A BERT-based Deep Learning Approach for Reputation Analysis in Social Media

Mohammad Wali Ur Rahman[1], Sicong Shao[1], Pratik Satam[2], Salim Hariri[1], Chris Padilla[3], Zoe Taylor[3], and Carlos Nevarez[3]

[1]Dept. of Electrical and Computer Engr, The University of Arizona, Tucson, AZ 85721 USA
[2]Dept. of Systems and Industrial Engr, The University of Arizona, Tucson, AZ 85721 USA
[3]Legendary Labs, Provo, UT 84604 USA

Abstract—Social media has become an essential part of the modern lifestyle, with its usage being highly prevalent. This has resulted in unprecedented amounts of data generated from users in social media, such as users' attitudes, opinions, interests, purchases, and activities across various aspects of their lives. Therefore, in a world of social media, where its power has shifted to users, actions taken by companies and public figures are subject to constantly being under scrutiny by influential global audiences. As a result, reputation management in social media has become essential as companies and public figures need to maintain their reputation to preserve their reputational capital. However, domain experts still face the challenge of lacking appropriate solutions to automate reliable online reputation analysis. To tackle this challenge, we proposed a novel reputation analysis approach based on the popular language model BERT (Bidirectional Encoder Representations from Transformers). The proposed approach was evaluated on the reputational polarity task using RepLab 2013 dataset. Compared to previous works, we achieved 5.8% improvement in accuracy, 26.9% improvement in balanced accuracy, and 21.8% improvement in terms of F-score.

Index Terms—Reputation polarity, artificial intelligence, natural language processing, machine learning, BERT, transformers, neural network, social media

I. INTRODUCTION

In recent years, Online Reputation Management (ORM) has become an integral part of marketing strategies for influential entities, such as large enterprises and public figures. ORM aims to help an entity build its public image, manage its reputation, mitigate the negative impact of potential actions, and reclaim its lost reputation. Social media posts, blogs and articles, and other opinionated content about the entity significantly impact ORM analysis. One of the most famous examples of ORMs recently is Cambridge Analytica due to their data scandal issue with Facebook [33].

To perform these duties for ORM, reputation management teams need to monitor texts, posts, and news regarding the target entity from online social media. A significant problem with traditional data sources (e.g., news articles, newspaper blogs, surveys, etc.) is that they usually present information from their own point of view, disregarding the opinions of most of the population. Besides, the traditional data source often causes a significant delay between the time an event occurs and the time news about such events is published or broadcasted [34]. Therefore, to overcome such limitations and collect public opinion regarding the entity of reputation in a direct and timely manner, ORM domain experts have focused more on online social media platforms [29, 30].

The population's opinions are expressed in texts posted on social media platforms [35]. Thus, instantaneous reactions regarding an event can be mined from online social media posts [36, 43]. Also, the rapid growth of social media platforms has enabled users to engage in online posts regarding different entities related to the events, which anyone can see and share on the social media platform [37, 38]. In this way, the posts can circulate among a large number of social media users and form a public opinion of the entity for ORM [31]. Therefore, ORM experts place a great deal of emphasis on analyzing social media posts based on reputation polarity. Reputation Polarity for ORM can be defined as follows: Whether the content posted on the Internet influences a target entity's reputation positively, negatively, or neutrally. It is important to distinguish between sentiments and reputation polarity derived from text, as pointed out in [1] and [20] because a negative sentiment text sample might actually prove to be a positive reputation text sample for a target entity and vice versa. As an example, consider the following text, *"Alas! Diego Maradona is no more! He will be missed by millions of fans."*, this text is a negative sentiment sample as it expresses sorrowfulness, but the reputation polarity of the text is positive for the entity "Diego Maradona". Thus, determining how text influences an entity's reputation is challenging.

Recent research has increasingly relied on pre-trained language models for different types of NLP tasks [17]. The advantage of pre-trained language models is that they have been trained on a large corpus of text data, which can significantly benefit NLP research. It is cheaper, faster, and easier to fine-tune such pre-trained models to perform downstream NLP tasks and achieve remarkable results with transfer learning. Hence, our approach is built upon BERT's modeling power, a popular pre-trained language model based on transformer-based deep learning [15]. Our study shows the feasibility of leveraging the contextual word representation derived from pre-trained BERT, along with a fine-tuning technique with extra reputation analysis data, to address reputation polarity task and achieve superior performance than previous works on RepLab 2013 dataset.

The rest of the paper is organized as follows: Section II

covers the related works, Section III discusses our BERT-based Reputation Analysis approach, Section IV presents our experimental results, and Section V concludes the paper.

