

mCPT at SemEval-2023 Task 3: Multilingual Label-Aware Contrastive Pre-Training of Transformers for Few- and Zero-shot Framing Detection

Markus Reiter-Haas^{*†}, Alexander Ertl^{*}, Kevin Innerebner, Elisabeth Lex
Graz University of Technology, Institute of Interactive Systems and Data Science
Sandgasse 36/III, 8010, Graz, Austria
reiter-haas@tugraz.at, ertl@student.tugraz.at
innerebner@student.tugraz.at, elisabeth.lex@tugraz.at

Abstract

This paper presents the winning system for the zero-shot Spanish framing detection task, which also achieves competitive places in eight additional languages. The challenge of the framing detection task lies in identifying a set of 14 frames when only a few or zero samples are available, i.e., a multilingual multi-label few- or zero-shot setting. Our developed solution employs a pre-training procedure based on multilingual Transformers using a label-aware contrastive loss function. In addition to describing the system, we perform an embedding space analysis and ablation study to demonstrate how our pre-training procedure supports framing detection to advance computational framing analysis.

1 Introduction

Approaches for computational framing detection are diverse (Ali and Hassan, 2022), as the framing concept itself is often just casually defined (Entman, 1993). Consequently, framing detection is challenging on its own, but also suffers from a lack of sufficient data (Kwak et al., 2020), especially in multilingual settings. The SemEval 2023 Task 3 Subtask 2 (Piskorski et al., 2023) aims at predicting 14 distinct media frames (Boydston et al., 2013) present within news articles in 9 languages. Due to label imbalances, as a result of the high dimension of the label space compared to the number of samples, traditional paradigms, e.g., per-label binary classification, do not apply well to the given setting without adaptations (Tarekegn et al., 2021).

We introduce *mCPT*, the label-aware Contrastive Pre-training of Transformers based on a multilingual encoder model (original team name on the leaderboard¹: *PolarIce*). We exploit two features

^{*}equal contribution

[†]corresponding author

¹<https://propaganda.math.unipd.it/semEval2023task3/SemEval2023testleaderboard.html>

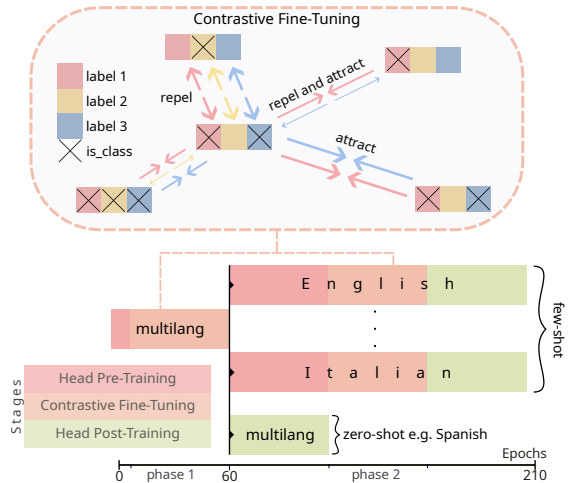


Figure 1: Our system performs label-aware contrastive fine-tuning (top). Embeddings of samples with similar labels are attracted, while they are repelled for dissimilar labels. The two-phase procedure (bottom) interleaves contrastive fine-tuning in both the multilingual and target language training.

of the task: (i) multi-label information and (ii) multilingual data for pre-training.

First, we leverage the label information by adopting a contrastive loss function, i.e., HeroCon-Loss (Zheng et al., 2022), for natural language processing that optimizes the embedding space with respect to the similarities of the label space. Therefore, samples with more similar labels occupy similar regions in the embedding space, whereas mostly dissimilar samples regarding their shared labels are pushed apart (refer to Figures 1 top and 2a).

Second, we design a custom two-phase procedure with multiple stages for multilingual training to maximize the available data (see Figure 1). In phase one, we train on all languages, while in phase two, we further fine-tune the model on the target language if such data exist i.e. few-shot setting, or continue training on all languages if not i.e. zero-shot setting.

