Team UTB-NLP at FinancES 2023: Financial Targeted Sentiment Analysis Using a Phonestheme Semantic Approach

Juan Cuadrado[1], Elizabeth Martinez[1], Juan Carlos Martinez-Santos[1] and Edwin Puertas [1,*,†]

[1]Universidad Tecnologica de Bolivar, Faculty of Engineering, Cartagena de Indias 17013001, Colombia

Abstract

Sentiment analysis in the financial domain is a challenging task that plays a crucial role in understanding public opinion, monitoring market trends, and assessing the impact of news on economic agents. In this shared task, we address targeted sentiment analysis in the financial domain, focusing on identifying the main economic target in news headlines and determining the sentiment polarity towards such targets. We propose a methodology that combines transformer-based models and phonestheme embeddings to extract meaningful features from the text, which are then used in a support vector machine (SVM) classifier for sentiment classification. Our approach shows promising results, outperforming the baseline with an F1-score of 0.529229 in Task 1. This research contributes to financial sentiment analysis by addressing the complexity of financial language and considering multiple economic agents' perspectives.

Keywords

FinancES, Transformers, Embeddings, Phonestheme, Sentiment Analysis

1. Introduction

The management of financial data and sentiment analysis in the economic domain has gained increasing attention in recent years due to the growing availability of online financial information [1, 2]. Extracting insights from financial news headlines has become vital for understanding market trends, making informed decisions, and predicting market behavior [3, 4]. Furthermore, targeted sentiment analysis, specifically focused on identifying the sentiment polarity toward specific economic targets, plays a crucial role in comprehending public opinion and its impact on the economy [5, 6].

While previous research in sentiment analysis has mainly focused on general domains, such as social media or product reviews, the financial field poses unique challenges due to its complex language, subjectivity, and dependency on the context [7, 8]. Financial terms are deeply rooted

IberLEF 2023, September 2023, Jaén, Spain

*Corresponding author.

†These authors contributed equally.

jflechas@utb.edu.co (J. Cuadrado); eayala@utb.edu.co (E. Martinez); jcmartinezs@utb.edu.co (J. C. Martinez-Santos); epuerta@utb.edu.co (E. P.)

0000-0002-8226-1372 (J. Cuadrado); 0000-0001-6592-347X (E. Martinez); 0000-0003-2755-0718 (J. C. Martinez-Santos); 0000-0002-0758-1851 (E. P.)

in the social, economic, and legal context, making sentiment analysis more intricate. Moreover, the same word or expression can have different connotations depending on the context in which it is used [9, 10, 11].

In this context, the shared task at hand, organized as part of IberLEF 2023 [12], addresses the challenges of targeted sentiment analysis in the financial domain [13]. It provides a dataset comprising news headlines collected from reputable financial and economic newspapers, with manual annotations of the target entity and sentiment polarity toward the target, companies, and consumers [14, 15]. While previous shared tasks have focused on sentiment analysis in other domains, none has specifically targeted the financial field.

For this paper, we participated in Task 1: Financial targeted sentiment analysis proposed by IberLEF 2023. This task involves identifying the main economic target from headlines of financial news and determining the sentiment polarity (positive, neutral, or negative) towards the identified target in the processed text. The evaluation measures for this task include precision, recall, and F1-score, with systems being ranked based on the arithmetic mean of the target F1-score and sentiment classification macro-F1.

To tackle this task, we propose an approach of combining transformer-based models and phonestheme embeddings [16, 17]. The transformer captures contextual information and semantic meaning, while the phonestheme embeddings provide a phonetic representation of the words, allowing us to capture the nuances of financial language. Furthermore, we employ a support vector machine (SVM) classifier to classify the sentiment polarity toward the identified targets [18].

Our experimental results demonstrate the effectiveness of our methodology, outperforming the baseline with a significant improvement in Task 1, achieving an F1-score of 0.529229. This research contributes to advancing sentiment analysis in the financial domain, providing valuable insights into market sentiment and enhancing decision-making processes. Our work is motivated by the goal of addressing the complexities of sentiment analysis in the financial field, and we draw inspiration from previous studies on polarity, emotion, and user statistics analysis, which have proven valuable in related tasks such as detecting fake profiles on Twitter [19, 20]

2. Related work

In this study [21], sentiment analysis in economic news headlines is explored for predicting stock value changes. The analysis compares four sentiment analysis tools, including BERT and a Recurrent Neural Network (RNN), with BERT and RNN demonstrating superior performance in accurately determining emotional values. These findings significantly impact understanding of the correlation between emotions and stock market fluctuations.

