

IITK at SemEval-2024 Task 4: Hierarchical Embeddings for Detection of Persuasion Techniques in Memes

Shreenaga Chikoti Shrey Mehta Ashutosh Modi
Indian Institute of Technology Kanpur (IIT Kanpur)
chikoti20@iitk.ac.in
ashutoshm@cse.iitk.ac.in

Abstract

Mememes are one of the most popular types of content used in an online disinformation campaign. They are primarily effective on social media platforms since they can easily reach many users. Mememes in a disinformation campaign achieve their goal of influencing the users through several rhetorical and psychological techniques, such as causal oversimplification, name-calling, and smear. The SemEval 2024 Task 4 *Multilingual Detection of Persuasion Technique in Memes* on identifying such techniques in the mememes is divided across three sub-tasks: (1) Hierarchical multi-label classification using only textual content of the mememe, (2) Hierarchical multi-label classification using both, textual and visual content of the mememe and (3) Binary classification of whether the mememe contains a persuasion technique or not using its textual and visual content. This paper proposes an ensemble of Class Definition Prediction (CDP) and hyperbolic embeddings-based approaches for this task. We enhance mememe classification accuracy and comprehensiveness by integrating HypEmo’s hierarchical label embeddings (Chen et al., 2023) and a multi-task learning framework for emotion prediction. We achieve a hierarchical F1-score of 0.60, 0.67, and 0.48 on the respective sub-tasks.

1 Introduction

Mememes are popular among people of all age groups today through different social media platforms (Keswani et al., 2020; Singh et al., 2020). These mememes help people know about the trends around them and can influence their decisions. Mememes are one of the popular modes for spreading disinformation among people (examples in Figure 1), as studies have suggested that people tend to believe what they see frequently in such mememes spread over the internet (Moravec et al., 2018). As evidenced by research (Shu et al., 2017) during the 2016 US Presidential campaign, nefarious actors, including



Figure 1: Sample set of mememes showing the multi-modal setting

bots, cyborgs, and trolls, leveraged mememes to evoke emotional reactions and propagate misleading narratives (Guo et al., 2020).

In this respect, SemEval-2024 Task 4 (Dimitrov et al., 2024) focuses on predicting the persuasive technique (from the visual and textual content) used in a mememe across four different languages: English, Arabic, North Macedonian and Bulgarian. The task is divided into three sub-tasks: (1) Hierarchical multi-label classification using only textual content of the mememe, (2) Hierarchical multi-label classification using both textual and visual content of the mememe and (3) Binary classification of whether the mememe contains a persuasion technique or not using its textual and visual content. The training data is provided for each sub-task but only in English. Taxonomy of various persuasion techniques (Figure 2) and their respective definitions are provided.

To address sub-task 1, we employed a dual approach involving definition-based modeling for each class and hierarchical classification using hyperbolic embeddings, as proposed in Chen et al. (2023). Based on hyperbolic embeddings, the

