Profiling Cryptocurrency Influencers With Few-shot Learning via Contrastive Learning

Notebook for PAN at CLEF 2023

Zengyao Li, Zhongyuan Han*, Jiyin Cai, Zhijian Huang, Shaohua Huang and Leilei Kong

Foshan University, Foshan, China

Abstract

We proposed a novel approach that leverages embedding augmentation based on contrastive learning to address the issue of few-shot learning in the task of analyzing cryptocurrency influencers.This representation learning technique enables the model to better distinguish between similar and dissimilar samples, thereby enhancing generalization in few-shot tasks.In our method, we treat sentences with the same label within a batch as positive samples while considering sentences with different labels as negative samples. We conduct additional pre-training to enhance the embedding representation capabilities of the encoder. Subsequently, we fine-tuned the model and used it for classification. Our method achieved the 4th official ranking on the official test dataset for this evaluation.The experimental results serve as compelling evidence of the effectiveness and utility of a contrastive learning-based approach in addressing the few-shot problem in analyzing cryptocurrency influencers. By employing this innovative method, we are able to effectively tackle the challenges associated with analyzing cryptocurrency influencers with limited labelled data.

Keywords

contrastive learning, few-shot learning, profiling cryptocurrency influencers

1. Introduction

Cryptocurrencies have gained significant popularity in recent years [1], with social media playing a crucial role in shaping public perceptions and influencing investment decisions. Understanding the characteristics and behaviours of cryptocurrency influencers on social media platforms can provide valuable insights for various stakeholders, including investors, researchers, and policymakers [2, 3].

In this study, we present a contrastive learning-based approach to address the PAN at CLEF 2023 task: Profiling Cryptocurrency Influencers with Few-shot Learning. The task comprises three sub-tasks, each targeting a specific aspect of cryptocurrency influencers. The first sub-task focuses on low-resource influencer profiling, with the objective of categorizing influencers into distinct classes based on their impact and reach. The second sub-task centres around

CLEF 2023: Conference and Labs of the Evaluation Forum, September 18–21, 2023, Thessaloniki, Greece
lzy1512192979@gmail.com (Z. Li); hanzhongyuan@gmail.com (Z. Han*); caijiyin0904@163.com (J. Cai); 627872946@qq.com (Z. Huang); 3390794416@qq.com (S. Huang); kongleilei@fosu.edu.cn (L. Kong)
0000-0001-8472-4150 (Z. Li); 0000-0001-8960-9872 (Z. Han*); 0009-0001-6418-908X (J. Cai); 0009-0002-5049-2093 (Z. Huang); 0009-0004-6659-0936 (S. Huang); 0000-0002-4636-3507 (L. Kong)

low-resource influencer interest identification, aiming to classify influencers' interests based on their Twitter posts. The third sub-task revolves around identifying influencers' intent in a low-resource context, where the goal is to discern the underlying purpose behind their tweets.

2. Related Work

The PAN committee organizing has initiated the Profiling task annually from 2013 to 2022. The PAN author profiling task spanned from 2013 to 2016 with the primary goal of determining the gender and age distribution of authors based on their documents [4, 5, 6, 7]. The following year, the PAN challenge expanded its scope beyond gender detection and incorporated language diversity identification as well [8]. In 2018, the task adopted a multimodal approach by providing participants with both textual and image data to detect the author's gender [9]. As the number of social media users surged in 2019, the PAN author profiling task shifted its focus to Twitter. In this instance, the objective involved not only bot detection but also the identification of gender among human authors [10]. Subsequently, the tasks in 2020, 2021, and 2022 revolved around identifying propagators of Twitter fake news, hate speech, and stereotype dissemination, respectively [11, 12].

Before 2018, researchers primarily focused on discussing text preprocessing and feature extraction methods for traditional classifiers. However, starting in 2018, there has been a shift towards the adoption of new deep-learning algorithms. With the introduction of the Transformer model [13], an increasing number of teams have opted to utilize the BERT model for author profiling tasks.

In the profiling task of the previous year, the team that achieved the top ranking employed a fusion model technique. They employed three distinct loss functions to fine-tune the Bertweet model and incorporated a CNN model based on BERT features. Additionally, they utilized soft voting to combine the predictions from these models, resulting in the highest performance. There are other teams that have also achieved remarkable results using ensemble models.