II. RELATED WORKS

Over the past decade, there have been a few notable contributions to reputation analysis in natural language processing area. RepLab is a well-known project regarding ORM competition, and identifying the reputation polarity of tweets has been one of the major tasks of RepLab [1]. The participants in the competition were provided with a corpus of tweets collected between June 1, 2012, to December 31, 2012. The tweets were about 61 entities from four domains: automotive, banking, universities, and celebrities. The tweets were manually annotated by thirteen annotators who were constantly supervised by ORM experts. Among the 11 teams that participated in the reputation polarity task, SZTE-NLP achieved the highest Accuracy (0.685) [2]. The features used for their approach are the number of positive and negative words in the text, determined by SentiWordNet [3] containing the features such as number of negation words, number of words with character repetitions, number of capitalized words, number and the polarity of acronyms, and bigrams. The features also contained whether the entity was mentioned in the tweet and the distances between the tokens and the mentions of the entity. Cossu et al. [4] used tf-idf in association with Gini purity criteria and SVM classifier for reputation analysis. Filgueiras et al. [5] used the number of words with positive and negative sentiment polarity, the number of emoticons with positive and negative polarity, punctuation marks, and the number of capitalized words as features. Most of the participating teams used different types of sentiment lexicons to derive the sentiment polarity of the tweets with varying kinds of textual selection approaches like concept count vectors, bag of words, and terms co-occurrence representation (TCOR), KL-divergence, etc. [6, 7, 8, 9]. In addtion, Giachanou et al. [12] used a lexicon augmentation approach to compute reputation scores.

Apart from the teams participating in the RepLab competitions, some other researchers also made their contribution in the area of Reputation Analysis. Alrubaian et al. used some user profile-specific features other than features derived from sentiment analysis [11]. Using these features, the authors employed Logistic Regression and Feature Rank Naïve Bayes models to classify the texts based on reputation polarity. The authors used their own manually annotated data for the assessment. Manaman et al. [13] used a similar set of features based on sentiment and n-grams. However, these methods, such as bag-of-words, have the major disadvantage of not always capturing the semantic relation. Texts constructed with the same words can have different meanings due to the context, but the bag-of-words method will consistently treat those words identically. N-gram-based methods solve this problem to some extent, but the dimensionality of N-gram features is usually very high, resulting in the risk of model overfitting.

III. BERT-BASED REPUTATION ANALYSIS APPROACH

In this section, we first review the RepLab 2013 dataset that focuses on monitoring the reputation of entities such as companies, organizations, and celebrities on Twitter. After that, we describe our BERT-based reputation analysis approach, including text preprocessing, pre-trained BERT model, and BERT-based reputation polarity.

A. Dataset

ORM monitoring involves constant scrutiny of online (especially social) media for information relating to the entity. The focus is on the opinions, news, and events of a given entity. It aims to detect any potential threat to its reputation early, that is, opinions and issues that could harm the entity's public image. This implies frequent examination of the latest online information. Social networks, especially Twitter, provide a key source of information in ORM monitoring. Therefore, RepLab 2013 dataset is selected for our research due to the fact that it focuses on monitoring company and individual reputations.

The RepLab 2013 dataset [1] contains tweets regarding 61 entities from four domains that are automotive, banking, universities, and celebrities. Table I provides the list of all the monitored Entities from the RepLab 2013 dataset, addressing the diversity of the dataset. Table II summarizes the description of the corpus, as well as the number of tweets for both training and test sets, and the distribution by language. Also, the distributions of training and testing samples across the four domains are shown in Table II. As we can see, the majority of the tweets are in the English language, while there are some tweets in Spanish as well.

The domains for reputation analysis were carefully selected for the following reasons. The automotive domain was selected for analyzing the entities where the reputation largely depends on the products. The banking domain was chosen as the representative of the entities where the transparency and ethical side of the activity play a huge role in determining their reputation. The university domain was considered due to these entities whose reputation largely depends on intangible products. Celebrities' domain was targeted because their reputation greatly depends on their character and services.
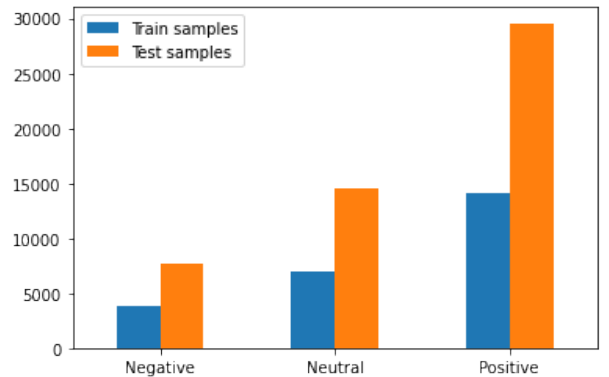

Figure 1: The reputation polarity class distribution in the dataset.

