Reshape or Update? Metric Learning and Fine-tuning for Low-Resource Influencer Profiling

Notebook for PAN at CLEF 2023

Roberto Labadie Tamayo[1], Areg Mikael Sarvazyan[1,2]

[1]Universitat Politècnica de València, València, Spain
[2]Symanto, Spain

Abstract
In these working notes, we present our contributions to the *"Profiling Cryptocurrency Influencers with Few-shot Learning"* shared task at PAN 2023 under the team name *pan23-dogecoin*, where we participated in all three subtasks. We focused on leveraging Metric Learning and Parameter-Efficient Fine-Tuning of Transformer-based models. We conducted extensive hyper-parameter search for both techniques, resulting in early-bird results of a tri-encoder model with 0.508 of macro-f1 for our best-performing system for subtask 1, and an application of PEFT techniques to Transformer-based language models which obtained 0.517, and 0.526 of macro-f1 for subtasks 2 and 3 respectively. We observed that our proposed metric learning-based system exhibited enhanced generalization capabilities with respect to the fine-tuning of Transformer-based models. Notably, despite the larger parameter size of some Transformer-based models, we obtained robust performance by utilizing smaller models pre-trained on domain-specific knowledge.

Keywords
Author Profiling, Metric Learning, Fine-tuning, Few-shot,

1. Introduction

Social media has become a fundamental part of our day-to-day lives. Through these services, millions of users connect and discuss various topics. Over the last decade, there has been an increasing interest in cryptocurrencies, providing the frame for the emergence of new digital communities sharing information. Consequently, understanding the behavior of these communities has become crucial for individuals invested in cryptocurrencies who wish to understand market trends, identify influential figures, and predict the success of new cryptocurrencies.

Numerous studies have explored the relationship between social media and the cryptocurrency ecosystem, highlighting the profound impact that online communities can have on the market [1, 2]. These studies have shown that influential figures, such as well-known cryptocurrency analysts, traders, and enthusiasts, can significantly influence the sentiment and adoption of specific cryptocurrencies through their social media activities. Positive or negative opinions

CLEF 2023: Conference and Labs of the Evaluation Forum, September 18-21, 2023, Thessaloniki, Greece
rlabtam@posgrado.upv.es (R. L. Tamayo); arsar1@inf.upv.es (A. M. Sarvazyan)
0000-0003-4928-8706 (R. L. Tamayo); 0009-0009-4055-1601 (A. M. Sarvazyan)

expressed by these individuals can sway public perception, impacting the value and market demand for different digital assets.

Furthermore, social media platforms serve as a vital source of information for individuals seeking to stay informed about the latest news, developments, and trends within the cryptocurrency space. Hashtags, dedicated forums, and communities surrounding cryptocurrencies on platforms such as Twitter provide a space for discussions, news dissemination, and the sharing of insights among community members.

All above form the basis of the shared task *"Profiling Cryptocurrency Influencers with Few-shot Learning"* [3] at PAN 2023 [4], where the primary objective is to profile Twitter users to determine their impact on cryptocurrency communities.

These working notes describe our participation in the shared task. More specifically, in Section 2 we review previous studies related to author profiling and few-shot learning before delving deeper into the tasks. In Section 3, we describe the subtasks and datasets provided by the organizers. Afterward, in Section 4 we introduce our approaches to tackle the challenge, and in Section 5 we describe our experimental setup and the achieved results. Finally, we draw some conclusions and present future lines of research in Section 6.

2. Related Work

When profiling users through semantic and stylometric features extracted from their textual constructions on a social media platform such as Twitter, one must consider that single message information is usually not enough for accurate profiling. Aggregating and contrasting this information from Twitter feeds is a matter explored along several shared tasks at PAN. These tasks have involved analyzing demographic characteristics of social media users [5, 6, 7, 8], such as gender, language variety, age, etc.

In recent years more complex tasks aiming at identifying personal traits and psychosocial behavior of users have been explored, ranging from profiling fake news [9] and hate speech [10] spreaders, to users using irony to target social groups according to stereotypical categorizations [11].