3. Model

3.1. Our Method

For this task, we propose a supervised contrastive learning method inspired by Simcse[14]. The method revolves around constructing groups of three sentences denoted as G = (x, x+, x-), where x represents the source sample, x+ denotes the positive sample with the same label, and x- represents the negative sample with a different label. During batch training, apart from using x- from the sentence group G1 as a negative sample, negative samples from other groups Gi (i!=1) within the same batch are also employed for training.

In figure 1, the subscript of S represents the category of the sentence. The purple line represents the source sample and the positive sample, where both samples belong to the same class label. Our objective is to minimize the semantic distance between them. In contrast, the green line represents the source sample and the negative sample, where these two samples have different class labels. Our objective is to maximize the semantic distance between them. In the

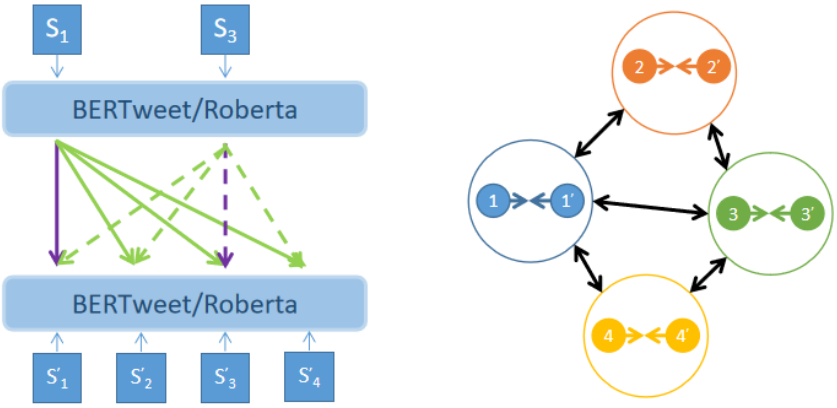

Figure 1: Contrastive Learning Framework

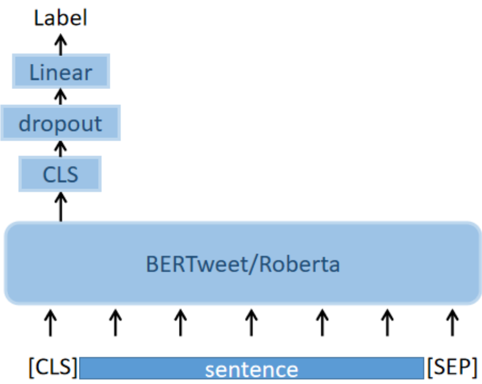

Figure 2: BERT-based model architecture

figure, the solid line and the dashed line depict the relationships between sentences within two batches. The training progress at this stage is evaluated using the Spearman coefficient, and the model with the highest coefficient is selected and further fine-tuned.

In the 4-classification task, Our goal is to obtain the feature space distribution shown in the right half of Figure 1 for the model. The samples from the same class exhibit close proximity, while samples from different classes are well-separated with significant distances between them.

3.2. Classifier

We selected the RoBERTa model and the BERTweet model as our baseline models. To serve as classifiers for prediction, we augmented the models with dropout and linear layers. Following the contrastive learning-based method described above, we conducted further pre-training of the baseline models. Subsequently, we connected the pre-trained models to a classifier for fine-tuning, ultimately obtaining the final model. The model architecture is illustrated in Figure 2.

4. Experiment

4.1. Dataset

This year's Pan evaluation task is a monolingual few-shot task aimed at identifying cryptocurrency influencers on Twitter. The task consists of three subtasks. Subtask 1: The training set for this subtask consists of five labels, with 32 users per label. Each label has a maximum of ten English tweets, resulting in a total of 160 users. The test set contains 220 users, with a maximum of 10 tweets per user. Subtask 2: This subtask includes five labels, with 64 users per label. Each label has one English tweet. The test set consists of 402 users, with one tweet per user. Subtask 3: This subtask involves four labels, with 64 users per label. Each label has one English tweet. The test set comprises 292 users, with one tweet per user. To process these data, we apply different processing strategies to each subtask as follows:

