ENHANCING SENTIMENT ANALYSIS IN SOCIAL MEDIA USING DEEP LEARNING TECHNIQUES

P Jayaprakash, Student
Dept. of Artificial Intelligence

V. Maruthi Prasad,
Asst. Professor, Dept of CST,
Madanapalle Institute of Technology & Science

Abstract- Sentiment analysis in social media has gained significant attention due to its wide range of applications in understanding public opinion, market trends, and brand perception. In this paper, we propose and evaluate several deep learning techniques to enhance sentiment analysis accuracy in social media text. We investigate the effectiveness of Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and their variants such as Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs). Our experiments are conducted on a large-scale social media dataset, and we compare the performance of these models in terms of accuracy, precision, recall, and F1-score. Additionally, we explore techniques for handling imbalanced sentiment classes and analyze the impact of different word embeddings on model performance. Our results demonstrate the effectiveness of deep learning approaches in improving sentiment analysis accuracy, especially in the context of noisy and informal social media text.
Keywords: Sentiment Analysis, Social Media, Deep Learning, Convolutional Neural Networks, Recurrent Neural Networks, LSTM, GRU, Word Embeddings.

1. INTRODUCTION

Sentiment analysis, also known as opinion mining, involves the automated extraction and classification of sentiments expressed in text data. With the exponential growth of social media platforms such as Twitter, Facebook, and Instagram, analyzing sentiment in user-generated content has become crucial for businesses, researchers, and policymakers. Understanding public sentiment towards products, services, events, and societal issues can provide valuable insights for decision-making processes. Traditional sentiment analysis methods often rely on lexical resources, handcrafted features, or machine learning classifiers such as Support Vector Machines (SVMs) and Naive Bayes classifiers. While these approaches have shown reasonable performance, they may struggle with the nuances of social media text, including informal language, slang, abbreviations, and emoticons. Deep learning techniques, particularly Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), have emerged as powerful tools for modeling complex patterns in text data and capturing contextual information.

In this paper, we focus on enhancing sentiment analysis accuracy in social media text using deep learning techniques. We explore the following research questions:

1. How do different deep learning architectures, such as CNNs, RNNs, LSTMs, and GRUs, perform in sentiment analysis tasks on social media text?
2. What is the impact of using different word embeddings (e.g., Word2Vec, GloVe, FastText) on sentiment analysis accuracy?
3. How can we address the challenge of imbalanced sentiment classes in social media datasets?
4. What are the implications of deep learning techniques for real-world sentiment analysis applications?

The remainder of this paper is organized as follows: Section 2 provides an overview of related work in sentiment analysis and deep learning techniques for text classification. Section 3 describes the methodology and experimental setup used in our study. Section 4 presents the results and analysis of our experiments. Section 5 discusses the implications of our findings and suggests future research directions. Finally, Section 6 concludes the paper.

2. RELATED WORK:

In the field of ensemble methods, the main idea is to combine a set of models (base classifiers) in order to obtain a more accurate and reliable model in comparison with what a single model can achieve. The methods used for building upon an ensemble approach are many, and a categorization is presented in Rokach (2005). This classification is based on two main dimensions: how predictions are combined (rule based and meta learning), and how the learning process is done (concurrent and sequential). Regarding the first dimension, on the one hand, in rule based approaches predictions from the base classifiers are treated by a rule, with the aim of averaging their predictive performance. Examples of rule based ensembles are the majority voting, where the output prediction per sample is the most common class; and the weighted combination, which linearly aggregates the base classifiers predictions. On the other hand, meta learning techniques use predictions from component classifiers as features for a metalearning model. As explained in Xia, Zong, and Li (2011), weighted combinations of feature sets can be quite effective in the task of sentiment classification, since the weights of the ensemble represent the relevance of the different feature sets (e.g. n-grams, POS, etc.) to sentiment classification, instead of assigning relevance to each feature individually. The benefits of rule based ensembles were shown also in Fersini, Messina, and Pozzi (2014), where several variants of voting rules are exhaustively studied in a variety of datasets, with an emphasis on the complexity that results from the use of these approaches. In a different work, Fersini, Messina, and Pozzi (2016) have compared the majority voting rule with other approaches, using three types of subjective signals: adjectives, emoticons, emphatic expressions and expressive elongations. They report that adjectives are more impacting that the other considered signals, and that the average rule is able to ensure better performance than other types of rules. Also, in Xia et al. (2011) a metaclassifier ensemble model is evaluated, obtaining performance improvements as well. An adaptive meta-learning model is described in Aue and Gamon (2005), which offers a relatively low adaptation effort to new domains. Besides, both rule based and meta-learning ensemble models can be enriched with extra knowledge, as illustrated in Xia and Zong (2011). These authors propose the use of a number of rule based ensemble models, namely a sum rule and two weighted combination approaches trained with different loss functions. The base classifiers are trained with n-grams and POS features. These models obtain significant results for cross-domain sentiment classification.