Our system performs competitively (top 10) on

all six few-shot (i.e., English, German, French, Italian, Polish, and Russian) and three zero-shot (i.e., Spanish, Greek, and Georgian) settings, beating the baselines on all languages. On Spanish, which is the only zero-shot language with a common language family and alphabet as the training languages, our system is the winning contribution with a Micro-F1 of 0.571 (compared to 0.120 of the baseline). Therefore, we argue that our system generalizes well to unseen data, even when no training data is available in similar target languages.

In sum, our contribution is three-fold²:

- C1 We adopt a multi-label contrastive loss function for natural language processing to optimize the embeddings of textual data.
- C2 We describe a two-phase multi-stage training procedure for multilingual scenarios with limited data, i.e., few- and zero-shot predictions.
- C3 We demonstrate the effectiveness of our winning system for framing detection supported by embedding and ablation studies.

2 Related Work

Framing Detection. According to Entman (1993), to frame is to select and emphasize some aspects of reality to encourage particular interpretations. That is, messages centered around a common topic may draw the receiver’s attention to distinct features, thus suggesting different courses of action, causal interpretations, etc. As such, computational framing detection requires natural language processing (NLP) methods that capture nuances of *how* content is presented rather than just *what* topic is present. Therein, studies focus on detecting vastly different conceptualizations of framing, such as blame frames (Shurafa et al., 2020), war frames (Wicke and Bolognesi, 2020), moral frames (Reiter-Haas et al., 2021), or media frames (Boydston et al., 2014; Kwak et al., 2020).

Regarding media frames, Boydston et al. (2013) identified a set of relevant frames that formed the basis for the media frame corpus (Card et al., 2015). Within this supervised frame detection scenario, Liu et al. (2019) indicate that Transformer-based approaches vastly outperform approaches using less powerful architectures such as LSTMs. As such, we also employ Transformer models with label-aware contrastive pre-training.

²Our code and model are publicly available at: <https://github.com/socialcomplab/semEval23-mcpt>

Supervised Contrastive Learning. Contrastive learning, originally mainly used in computer vision settings (e.g., Chopra et al., 2005), has recently found increased attention in the NLP research community due to its efficacy on tasks with limited amounts of data and its applicability to Transformer embeddings (e.g., Tunstall et al., 2022). The general concept of supervised contrastive learning (Khosla et al., 2020a) is that latent representations (or embeddings in NLP) of samples with the same labels should be close in embedding space, while samples with different labels should be further apart.

Su et al. (2022) and Zheng et al. (2022) have independently proposed contrastive learning methods for multi-label settings that weight similarities of samples by the similarity of their label vectors, i.e. hidden representations of samples with similar label vectors should be more similar than hidden representations of samples with less similar label vectors. Su et al. (2022) weight a Euclidean distance-based measure of embeddings by a normalized dot product of the label vectors. HeroCon (Zheng et al., 2022) generalizes supervised contrastive loss (Khosla et al., 2020b) and beats previous state-of-the-art contrastive learning paradigms in multi-label settings on multiple image data sets. Tunstall et al. (2022) introduce SETFIT, an algorithm for the data-efficient fine-tuning of sentence embeddings, primarily on binary labels. SETFIT first fine-tunes sentence embeddings in a contrastive manner before training a classification head.

We combine the idea of the contrastive pre-training stage from SETFIT and adopt HeroCon for NLP loss to improve performance on multi-label datasets.

3 Methods

At the core of our system lies a multilingual Transformer model with dense neural layers comprising the head. Contrastive fine-tuning is performed as part of a multi-stage training procedure.

3.1 Contrastive Fine-Tuning

Our contrastive fine-tuning objective (C1, Figure 1 top) is centered around the idea that embeddings of samples with similar labels should be close while embeddings of samples with very distinct label vectors should be distant. Following Zheng et al. (2022) for every batch and every class, we com-

Table 1: **Test set results on the official leaderboard** on Subtask 2, first few-shot (top) then zero-shot (bottom). The results are sorted by Micro-F1 of *mCPT*, i.e., our system performance on the target metric. Our system outperforms the *Base* on all languages, both on Micro-F1 and Macro-F1, with the majority of improvements being very significant[†]. Similarly, *mCPT* performs better than *SETFIT* on all Latin-based languages. Our winning contribution to Spanish is also significantly better than SETFIT, as well as the averaged Micro and Macro-F1 scores.