In the past, most proposed works applied traditional machine learning techniques to profile authors, relying mainly upon Support Vector Machines [12, 13], Logistic Regression [14] and Random Forest [15] over representations of the profiles as Bag of Words, n-grams and their term-frequency based vector representation combined with linguistic features.

However, given further developments in neural models and their ability to extract a more adequate vector representation of text, many works also introduced neural techniques to both model the set of tweets into a profile representation, and classify these profiles. These works, in many cases, did not exclude the use of linguistic features to combine them with architectures like Long Short Term Memory neural networks [16, 17], Convolutional Neural Networks [18], and Transformer-based models [19, 20, 21].

In this year's edition of PAN, the complex task of handling psychosocial aspects of subjects is studied in a low-resource scenario, where the amount of information about users belonging to each class is very limited. The latter makes the challenge of profiling users a few-shot task. The field of few-shot learning has gained significant attention as a promising approach for

Table 1
Dataset statistics for each subtask. UPC denotes Users Per Class, TPU denotes Tweets per User.

Subtask		Classes	UPC	TPU
(i)	*Profiling*	5	32	10
(ii)	*Interest Id.*	5	64	1
(iii)	*Intent Id.*	4	64	1

addressing the challenge of limited data availability. Few-shot learning aims to develop models that can learn from only a few examples per class, mimicking the human ability to generalize and recognize new concepts with limited exposure. One notable advancement in few-shot learning is the introduction of Large Language Models (LLMs), such as OpenAI's GPT-3 [22]. These models have been leveraged to tackle various natural language processing (NLP) tasks unseen at training time, including text classification, sentiment analysis, and language generation, by using strategies such as in-context learning [23].

Nevertheless, to obtain competitive results, these few-shot techniques require models with considerably large context lengths and usually perform better as they are scaled to larger number of learned parameters. Hence, appealing to models whose training relies on prior knowledge about the similarity between the instances of the corpus, such as Bi-Encoder [24] and Tri-Encoder networks [25], constitute a viable solution for many low resources tasks in terms of computational resources and performance [26], especially for author profiling tasks [27].

3. Data

We participate in the three subtasks proposed in the *"Profiling Cryptocurrency Influencers with Few-shot Learning"* shared task: (i) low-resource influencer profiling, (ii) low-resource influencer interest identification, and (iii) low-resource influencer intent identification. We only employ the data provided by the task organizers [28], whose statistics are presented in Table 1. Here we clearly observe that subtask 1 corresponds to the regular schema of a profiling task, whereas the remaining two subtasks can also be framed as text classification tasks instead of author profiling ones, given that for each author, we only have a single tweet instance.

4. Submitted Systems

The systems described in these working notes rely on the use of Large Language Models (LLMs), using them in two different approaches: (i) as text encoders to find relationships in a high-dimensional embedding space and (ii) as backbones to be fine-tuned for the three downstream subtasks. We frame the first subtask as an author profiling task, treating the remaining ones as classification tasks. To this end, we acquire textual definitions of the target classes to characterize a type of author profile attending to the scale in which their tweets influence the cryptocurrency community (see in Appendix A). We leverage the definitions for each category independently of the actual task of profiling, taking into account the behavior of a Twitter user[1]. While these definitions do not reflect the writing style or content of a tweet, we appealed to the capability of LLMs to determine useful relationships from these descriptions such that they can discriminate between classes in a zero-shot setting.

4.1. Metric Learning

According to previous works on author profiling and few-shot learning [29, 26, 30], our first approach consisted in experimenting with various configurations of bi-encoder and tri-encoder networks.

For this, we opted for a pre-trained Transformer-based [31] language model as a backbone to obtain profile-level embeddings, and learning a transformation in a latent space. Here, we optimize such that samples from the same class are closer together[2], whereas samples from different classes are farther apart. Finally, an instance of a user profile is classified into the class of its closest prototype, i.e., an instance from the training set or user definition.