Building upon this, Zhang et al. [22] propose a novel approach to fine-grained financial sentiment analysis (FSA). Their regression model integrates corpus-level statistics obtained through an autoencoder with semantic features, enhancing sentiment orientation prediction for financial texts. Experimental results demonstrate the method's effectiveness, showcasing significant improvements compared to baseline models without additional computational overhead. Furthermore, the study highlights the importance of considering neglected signs in FSA and their impact on sentiment score prediction.

Moreover, Araci et al. [23] introduce FinBERT, a domain-specific language model based on BERT, tailored for financial sentiment analysis. By leveraging pre-trained language models and fine-tuning financial data, FinBERT outperforms state-of-the-art machine learning methods, even with limited labeled data. The research emphasizes the effectiveness of incorporating pre-trained language models for understanding financial sentiment.

Regarding social media's impact on financial indices, Valle et al. [24] investigate the relationship between Twitter posts and financial markets during the pandemic. Through sentiment analysis of Twitter data and analysis of global economic data, the study identifies a significant influence of Twitter polarity on stock market behavior. Furthermore, by utilizing a lexicon-based approach and shifted correlation analysis, the research uncovers hidden correlations, highlighting the role of social media in shaping financial decision-making during the pandemic.

Another study by [3] examines public sentiment toward the Colombian presidential debate through sentiment analysis on Twitter. By focusing on tweets related to three candidates, the study utilizes the BERT model and achieves an accuracy of 76% and an F1 score of 85%. This research provides insights into the public sentiment toward presidential candidates during debates.

Puertas' doctoral thesis [25] aims to improve polarity interpretation in written texts by incorporating phonetic and emotional elements at various linguistic levels. By analyzing microblogging sources, the study investigates the contribution of phonetic and emotional features in predicting polarity. Results demonstrate that combining lexical, semantic, phonetic, and emotional elements achieves an impressive F1 measure of approximately 80% for polarity detection.

Furthermore, Puertas et al. [26] propose an approach integrating Design Science Research with Natural Language Processing (NLP), Computational Linguistics (CL), and Artificial Intelligence (AI) techniques for detecting sociolinguistic features in digital social networks. Through a case study that analyzes the semantic values of Twitter accounts belonging to Colombian universities, the research successfully identifies sociolects, uncommon words, and community-specific vocabulary. This finding emphasizes the efficacy of their approach.

Last, Pan et al. [27] evaluate transformer models for financial targeted sentiment analysis in Spanish. They address the challenges of extracting the main economic target and determining sentiment polarity from financial texts, providing a valuable corpus of Spanish financial tweets and news headlines. The evaluation of different Spanish-specific large language models demonstrates the performance of MarIA and BETO. This research contributes to developing effective sentiment analysis techniques for financial data in Spanish.

3. Methodology

In this study, we address the specific requirements of the reviewer while incorporating additional relevant information. The revised methodology focuses solely on meeting the reviewer's request without directly referencing their comments.

We start by using a dataset comprising Spanish news headlines obtained from reputable digital newspapers specializing in economic, financial, and political news. To ensure the relevance of the data, we conduct a two-stage filtering process targeting sections that contain economic-related content, eliminating irrelevant headlines. Afterwards, the dataset undergoes manual annotation by three committee members who assign sentiment polarity to three entities: target, companies, and consumers. This results in a dataset of approximately 6,000 to 8,000 headlines, which we further divide into training and test sets for the shared task.

To perform sentiment analysis, we establish the following pipeline. First, the dataset undergoes a preprocessing stage where we remove stopwords and punctuation. Next, we extract features by combining phonestheme embeddings with the RoBERTa Transformer model [28]. These features are then combined using the join features technique. To ensure model stability, we apply regularization techniques. Finally, we employ a Support Vector Machine (SVM) model for sentiment classification.

To extract phonestheme embeddings, we utilize the CESS-ESP corpus for Spanish and the Brown Corpora for English. These corpora provide standard syntactic and lexical structures necessary for this study. Additionally, we use the NRC Valence, Arousal, and Dominance lexicon, which is created through manual annotation using Best-Worst Scaling. This lexicon includes a list of over 20,000 English words and their respective translations in Spanish, along with Valence (V), Arousal (A), and Dominance (D) scores. These scores range from 0 to 1, representing fine-grained real-value annotations. These resources play a crucial role in training the phonetic embeddings.