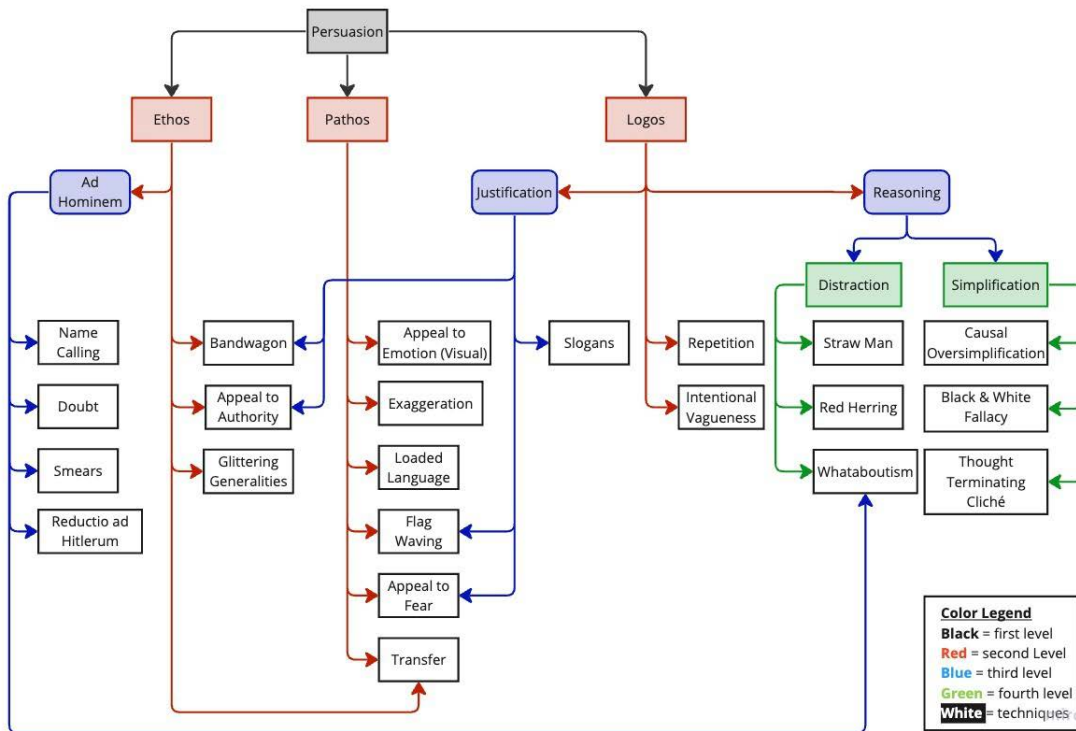


Figure 2: Taxonomy of persuasion techniques for sub-task 2

method facilitates a nuanced classification of persuasion techniques by leveraging hierarchical structures. The incorporation of definition-based modeling allows for a dataset-agnostic approach, enhancing the precision of classification without reliance on hierarchical structures.

For sub-task 2, we augmented our methodology by integrating CLIP embeddings (Radford et al., 2021) to capture essential features from memes’ textual and visual components. This fusion of textual and visual information enables a more comprehensive analysis of meme content.

In addressing sub-task 3, we adopted an ensemble approach, leveraging transfer learning from both the DistilBERT (Sanh et al., 2019) and CLIP embeddings (Radford et al., 2021). This ensemble technique enhances the robustness and effectiveness of our classification system by amalgamating insights from both pre-trained models. We release the code via GitHub.¹

2 Background

The goal of propaganda is to enhance people’s mindsets (Singh et al., 2020), especially at the time of elections, where the trends in the media influ-

ence the votes of the people (Shu et al., 2017). Propaganda uses psychological and rhetorical techniques to serve its purpose. Such methods include using logical fallacies and appealing to the audience’s emotions. Logical fallacies are usually hard to spot since the argumentation, at first sight, might seem correct and objective. However, a careful analysis shows that the conclusion cannot be drawn from the premise without misusing logical rules (Gupta and Sharma, 2021). Another set of techniques uses emotional language to induce the audience to agree with the speaker only based on the emotional bond that is being created, provoking the suspension of any rational analysis of the argumentation (Szabo, 2020).

Corpora development has been instrumental in advancing deception detection methodologies. Rashkin et al. (2017) introduced the TSHP-17 corpus, providing document-level annotation across four classes: trusted, satire, hoax, and propaganda. However, their study on the classification task revealed limitations in the generalizability of n-gram-based approaches. Building on this, Barrón-Cedeno et al. (2019) contributed the QProp corpus, which specifically targeted propaganda detection, employing a binary classification scheme of propaganda versus non-propaganda. Similarly, Habernal

¹<https://github.com/Exploration-Lab/IITK-SemEval-2024-Task-4-Persuasion-Techniques>




Figure 3: The frequency Distribution of Labels in the training dataset

et al. (2018) developed a corpus annotated with fallacies, including *ad hominem* and *red herring*, directly relevant to propaganda techniques.

BERT-based variants have emerged as promising methodologies for classification tasks in tandem with corpus development. Yoosuf and Yang (2019) proposed a fine-tuning approach post-world-level classification using BERT, while Fadel et al. (2019) presented a pre-trained ensemble model integrating BiLSTM, BERT, and RNN components. Further extending the capabilities of BERT, Costa et al. (2023) advocated for a multilingual setup, employing translation to English before utilizing RoBERTa. Additionally, Teimas and Saias (2023) proposed a hybrid technique combining CNN with DistilBERT for improved detection accuracy.