Table I: The Monitored Entities from Different Domains in Our Reputation Polarity Prediction Task

Automobile		Banking	Universities	Celebrities	
BMW	Chrysler	RBS bank	Harvard University	Adele	Coldplay
Audi	Subaru	Barclays	Stanford University	Alicia Keys	Lady Gaga
Volvo	Ferrari	HSBC	The University of California, Berkeley	The Beatles	Madonna
Toyota	Bentley	Bank of America	MIT	Led Zeppelin	Jennifer Lopez
Volkswagen	Porsche	Wells Fargo	Princeton University	Aerosmith	Justin Bieber
Honda	Yamaha	PNC	Columbia University	Bon Jovi	Shakira
Nissan	KIA	Capital One	Yale University	U2	PSY
Fiat	Ford	Banco Santander	The Johns Hopkins University	AC/DC	The Script
Suzuki	Jaguar	Bankia	New York University	The Wanted	Whitney Houston
Mazda	Lexus	BBVA	University of Oxford	Maroon 5	Britney Spears
		Goldman Sachs			

Table II: The Description of RepLab 2013 Dataset

	Automobile	Banking	Universities	Celebrities
No. of Entities	20	11	10	20
Train Samples	15,123	7,774	6,960	15,822
Test Samples	31,785	16,621	14,944	33,498
Total Samples	46,908	24,395	21,904	49,320
English Tweets	38,614	16,305	20,342	38,283
Spanish Tweets	8,294	8,090	1,562	11,037

Each tweet is manually annotated in three polarity classes: POSITIVE reputation, NEGATIVE reputation, and NEUTRAL reputation. In this research, all of the tweets marked as "Unrelated" are ignored as per RepLab 2013 instructions. Some of the tweets in the dataset are not available because the user have deleted the tweet or changed the Twitter account privacy settings. The distribution of resulting tweets from each polarity class is presented by Figure 1.

B. Pre-trained BERT Model

A key innovation in contextualized representation learning is BERT [15]. This work shows that even though the word embedding layer (of a traditional deep learning model for NLP task) is trained from large language corpora, training a variety of neural network architectures that encode contextual representations only based on the limited supervised data on end NLP tasks is still not sufficient. As opposed to ELMo [14] and ULMFiT [44], which aim to add more features for a particular architecture based on human understanding of the end NLP task, BERT uses a fine-tuning process that almost does not require a specific architecture for each end NLP task. This is a desirable outcome for automated reputation analysis, as an artificial intelligence-based reputation management system aims to minimize the use of prior human knowledge in the design of the model. Instead, such knowledge should be learned through data.

BERT provides two parameter-intensive settings that are BERT-Base and BERT-Large. BERT-Base has 768 hidden dimensions, 12-layer transformer blocks, and 12 attention heads, with 110 million parameters, while BERT-Large has 1024 hidden dimensions, 24-layer transformer blocks, and 16 attention heads, with 340 million parameters [15]. Pre-training with BERT models involves two major methods: masked language modeling and next sentence prediction. The masked language modeling method randomly masks 15% of tokens in the input, and the goal is to predict the masked token according to its context. In the next sentence prediction, pairs of sentences are used as input, and BERT predicts if the second sentence in the pair is the following sentence. Using these two methods, BERT can build highly effective bidirectional contextual representations in order to advance the performance for a variety of NLP tasks. Hence, due to the success of BERT models in NLP, we propose a BERT-based reputation polarity approach.

C. Text Pre-processing

Text preprocessing is an essential step for reputation polarity tasks. It transforms the text in social media into a more digestible format that is more conducive to machine learning models. With this process, tweets from RepLab 2013 dataset are cleaned and prepared for training models. Specifically, we remove extra spaces, special characters, emojis, URL links, and punctuation marks.

D. BERT-based Reputation Polarity

We address text reputation polarity prediction in social media, which is the core task of ORM. Reputation polarity determines whether a text's content has positive, negative, or neutral implications for the monitored entity's reputation. The traditional approaches to this task include using sentiment lexicons and linguistic features [1]. However, they struggled in social media due to its noisy nature and, most importantly, the fact that reputation polarity is not only expressed in sentiment-bearing words but also impacted by other word usage related to reputation and context. Therefore, the BERT-based approach that can learn discriminative contextual word representation, has excellent potential for this task. Thus, we propose a BERT-based deep learning Approach for building the reputation polarity classifier. We address predicting reputation polarity as a supervised task due to the label being available in the training data. PyTorch [16] platform and Transformers library from HuggingFace [39] are used for the implementation of our approach.