Language	# Samples Train/Dev/Test	Micro-F1			Macro-F1			Position	
		mCPT	SETFIT	Base	mCPT	SETFIT	Base	#	Teams
German (\mathcal{G} , L)	132 / 45 / 50	.622*	.549	.487	.564*	.492	.418	6	/19
Polish (\mathcal{S} , L)	145 / 49 / 47	.597	.584	.594	.555	.542	.532	9	/19
Italian (\mathcal{R} , L)	227 / 76 / 61	.584*	.502	.486	.469**	.371	.372	5	/19
English (\mathcal{G} , L)	433 / 83 / 54	.535*	.469*	.350	.482*	.409*	.274	5	/23
French (\mathcal{R} , L)	158 / 53 / 50	.469*	.463*	.329	.429*	.419*	.276	9	/19
Russian (\mathcal{S})	143 / 48 / 72	.409*	.421*	.230	.367**	.258	.218	5	/18
<i>Spanish</i> (\mathcal{R} , L)	– / – / 30	.571**	.418*	.120	.455**	.305*	.095	1	/17
<i>Greek</i>	– / – / 64	.516*	.427	.345	.410*	.338*	.057	7	/16
<i>Georgian</i>	– / – / 29	.400*	.404*	.260	.291	.384*	.251	9	/16
Summary	1238 / 354 / 457	.523**	.471*	.356	.447**	.391*	.277	6.2	/18.4

^{mCPT} Our system; ^{SF} SETFIT Transformer model; ^{Base} Challenge Baseline (n-grams count + SVC);

[†] We assume a normal approximation interval on a binomial distribution 99.5% confidence level ($z = 2.81$) concerning the number of labels as proxy. We will update the table with a statistical test on the samples once the test labels are released.

* Significant improvement outside the confidence interval compared to Base; ** also over SETFIT; ¹ Winner;

Bold Best performance; ^{Italic} Zero-Shot Language; ^G Germanic; ^S Slavic; ^R Romance; ^L Latin alphabet;

pute the similarity between positive samples, i.e., samples that are of that class, and all others. As such, samples may both repel and attract each other within different classes yet do neither if they are both negative.

Our loss function is a linear combination of two terms: A binary cross entropy term \mathcal{L}_{BCE} that jointly optimizes the head and body in the contrastive fine-tuning stage and a contrastive term \mathcal{L}_{CON} :

$$\mathcal{L} = \mathcal{L}_{BCE} + \alpha \mathcal{L}_{CON} \quad (1)$$

where α is a weighting hyperparameter. The contrastive loss is given by:

$$\mathcal{L}_{CON} = \frac{1}{|C|} \sum_{c \in C} -\mathbf{E}_{X_i, X_j \in \mathcal{P}(c)} \left[\log \frac{\sigma_{ij} f(X_i, X_j)}{\delta_{ij}} \right] \quad (2)$$

where C is the set of all classes (e.g. *Economic*), $\mathcal{P}(c)$ is the set of all positive samples i.e., all embeddings X_i that are of class c , and $f(\cdot, \cdot)$ is the cosine similarity measure between embeddings. The loss is normalized by:

$$\delta_{ij} = \frac{\sigma_{ij} f(X_i, X_j) + \sum_{X_k \in \mathcal{N}(c)} \gamma_{ik} f(X_i, X_k)}{|\mathcal{N}(c)| + 1} \quad (3)$$

where $\mathcal{N}(c)$ is the set of all negative samples and σ_{ij} and γ_{ik} are given by:

$$\sigma_{ij} = 1 - d(Y_i, Y_j) / |C|, \quad \gamma_{ik} = d(Y_i, Y_k)$$

respectively where d describes the Hamming distance between label vectors Y_i .

3.2 Training Procedure

We follow a two-phase training procedure illustrated in Figure 1 (bottom), focusing on optimally utilizing the available data, which in our case are the six languages of subtask 2 ($C2$). Inspired by Tunstall et al. (2022), we first optimize the embedding space of a (multilingual) Transformer model. Herein, our approach makes the assumption that the embeddings, regardless of language, possess mutual information given similar labels (Zheng et al., 2022). While the embedding space may be improved in this manner, we fine-tune on the target language to further improve the performance.