In the contrastive training of this metric learning approach, we explored both a bi-encoder strategy and a triplet-based approach. A user's profile, `[AUTHOR-PROFILE]`, was constructed by concatenating their tweets with the a newly introduced `[SEP-TWEET]` token, embedding the resulting text as follows:

Represent the cryptocurrency Tweet posts for profiling their authors into the classes nano micro macro mega non-influencer; Input: `[AUTHOR-PROFILE]`

We also took into account the definitions, `[USER-DEFINITION]`, mentioned in Section 4 as elements of our metric space when building the pairs (or triplets) for training. These definitions were considered as instances belonging to their corresponding class. To obtain the embeddings of definitions, we used the following prompting scheme:

Represent the cryptocurrency influencers definition for comparing it with the representation of their tweets; Input: `[USER-DEFINITION]`

Finally, we introduced an online computation of the training pairs and triplets by recomputing them after a fixed number of epochs to stress the learned transformation function towards higher inter-class sparsity. For contrastive pairs we approach the construction of negative examples under a hard criterion by selecting pairs of elements from different classes that are closer in the newly-learned embedding space. Similarly for triplets, given an anchor example, we select five random examples from its own class as positive points, and get the five closest examples from the complementary classes as negatives, also ensuring a balanced distribution.

4.2. Fine-tuning

In addition to the metric learning strategy described in Section 4.1, we explored a straightforward classification approach by fine-tuning a language model for each subtask. For this, we adapt a pre-trained Transformer-based text encoder to each subtask using the state-of-the-art parameter-efficient fine-tuning technique LoRA [32]. This simpler approach allows us to avoid

[1]Extracted and modified from https://greenhousemag.com/
[2]Using the squared euclidean distance.

the construction of contrastive examples, which can introduce some undesired bias into the learning process even when we compute the triplets and pairs dynamically. Nevertheless, this kind of language model is limited when dealing with long context sequences, forcing us to truncate the user profile to 512 tokens. This means that most profile sequences will be truncated, as can be seen in Figure 1.

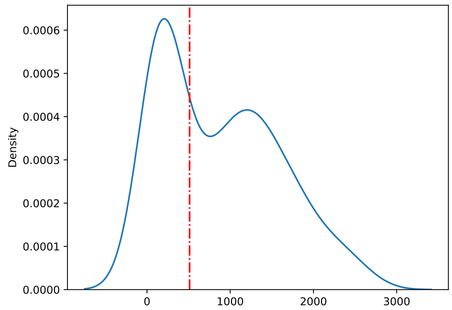

Figure 1: Length distribution for the profile sequences in the training set. The vertical red line is the cutoff at 512 tokens.

5. Experiments and Results

To evaluate the performance of our models in the development process, we employed a 4-fold cross-validation strategy while preserving class balance. No text pre-processing was carried out for any approach because we were looking for an end-to-end methodology that avoids conditioning the learning process of the models on hand-crafted heuristics. Following the task organizers' decision, we report mean macro-f1 between the validation results in each fold, as well as the macro-f1 of the private test set in our early-bird submissions.

5.1. Metric Learning

The bi-encoder approach was only used for the first subtask. This subtask was the only one among all the three where a simple transformer-based sentence encoder described in Section 4.2 was unable to handle the complete profile sequences, which we address by the use of INSTRUCTOR [33] as the profile encoder. However, given its maximum context length of 2048, truncation was still necessary.

The profile encoder was kept frozen; only the transformation layer was learned using Adam [34] with learning rates ranging from 1e-4 to 1e-1. The dimension of this latent space and the frequency of pair or triplet recalculation was studied, in addition to whether the definitions as prototypes should be included to train the network and classify new samples. To accomplish this tuning process, we relied on Bayesian optimization as implemented in Optuna [35].

During the training, we used two loss functions, contrastive [36] or triplet loss [37] for bi-encoder and triplet-based strategies, respectively. In Table 2 we present the results of the best combination involving both perspectives under an online example construction schema, which involved updating the training examples after every five epochs.

Table 2
Mean macro-f1 in a 4-fold cross-validation setting for metric learning in subtask 1. *Def.* refers to using definitions as anchors, *Ins.* refers to employing also including all training instances as anchors.

Strategy	Macro-F1
Contrastive + Ins.	0.332
Contrastive + Def. + Ins.	0.356
Triplet + Ins.	0.449
Triplet + Def. + Ins.	**0.461**

From this table, we observe how introducing the definitions from Section 4 improves the results obtained by both pairs and triplet-based systems. Additionally, providing positive and negative pairs, i.e., using triplet loss, results in considerable improvements in our metric learning approach.