For training the phonestheme embeddings, we utilize the CESS-ESP corpus containing 188,650 syntactically annotated Spanish words and the Brown Corpora, which contains approximately one million syntactically annotated English words. These corpora are instrumental in capturing the syntactic and semantic properties necessary for the embeddings.

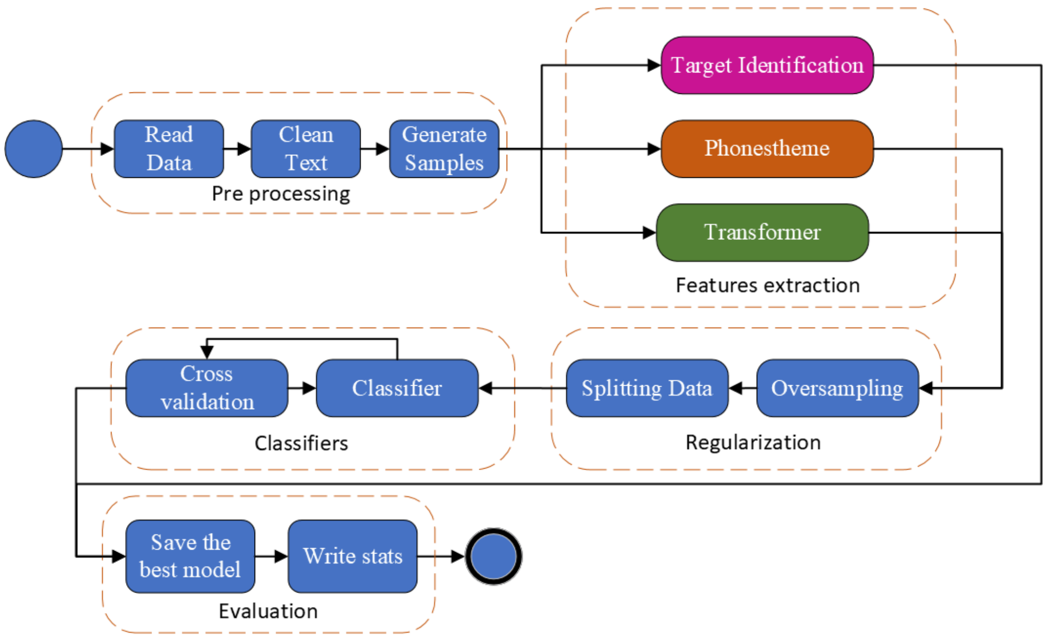

Figure 1: System Pipeline

The methodology also includes the extraction of phonetic elements. This process involves specific tasks such as extracting phonesthemes, phonetic frequencies, and phoneme sequences. These tasks are carried out using the trained embeddings and various linguistic techniques, such as syllable extraction and encoding in the International Phonetic Alphabet (IPA).

Additionally, we analyze the preprocessed data for target identification to identify words associated with Nsubj dependencies. We then identify chunks of words and search for Nsubj words within these chunks. The identified pieces serve as the targets.

4. System Overview

This section presents the predictive model utilized as a solution for Task 1 of the FinancEs Iberlerf 2023, which involves identifying the main economic target in a news headline and determining the sentiment polarity toward the identified target in the text.

To effectively tackle this task, our proposed model follows a systematic five-step approach 1. Firstly, we preprocess the raw data, ensuring proper formatting and readiness for subsequent analysis. This involves removing stopwords, handling punctuation, and standardizing the text.

Next, meaningful features are extracted from the preprocessed data, capturing essential information related to sentiment analysis. Various techniques, including advanced language models and contextual embeddings, are employed to capture semantic meaning and contextual understanding.

In the third step, we combine the extracted features to comprehensively represent the data. This integration aims to capture synergistic effects and enable our model to leverage a broader range of information for more informed predictions.

To enhance model stability and prevent overfitting, regularization techniques are applied in the fourth step. These techniques, such as dropout and weight decay, contribute to generalizing the model's predictions and improving its robustness.