Phase one consists of two separate stages: *head pre-training* and *contrastive fine-tuning*. For phase two, the head is re-used for zero-shot settings, while discarded and randomly re-initialized for few-shot settings. In the former case, we conduct the *post-training* stage on all languages, while in the latter it is essential to *pre-train* the head on the target language before proceeding to the *contrastive fine-tuning* and *head post-training* stages. Both the *head pre-training* and *head post-training* stages only compute the binary cross entropy term \mathcal{L}_{BCE} ,




Figure 2: *Effect of the loss function on the embedding space.* (a) Shows the repositioning of randomly generated samples (both embeddings and labels) in two-dimensional space. The contrast loss function on its own increases the cosine similarity of latent representations with similar labels, while decreasing the similarity of representations with different labels. Note the positioning of trained embeddings with identical labels along lines drawn from the origin. (b) Without contrastive pre-training, the pairs of embeddings in the English dev set are similar regardless of their label distance. (c) After 50 epochs, the embedding cosine similarity reflects the Hamming distance of the labels.

while simultaneously leaving the body unchanged, i.e., frozen. While the stages are identical, the rationale for each of them is very distinct: As the head is randomly initialized, we first *pre-train* it to avoid high gradients in the subsequent *contrastive fine-tuning* stage. In contrast, *post-training* allows the head to better fit the fine-tuned embeddings.

4 Experiments

We now present the results of our *mCPT* system, supported by embedding space and ablation studies (C3). We base *mCPT* on the multilingual³ sentence Transformer model *paraphrase-multilingual-MiniLM-L12-v2* (Reimers and Gurevych, 2019) and demonstrate that competitive results can be achieved with a relatively small amount of parameters, i.e., 117M parameters (Wang et al., 2020), given a training method tailored to the task. The model was chosen for its sentence embedding performance on multiple languages and its small size relative to similar state-of-the-art multilingual Transformer models.

Our model architecture comprises mean-pooling, no normalization of embeddings, a dense head with one hidden layer of size 256, and a dropout of 0.5. We train the model with separate learning rates for the classification head ($1e-3$) and the body ($2e-5$), a weighting parameter α of 0.01, a batch size of

³The base model was trained on 50+ languages including all nine of the shared task, thus being suitable for the problem.

26 for 10 and 50 epochs for *head pre-training* and *contrastive fine-tuning* respectively in phase one (more details in Appendix A).

Baseline Models. We compare the performance of our system against two baselines. First, we consider the results of the official baseline *Base* (i.e., n-grams and support vector classification; Piskorski et al., 2023). Second, we compare against SETFIT (Tunstall et al., 2022) with the same base encoder as ours, on the post-challenge test set (details in Appendix B).

4.1 Main Results

mCPT performs better on Latin alphabets (marked by *L*) in both few- and zero-shot settings, and improves upon the two baselines (as presented in Table 1). In Slavic languages (*S*), Polish is second-best in terms of Micro-F1 (0.597), but only slightly outperforms the baseline, whereas the improvement on Russian is very significant but only achieves the second-lowest score of 0.409. In comparison, both Germanic languages (*G*), German and English perform well, where we also have our highest overall Micro-F1 of 0.622 for German, but also a high baseline of 0.487. Although we find significant improvement on Greek and Georgian (which are zero-shot languages that do not share a major branch with any other language) over the baseline, both perform poorly in terms of Micro-F1 (i.e., Georgian having the lowest Micro-F1 of

0.400). Hence, we suspect that not enough information from the other languages could be transferred. Finally, in the Romance languages (\mathcal{R}), Italian and Spanish perform well, while French with a Micro-F1 of 0.516 performs lower than Greek.

The performance of Spanish is especially noteworthy, as it is the only zero-shot language that shares an alphabet as well as a language family with the training data languages. Therefore, we argue that our winning performance on Spanish (Micro-F1 of 0.571 compared to 0.12 of the baseline) stems from the fact that the knowledge was successfully transferred from the other languages to the zero-shot setting. This is further supported by *SETFIT*, which improves upon the baseline but shows lower performance on Latin-based languages.