5.2. Fine-tuning

When fine-tuning language models, we compared their performance by also taking into account the size of the models due to hardware restrictions in the submission platform. For those models which allow it, we studied whether employing low-rank adapters outperforms conventional fine-tuning.

In all the cases, the tuned parameters were optimized with RMSprop [38]. Particularly when applying conventional fine-tuning of the whole network, we apply an increasing learning rate from shallower layers to deeper ones [39], starting from 1e-5 and increasing it on each layer with a factor of 0.1 units.

We experimented with full fine-tuning and LoRA as implemented in the HuggingFace ecosystem [40] for three language models: CryptoBERT[3], and two BLOOM [41] models, BLOOM-1b1[4] and BLOOM-7b1[5]. Table 3 shows the evaluation results after fine-tuning these models.

Table 3
Mean macro-f1 in a 4-fold cross-validation setting for fine-tuned models in all subtasks.

Strategy	Subtask 1	Subtask 2	Subtask 3
CryptoBERT	0.554	0.405	0.647
CryptoBERT + LoRA	**0.68**	0.514	**0.608**
BLOOM-1b1	0.610	0.504	0.551
BLOOM-1b1 + LoRA	0.623	**0.528**	0.588
BLOOM-7b1 + LoRA	0.651	0.520	0.608

From these results, we observe that applying LoRA outperforms the full model fine-tuning for CryptoBERT and BLOOM-1b1. Given BLOOM-7b1 model size, we were unable to experiment

[3]tinyurl.com/cryptobert
[4]tinyurl.com/bloom-1b1
[5]tinyurl.com/bloom-7b1

with full fine-tuning. We hypothesize that given the number of examples in this low-resource dataset, catastrophic forgetting of learned information in the pre-training phase [42] may become more critical given the bias imposed towards such a small training set.

Additionally, while BLOOM models exhibit many more parameters than CryptoBERT, given the domain in which the latter has been pretrained and finetuned, this model achieves better performance in two of the three tasks. Finally, we note that for subtask 1 all fine-tuned models outperform our metric learning approach.

5.3. Official Results

Regarding the official submissions to the *"Profiling Cryptocurrency Influencers with Few-shot Learning"* task in early-bird stages, we performed a majority vote for the models obtained by finetuning each fold from the cross-validation process. In this way, we submitted one run for subtask 1 involving the use of the best triplet-based model, obtaining a macro-f1 of 0.508 in the test set. We also submitted runs using CryptoBERT-LoRA obtaining 0.465, 0.517 and 0.526 of macro-f1 for subtasks 1, 2 and 3 respectively.

Interestingly, in contrast with the results from our cross-validation scheme, the tri-encoder system obtains more robust results for subtask 1, which is aligned with our hypothesis related to its capacity of processing long profile sequences. In addition to this, we believe the fact that it only learns a transformation of the representation space, where objects from different classes are distant, and that the constructed training examples grow in cubic orders, making tri-encoder more robust and generalizable than fine-tuned language models. We find this is an important lesson: a robust cross-validation scheme should be one of the topmost priorities, especially when participating in competitions with private test sets.

6. Conclusions

In these working notes, we describe our submissions to the *"Profiling Cryptocurrency Influencers with Few-shot Learning"* shared task at PAN 2023. We participated in all three subtasks, experimenting with bi-encoders and PEFT of transformer-based models. Our best-performing systems achieved macro-F1 measurements of 0.508, 0.517, and 0.526 for tasks 1, 2, and 3 in the early-bird stages, respectively, based on the official test set. From our study, we observed that the proposed metric learning-based system demonstrated superior generalization capabilities with respect to fine-tuned language models. Additionally, despite the large parameter size of some Transformer-based models, we obtained more robust performance by employing models pre-trained on knowledge aligned with the specific task domain.

While the primary focus of this task was to evaluate machine learning models' ability to learn in low-resource scenarios, for future work, we intend to investigate how augmenting the training examples using methods such as back-translation impacts the systems' generalization capabilities. Furthermore, considering the decent results achieved by transformer-based models, we plan to gather a more informative characterization of the classes in terms of writing style and content. This characterization will help us select the elements that belong to the final profile sequence by contrasting them with these defined characteristics. Subsequently, a condensed and reduced version of the profile will be fed into the discriminative models.