Figure 2 describes BERT-based reputation polarity architecture. As we can see, the BERT uses the token representation vectors E_i as input. For each token, E_i is the sum of three representation vectors, including a position embedding vector, a sentence vector, and a typical word embedding vector. The position embeddings provide the model with information about the token's location within the sentence since transformers lack this information. The sentence vector is used when more context is needed than a single sentence. A typical word
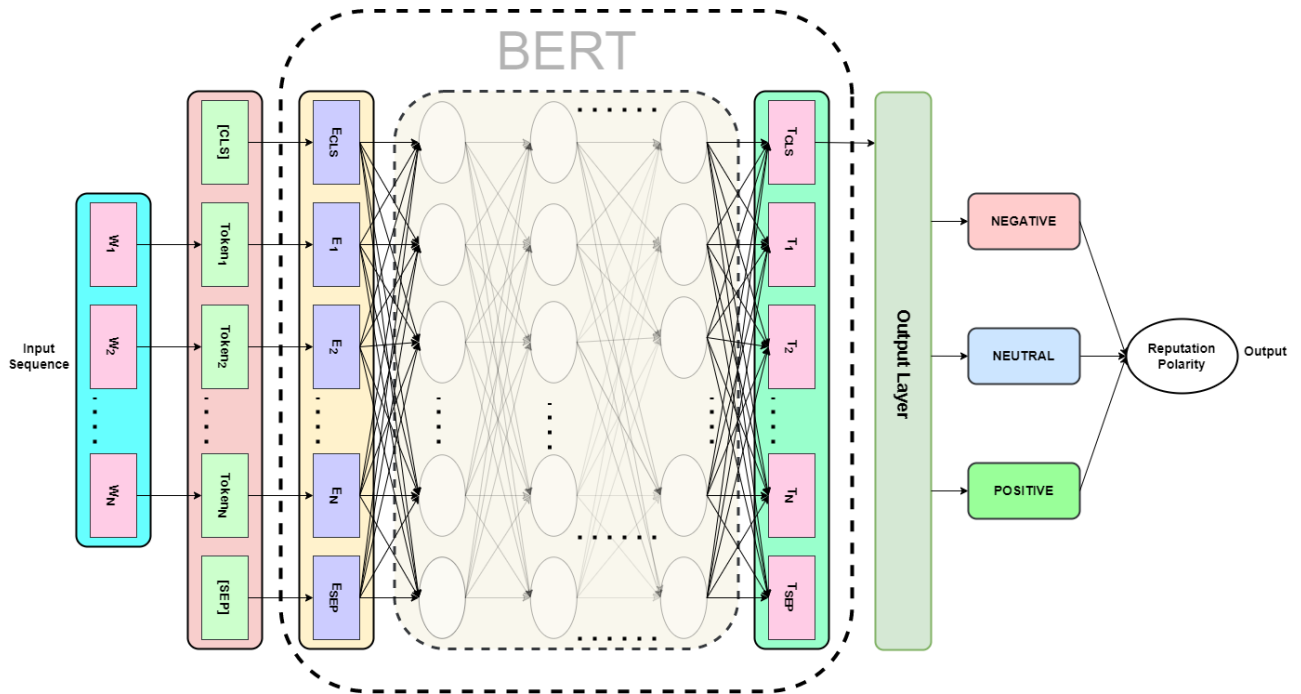
Figure 2: BERT architecture for reputation polarity task.

embedding provides vector representations for words in the context. Each input token E_i is represented by representation T_i in the layers of transformer blocks by considering all token representations within the sentence, addressing capability for capturing long-term dependencies of transformer-based models.

The BERT augments input sentence with adding [CLS] token and [SEP] token. The [CLS] token contains the embedding for special classification, and the [SEP] token is utilized for separating segments. For our reputation polarity task, BERT uses the last hidden state h from the first token [CLS] to represent the whole input sequence. The softmax classifier is added on the top of BERT to predict the probability of reputation polarity class c :

$$p(c|h) = softmax(Vh) \qquad (1)$$

where V is the parameter matrix for specific reputation polarity prediction task. We perform fine-tuning on all the parameters of BERT and parameter matrix V jointly learned during the fine-tuning on the reputation polarity training data.

As we can see from Figure 1, class-imbalance data occurs in our reputation polarity prediction task. The data of the minority class is significantly less than the data of the majority class. The prediction performance may be suboptimal without addressing class imbalance. To deal with this problem, a cost-based strategy [40] is employed because it can be seamlessly integrated into our BERT-based reputation polarity model. More specifically, the weight is attributed to each training sample during the training, and larger weights are applied to the class with smaller amounts, which are the training samples with higher cost if they are classified incorrectly. In contrast, lower weights are assigned to the class with larger amounts,

which are the training samples with lower costs if they are wrongly classified.

IV. EXPERIMENTATION RESULTS

A. Experiment Settings

To demonstrate the superiority of our approach, we evaluated it through RepLab 2013 dataset. As the dataset is multilingual and contains tweets from both English and Spanish Languages, we conduct the experiments by using three different pre-trained BERT models [39]: (1) bert-base-multilingual-uncased model: BERT-Base multilingual model pretrained on the top 104 languages, (2) bert-base-uncased model: BERT-Base monolingual model pretrained on English, and (3) bert-large-uncased: BERT-Large monolingual model pretrained on English. Then three different combinations of these models are formed, namely Mode 1, Mode 2, and Mode 3.