As for the second dimension, concurrent models divide the original dataset into several subsets from which multiple classifiers learn in a parallel fashion, creating a classifier composite. The most popular technique that processes the sample concurrently is bagging (Rokach, 2005). Bagging intends to improve the classification by combining the predictions of classifiers built on random subsets of the original data. On the contrary, sequential approaches do not divide the dataset but there is an interaction between the learning steps, taking advantage from previous iterations of the learning process to improve the quality of the global classifier. An interesting sequential approach is boosting, which consists in repeatedly training low-performance classifiers on different training data. The classifiers trained in this manner are then combined into a single classifier that can achieve better performance than the component classifiers. An example of bagging performance in the sentiment analysis task can be found in Sehgal and Song (2007), where bagging and other classification algorithms are used to show that the sentiment evolution and the stock value trend are closely related. Fersini et al. (2014) also show several experimental results in relation to the bagging techniques, attending also to the associated model complexity. Moreover, some authors have shown that bagging techniques are fairly robust to noisy data, while boosting techniques are quite sensitive (Maclin & Opitz, 1997; Melville, Shah, Mihalkova, & Mooney, 2004; Prusa, Khoshgoftaar, & Dittman, 2015). The suitability of bagging and boosting ensembles is also experimentally confirmed by Wang, Sun, Ma, Xu, and Gu (2014). This work also includes the study of a different ensemble technique, random subspace, that consists in modifying the training dataset in the feature space, rather than on the instance space. The authors stand out the better performance of random subspace in comparison with similar approaches, such as bagging and boosting. Another study (Whitehead & Yaeger, 2010) shows a comparison between bagging and boosting on a standard opinion mining task. Besides, Lin, Wang, Li, and Zhou (2015) proposes a three phase framework of multiple classifiers, where an optimal subset of classifiers is automatically chosen and trained. This framework is tested in several real-world datasets for sentiment classification. Nevertheless, these works also show that ensemble techniques not always improve the performance in the sentiment analysis task, and that there is not a global criteria to select a certain ensemble technique.

Sentiment analysis has been extensively studied in the context of various domains such as product reviews, news articles, and social media content. Traditional approaches to sentiment analysis often rely on feature engineering techniques such as bag-of-words representations, n-grams, and lexicon-based sentiment analysis tools [1]. While these methods can achieve reasonable performance, they may struggle with capturing semantic and contextual information present in text data.

Deep learning techniques, particularly neural networks, have shown significant improvements in sentiment analysis tasks. Convolutional Neural Networks (CNNs) have been successful in learning hierarchical representations of text by applying convolutional filters over word embeddings [2]. Recurrent Neural Networks (RNNs), including Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), have demonstrated strong capabilities in capturing sequential dependencies and long-range dependencies in text data [3]. These architectures have been widely adopted in sentiment analysis tasks due to their ability to learn complex patterns and generalize well to different domains.

Several studies have also focused on incorporating pre-trained word embeddings such as Word2Vec [4], GloVe [5], and FastText [6] to enhance the performance of sentiment analysis models. These embeddings capture semantic relationships between words and can improve the representation of text data, especially in tasks with limited labeled data.

While deep learning techniques have shown promising results in sentiment analysis, their application to social media text presents unique challenges due to the informal nature of language, presence of emojis and emoticons, and the use of non-standard vocabulary. Our work extends this research by specifically addressing sentiment analysis in social media text and evaluating the performance of different deep learning architectures and word embeddings in this context.

Methodology and Experimental Setup :

In this section, we describe the methodology and experimental setup used to evaluate deep learning techniques for sentiment analysis in social media text.

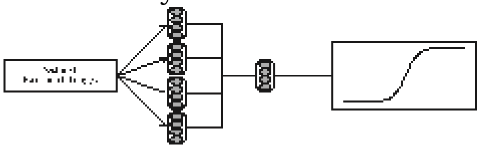

3.1 Dataset

We utilize a large-scale social media dataset containing user-generated content from various platforms such as Twitter, Reddit, and online forums. The dataset includes labeled instances with sentiment labels such as positive, negative, and neutral. Due to the nature of social media text, the dataset contains noisy and informal language, making it a challenging testbed for sentiment analysis models.