Acknowledgments

This work has been partially developed with the support of valgrAI - Valencian Graduate School and Research Network of Artificial Intelligence and the Generalitat Valenciana, and co-funded by the European Union.

References

[1] H. Omar, L. A. Lasrado, Uncover social media interactions on cryptocurrencies using social set analysis (ssa), Procedia Computer Science 219 (2023). doi:https://doi.org/10.1016/j.procs.2023.01.277, cENTERIS – International Conference on ENTERprise Information Systems / ProjMAN – International Conference on Project MANagement / HCist – International Conference on Health and Social Care Information Systems and Technologies 2022.

[2] M. Ortu, S. Vacca, G. Destefanis, C. Conversano, Cryptocurrency ecosystems and social media environments: An empirical analysis through hawkes' models and natural language processing, Machine Learning with Applications 7 (2022). doi:https://doi.org/10.1016/j.mlwa.2021.100229.

[3] M. Chinea-Rios, I. Borrego-Obrador, M. Franco-Salvador, F. Rangel, P. Rosso, Profiling Cryptocurrency Influencers with Few shot Learning at PAN 2023, in: CLEF 2022 Labs and Workshops, Notebook Papers, 2023.

[4] J. Bevendorff, I. Borrego-Obrador, M. Chinea-Ríos, M. Franco-Salvador, M. Fröbe, A. Heini, K. Kredens, M. Mayerl, P. Pęzik, M. Potthast, F. Rangel, P. Rosso, E. Stamatatos, B. Stein, M. Wiegmann, M. Wolska, , E. Zangerle, Overview of PAN 2023: Authorship Verification, Multi-Author Writing Style Analysis, Profiling Cryptocurrency Influencers, and Trigger Detection, in: A. Arampatzis, E. Kanoulas, T. Tsikrika, A. G. Stefanos Vrochidis, D. Li, M. Aliannejadi, M. Vlachos, G. Faggioli, N. Ferro (Eds.), Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the Fourteenth International Conference of the CLEF Association (CLEF 2023), Lecture Notes in Computer Science, Springer, 2023.

[5] F. Rangel, P. Rosso, B. Verhoeven, W. Daelemans, M. Potthast, B. Stein, Overview of the 4th Author Profiling Task at PAN 2016: Cross-Genre Evaluations, in: K. Balog, L. Cappellato, N. Ferro, C. Macdonald (Eds.), Working Notes Papers of the CLEF 2016 Evaluation Labs, volume 1609 of *Lecture Notes in Computer Science*, 2016.

[6] F. Rangel, P. Rosso, M. Potthast, B. Stein, Overview of the 5th Author Profiling Task at PAN 2017: Gender and Language Variety Identification in Twitter, in: L. Cappellato, N. Ferro, L. Goeuriot, T. Mandl (Eds.), Working Notes Papers of the CLEF 2017 Evaluation Labs, volume 1866 of *CEUR Workshop Proceedings*, 2017.

[7] F. Rangel, M. Montes-y-Gómez, M. Potthast, B. Stein, Overview of the 6th Author Profiling Task at PAN 2018: Cross-domain Authorship Attribution and Style Change Detection, in: L. Cappellato, N. Ferro, J.-Y. Nie, L. Soulier (Eds.), CLEF 2018 Evaluation Labs and