Finally, a suitable algorithm, such as a support vector machine (SVM), is employed in the fifth step to classify the regularized data and determine the sentiment polarity toward the identified target. The SVM classifier leverages the extracted features and the integrated representation to make predictions based on learned patterns.

The model's performance evaluation is based on F1 metrics, which comprehensively assess sentiment polarity and target identification. The F1 score balances precision and recall, providing an accurate measure of the model's ability to capture the correct sentiment polarity and identify the main economic target within the text.

4.1. Data Description

The dataset utilized for this task consists of Spanish news headlines collected from specialized digital newspapers focusing on economic, financial, and political news. These newspapers, including Expansión, El Economista, Modaes, and El Financiero, are based in various Spanish-speaking countries. The challenge provides the dataset with pre-existing labels, indicating the target entity and sentiment polarity across three dimensions: target, companies, and consumers. These labels were assigned to each headline, classifying them as positive, neutral, or negative concerning the respective entities. The dataset underwent a two-stage filtering process to ensure its relevance, including identifying economic-related content from specific subsections of the newspapers and removing headlines not within the financial domain. The resulting dataset, composed of 6,000 to 8,000 news headlines, will be used for the shared tasks, with training and test sets released in an 80%-20% ratio.

4.2. Pre-processing

During the pre-processing stage, the raw data is subjected to various techniques to ensure its suitability for subsequent analysis. We employed several standard pre-processing methods, including removing stopwords, handling punctuation marks, and standardizing the text.

Stopword removal involves eliminating commonly occurring words that do not carry significant meaning or contribute to the analysis. Punctuation handling focuses on managing punctuation marks, such as commas, periods, and question marks, to maintain proper sentence structure and readability. Finally, text standardization involves converting the text to a consistent format, such as lowercase or removing diacritical marks, to promote uniformity and facilitate practical analysis and comparison.

By applying these pre-processing techniques, we cleaned the data, reduced noise, and achieved consistency, ensuring the data was ready for feature extraction and subsequent modeling stages.

4.3. Feature Extraction

We employ three feature extraction methods to capture the relevant information from the pre-processed data. Each approach focuses on extracting specific characteristics that contribute to sentiment analysis.

4.3.1. Target Identification

The first method, target identification, is crucial in sentiment analysis. It involves identifying the economic target or objective from the pre-processed text. To achieve this, we utilize the Spacy library to analyze the syntactic structure of the text. Specifically, we search for words or chunks associated with the Nsubj dependency, representing the subject of a sentence or phrase. By locating the first Nsubj word encountered, we identify the target entity around which the sentiment analysis revolves. This method allows us to pinpoint the critical element of interest in the text and analyze sentiment about it.

4.3.2. Phonestheme Features

The second method, phonestheme feature extraction, focuses on capturing phonetic representations of words. By considering the phonetic structure of words, we aim to charge additional nuances of financial language that contribute to sentiment analysis. Phonesthemes represent the smallest sound units in a language, and by leveraging pre-trained Word2Vec models specifically trained on phonesthemes, we obtain vector representations of these phonetic units. These phonestheme features provide a unique perspective on the text, enabling the model to capture phonetic patterns and further enhance its ability to analyze sentiment in financial texts. In addition, by incorporating phonetic information, we can uncover subtle variations in pronunciation that may carry sentiment implications.

4.3.3. Transformer Features

The third method, transformer feature extraction, utilizes a transformer-based model, specifically RoBERTa, to capture contextual information and semantic meaning from the pre-processed text. Transformers are powerful deep-learning models that capture complex contextual relationships within a text. By processing the pre-processed text through the RoBERTa transformer, we obtain contextualized embeddings that encode rich semantic and syntactic information. These transformer features play a crucial role in sentiment analysis, allowing the model to grasp the nuanced meaning and sentiment expressed in the text. The contextual information captured by the transformer enhances the model's understanding of sentiment-related concepts, idiomatic expressions, and subtle linguistic nuances in financial texts.

4.4. Feature Union

In this step, we combine the phonesthemes and transformer features to create a unified feature set for regularization. We achieved the combination by using the append() method, which allows us to merge the extracted features into a single dataset. By consolidating these different features, we create a comprehensive representation of the text data, capturing diverse aspects that contribute to sentiment analysis.

4.5. Regularization

Once the features have been extracted and combined, the next step is to address the class imbalance and prepare the data for modeling. We employ a two-step process to achieve this: class balancing using SMOTE (Synthetic Minority Over-sampling Technique) and data partitioning through k-fold cross-validation.