- `Base Scheme 1(subtask 1)`
  Text splicing (multiple tweets from the same user are spliced, when a tweet exceeds 510 lengths, it will be spliced into a new one)
  Multiple line breaks are replaced with one, and multiple spaces are also replaced with one Convert emoticons to English text
- `Base Scheme 2 (subtask 2, subtask3)`
  Multiple line breaks are replaced with one, and multiple spaces are also replaced with one Convert emoticons to English text
- `Addtional Scheme 1 (all subtasks)`
  Note $ and # (add currency to vocabulary)
- `Addtional Scheme 2 (subtask2, subtask3)`
  Pairwise combination with label text for data augmentationStyle of author name (chinese)

We perform various experiments by combining fundamental strategies with additional strategies, and we will analyze the specific results in Section 5.

4.2. Experimental setting

In order to train the model and evaluate its effectiveness, a 5-fold cross-validation method was employed during the training phase. The optimal hyperparameters that were found are shown in table 1. In this experiment, the models we ultimately submitted were obtained by fine-tuning the hyperparameters in Table 1. Additionally, during the fine-tuning stage of the pre-trained model, we utilized cross-entropy as the loss function for model. The entire experiment was conducted using the PyTorch framework.

Table 1
Hyperparameter Settings

Model	lr	Batch size	Epoch	Optimizer	Scheduler
Bertweet-base	2e-5	8	50	AdamW	get_linear_schedule_with_warmup
Roberta-base	5e-5	8	50	AdamW	get_linear_schedule_with_warmup

Table 2
Comparison of Macro-F1 between Different Models of Subtask1

Model	1	2	3	4	5	Avg	Test
Bertweet-base	50.2857	54.7403	55.1313	47.6398	48.5000	51.2594	/
Roberta-base	59.3073	51.2732	57.8632	53.2467	59.8787	56.3138	/
Bertweet-base + CL	60.0733	53.3427	59.7034	56.2397	60.5312	57.9781	55.1

5. Result

This section introduces the experimental results.

Table 2 exhibits the outcomes obtained through the utilization of 5-fold cross-validation on the training dataset of subtask 1, along with the results obtained on the test dataset. The model we finally submitted was trained using the contrastive learning method, and it achieved a score of 55.1, placing us in the 7th position.

Table 3 showcases the results and test set outcomes derived from 5-fold cross-validation performed on the training set of subtask 2. In addition to using contrastive learning, we also conducted experiments using AWP adversarial training and the combination strategy. For the final model submission, we included both the contrastive learning model and an ensemble model that combines all three methods. As shown in the figure, the contrastive learning model achieved a higher score compared to the ensemble model. Based on our score, we ultimately ranked 5th in this subtask.

In Table 4, you can find the results and test set outcomes obtained through 5-fold cross-validation on the training set of subtask 3.Regarding subtask 3, our submitted model utilized the same approach as for subtask 2, and the experimental results ultimately demonstrated that using the contrastive learning method alone outperformed the performance of the ensemble model. In subtask 3, we ultimately achieved 11th place. Our speculation is that there may be a conflict between adversarial training and contrastive learning methods, leading to the observed decrease in model performance. Our next step will be to investigate the underlying reasons behind this discrepancy to understand the root cause.

6. Conclusion

To address the task of analyzing cryptocurrency influencers proposed by PAN2023, we propose in this paper an approach of contrastively learned fusion models. Additionally, we also explored the performance of adversarial training in this task. We discovered that there may be a certain

Table 3
Comparison of Macro-F1 between Different Models of Subtask2

Model	1	2	3	4	5	Avg	Test
Bertweet-base	60.3595	70.6666	64.3458	82.6203	66.2257	68.8436	/
Roberta-base	64.9344	70.5991	54.0474	68.6010	70.2701	65.6904	/
Bertweet-base + combination	54.9411	73.8650	58.9674	76.6224	65.6087	66.0009	/
Bertweet-base + AWP	74.5034	72.1264	60.4984	78.1898	67.0887	70.4813	/
Bertweet-base + CL	73.9417	73.7750	62.4980	79.7738	67.32908	71.4635	61.63
Ensemble model	/	/	/	/	/	/	55.30