In Mode 1, we fine-tuned the *bert-base-multilingual-uncased* model (i.e., BERT Multilingual Base model) on the entire training dataset (in both English and Spanish texts) and evaluated the model's performance on the test dataset. *bert-base-multilingual-uncased* tokenizer is used for text tokenization. With this setting, Mode 1 can perform reputation polarity prediction on both English and Spanish tweets with a single multilingual model.

In Mode 2, we first identified the language of the tweets using Twitter API [42]. If the language of the training tweets were English, the tweets were first tokenized with the *bert-large-uncased* tokenizer. Afterwards, they were passed on to the pre-trained *bert-large-uncased* model (i.e., BERT Large monolingual English model) for fine-tuning on English data to create a BERT model that can detect reputation polarity focusing on English. On the other hand, if the language of
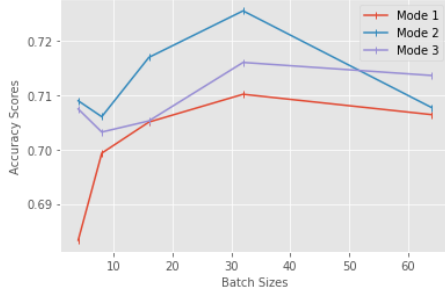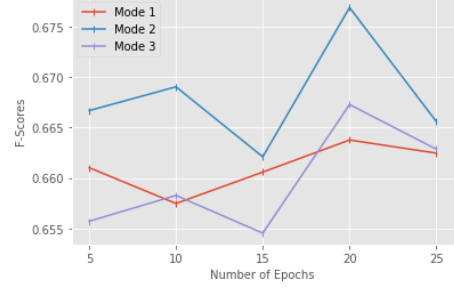
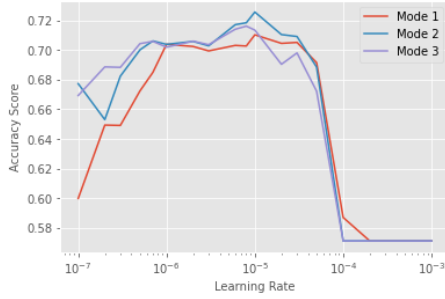(a) Accuracy Scores of different modes across different batch sizes
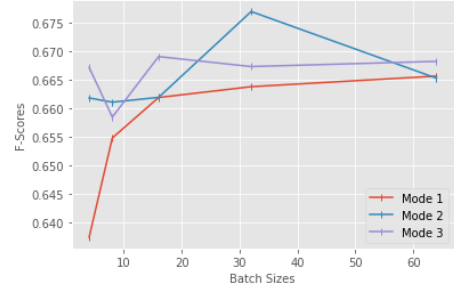
(b) F-Scores of different modes across different batch sizes
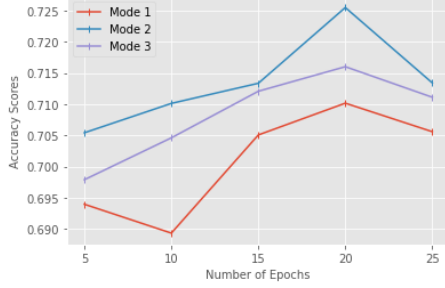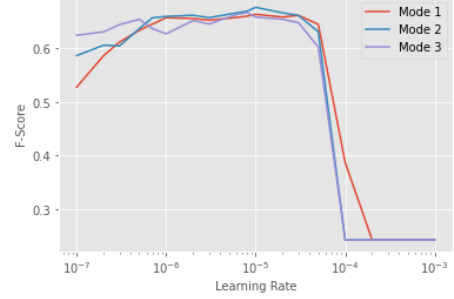

(c) Accuracy Scores of different modes across different epochs

(d) F-Scores of different modes across different epochs

(e) Accuracy Scores of different modes across different learning rates

(f) F-Scores of different modes across different learning rates

Figure 3: Performance of our approach with different modes.

Table III: The Best Experiment Results of Our BERT-based Reputation Analysis Approach with Different Modes

Model	Epochs	Learning_rate	Batch_size	Precision	Recall	F-Score	Accuracy	Training time in minutes	Test time in minutes
BERT (Mode 1)	20	$1 \cdot 10^{-5}$	32	0.68	0.65	0.66	0.71	97.1	3.2
BERT (Mode 2)	20	$1 \cdot 10^{-5}$	32	0.69	0.67	0.67	0.73	147.9	3.5
BERT (Mode 3)	20	$8 \cdot 10^{-6}$	32	0.69	0.65	0.67	0.72	108.9	3.6

the training tweets were Spanish, they were tokenized with a *bert-base-multilingual-uncased* tokenizer, and then passed on to the *bert-base-multilingual-uncased* model for fine-tuning on Spanish data to build a BERT model that can predict reputation polarity focusing on Spanish. For the evaluation, the test tweets were passed to either of the models depending on their language (English or Spanish). If the language of the test tweet is English, the reputation polarity of the tweet was predicted using the fine-tuned BERT large model. On the contrary, if the identified language of the test tweet was Spanish, the multilingual BERT base model was used for reputation polarity prediction.