In the class balancing step, we account for the disproportionate distribution of polarities in the dataset. Then, SMOTE is applied to generate synthetic samples for the minority class, thereby equalizing the representation of different sentiment categories. This technique helps prevent bias toward the majority class and ensures a more robust and unbiased model training.

Following class balancing, we partitioned the dataset into training and testing subsets using k-fold cross-validation. This method divides the data into k equally sized folds, where each fold serves as a testing set once while we use the remaining k-1 folds for training. By repeatedly cycling through the folds, we obtain more reliable performance estimates and reduce the impact of data variability.

4.6. Classifier

We implemented several classifiers and evaluated their performance using the f1-score metric. Table 1 presents a detailed comparison of the classifiers' performance.

Table 1
Classifiers Results

Name	Model	Accuracy	Precision	Recall	F1
RF	Random Forest	0.68	0.69	0.68	0.68
DT	Decision Tree	0.59	0.59	0.59	0.59
NB	Gaussian Naive Bayes	0.58	0.67	0.58	0.54
SVM	Support Vector Machine	0.69	0.70	0.69	0.69
MLP	MLP Classifier	0.66	0.66	0.66	0.66
KNN	KNeighbors	0.65	0.67	0.65	0.65

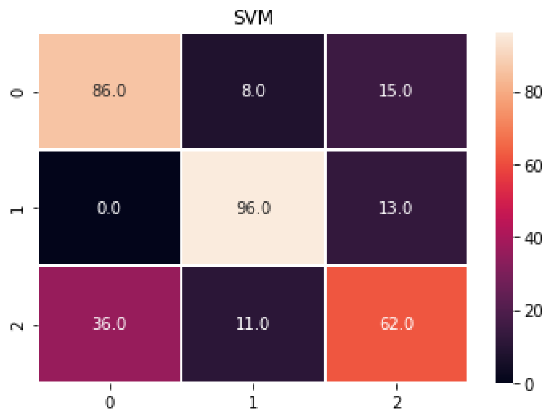

Figure 2: SVM Confusion Matrix.

We selected the Support Vector Machine (SVM) classifier for our sentiment analysis task after the evaluation. SVM is renowned for handling high-dimensional feature spaces and accurately classifying data points into different classes. We aim to achieve optimal sentiment analysis results by leveraging its robustness and flexibility.

Table 1 presents each classifier's performance metrics, including accuracy, precision, recall, and f1-score. Based on these quantitative results, the SVM classifier achieved the highest f1-score, indicating its superior performance in sentiment analysis for our task.

Next, we provide the confusion matrix for the SVM training results presented below in Figure 2, providing a better understanding of the classifier's behavior. Specifically, we observed that the classifier exhibits a higher misclassification rate for the positive polarity label. It indicates a tendency to predict positive sentiments as unfavorable, resulting in a false negative error. The classifier correctly predicts the positive polarity label only 62% of the time.

4.7. Evaluation

The test dataset was read and processed during the evaluation stage. Then, we evaluated the extracted features for Task 1 using the implemented classifier and recorded the results.

5. Experimental Setup

The methodology employed in this study comprises several steps. Firstly, the data undergoes preprocessing to cleanse and standardize the text. Subsequently, we extracted meaningful features from the text, including target identification, phonestheme features, and transformer features. Next, we combined these extracted features into a unified feature set. Next, regularization techniques are applied to address the class imbalance, utilizing the SMOTE technique for balancing sentiment polarities. Last, we employed the Support Vector Machine (SVM) model for classification and performance evaluation using the F1 score.

6. Results

We compared our polarity sentiment analysis and target identification performance with the baseline approach. We used the following metrics to evaluate the performance: F1 score for sentiment polarity (Task 1) and F1 score for target identification. The results obtained are presented in Table 2.

Table 2
Ranking Results Iberlef FinancEs 2023

F1 Task1	Target	F1 Target Sentiment
0.52923	0.41008	0.64838

Team UTB-NLP, led by user eapuerta, participated in Task 1 of the competition, focusing on financial targeted sentiment analysis. With seven submissions, our team achieved an impressive 6th position out of all participants. Our methodology involved preprocessing the dataset using transformers, phonestheme embeddings, and feature extraction and regularization techniques. Finally, we implemented an SVM classifier for sentiment classification.