Exploring multimodal content, Glenski et al. (2019) delved into multilingual multimodal deception detection, mainly focusing on hateful memes. Leveraging visual and textual content, they utilized fine-tuning techniques with state-of-the-art models like ViLBERT and VisualBERT and transfer learning-based approaches (Gupta et al., 2021).

3 Data Description

The competition consisted of two different phases mainly the development phase which we refer to as the development set and for the development phase we were provided the training and validation sets for benchmarking our models

All three sub-tasks have different sets of memes split across training, validation and Development sets as shown in Table 2. We have also plotted the Distribution of the labels across the Figure 3 training data and the Figure 4 validation data.

Our analysis used a dictionary to map various rhetorical techniques to numerical values for plotting. This dictionary is as follows:

Persuasion Technique	Number mapped to
Presenting Irrelevant Data (Red Herring)	0
Bandwagon	1
Smears	2
Glittering generalities (Virtue)	3
Causal Oversimplification	4
Whataboutism	5
Loaded Language	6
Exaggeration/Minimisation	7
Repetition	8
Thought-terminating cliché	9
Name calling/Labeling	10
Appeal to authority	11
Black-and-white Fallacy/Dictatorship	12
Obfuscation, Intentional vagueness, Confusion (Straw Man)	13
Reductio ad hitlerum	14
Appeal to fear/prejudice	15
Misrepresentation of Someone's	16
Position (Straw Man)	17
Flag-waving	18
Slogans	19
Doubt	19

Table 1: Dictionary Mapping for different persuasion techniques for Subtask 1




Figure 4: The frequency Distribution of Labels in the validation dataset

4 System overview

The proposed system for all the sub-tasks involves task-specific modifications made to the BERT model and earlier proposed works including CLIP Model (Radford et al., 2021), Class Definition based Emotion Predictions (Singh et al., 2021, 2023) and HypEmo model (Chen et al., 2023) (described below).

4.1 Data Pre-processing

To ensure consistency and standardization, we begin by pre-processing the text. This involves removing newline characters, commas, numerical values, and other special characters. Additionally, the entire text is converted to lowercase. In our approach, we leverage the Development (Dev) and Training sets, focusing solely on samples containing non-zero classes.



Figure 5: The meme sarcastically suggests that individuals who oppose Trump are being unfairly equated with terrorists, highlighting the absurdity of such comparisons. Two persuasion techniques are used: (i) *Loaded Language*, and (ii) *Name calling* that can be inferred from the text and the visual content.

Sub-task	Train Data	Validation Data	Development Data
Sub-task 1	7000	500	1000
Sub-task 2	7000	500	1000
Sub-task 3	1200	300	500

Table 2: Distribution of data across sub-tasks

4.2 Sub-task 1: Hierarchical Multi-label Text Classification

We present a novel approach to meme classification, drawing upon the methodologies of two key frameworks: HypEmo and a multi-task learning model focused on emotion definition modeling.

HypEmo (Chen et al., 2023) utilizes pre-trained label hyperbolic embeddings to capture hierarchical structures effectively, particularly in tree-like formations. Initially, the hidden state of the [CLS] token from the RoBERTa backbone model is projected using a Multi-Layer Perceptron (MLP). Subsequently, an exponential map is applied to project it into hyperbolic space. The distance from pre-trained label embeddings is the weight for the cross-entropy loss function, enhancing the model’s sensitivity to label relationships.

To implement the HypEmo architecture, we transform the Directed Acyclic Graph (DAG) (Figure 2) into a tree structure. This involves duplicating children with multiple parents, resulting in distinct embeddings for each label. For example, a sentence with various labels is converted into separate samples, each assigned one label. Utilizing the Poincaré hyperbolic entailment cones model (Ganea et al., 2018) with 100 dimensions, the con-




Figure 6: Proposed architecture for sub-task 2

structed tree undergoes training, with predictions generated via softmax. Peaks are identified through Z-score analysis associated with each class, with thresholds set accordingly.

Singh et al. (2021, 2023) have introduced a complementary approach focusing on emotion prediction through a multi-task learning framework. This model incorporates auxiliary tasks, including masked language modeling (MLM) and class definition prediction, to enhance the understanding of emotional concepts. In our setup, class definitions are merged using a [SEP] token, with the model trained to predict whether the conjoined definition matches the actual definition. Binary cross-entropy loss is employed for this task, along with MLM for fine-tuning the model. Additionally, binary cross-entropy loss is used for each class during training. We utilize class definitions provided by the meme classification competition for the auxiliary task of class-definition prediction.