Workshop – Working Notes Papers, 10-14 September, Avignon, France, CEUR-WS.org, 2018.
[8] F. Rangel, P. Rosso, Overview of the 7th Author Profiling Task at PAN 2019: Bots and Gender Profiling, in: L. Cappellato, N. Ferro, D. Losada, H. Müller (Eds.), CLEF 2019 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2019.
[9] F. Rangel, A. Giachanou, B. Ghanem, P. Rosso, Overview of the 8th Author Profiling Task at PAN 2020: Profiling Fake News Spreaders on Twitter, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névéol (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020.
[10] F. Rangel, P. Rosso, G. L. D. L. P. Sarracén, E. Fersini, B. Chulvi, Profiling Hate Speech Spreaders on Twitter Task at PAN 2021, in: G. Faggioli, N. Ferro, A. Joly, M. Maistro, F. Piroi (Eds.), CLEF 2021 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2021.
[11] R. Ortega-Bueno, B. Chulvi, F. Rangel, P. Rosso, E. Fersini, Profiling Irony and Stereotype Spreaders on Twitter (IROSTEREO). Overview for PAN at CLEF 2022., in: G. Faggioli, N. Ferro, A. Hanbury, M. Potthast (Eds.), CLEF 2022 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2022.
[12] A. Basile, G. Dwyer, M. Medvedeva, J. Rawee, H. Haagsma, M. Nissim, N-GrAM: New Groningen Author-profiling Model—Notebook for PAN at CLEF 2017, in: L. Cappellato, N. Ferro, L. Goeuriot, T. Mandl (Eds.), CLEF 2017 Evaluation Labs and Workshop – Working Notes Papers, 11-14 September, Dublin, Ireland, CEUR-WS.org, 2017.
[13] J. Pizarro, Using N-grams to detect Bots on Twitter, in: L. Cappellato, N. Ferro, D. E. Losada, H. Müller (Eds.), Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum, Lugano, Switzerland, September 9-12, 2019, volume 2380 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2019.
[14] M. Martinc, I. Škrjanec, K. Zupan, S. Pollak, PAN 2017: Author Profiling - Gender and Language Variety Prediction—Notebook for PAN at CLEF 2017, in: L. Cappellato, N. Ferro, L. Goeuriot, T. Mandl (Eds.), CLEF 2017 Evaluation Labs and Workshop – Working Notes Papers, 11-14 September, Dublin, Ireland, CEUR-WS.org, 2017.
[15] F. Johansson, Supervised Classification of Twitter Accounts Based on Textual Content of Tweets, in: L. Cappellato, N. Ferro, D. E. Losada, H. Müller (Eds.), Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum, Lugano, Switzerland, September 9-12, 2019, volume 2380 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2019.
[16] H. Giglou, J. Razmara, M. Rahgouy, M. Sanaei, LSACoNet: A Combination of Lexical and Conceptual Features for Analysis of Fake News Spreaders on Twitter—Notebook for PAN at CLEF 2020, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névéol (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020.
[17] R. Labadie-Tamayo, D. Castro-Castro, R. Ortega-Bueno, Fusing Stylistic Features with Deep-learning Methods for Profiling Fake News Spreader—Notebook for PAN at CLEF 2020, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névéol (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020.
[18] M. Siino, E. D. Nuovo, I. Tinnirello, M. L. Cascia, Detection of Hate Speech Spreaders using Convolutional Neural Networks, in: G. Faggioli, N. Ferro, A. Joly, M. Maistro, F. Piroi (Eds.), CLEF 2021 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2021.
[19] H. Cao, Z. Han, Z. Mo, Z. Li, Z. Xiao, Z. Li, L. Kong, A Multi-Model Voting Ensemble