The results showcased a remarkable F1 score of 0.52923, indicating our team's proficiency in identifying the main economic target from financial news headlines and determining the sentiment polarity toward the target in the processed text. Our successful performance highlights the effectiveness of our methodology in addressing sentiment analysis challenges within the financial domain, making a significant contribution to the advancement of this field.

7. Conclusion

In conclusion, this study has successfully addressed the challenges of targeted sentiment analysis in the financial domain, providing valuable insights into market sentiment and its economic impact. Through the analysis of a carefully curated dataset of news headlines from reputable financial and economic newspapers, we have developed an effective approach that combines transformer-based models and phonestheme embeddings. This approach captures contextual information, semantic meaning, and the nuances of financial language, leading to improved sentiment analysis results.

Our experimental results demonstrate the superiority of our methodology, surpassing the baseline with a significant improvement in Task 1 and achieving an impressive F1-score of 0.529229. These findings contribute to advancing sentiment analysis in the financial domain and offer valuable tools for understanding public opinion, making informed decisions, and predicting market behavior.

To further enhance the applicability and impact of our research, future work should focus on providing a more comprehensive and detailed description of the dataset used in this study. This would involve describing the collection process, annotation methodology, and the characteristics of the included news headlines. By providing this information, researchers will have a deeper understanding of the dataset and its potential applications, facilitating further advancements in targeted sentiment analysis within the financial domain.

Furthermore, it is important to recognize the limitations of this study. Factors such as the dataset size and specific challenges encountered during the analysis may have influenced the results to some extent. To address these limitations, future research should explore strategies for expanding the dataset and consider alternative approaches, such as incorporating domain-specific knowledge or exploring other types of embeddings.

The implications of accurate sentiment analysis in the financial domain are significant, enabling investors, financial institutions, and policymakers to gain deeper insights into market trends, improve investment strategies, and manage financial risks more effectively. Additionally, the proposed methodology has the potential for broader applications, such as financial risk assessment, market forecasting, and sentiment-driven trading strategies.

Acknowledgments

Thanks to the master's degree scholarship program in engineering at the Universidad Tecnológica de Bolivar (UTB) in Cartagena, Colombia.

References

[1] L. Putterman, Microeconomics: Behavior, institutions, and evolution, 2005.

[2] M. Hasan, J. Popp, J. Oláh, et al., Current landscape and influence of big data on finance, Journal of Big Data 7 (2020) 1–17.

[3] E. Puertas, J. C. Martinez-Santos, P. Andrés Pertuz-Duran, Presidential preferences in colombia through sentiment analysis, in: 2022 IEEE ANDESCON, 2022, pp. 1–5. doi:10.1109/ANDESCON56260.2022.9989700.

[4] A. S. George, A. H. George, T. Baskar, A. G. Martin, Human insight ai: An innovative technology bridging the gap between humans and machines for a safe, sustainable future, Partners Universal International Research Journal 2 (2023) 1–15.

[5] X. Man, T. Luo, J. Lin, Financial sentiment analysis (fsa): A survey, in: 2019 IEEE International Conference on Industrial Cyber Physical Systems (ICPS), IEEE, 2019, pp. 617–622.

[6] L. Zhao, L. Li, X. Zheng, J. Zhang, A bert based sentiment analysis and key entity detection

