approach outperformed all three baselines in terms of all metrics by a significant margin of at least 5%. Moreover, the proposed approach is computationally inexpensive because it represents each document as a $196-$dimensional feature vector.

5 Conclusion

In this paper, we have presented an enhanced approach for emotion classification by leveraging semantic and linguistic-based feature representations. To transform each document for the required representation, we have used linguistic and semantically-similar features-based vector representations, i.e., linguistic-based features, semantically-similar features, and semantically-similar linguistic features. The empirical analysis conducted on the GoEmotions dataset demonstrates that classification models significantly outperform baseline approaches by a substantial margin of 5% when employing feature vectors derived from the fusion of pre-trained word embedding and lexicon knowledge.

References

Ameeta Agrawal, Aijun An, and Manos Papagelis. 2018. Learning Emotion-enriched Word Representations. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 950–961, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Cecilia Ovesdotter Alm, Dan Roth, and Richard Sproat. 2005. Emotions from Text: Machine Learning for Text-based Emotion Prediction. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 579–586, Vancouver, British Columbia, Canada. Association for Computational Linguistics.

Anjali Bhardwaj and Muhammad Abulaish. 2022. MINDS: A Multi-label Emotion and Sentiment Classification Dataset Related to COVID-19. *WNLPe-Health 2022: The First Workshop on Context-aware NLP in eHealth*, pages 3–13.

Joan C Borod. 2000. *The Neuropsychology of Emotion*. Affective Science.

Bernhard E Boser, Isabelle M Guyon, and Vladimir N Vapnik. 1992. A Training Algorithm for Optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152. Association for Computing Machinery.

Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan S. Cowen, Gaurav Nemade, and Sujith Ravi. 2020. GoEmotions: A Dataset of Fine-Grained Emotions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL*, pages 4040–4054. Association for Computational Linguistics.

Paul Ekman. 1999. Basic Emotions. *Handbook of cognition and emotion*, 98(16):45–60.

Manaal Faruqui, Jesse Dodge, Sujay Kumar Jauhar, Chris Dyer, Eduard Hovy, and Noah A. Smith. 2015. Retrofitting Word Vectors to Semantic Lexicons. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1606–1615, Denver, Colorado. Association for Computational Linguistics.

Yoav Freund, Robert Schapire, and Naoki Abe. 1999. A Short Introduction to Boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780):1612.

Margherita Grandini, Enrico Bagli, and Giorgio Visani. 2020. Metrics for Multi-class Classification: an Overview. *arXiv preprint arXiv:2008.05756*.

Xia Hu, Jiliang Tang, Huiji Gao, and Huan Liu. 2013. Unsupervised Sentiment Analysis with Emotional Signals. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 607—-618, New York, NY, USA. Association for Computing Machinery.

Zongxi Li, Xinhong Chen, Haoran Xie, Qing Li, and Xiaohui Tao. 2020. Emochannelattn: Exploring Emotional Construction Towards Multi-Class Emotion Classification. In *2020 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, pages 242–249. IEEE.

Saif M. Mohammad. 2018a. Obtaining Reliable Human Ratings of Valence, Arousal, and Dominance for 20,000 English Words. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL*, volume 1, pages 174–184, Melbourne, Australia. Association for Computational Linguistics.

Saif M. Mohammad. 2018b. Word Affect Intensities. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC*, Miyazaki, Japan. European Language Resources Association (ELRA).

Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a Word-Emotion Association Lexicon. *Computational Intelligence*, 29(3):436–465.

Andrius Mudinas, Dell Zhang, and Mark Levene. 2012. Combining Lexicon and Learning based Approaches for Concept-Level Sentiment Analysis. In *Proceedings of the first international workshop on issues of sentiment discovery and opinion mining*, pages 1–8.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 1532–1543, Doha, Qatar. ACL.

Soujanya Poria, Alexander F. Gelbukh, Erik Cambria, Amir Hussain, and Guang-Bin Huang. 2014. Emosenticspace: A novel framework for affective common-sense reasoning. *Knowledge-Based Systems*, 69:108–123.

Jason DM Rennie. 2001. Improving Multi-class Text Classification with Naive Bayes.

Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly Voll, and Manfred Stede. 2011. Lexicon-Based Methods for Sentiment Analysis. *Computational linguistics*, 37(2):267–307.

Duy-Tin Vo and Yue Zhang. 2015. Target-Dependent Twitter Sentiment Classification with rich Automatic Features. In *Twenty-fourth international joint conference on artificial intelligence*, pages 1347—1353.

Wiebke Wagner. 2010. *Steven Bird, Ewan Klein and Edward Loper: Natural Language Processing with Python, Analyzing Text with the Natural Language Toolkit*, volume 44. O'Reilly Media, Inc., Beijing.

Zi-Qiang Wang, Xia Sun, De-Xian Zhang, and Xin Li. 2006. An Optimal SVM-Based Text Classification

Algorithm. In *2006 International Conference on Machine Learning and Cybernetics*, pages 1378–1381. IEEE.

Nesar Ahmad Wasi and Muhammad Abulaish. 2020. Document-level sentiment analysis through incorporating prior domain knowledge into logistic regression. In *Proceedings of the 19th IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, pages 969–974. IEEE.

Nesar Ahmad Wasi and Muhammad Abulaish. 2024. SKEDS — An External Knowledge Supported Logistic Regression approach for Document-level Sentiment Classification. *Expert Systems with Applications*, 238:121987.

Ali Yadollahi, Ameneh Gholipour Shahraki, and Osmar R Zaiane. 2017. Current State of Text Sentiment Analysis from Opinion to Emotion Mining. *ACM Computing Surveys (CSUR)*, 50(2):1–33.

Sayyed M Zahiri and Jinho D Choi. 2018. Emotion Detection on TV Show Transcripts with Sequence-based Convolutional Neural Networks. In *Workshops at the thirty-second aaai conference on artificial intelligence*.