Finally, we merge the predictions generated by both models (HypEmo, Fine-grained class-definition based model) to compute the final predictions. This integrated approach aims to leverage the strengths of each framework, enhancing the accuracy and comprehensiveness of meme classification outcomes.

4.3 Sub-task 2: Hierarchical Multi-label Text and Image Classification

We model this sub-task by experimenting with using an ensemble of HypEmo (Chen et al., 2023)

and the class definition-based multi-task learning model (Singh et al., 2021, 2023) for the textual content of the meme and using the CLIP model (Radford et al., 2021) embeddings for extracting the relevant features from the visual content of the meme. We construct a similar DAG structure for sub-task 1 and generate the hyperbolic embeddings. The image embeddings obtained from the CLIP model are concatenated with the embeddings generated for the textual contents before sending the combined feature vector for training. Then, the model is trained, and the predictions are generated using the softmax activation function. The Z-score analysis is done on the resulting predictions to make the classification, similar to task 1. An overview of the architecture of the modified HypEmo model is shown in Figure 7.




Figure 7: Proposed architecture for sub-task 3

shown below:

$$L(\mathbf{x}, \mathbf{y}) = -\frac{1}{N} \sum_1^N (w * y_i * \log(x_i) + (1 - w) * (1 - y_i) * \log(1 - x_i))$$

$$w = \frac{1}{f}(K - f)$$

where N is the batch size, i is the index of the i^{th} batch element, f is the frequency of the positive class, \mathbf{x} is the output of the last *sigmoid* layer, \mathbf{y} is the vector of the ground truth labels, and K is the total size of the training dataset. Finally, by choosing the one with a higher probability, we use the output probabilities of the final *sigmoid* layer to predict whether a persuasion technique is present in the meme.

5 Experimental setup

5.1 Implementation Details

We have used the official PyTorch implementation (Paszke et al., 2019) for implementing all the models across sub-tasks. We have used the HypEmo² model and the Class Definition Prediction (CDP)³ model for generating the hyperbolic embeddings and class-definition based features of the textual contents, respectively and the CLIP⁴ mainly the 'clip-ViT-B-32' model for generating embeddings for the visual features of the meme. Some portions of the test set have languages other than English for

²HypEmo, <https://github.com/dinobby/hypemo>

³CDP, <https://github.com/Exploration-Lab/FineGrained-Emotion-Prediction-Using-Definitions>

⁴CLIP, <https://github.com/openai/CLIP>

4.4 Sub-task 3: Binary Text and Image Classification

In this task, we must classify whether a meme contains a persuasion technique based on its textual and visual content. We use the pre-trained BERT_{BASE} model (Devlin et al., 2019) and the Convolution Neural Network (CNN) (O'Shea and Nash, 2015) layer to extract the features from the text and image, respectively. We attach a feed-forward [CLS] token embedding along with two linear layers connected by the *sigmoid* activation function in between, which generates the sentence embeddings corresponding to the textual content in the meme. We use a network of four CNN layers connected through the ReLU activation function, which progressively extracts features from the input image. Max pooling layers are used to down-sample the feature maps, increasing robustness to minor variations. The resultant image embeddings are concatenated with the sentence embeddings, and a linear classifier is applied to the combined feature vector with the *sigmoid* activation function. We use the binary cross-entropy loss function to train the model and tune the hyperparameters on the validation set. An overview of the model architecture is shown with an example in Figure 8.

Since the training data is in a 2:1 ratio for the "persuasive" (positive, labeled as 1) and "not-persuasive" (negative, labeled as 0) class, which leads to an imbalance in the dataset, we use the weighted binary cross entropy loss function as

testing purposes. Since the models described earlier were trained in English, we translated the non-English data into English language using the implementation of the OPUS-MT model (Tiedemann and Thottingal, 2020) from the HuggingFace⁵ library and inference was done on the translated text. We created an ensemble of classes predicted by all the models and took a union of the predicted labels to produce the final predicted set of labels to which the meme belonged.

We have used the data in the same ratio provided in the task to train the models. We combine the train validation dataset for training in each subtask and test it in the four languages.