4.2 Embedding Space Analysis

Figure 2 demonstrates how our contrastive training procedure optimizes the embeddings of the Transformer body. Figure 2a exemplarily shows the repositioning of samples due to the loss function in two-dimensional space. Observe how the trained labels (opaque) align, with the blue-yellow label between the two blue-only labels and two yellow-only labels, while simultaneously pushing the red labels to the side (more detailed analysis in Appendix C). For the analysis of the high-dimensional embeddings (384) and label spaces (14) on the dev set, we use boxen plots concerning embedding cosine similarity for all pairwise samples within a given Hamming distance.

Regarding the pre-trained base model on English (Figure 2b), we find a suboptimal correlation with $R^2 = 0.005$ and $\beta = -0.009$. In comparison, after contrastively training the model, the correlation becomes much more pronounced, i.e., $R^2 = 0.241$ and $\beta = -0.079$ for English as shown in Figure 2c. Furthermore, the spread of pairwise embedding similarity distribution increases, as a greater amount of samples become dissimilar to each other, especially for higher label distance. Thus, we conclude that our system leads to higher utilization of the available embedding space, which in turn boosts performance. Appendix D contains the remaining languages.

Finally, we want to emphasize that our data set has no perfectly dissimilar label vector pairs, which would make hard negative mining approaches (Gao et al., 2021) infeasible, e.g., for using plain con-

Table 2: Ablation study (top) and the proposed contrastive sampling extension (bottom) on the dev set. In general, we observe that *mCPT* performs best with all components, i.e., pre-training (PT), contrastive loss (\mathcal{L}_{CON}), and end-to-end training (E2E), enabled. Contrast sampling (CS) suggest further improvements.

Model	en	it	ru	fr	ge	po
mCPT	.682	.585	.520	.570	.561	.636
- PT	.681	.545	.475	.563	.583	.616
- \mathcal{L}_{CON}	.657	.521	.436	.524	.570	.645
- E2E	.629	.519	.500	.535	.586	.633
mCPT+CS	.688	.590	.519	.575	.591	.638

trastive loss (Chopra et al., 2005).

4.3 Ablation Study and Extension

Table 2 indicates the effectiveness of our combined training approach *mCPT*. From *mCPT*, we remove components iteratively, first removing the multilingual pre-training phase, then the contrastive term (see Equation 2), and finally, end-to-end training leaving only a trained classification head with no embedding fine-tuning. Comparing the results of the ablation study with those of Table 1 it is interesting to note that our approach works best on languages with lower scores. A hypothesis is that the out-of-the-box Transformer embeddings already fit the data well and that *mCPT* is not able to improve upon the already strong baseline. Finally, we find that adding a contrast sampling extension could further improve the results (see Appendix E).

5 Conclusion

In this paper, we describe our system (*mCPT*) for the framing detection shared task (Piskorski et al., 2023). We introduce an approach based on a label-aware contrastive loss and training procedure for Transformers to deal with the challenges of multilingual multi-label prediction with few or even zero samples. The generalization ability of our system is demonstrated by providing the winning contribution for the Spanish framing detection subtask where no training samples were available⁴. Hence, we believe that our system is a notable advancement for computational framing research.

⁴We refer to Appendices F and G for discussions on limitations and ethical considerations, respectively.

Acknowledgements

We would like to thank the anonymous reviewers for taking the necessary time and effort to review the manuscript and appreciate their universal positive feedback.