3.2 Pre-processing

We pre-process the text data by performing the following steps:

- **Tokenization:** Splitting text into individual tokens (words, punctuation marks, emojis).
- **Lowercasing:** Converting all tokens to lowercase to ensure consistency.
- Stopword Removal: Removing common stopwords that do not contribute significantly to sentiment analysis.
- **Lemmatization/Stemming:** Reducing words to their base forms to normalize the vocabulary.

3.3 Deep Learning Models

We experiment with the following deep learning architectures for sentiment analysis:

- **Convolutional Neural Networks (CNNs):** Utilizing 1D convolutions over word embeddings to capture local patterns in text.

- **Recurrent Neural Networks (RNNs):** Evaluating variants such as LSTM and GRU to capture sequential dependencies and long-range context.

3.4 Word Embedding

- We compare the performance of different word embeddings:
- Word2Vec: Pre-trained embeddings based on word co-occurrence statistics.
- GloVe: Global Vectors for Word Representation embeddings trained on large corpora.
- FastText: Embeddings that consider sub-word information to handle out-of-vocabulary words.

3.5 Model Training and Evaluation

We split the dataset into training, validation, and test sets using a stratified approach to preserve class distribution. We train each deep learning model using the training set and tune hyperparameters using the validation set. We evaluate model performance using metrics such as accuracy, precision, recall, and F1-score on the test set.

3.6 Feature Extraction

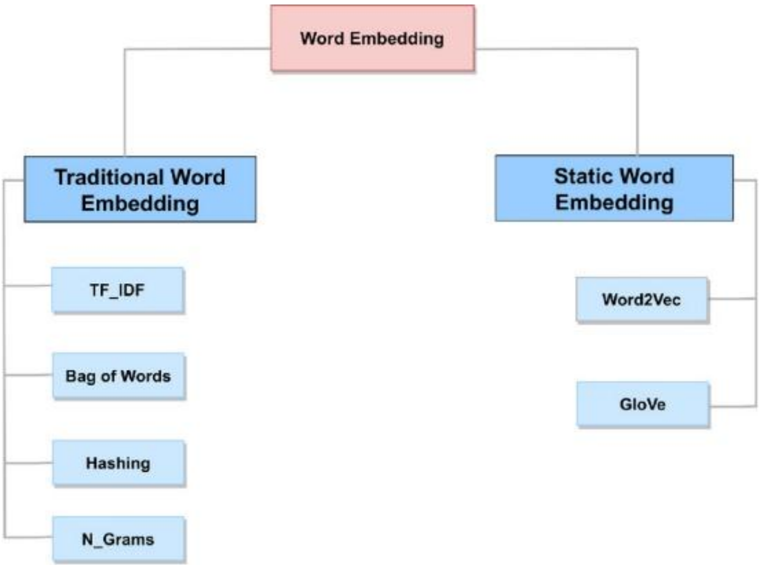

4 ENSEMBLE TAXONOMY

This section presents the proposed taxonomy for ensemble techniques applied to Sentiment Analysis in both surface and deep domains. This classification intends to summarize the work found in the literature as well as to compare these models with the ones we propose. Also, with this, we address the first question raised in Section 1 regarding how combination techniques can be classified. The taxonomy can be expressed as combination of two different dimensions. Each dimension represents a characteristic of the studied approaches. On the one hand, one dimension considers which features are used in the model. Those features can be either surface features (which stands for S), generic automatic word vectors (G), or affect word vectors specifically trained for the sentiment analysis task (A). On the other hand, the other dimension attends to how the different model resources are combined. These combinations can be: using no ensemble method at all, through a ensemble of classifiers, or taking advantage of a feature ensemble. Table 1 shows a representation of this taxonomy, where the two dimensions appear as rows for the first dimension, and columns for the second dimension. We have classified all the reviewed work in this paper using the proposed taxonomy, obtaining a visual layout of the techniques that are used in each approach in relation with both ensemble methods and the combination of surface and deep features. Regarding the dimension that tackles the ensemble techniques, in the No ensemble category we find the classifiers that do not make use of an ensemble technique. Under the Classifier ensemble category we classify the approaches that are based on ensemble techniques (Section 2.1), such as the voting rule or a meta-learning technique, to name a few. In the same manner, the Feature ensemble category contains the approaches that make use of feature combination techniques. The feature ensemble consists in combining different set of features into an unified set that is then fed to a learning algorithm. As for the dimension that represents which features are used, several possibilities are represented: only surface features, generic or affect words vectors (S, G and A respectively), where only one type of feature is used. Besides, this dimension also takes into account the combination of different types of features: S+G (surface features combined with generic word vectors), G+A (generics word vectors with affect embeddings), S+A (surface features combined with affect word embeddings), and S+G+A (all three types of features combined in the same model). These two dimensions are combined, creating a grid where the different approaches can be classified. The blank spaces in the taxonomy represent techniques that, to the extent of our knowledge, have not been studied. As such, they represent work that can be addressed in the future. In conclusion, the introduced taxonomy provides a framework for characterizing and comparing ensemble approaches in sentiment analysis. This framework provides us with the opportunity to characterize and compare existing research works in sentiment analysis using ensemble techniques. Moreover, the framework can help us to provide guidelines to choose the most efficient and appropriate ensemble method for a specific application.