Mode 3 was trained in a manner similarly to Mode 2. Here we also have different BERT model for each language.

However, for the English training tweets, we have used them for fine-tuning *bert-base-uncased* model (i.e., BERT Base monolingual English model) instead of *bert-large-uncased* model. During the evaluation, the reputaion polarity of the English test tweets were predicted using the *bert-base-uncased* model, which focuses on reputation polarity in English. While, for the fine-tuning and evaluation of spanish tweets, *bert-base-multilingual-uncased* model in Mode 2 was still used.

For the training procedure, we have used an Adam optimizer with weight decay from Hugging Face recommended in [18]. In addition, we have used a linear scheduler with no warm-up steps. For the other hyperparameters, the original BERT authors have made some recommendations for fine-tuning tasks [15]. Apart from the number of epochs, we have mainly

used the recommended hyperparameters for our classification tasks. We perform hyperparameter search for our fine-tuning process with the following setting. The learning rates were experimented in the range from 10^{-7} to 10^{-3}. The number of epoch is selected in {5, 10, 15, 20, 25}, and the number of batch sizes is selected in {4, 8, 16, 32, 64}. For the output layer, the activation function we have used are either SoftMax or LogSoftmax. Categorical cross-entropy loss function is used during the training. The cost-based strategy is employed during the training to address the class imbalance problem. The experiments were conducted with Google Colab Pro Plus [41].

B. Results

We report the best results for all the hyperparameter combinations and plotted the graphs in Figure 3. The evaluation metrics selected for our visualization results are F-score and Accuracy because they are the official evaluation metrics used for the RepLab 2013 competition for the Reputation polarity detection task. Table III showcases our experiment runs for the three different modes of our BERT models. From the graphs in Figure 3, we can see, that Mode 2 gives us the best results, achieving the highest accuracy of 73% with an F-score of 67%. Besides, our BERT model in Mode 3 achieves a 72% accuracy and Mode 1 achieves 71% accuracy. Figure 3 show the changes in Accuracy and F-score respectively for 3 different modes with different hyperparameters.

In Figures 3(a) and 3(b), we plotted the best results obtained with different batch sizes. We observed that batch size of 32 yield best results. In Figures 3(c) and 3(d), we plotted the best accuracy and best F-Score obtained for different epoch numbers mentioned, and for our experiments, all the mode performed best when the number of epochs was 20. Figures 3(e) and 3(f) shows the prediction performance with different learning rates, ranging from 10^{-7} to 10^{-3}. As we can see, all the modes obtain better prediction with lower learning rates, ranging from 10^{-6} to 10^{-5}. We also noticed that the models' performance starts to drop with higher learning rates. Besides, we observed that the accuracy of both models remains unchanged from 57% in very high learning rates mainly because the models fail to classify text as NEGATIVE or NEUTRAL, while they classify texts as POSITIVE, which is the majority class from the training set. From the figures, we also notice that both Mode 3 and Mode 2 attain better overall results than Mode 1.

Table IV showcases how our appropach performs compared to other approaches that classify texts based on Reputation Polarity. All the approches listed in this table use the RepLab 2013 dataset comprised of all 61 entities for their evaluation [1, 2, 4, 5, 6, 7, 8, 9, 10, 20] under the same training and testing sets. The best results reported for each team were presented. All of the teams participating in the RepLab 2013 competition reported their Accuracy and F-scores as per the official RepLab 2013 evaluation criteria. On the other hand, Peetz et al. [20] evaluated their systems in terms of F-score and Balanced Accuracy since the dataset is imbalanced. It is plain to see that our BERT-based approaches outperform all

Table IV: Comparison of Different Approaches in Reputation Polarity Task

Teams	Accuracy	BAC	F-score
BERT(Mode 2)	**0.73**	**0.66**	**0.67**
BERT(Mode 3)	0.72	0.66	0.67
BERT(Mode 1)	0.71	0.65	0.66
Peetz et al. [20]	-	0.52	0.55
SZTE-NLP [2]	0.69	-	0.38
LIA [4]	0.65	-	0.19
Popstar [5]	0.64	-	0.37
UAMCLYR [8]	0.62	-	0.29
UNED ORM [9]	0.62	-	0.15
DIUE [7]	0.55	-	0.25
Volvam [6]	0.54	-	0.26
DAEDALUS [10]	0.44	-	0.34

the other approaches in every evaluation metric. For example, in terms of Accuracy, the BERT model in Mode 2 performs better by 4% than the best-reported Accuracy by SZTE-NLP, while the BERT models in Mode 3 and Mode 1 has 3% and 2% better Accuracy respectively. On the other hand, the Mode 2 and Mode 3 models surpass Peetz et al. by 14% in terms of Balanced Accuracy and 12% in terms of F-Score. Meanwhile, the BERT model in Mode 1 surpasses Peetz et al. by 13% in terms of Balanced Accuracy and and 11% in terms of F-Score. These results address that the BERT successfully learns contextual word representation for reputation analysis from a sequence, which other methods fail to achieve. Therefore, we validate the superiority of the BERT-based approach for reputation polarity.