References

- Mohammad Ali and Naemul Hassan. 2022. [A survey of computational framing analysis approaches](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9335–9348, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Amber E Boydston, Dallas Card, Justin Gross, Paul Resnick, and Noah A Smith. 2014. Tracking the development of media frames within and across policy issues.
- Amber E Boydston, Justin H Gross, Philip Resnik, and Noah A Smith. 2013. Identifying media frames and frame dynamics within and across policy issues. In *New Directions in Analyzing Text as Data Workshop, London*.
- Dallas Card, Amber Boydston, Justin H Gross, Philip Resnik, and Noah A Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. [Learning a similarity metric discriminatively, with application to face verification](#). In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546. IEEE.
- Robert M Entman. 1993. Framing: Towards clarification of a fractured paradigm. *McQuail's reader in mass communication theory*, 390:397.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [Simcse: Simple contrastive learning of sentence embeddings](#). *arXiv preprint arXiv:2104.08821*.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020a. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020b. [Supervised contrastive learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc.
- Haewoon Kwak, Jisun An, and Yong-Yeol Ahn. 2020. [A systematic media frame analysis of 1.5 million new york times articles from 2000 to 2017](#). In *12th ACM Conference on Web Science*, pages 305–314.
- Qisheng Liao, Meiting Lai, and Preslav Nakov. 2023. Marseclipse at semeval-2023 task 3: Multi-lingual and multi-label framing detection with contrastive learning. In *Proceedings of the 17th International Workshop on Semantic Evaluation, SemEval 2023*, Toronto, Canada.
- Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. [Detecting frames in news headlines and its application to analyzing news framing trends surrounding us gun violence](#). In *Proceedings of the 23rd conference on computational natural language learning (CoNLL)*, pages 504–514.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. Semeval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multi-lingual setup. In *Proceedings of the 17th International Workshop on Semantic Evaluation, SemEval 2023*, Toronto, Canada.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Markus Reiter-Haas, Simone Kopeinik, and Elisabeth Lex. 2021. [Studying moral-based differences in the framing of political tweets](#). In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, pages 1085–1089.
- Chereen Shurafa, Kareem Darwish, and Wajdi Zaghuan. 2020. [Political framing: Us covid19 blame game](#). In *International Conference on Social Informatics*, pages 333–351. Springer.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tiejun Liu. 2020. MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33:16857–16867.
- Xi'ao Su, Ran Wang, and Xinyu Dai. 2022. [Contrastive learning-enhanced nearest neighbor mechanism for multi-label text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 672–679, Dublin, Ireland. Association for Computational Linguistics.
- Adane Nega Tarekegn, Mario Giacobini, and Krzysztof Michalak. 2021. [A review of methods for imbalanced multi-label classification](#). *Pattern Recognition*, 118:107965.

Lewis Tunstall, Nils Reimers, Unso Eun Seo Jo, Luke Bates, Daniel Korat, Moshe Wasserblat, and Oren Pereg. 2022. [Efficient few-shot learning without prompts](#). *arXiv preprint arXiv:2209.11055*.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788.

Philipp Wicke and Marianna M Bolognesi. 2020. [Framing covid-19: How we conceptualize and discuss the pandemic on twitter](#). *PLoS one*, 15(9):e0240010.

Ben Wu, Olesya Razuvayevskaya, Freddy Heppell, João A. Leite, Carolina Scarton, Kalina Bontcheva, and Xingyi Song. 2023. Sheffieldveraai at semeval-2023 task 3: Mono and multilingual approaches for news genre, topic and persuasion technique classification. In *Proceedings of the 17th International Workshop on Semantic Evaluation, SemEval 2023*, Toronto, Canada.

Lecheng Zheng, Jinjun Xiong, Yada Zhu, and Jingrui He. 2022. [Contrastive learning with complex heterogeneity](#).

A Training Environment

We performed the main experiments on the Kaggle platform (www.kaggle.com) with the P100 graphics card. We chose a free platform for the computation to demonstrate that our system is tailored towards the task at hand and is accessible for everybody, rather than relying on large amounts of computational resources. We empirically selected the hyperparameters to fit the platform. Herein, we chose a batch size of 26 which optimally utilizes the available GPU memory. The multilingual pre-training takes approximately 1.5 hours, while the language-specific fine-tuning takes 1 hour each.

B SETFIT Parameters

We choose *SETFIT* as it is similar in concept to our system, i.e., contrastive learning for Transformers, but not aligned with the shared task, i.e., does not explicitly consider multi-label problems. Hence, the comparison demonstrates how our system is an improvement over established approaches in this setting and emphasizes that the adaptations of the contrastive loss and training procedure are indeed beneficial. We report the results of *SETFIT* on the post-challenge leaderboard without further adaption after the initial submissions for fair comparisons on the test set.