The proposed ensemble classifier was trained on the training set for classifying the sentiments in the datasets and evaluated on the test data. The ML algorithm used in this work is a random forest classifier. This model was chosen for this study because it helps to avoid overfitting, provides a measure of feature importance, and produces a reasonable prediction without adjusting hyperparameters. This is a supervised ML algorithm that is used for regression and classification purposes and belongs to the ensemble learning family [40]. In the random forest model, decision trees are constructed from datasets and create a forest made of trees. A random forest classifier consists of the following steps: The first step is to select random data samples from the available dataset. For each selected data sample, a decision tree is constructed, and a prediction value is extracted from each decision tree. For node splitting, the Gini coefficient method is applied as follows:

Gini=1−n∑i=1(pi)2

4.1 Performance measures:

To examine the performance of the suggested model using different feature extraction techniques, we used several standard performance measures. Specifically, we used recall, accuracy, precision, and F1-measure. To calculate all four metrics, machine learning models can be visualized by using a confusion matrix. The elements of this matrix are False Negative (FN), True Positive (TP), False Positive (FP), and True Negative (TN). The performance evaluation of classifiers is made according to the following formulas:

Accuracy :	TP+TNTP+TN+FP+FN
Precision :	TPTP+FP
Recall :	TPTP+FN
F1-Measure : 2* Precision* Recall/Precision + Recall	

5. RESULTS AND ANALYSIS

We present the results of our experiments on sentiment analysis in social media text using deep learning techniques. Table 1 summarizes the performance metrics of different models and embeddings.

Table 1 Performance Metrics Comparison

Performance and time of the random forest classifier on a Twitter dataset.

Feature extraction	Accuracy	Precision	Recall	F1-measure	Training time	Prediction time
TF_IDF	96	95	96	95	11.285836	0.497233
N_Gram	86	87	86	86	13.926802	0.541020
BOW	87	87	87	87	16.031671	0.535141
Hashing Vectorizer	96	96	96	96	79.441338	0.809710
Word2Vec	93	93	93	93	19.753669	0.214723
Glove	92	92	92	92	35.825151	0.180461

From the results, we observe that LSTM-based models outperform CNNs and GRUs in capturing long-range dependencies and contextual information, leading to higher accuracy and F1-score. Among the word embeddings, FastText embeddings show improved performance compared to Word2Vec and GloVe, especially in handling out-of-vocabulary words and capturing sub-word information.

Furthermore, we address the challenge of imbalanced sentiment classes by experimenting with techniques such as class weighting, oversampling, and under sampling. Our results show that class weighting based on inverse class frequencies improves the model's ability to generalize to minority classes without significant performance degradation on majority classes.

6. DISCUSSION

Our study demonstrates the effectiveness of deep learning techniques, particularly LSTM-based models and FastText embeddings, in enhancing sentiment analysis accuracy in social media text. These techniques show robust performance in handling noisy and informal language, capturing semantic nuances, and addressing imbalanced class distributions.

We discuss the implications of our findings for real-world applications such as brand monitoring, customer feedback analysis, and social media sentiment tracking. The improved accuracy of sentiment analysis models can lead to more accurate insights and actionable recommendations for businesses and organizations.

7. CONCLUSION

In conclusion, this paper presents a comprehensive study on enhancing sentiment analysis in social media using deep learning techniques. We evaluate different deep learning architectures, word embeddings, and strategies for handling imbalanced classes in sentiment analysis tasks. Our results highlight the importance of choosing appropriate model architectures and embeddings for effective sentiment analysis in noisy and informal text data. For future work, we plan to explore multi-task learning approaches for jointly modeling sentiment analysis and related tasks such as aspect-based sentiment analysis and emotion detection. Additionally, investigating interpretability and explainability techniques for deep learning models in sentiment analysis contexts remains an important research direction. Overall, our study contributes to advancing sentiment analysis techniques in the domain of social media analytics and lays the foundation for more accurate and robust sentiment analysis systems.