V. CONCLUSION

Over the past few years, BERT has shown its superior performance of text representation learning in many different NLP tasks [21, 22, 23, 24, 25]. However, BERT's application to classify texts based on reputation polarity has not yet been validated. The main contribution of this paper is to shed light on BERT's capability in reputation analysis. In this research, we have fine-tuned pre-trained BERT models to classify texts based on reputation polarity prediction using a standard dataset dedicated to reputation research. The experimental results show that our BERT-based approach can achieve a 5.8% improvement in Accuracy in detecting reputation polarity, 26.9% improvement in Balanced Accuracy, and a 21.8% improvement in terms of F-score. We should note that it is possible further to improve the performance with more hyper-parameter tuning. Besides, the modified versions of BERT such as RoBERTa, AlBERT, distilBERT [26, 27, 28] may further improve the prediction performance for reputation analysis, We leave these possibilities for future works.

ACKNOWLEDGMENT

REFERENCES

[1] E. Amigó, J. Carrillo de Albornoz, I. Chugur, A. Corujo, J. Gonzalo, T. Martín, M. Edgar, M. de Rijke & D. Spina, "Overview of replab 2013: Evaluating online reputation monitoring systems", In *International conference of the cross-language evaluation forum for european languages, Springer, Berlin, Heidelberg*, 2013.

[2] V. Hangya, & R. Farkas, "Filtering and polarity detection for reputation management on tweets", In *CEUR WORK-SHOP PROCEEDINGS*, 2013.

[3] A. Esuli & F. Sebastiani, "SentiWordNet: A publicly available lexical resource for opinion mining", In *LREC '06*, 2006.

[4] J. V. Cossu, B. Bigot, L. Bonnefoy, M. Morchid, X. Bost, G. Senay, R. Dufour, V. Bouvier, J. Torres-Moreno & M. El-Bèze, "Lia@ replab 2013", In *RepLab: An evaluation campaign for Online Reputation Management Systems, Fourth International Conference of the CLEF initiative*, 2013.

[5] J. Filgueiras, & S. Amir, "POPSTAR at RepLab 2013: Polarity for Reputation Classification", In *CLEF (Working Notes)*, 2013.

[6] A. Mosquera, J. Fernandez, J. M. Gomez, P. Martınez-Barco, & P. Moreda, "Dlsi-Volvam at RepLab 2013: Polarity classification on Twitter data", In *CLEF '13*, 2013.

[7] J. Saias, "In search of reputation assessment: experiences with polarity classification in RepLab 2013", In *CLEF '13*, 2013.

[8] E. Villatoro-Tello, C. Rodríguez-Lucatero, C. Sánchez-Sánchez, & A. P. López-Monroy, at "RepLab 2013: Profiling Task", In *CLEF (Working Notes)*, 2013.

[9] A. Castellanos, J. Cigarrán, & A. García-Serrano, "Modelling Techniques for Twitter Contents: A step beyond classification based approaches", In *CLEF (Working Notes)*, 2013.

[10] J. Villena-Román, S. Lana-Serrano, C. Moreno, J. García-Morera, & J. C. G. Cristóbal. "DAEDALUS at RepLab 2012: Polarity Classification and Filtering on Twitter Data", In *CLEF (Online Working Notes/Labs/Workshop)*, vol. 60, 2012.

[11] M. Alrubaian, M. Al-Qurishi, M. Al-Rakhami, M. M. Hassan, & A. Alamri, "Reputation-based credibility analysis of Twitter social network users", *Concurrency and Computation: Practice and Experience*, 2013.

[12] A. Giachanou, J. Gonzalo, & F. Crestani, "Propagating sentiment signals for estimating reputation polarity", *Information Processing & Management*, 2019.

[13] H. S. Manaman, S. Jamali, & A. AleAhmad, "Online reputation measurement of companies based on user-generated content in online social networks", *Computers in Human Behavior*, 2016.

[14] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, & L. Zettlemoyer, "Deep contextualized word representations", *arXiv*, 2018.

[15] J. Devlin, M. W. Chang, K. Lee, & K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding". *arXiv*, 2018.