5.2 Evaluation Metrics

Sub-tasks 1 and 2 depend on a hierarchy, as shown in Figure 2. Hierarchical-F1 (Kiritchenko et al., 2006) is used as the evaluation metric for these two sub-tasks. In these two, the gold label is always a leaf node of the DAG, considering the hierarchy in Figure 2 as a reference. However, any node of the DAG can be a predicted label with:

- If the prediction is a leaf node and it is the correct label, then a full reward is given. For example, *Red Herring* is predicted and is the gold label.
- If the prediction is NOT a leaf node and an ancestor of the correct gold label, then a partial reward is given (the reward depends on the distance between the two nodes). For example, if the gold label is *Red Herring* and the predicted label is *Distraction* or *Appeal to Logic*.
- If the prediction is not an ancestor node of the correct label, then a null reward is given. For example, if the gold label is *Red Herring* and the predicted label is *Black and White Fallacy* or *Appeal to Emotions*.

Sub-task 3 uses macro-F1 as the evaluation metric for the binary classification task. This ensures equal importance to the "persuasion technique present" and "no persuasion technique" classes, regardless of potential data imbalance.

6 Results

We conducted several experiments across all the sub-tasks, and the detailed information can be seen

⁵OPUS-MT, <https://huggingface.co/Helsinki-NLP/opus-mt-bg-en>

Technique	Arabic	Bulgarian	North Macedonian
Presenting Irrelevant Data (Red Herring)	0.	0.	0.
Bandwagon	0.	0.	0.
Smears	0.67	0.84	0.90
Glittering generalities (Virtue)	0.29	0.10	0.
Causal Oversimplification	0.	0.	0.
Whataboutism	0.	0.05	0.
Loaded Language	0.41	0.62	0.37
Exaggeration/Minimisation	0.	0.	0.
Repetition	0.50	0.34	0.
Thought-terminating cliché	0.	0.19	0.
Name calling/Labeling	0.44	0.45	0.49
Appeal to authority	0.	0.30	0.31
Black-and-white Fallacy/Dictatorship	0.	0.06	0.
Obfuscation, Intentional vagueness, Confusion (Straw Man)	0.	0.	0.
Reductio ad hitlerum	0.	0.	0.
Appeal to fear/prejudice	0.04	0.21	0.1
Misrepresentation of Someone's Position (Straw Man)	0.	0.	0.
Flag-waving	0.	0.33	0.
Slogans	0.	0.43	0.16
Doubt	0.	0.15	0.11
Transfer	0.	0.48	0.61
Appeal to (Strong) Emotions	0.	0.18	0.09

Table 3: Macro F1 scores for different persuasion classes for the given languages for Subtask 2

Technique	Arabic	Bulgarian	North Macedonian
Presenting Irrelevant Data (Red Herring)	0.	0.	0.
Bandwagon	0.	0.	0.
Smears	0.33	0.18	0.17
Glittering generalities (Virtue)	0.	0.07	0.
Causal Oversimplification	0.	0.	0.
Whataboutism	0.	0.08	0.
Loaded Language	0.39	0.62	0.55
Exaggeration/Minimisation	0.11	0.	0.
Repetition	0.40	0.36	0.
Thought-terminating cliché	0.	0.28	0.
Name calling/Labeling	0.39	0.58	0.54
Appeal to authority	0.	0.38	0.22
Black-and-white Fallacy/Dictatorship	0.	0.04	0.
Obfuscation, Intentional vagueness, Confusion (Straw Man)	0.	0.	0.
Reductio ad hitlerum	0.	0.	0.
Appeal to fear/prejudice	0.	0.05	0.
Misrepresentation of Someone's Position (Straw Man)	0.	0.	0.
Flag-waving	0.	0.29	0.
Slogans	0.	0.37	0.04
Doubt	0.25	0.16	0.1

Table 4: Macro F1 scores for different persuasion classes for the given languages for Subtask 1

in Table 3, Table 4, Table 5, Table 8 and Table 9.

For Task 1, we started experimenting with the BERT and RoBERTa models, achieving a hierarchical F1 score of 0.55 and 0.60 on the test set of the English language. But, in this approach, we did not take the hierarchy and the definitions of the classes into consideration. We tried to accommodate that using the combination of HypEmo and CDP models.