Table 4
Comparison of Macro-F1 between Different Models of Subtask3

Model	1	2	3	4	5	Avg	Test
Bertweet-base	82.7838	74.7717	78.6280	70.4873	82.1657	77.7673	/
Roberta-base	80.8509	77.1247	76.7732	71.8944	82.3757	77.8038	/
Bertweet-base + combination	84.6768	72.9624	78.1677	69.5625	84.3846	77.9508	/
Bertweet-base + AWP	84.5232	76.8985	80.6422	70.7047	88.3703	80.2278	/
Bertweet-base + CL	82.4770	78.6378	84.4551	71.3504	88.2307	81.0302	57.62
Ensemble model	/	/	/	/	/	/	53.72

contradiction between adversarial training and contrastive learning, resulting in a decrease in performance when both methods are used simultaneously. Ultimately, our proposed method achieved the 4th position in the official ranking.

7. Acknowledgments

This work is supported by the Natural Science Foundation of Guangdong Province, China (No.2022A1515011544).

References

[1] R. Sawhney, S. Agarwal, V. Mittal, et al., Cryptocurrency bubble detection: A new stock market dataset, financial task & hyperbolic models, in: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2022, pp. 5531–5545.

[2] M. Chinea-Rios, I. Borrego-Obrador, M. Franco-Salvador, et al., Profiling cryptocurrency influencers with few shot learning at pan 2023, in: CLEF 2022 Labs and Workshops, Notebook Papers, 2023.

[3] J. Bevendorff, M. Chinea-Ríos, M. Franco-Salvador, et al., Overview of pan 2023: Authorship verification, multi-author writing style analysis, profiling cryptocurrency influencers, and trigger detection, in: European Conference on Information Retrieval, Springer, 2023, pp. 518–526.

[4] F. Rangel, P. Rosso, M. Koppel, et al., Overview of the author profiling task at pan 2013, in: CLEF conference on multilingual and multimodal information access evaluation, CELCT, 2013, pp. 352–365.
[5] F. Rangel, P. Rosso, I. Chugur, et al., Overview of the 2nd author profiling task at pan 2014, in: CLEF 2014 Evaluation Labs and Workshop Working Notes Papers, Sheffield, UK, 2014, 2014, pp. 1–30.
[6] F. Rangel, P. Rosso, M. Potthast, et al., Overview of the 3rd author profiling task at pan 2015, in: CLEF, 2015, pp. 1–8.
[7] F. Rangel, P. Rosso, B. Verhoeven, et al., Overview of the 4th author profiling task at pan 2016: cross-genre evaluations, in: Working Notes Papers of the CLEF 2016 Evaluation Labs. CEUR Workshop Proceedings/Balog, Krisztian [edit.]; et al., 2016, pp. 750–784.
[8] F. Rangel, P. Rosso, M. Potthast, et al., Overview of the 5th author profiling task at pan 2017: Gender and language variety identification in twitter, Working notes papers of the CLEF (2017) 1613–0073.
[9] F. Rangel, P. Rosso, M. Montes-y Gómez, et al., Overview of the 6th author profiling task at pan 2018: multimodal gender identification in twitter, Working notes papers of the CLEF (2018) 1–38.
[10] F. Rangel, P. Rosso, Overview of the 7th author profiling task at pan 2019: bots and gender profiling in twitter, Working notes papers of the CLEF 2019 evaluation labs 2380 (2019) 1–7.
[11] F. Rangel, A. Giachanou, B. H. H. Ghanem, et al., Overview of the 8th author profiling task at pan 2020: Profiling fake news spreaders on twitter, in: CEUR workshop proceedings, volume 2696, Sun SITE Central Europe, 2020, pp. 1–18.
[12] F. Rangel, G. L. d. l. Peña-Sarracén, M. A. Chulvi-Ferriols, et al., Profiling hate speech spreaders on twitter task at pan 2021, in: Proceedings of the Working Notes of CLEF 2021, Conference and Labs of the Evaluation Forum, Bucharest, Romania, September 21st to 24th, 2021, CEUR, 2021, pp. 1772–1789.
[13] A. Vaswani, N. Shazeer, N. Parmar, et al., Attention is all you need, Advances in neural information processing systems 30 (2017).
[14] T. Gao, X. Yao, D. Chen, Simcse: Simple contrastive learning of sentence embeddings, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 6894–6910.