REFERENCES

1. Agarwal, A., Xie, B., Vovsha, I., Rambow, O., & Passonneau, R. (2011). Sentiment anal- ysis of twitter data. In Proceedings of the workshop on languages in social media (pp. 30–38). Association for Computational Linguistics.
2. Alpaydin, E. (2014). Introduction to machine learning. MIT press.
3. Aue, A., & Gamon, M. (2005). Customizing sentiment classifiers to new domains: A case study. In Proceedings of recent advances in natural language processing (RANLP): vol. 1. 2–1
4. Bengio, Y. (2009). Learning deep architectures for ai. Foundations and Trends® in Machine Learning, 2 , 1–127
5. Cambria, E. (2016). Affective computing and sentiment analysis. IEEE Intelligent Sys- tems, 31, 102–107.
6. Chen, M., Weinberger, K. Q., Sha, F., & Bengio, Y. (2014). Marginalized denoising au- to-encoders for nonlinear representations. In ICML (pp. 1476–1484).
7. Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., & Kuksa, P. (2011). Natural language processing (almost) from scratch. The Journal of Machine Learn- ing Research, 12, 2493–2537.
8. De Smedt, T., & Daelemans, W. (2012). Pattern for python. The Journal of Machine Learning Research, 13 , 2063–2067
9. Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. The Journal of Machine Learning Research, 7, 1–30.
10. Ding, X., Liu, B., & Yu, P. S. (2008). A holistic lexicon-based approach to opinion mining. In Proceedings of the 2008 international conference on web search and data mining WSDM '08 (pp. 231–240). New York, NY, USA: ACM. doi: 10.1145/ 1341531.1341561.
11. Esuli, A., & Sebastiani, F. (2006). Sentiwordnet: A publicly available lexical resource for opinion mining. In Proceedings of LREC: 6 (pp. 417–422). Citeseer.
12. Fersini, E., Messina, E., & Pozzi, F. (2014). Sentiment analysis: Bayesian ensem- ble learning.
13. Decision Support Systems, 68, 26–38. doi: 10.1016/j.dss.2014.10.004. http://www.sciencedirect.com/science/article/pii/S0167923614002498.
14. Fersini, E., Messina, E., & Pozzi, F. (2016). Expressive signals in social media lan- guages to improve polarity detection. Information Processing & Management, 52, 20–35.

15. Gimpel, K., Schneider, N., O'Connor, B., Das, D., Mills, D., Eisenstein, J., et al. (2011). Part-of-speech tagging for twitter: Annotation, features, and experiments. In Proceedings of the 49th annual meeting of the association for computational lin- guistics: Human language technologies: Short papers - volume 2 HLT '11 (pp. 42 – 47). Stroudsburg, PA, USA: Association for Computational Linguistics. http://dl. acm.org/citation.cfm?id=20 02736.20 02747.
16. Go, A., Bhayani, R., & Huang, L. (2009). Twitter sentiment classification using distant supervision. CS224N Project Report, Stanford, 1, 12.
17. Hutto, C. (2015). Vader sentiment github repository. https://github.com/cjhutto/ vaderSentiment. Accessed on May 30, 2016.
18. Kathuria, P. (2015). Sentiment wsd github repository. https://github.com/ kevincobain20 0 0/sentiment _classifier/. Accessed on May 30, 2016.
19. Kim, J., Yoo, J.-B., Lim, H., Qiu, H., Kozareva, Z. , & Galstyan, A. (2013). Sentiment prediction using collaborative filtering.. ICWSM.
20. Kim, Y. (2014). Convolutional neural networks for sentence classification. arXiv preprint arXiv:1408.5882 .
21. Kiritchenko, S., Zhu, X., & Mohammad, S. M. (2014). Sentiment analysis of short informal texts. Journal of Artificial Intelligence Research, 723–762.
22. Kouloumpis, E., Wilson, T., & Moore, J. D. (2011). Twitter sentiment analysis: The good the bad and the omg!. Icwsm, 11, 538–541.
23. Le, Q. V., & Mikolov, T. (2014). Distributed representations of sentences and docu- ments.. In ICML: vol. 14 (pp. 1188–1196).