For the HypEmo model, the model was trained to prioritize higher-level labels in the Directed Acyclic Graph (DAG). During this process, we explored two options: eliminating children when the model predicted the parent label and retaining the children. We observed a significant impact on the hierarchical F1 score, with the first formulation

Language	Base F1	Hierarchical F1	Hierarchical Precision	Hierarchical Recall
English	0.37	0.60	0.53	0.69
Arabic	0.37	0.42	0.32	0.60
Bulgarian	0.28	0.48	0.40	0.62
North-Macedonian	0.30	0.41	0.33	0.56

Table 5: Hierarchical-F1 scores computed across four languages of the test set for sub-task 1. Base F1 score here is the Baseline F1 score

Language	Baseline F1	Hierarchical F1	Hierarchical Precision	Hierarchical Recall
English	0.44	0.67	0.67	0.67
Arabic	0.57	0.53	0.50	0.57
Bulgarian	0.50	0.65	0.66	0.63
North-Macedonian	0.55	0.67	0.72	0.62

Table 6: Hierarchical-F1 scores calculated for four languages within the test set for sub-task 2, with Base-F1 denoting the Baseline F1 score depicted on the leaderboard

yielding 0.45 F1 and the second approach resulting in 0.59 on the test set. We also tried to predict the labels utilizing only the definitions of the classes, using the CDP model, which yielded a hierarchical F1 score of 0.57 and 0.59 on the dev set and the test set, respectively.

For constructing an ensemble, one approach considered concatenating embeddings or softmax predictions from both models for further classification using a neural network. However, this approach was not viable due to limited samples for generalization. The most effective model emerged from utilizing the ensemble with fine-tuning of hyperparameters. Combining predictions from both models yielded a hierarchical F1 score of 0.60.

Table 8 shows that the best generalizability across all tasks is achieved via the HypEmo + CDP(Union) for subtask1.

For sub-task 2, we trained the model from scratch after including the two labels used in the ensemble used in sub-task 1 and changed the feature embeddings being trained by considering the features from the visual content. However, as you can see, there is very little to no difference between the results using CLIP and not using CLIP. We can also see that, unlike the first subtask, they perform better due to more data.

We can see the F1-score analysis tables for each subtask, i.e., in Table 4, Table 3 for subtask1 and subtask2.

For sub-task 3, we trained the model on an en-

Model	English	Arabic	Bulgarian	North-Macedonian
BERT	0.55	0.39	0.40	0.36
RoBERTa	0.60	0.37	0.45	0.38
HypEmo	0.55	0.43	0.42	0.39
CDP	0.59	0.40	0.48	0.43
HypEmo + CDP (Union)	0.60	0.42	0.48	0.41

Table 7: Hierarchical-F1 scores calculated for four languages within the test set for sub-task 1 across different models

Model	English	Arabic	Bulgarian	North-Macedonian
HypEmo (Without CLIP)	0.63	0.511	0.58	0.63
HypEmo (With CLIP)	0.63	0.49	0.59	0.62
CDP	0.64	0.51	0.62	0.65
HypEmo + CDP (Union)	0.67	0.53	0.65	0.67

Table 8: Hierarchical-F1 scores calculated for four languages within the test set for sub-task 2 across different models

Language	Base F1	Macro-F1
English	0.25	0.49
Arabic	0.23	0.47
North Macedonian	0.09	0.49
Bulgarian	0.16	0.48

Table 9: Macro-F1 scores computed across 4 languages of the test set for sub-task 3.

Sub-task	Ranking
English-Sub-task1	21
English-Sub-task2	10
English-Sub-task3	19
Bulgarian-Sub-task1	14
Bulgarian-Sub-task2	8
Bulgarian-Sub-task3	11
North Macedonian-Sub-task1	13
North Macedonian-Sub-task2	7
North Macedonian-Sub-task3	7
Arabic-Sub-task1	4
Arabic-Sub-task2	6
Arabic-Sub-task3	13

Table 10: Leaderboard position of our team in the competition in each sub-task

semble of BERT and CNN models to consider the textual and visual features. It was seen that the ensemble performs just slightly better than using the BERT model, that is, considering only the textual cues. Visual cues are considered significantly when persuasion techniques like *Smears* are used, as seen in sub-task 2. For the rest of the persuasion techniques, the visual cues were seen not to make a significant impact on the classification task. On the gold labels of the dev set, the ensemble gave a macro-F1 score of 0.67, which is a slight im-

provement from the BERT model, which showed a macro-F1 score of 0.63 on the dev set.