We mimic the parameters setting where applicable while preserving the standard training procedure to maximize comparability. We first contrastively train the body for 10 epochs before training the full model end-to-end for 50 epochs with a batch size of 26 with learning rates of $1e - 3$ and $2e - 5$ for the head and the body respectively. The body-then-end-to-end procedure was suggested by the usage guide. The training runtime is approximately 10 hours on the Kaggle platform (which again was chosen for a fair comparison). Initially, we experimented with SETFIT in the challenge period but decided to submit our presented system instead.

C Repositioning of Samples

In Figure 2a, we show the effect of the loss function and how the label and embedding space are intertwined. Specifically, we demonstrate how randomly generated embeddings in two-dimensional space with three-dimensional label vectors shift towards more optimal positions after applying the contrastive loss function. Accordingly, the initially random positions of embeddings with equivalent labels end up on straight lines drawn from the origin. The observed effect is a direct consequence of similar label vectors attracting and opposite labels repelling each other. Moreover, a partial label similarity with two distinct groups ends between those groups, as a result of both forces being active. For instance, consider the line from top left to bottom right: blue-only labels become attracted, and repel yellow-only labels, while the sample with blue and yellow lies in between. Due to the resulting positioning, the embedding space becomes disentangled leading to an increase in linear separability, which benefits classifiers such as our differentiable head.

D Embeddings and Labels Correlation

Here, we present the remaining languages for the embedding space analysis in Figure 3. When considering the plots before contrastive training is performed, the correlations between embedding spaces and label spaces are very weak, as well as most embeddings being very similar regardless of their label distance. Hence, virtually no pair of samples are opposite to each other, i.e., has a noteworthy negative cosine similarity. Conversely, all pairs are similar to a certain extent. We argue that in these given embedding spaces, it is challenging