approach for online financial texts, in: 2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design (CSCWD), IEEE, 2021, pp. 1233–1238.
[7] K. Dashtipour, M. Gogate, J. Li, F. Jiang, B. Kong, A. Hussain, A hybrid persian sentiment analysis framework: Integrating dependency grammar based rules and deep neural networks, Neurocomputing 380 (2020) 1–10.
[8] R. Guarasci, S. Silvestri, G. De Pietro, H. Fujita, M. Esposito, Assessing bert's ability to learn italian syntax: A study on null-subject and agreement phenomena, Journal of Ambient Intelligence and Humanized Computing 14 (2023) 289–303.
[9] J. A. García-Díaz, F. García-Sánchez, R. Valencia-García, Smart analysis of economics sentiment in spanish based on linguistic features and transformers, IEEE Access 11 (2023) 14211–14224.
[10] K. Ishizuka, K. Nakata, Text mining for factor modeling of japanese stock performance, in: 2021 IEEE 8th International Conference on Industrial Engineering and Applications (ICIEA), IEEE, 2021, pp. 538–542.
[11] A. Gutiérrez-Fandiño, P. Kolm, J. Armengol-Estapé, et al., Fineas: Financial embedding analysis of sentiment, arXiv preprint arXiv:2111.00526 (2021).
[12] S. M. Jiménez-Zafra, F. Rangel, M. Montes-y Gómez, Overview of IberLEF 2023: Natural Language Processing Challenges for Spanish and other Iberian Languages, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2023), co-located with the 39th Conference of the Spanish Society for Natural Language Processing (SEPLN 2023), CEUR-WS.org, 2023.
[13] J. A. García-Díaz, Almela, F. García-Sánchez, G. Alcaráz Mármol, M. J. Marín-Pérez, R. ValenciaGarcía, Overview of FinancES at IberLEF 2023: Financial Targeted Sentiment Analysis in Spanish, Procesamiento del Lenguaje Natural 71 (2023).
[14] J. A. García-Díaz, M. P. Salas-Zárate, M. L. Hernández-Alcaraz, R. Valencia-García, J. M. Gómez-Berbís, Machine learning based sentiment analysis on spanish financial tweets, in: Trends and Advances in Information Systems and Technologies: Volume 1 6, Springer, 2018, pp. 305–311.
[15] F. Sun, A. Belatreche, S. Coleman, T. M. McGinnity, Y. Li, Pre-processing online financial text for sentiment classification: A natural language processing approach, in: 2014 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr), IEEE, 2014, pp. 122–129.
[16] E. Puertas, J. C. Martinez-Santos, Phonetic detection for hate speech spreaders on twitter, volume 2936, 2021.
[17] E. Puertas, Embedding of phonestheme in spanish, 2020. URL: https://doi.org/10.5281/zenodo.4299242. doi:10.5281/zenodo.4299242.
[18] A. M. Husein, B. Sipahutar, R. Dashuah, E. Hutauruk, Sentiment analysis od face to face school policy on twitter social media with support vector machine (svm), Sinkron: jurnal dan penelitian teknik informatika 8 (2023).
[19] L. G. Moreno-Sandoval, E. A. P. Del Castillo, A. P. Quimbaya, J. A. Alvarado-Valencia, Assembly of polarity, emotion and user statistics for detection of fake profiles., in: CLEF (Working Notes), 2020.
[20] C. Hays, Z. Schutzman, M. Raghavan, E. Walk, P. Zimmer, Simplistic collection and labeling practices limit the utility of benchmark datasets for twitter bot detection, arXiv preprint

arXiv:2301.07015 (2023).
[21] L. Nemes, A. Kiss, Prediction of stock values changes using sentiment analysis of stock news headlines, Journal of Information and Telecommunication 5 (2021) 375–394.
[22] H. Zhang, Z. Li, H. Xie, R. Y. Lau, G. Cheng, Q. Li, D. Zhang, Leveraging statistical information in fine-grained financial sentiment analysis, World Wide Web 25 (2022) 513–531.
[23] D. Araci, Finbert: Financial sentiment analysis with pre-trained language models, arXiv preprint arXiv:1908.10063 (2019).
[24] D. Valle-Cruz, V. Fernandez-Cortez, A. López-Chau, R. Sandoval-Almazán, Does twitter affect stock market decisions? financial sentiment analysis during pandemics: A comparative study of the h1n1 and the covid-19 periods, Cognitive computation (2022) 1–16.
[25] E. A. Puertas Del Castillo, et al., Análisis de elementos fonéticos y elementos emocionales para predecir la polaridad en fuentes de microblogging (2023).
[26] E. Puertas, L. G. Moreno-Sandoval, J. Redondo, J. A. Alvarado-Valencia, A. Pomares-Quimbaya, Detection of sociolinguistic features in digital social networks for the detection of communities, Cognitive Computation 13 (2021) 518–537.
[27] R. Pan, J. A. García-Díaz, F. Garcia-Sanchez, R. Valencia-García, Evaluation of transformer models for financial targeted sentiment analysis in spanish, PeerJ Computer Science 9 (2023) e1377.
[28] J. M. Pérez, D. A. Furman, L. A. Alemany, F. Luque, Robertuito: a pre-trained language model for social media text in spanish, arXiv preprint arXiv:2111.09453 (2021).