[16] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, & A. Desmaison, "Pytorch: An imperative style, high-performance deep learning library", *Advances in neural information processing systems*, 32, 2019

[17] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Lauf, M. Funtowicz & J. Davison, "Transformers: State-of-the-art natural language processing", In *Conference on empirical methods in natural language processing*, 2020.

[18] I. Loshchilov & F. Hutter, "Fixing weight decay regularization in adam", *arXiv*, 2018.

[19] Sun, C., Qiu, X., Xu, Y., & Huang, X. (2019, October). How to fine-tune bert for text classification?. In China national conference on Chinese computational linguistics (pp. 194-206). Springer, Cham.

[20] M. H. Peetz, M. de Rijke, & R. Kaptein, "Estimating reputation polarity on microblog posts", *Information Processing & Management*, 2016, 52(2), 193-216.

[21] W. Yang, Y. Xie, A. Lin, X. Li, L. Tan, K. Xiong, M. Li, & J. Lin. "End-to-end open-domain question answering with bertserini", *arXiv preprint arXiv:1902.01718*, 2019.

[22] W. Shi, & V. Demberg, "Next sentence prediction helps implicit discourse relation classification within and across domains", In *conference on empirical methods in natural language processing* and the *international joint conference on natural language processing (EMNLP-IJCNLP)*, 2019.

[23] Z. Gao, A. Feng, X. Song & X. Wu, "Target-dependent sentiment classification with BERT", *Ieee Access*, 2019.

[24] H. Xu, B. Liu, L. Shu & P. S. Yu, "BERT post-training for review reading comprehension and aspect-based sentiment analysis" *arXiv*, 2019.

[25] C. Sun, L. Huang, & X. Qiu, "Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence", *arXiv*, 2019.

[26] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, & V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach", *arXiv preprint arXiv:1907.11692*, 2019.

[27] V. Sanh, L. Debut, J. Chaumond, & T. Wolf, "Distil-BERT, a distilled version of BERT: smaller, faster, cheaper and lighter", *arXiv*, 2019.

[28] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma & R. Soricut, "Albert: A lite bert for self-supervised learning of language representations", *arXiv*, 2019.

[29] N. Bilton, "The Growing Business of Online Reputation Management", *New York Times*, 2011.

[30] D. Proserpio, G. Zervas. "Online Reputation Management: Estimating the Impact of Management Responses on Consumer Reviews", *Marketing Science*, 36(5), 2017.

[31] M. Bruhn, V. Schoenmueller, & D. B. Schäfer, "Are social media replacing traditional media in terms of brand equity creation?", *Management research review*, 2012.

[32] Clark, K., Khandelwal, U., Levy, O., & Manning, C. D. What does bert look at? an analysis of bert's attention. arXiv, 2019.

[33] J. Isaak & M. J. Hanna, "User data privacy: Facebook, Cambridge Analytica, and privacy protection", *Computer*,

51(8), 56-59, 2018.

[34] B. Wu, & H. Shen, "Analyzing and predicting news popularity on Twitter", *International Journal of Information Management*, 35(6), 2015.

[35] S. Shao, C. Tunc, A. Al-Shawi, & S. Hariri, "One-Class Classification with Deep Autoencoder Neural Networks for Author Verification in Internet Relay Chat," in *IEEE/ACS International Conference on Computer Systems and Applications (AICCSA)*, 2019.

[36] S. Shao, C. Tunc, A. Al-Shawi, & S. Hariri, "Automated Twitter Author Clustering with Unsupervised Learning for Social Media Forensics," in *IEEE/ACS International Conference on Computer Systems and Applications (AICCSA)*, 2019.

[37] S. Shao, C. Tunc, A. Al-Shawi, & S. Hariri. "An Ensemble of Ensembles Approach to Author Attribution for Internet Relay Chat Forensics," *ACM Transactions on Management Information Systems (TMIS)*, 2020.

[38] S. Shao. "Machine Learning-Based Author Identification for Social Media Forensics," *Doctoral Dissertation, The University of Arizona*, 2021.

[39] Hugging Face. BERT-Hugging Face. Retrieved 2022 from *https://huggingface.co/docs/transformers/*.

[40] Z.H. Zhou & X.Y. Liu, "Training cost-sensitive neural networks with methods addressing the class imbalance problem," *IEEE Transactions on knowledge and data engineering*, 2005.

[41] Google. Google-Colaboratory. Retrieved 2022 from *https://colab.research.google.com/*.

[42] Twitter. Twitter-API. Retrieved 2022 from *https://developer.twitter.com/en/docs/twitter-api*.

[43] J. Bernard, S. Shao, C. Tunc, H. Kheddouci, and S. Hariri, "Quasi-cliques analysis for IRC channel thread detection," in *International Conference on Complex Networks and their Applications*, 2018.

[44] J. Howard and Sebastian Ruder, "Universal language model fine-tuning for text classification," *arXiv*, 2018.