Classifier based on BERT for Profiling Irony and Stereotype Spreaders on Twitter, in: G. Faggioli, N. Ferro, A. Hanbury, M. Potthast (Eds.), CLEF 2022 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2022.
[20] R. Labadie Tamayo, D. Castro Castro, R. Ortega Bueno, Deep Modeling of Latent Representations for Twitter Profiles on Hate Speech Spreaders Identification Task—Notebook for PAN at CLEF 2021, in: G. Faggioli, N. Ferro, A. Joly, M. Maistro, F. Piroi (Eds.), CLEF 2021 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2021.
[21] C. Bagdon, Profiling Spreaders of Hate Speech with N-grams and RoBERTa—Notebook for PAN at CLEF 2021, in: G. Faggioli, N. Ferro, A. Joly, M. Maistro, F. Piroi (Eds.), CLEF 2021 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2021.
[22] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, volume 33, Curran Associates, Inc., 2020, pp. 1877–1901. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.
[23] Q. Dong, L. Li, D. Dai, et al., A survey on in-context learning, 2023.
[24] G. Koch, R. Zemel, R. Salakhutdinov, et al., Siamese neural networks for one-shot image recognition, in: ICML deep learning workshop, volume 2, Lille, 2015.
[25] E. Hoffer, N. Ailon, Deep metric learning using triplet network, in: A. Feragen, M. Pelillo, M. Loog (Eds.), Similarity-Based Pattern Recognition, Springer International Publishing, Cham, 2015, pp. 84–92.
[26] T. Mueller, G. Pérez-Torró, M. Franco-Salvador, Few-shot learning with siamese networks and label tuning, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2022.
[27] M. Chinea-Rios, T. Müller, G. L. D. la Pena Sarrac'en, F. Rangel, M. Franco-Salvador, Zero and few-shot learning for author profiling, in: International Conference on Applications of Natural Language to Data Bases, 2022.
[28] F. Rangel, M. Chinea-Rios, M. Franco-Salvador, P. Rosso, PAN23 Profiling Cryptocurrency Influencers with Few-shot Learning, 2023. URL: https://doi.org/10.5281/zenodo.7701748. doi:10.5281/zenodo.7701748.
[29] M. Chinea-Rios, T. Müller, G. L. De la Peña Sarracén, F. Rangel, M. Franco-Salvador, Zero and few-shot learning for author profiling, in: P. Rosso, V. Basile, R. Martínez, E. Métais, F. Meziane (Eds.), Natural Language Processing and Information Systems, Springer International Publishing, Cham, 2022, pp. 333–344.
[30] C. Saedi, M. Dras, Siamese networks for large-scale author identification, Computer Speech & Language 70 (2021).
[31] A. Vaswani, N. M. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: NIPS, 2017.
[32] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, W. Chen, Lora: Low-rank adaptation of large language models, in: International Conference on Learning Representations, 2021.

[33] H. Su, J. Kasai, Y. Wang, et al., One embedder, any task: Instruction-finetuned text embeddings, arXiv preprint arXiv:2212.09741 (2022).
[34] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Y. Bengio, Y. LeCun (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. URL: http://arxiv.org/abs/1412.6980.
[35] T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: A next-generation hyperparameter optimization framework, in: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining, 2019.
[36] R. Hadsell, S. Chopra, Y. LeCun, Dimensionality reduction by learning an invariant mapping, in: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), volume 2, 2006. doi:10.1109/CVPR.2006.100.
[37] F. Schroff, D. Kalenichenko, J. Philbin, Facenet: A unified embedding for face recognition and clustering, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015.
[38] G. Hinton, N. Srivastava, K. Swersky, Lecture 6a overview of mini–batch gradient descent, Coursera Lecture slides https://class. coursera. org/neuralnets-2012-001/lecture,[Online] (2012).
[39] J. Howard, S. Ruder, Universal language model fine-tuning for text classification, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, 2018. doi:10.18653/v1/P18-1031.
[40] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, A. Rush, Transformers: State-of-the-art natural language processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, Online, 2020, pp. 38–45. URL: https://aclanthology.org/2020.emnlp-demos.6. doi:10.18653/v1/2020.emnlp-demos.6.
[41] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, et al., BLOOM: A 176b-parameter open-access multilingual language model, CoRR abs/2211.05100 (2022). doi:10.48550/arXiv.2211.05100.
[42] T. He, J. Liu, K. Cho, M. Ott, B. Liu, J. Glass, F. Peng, Analyzing the forgetting problem in the pretrain-finetuning of dialogue response models, 2021. arXiv:1910.07117.

A. Textual definitions of target classes

In Table 4 we present the definitions of the target classes for subtask 1 used in our metric learning approach. These definitions intend to capture information of types of Twitter accounts attending to the scale of their impact in social media.

Table 4
Definitions of target classes for subtask 1.

Class	Definition
mega	An influencer with at least 1 million followers on at least one social media platform, blog, podcast, or other Internet communication channel.
macro	A creator with more than 100,000 followers and at least a 3% engagement rate, often with a title such as celebrity, TV personality, athlete, or thought leader in their community.
micro	A social media user with between 1,000 and 100,000 followers, who has a larger social media presence than a normal person, but smaller than a celebrity.
nano	An everyday social media user with 100 to 10,000 followers, who shares typical content such as photos of family, friends, cat videos, and memes.
non-influencer	Someone who does not have the same level of authority or following within a niche and does not have the same ability to influence their audience.