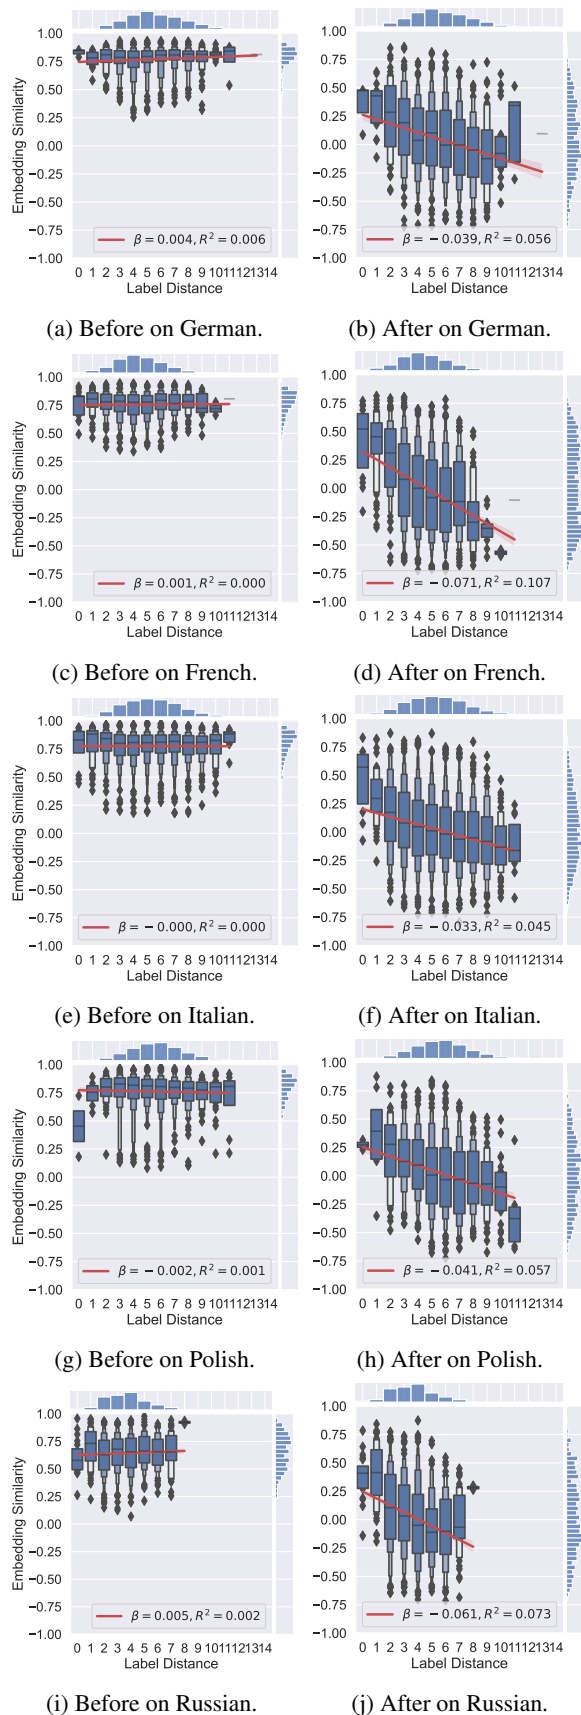


Figure 3: *Effect of the loss function on the embedding space.* Evaluated on various language dev sets.

for a classifier to learn a discriminative function. Moreover, due to the suboptimal positioning of embeddings, a substantial amount of the embedding space does not contribute to the prediction performance, thus wasting the model’s potential expressiveness. Hence, the interpretation from Section 4.2 can be directly applied to the five other languages, as the effect is the same (although differently pronounced). For instance, models of German and Russian even have a slight upward slope when applied without contrastive training. Hence, the negative slope and regression fit increases with contrastive training, as intended and expected.

E Contrast Sampling Extension.

Adding a custom *contrast sampler* which ensures that at least one sample from every class is present per batch further improves consistency as well as performance. Due to the nature of the contrastive objective, it is imperative that every batch contain negative as well as positive pairs of samples for every class. This is not guaranteed by sampling randomly, especially if the label distribution is imbalanced. As illustrated by Table 2 it outperforms pure *mCPT* in five out of six languages while coming in second by a small margin in Russian. We attribute this largely to the variance introduced by using a relatively small batch size of 26 compared to the number of labels (14). This variance may lead to undesirable gradient updates in some iterations when batches contain label distributions that are not representative.

F Limitations

We recognize three main limitations of our work, which are distinct in their aspect.

First, the *performance limitation*; while our system has competitive results across the board, it only performs best in one of the nine languages on the leaderboard. In comparison, team *MarsEclipse* (Liao et al., 2023), which also focused on the framing detection subtask, wins all six few-shot languages and performs second on two of three (i.e., Greek and Georgian) zero-shot languages. They only perform worse at Spanish (6th), which is opposite to our placement. Team *SheffieldVeraAI* (Wu et al., 2023), who also participated in the other two subtasks regarding news genre and persuasion technique detection, perform well across the board and wins the Greek and Georgian framing detection tasks. Hence, our system occupies the niche

of zero-shot prediction when trained with similar languages (i.e, in our case Latin-based).

Second the *technical limitation*, our system was trained using a small multilingual model as we aimed towards adapting Transformer pre-training for the multi-label challenge in particular rather than achieving the highest performance with computationally expensive training. However, as a consequence, we do not know how well our system scales to bigger models, such as MPNet (Song et al., 2020), and plan to address this limitation in future work.

Third the *task setting limitation*, we want to emphasize a potential limitation resulting from the shared task setting. Ali and Hassan (2022) argue that the specified labels in the media frame corpus (Card et al., 2015) revolve around topics (i.e., the *what*) rather frames (i.e., the *how*). As the same labels were adopted for the shared task, the conceptualizations of frames are expected to be similar to a certain extent. They thus would also affect the resulting models and predictions.

G Ethics Statement

We want to discuss three ethical considerations of our system. First, our system is based on pre-trained Transformers, which inherit *biases* from their training data. For the shared task, these biases are negligible, but are a concern for real-world applications. The second consideration relates to *fairness* concerns. The performance varies strongly between languages, with more researched languages typically resulting in better performance. We, thus, embrace the multilingual setting of the shared task with one-third zero-shot languages, but similarly achieved better performance in Latin-based languages. Third, our system leads to better detection of media frames, which is an important research direction. However, the system could in theory also be used in a disputed or even malicious manner, e.g., for *reframing* political statements. Hence, we do not advise specific applications of our system besides better framing detection.