7 Conclusion

Detection of persuasion techniques in memes is seen in a multi-modal setting in this task, but the significant features are drawn from the textual cues in the memes, which can be seen in the results of sub-tasks 1 and 2. The CLIP and other visual language models still need considerable development, and visual cues are helpful for only specific input-output pairs. Identifying whether a persuasion technique is present in the meme but does not apply to the multi-label classification task can be beneficial. Also, we have used a basic ensemble of the latest works in this area and modified them for task-specific requirements. Still, other complex architectures can be explored to get better results.

References

- Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Propopy: Organizing the news based on their propagandistic content. *Information Processing & Management*, 56(5):1849–1864.
- Chih-Yao Chen, Tun-Min Hung, Yi-Li Hsu, and Lun-Wei Ku. 2023. [Label-aware hyperbolic embeddings for fine-grained emotion classification](#).
- Nelson Filipe Costa, Bryce Hamilton, and Leila Kosseim. 2023. Clac at semeval-2023 task 3: Language potluck roberta detects online persuasion techniques in a multilingual setup. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 1613–1618.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Dimitar Dimitrov, Firoj Alam, Maram Hasanain, Abul Hasnat, Fabrizio Silvestri, Preslav Nakov, and Giovanni Da San Martino. 2024. Semeval-2024 task 4: Multilingual detection of persuasion techniques in memes. In *Proceedings of the 18th International Workshop on Semantic Evaluation, SemEval 2024*, Mexico City, Mexico.
- Ali Fadel, Ibraheem Tuffaha, and Mahmoud Al-Ayyoub. 2019. Pretrained ensemble learning for fine-grained propaganda detection. In *Proceedings of the second workshop on natural language processing for internet freedom: censorship, disinformation, and propaganda*, pages 139–142.
- Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic entailment cones for learning hierarchical embeddings. In *International Conference on Machine Learning*, pages 1646–1655. PMLR.
- Maria Glenski, Ellyn Ayton, Josh Mendoza, and Svitlana Volkova. 2019. Multilingual multimodal digital deception detection and disinformation spread across social platforms. *arXiv preprint arXiv:1909.05838*.
- Bin Guo, Yasan Ding, Lina Yao, Yunji Liang, and Zhiwen Yu. 2020. The future of false information detection on social media: New perspectives and trends. *ACM Computing Surveys (CSUR)*, 53(4):1–36.
- Kshitij Gupta, Devansh Gautam, and Radhika Mamidi. 2021. [Volta at semeval-2021 task 6: Towards detecting persuasive texts and images using textual and multimodal ensemble](#).
- Vansh Gupta and Raksha Sharma. 2021. [NLPITR at SemEval-2021 task 6: RoBERTa model with data augmentation for persuasion techniques detection](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 1061–1067, Online. Association for Computational Linguistics.
- Ivan Habernal, Patrick Pauli, and Iryna Gurevych. 2018. Adapting serious game for fallacious argumentation to german: Pitfalls, insights, and best practices. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Vishal Keswani, Sakshi Singh, Suryansh Agarwal, and Ashutosh Modi. 2020. [IITK at SemEval-2020 task 8: Unimodal and bimodal sentiment analysis of Internet memes](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1135–1140, Barcelona (online). International Committee for Computational Linguistics.
- Svetlana Kiritchenko, Richard Nock, and Fazel Famili. 2006. [Learning and evaluation in the presence of class hierarchies: Application to text categorization](#). volume 4013, pages 395–406.
- Patricia Moravec, Randall Minas, and Alan Dennis. 2018. [Fake news on social media: People believe what they want to believe when it makes no sense at all](#). *SSRN Electronic Journal*.
- Keiron O’Shea and Ryan Nash. 2015. [An introduction to convolutional neural networks](#).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#).

- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 2931–2937.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*, 19(1):22–36.
- Gargi Singh, Dhanajit Brahma, Piyush Rai, and Ashutosh Modi. 2021. [Fine-grained emotion prediction by modeling emotion definitions](#). In *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 1–8, Los Alamitos, CA, USA. IEEE Computer Society.
- Gargi Singh, Dhanajit Brahma, Piyush Rai, and Ashutosh Modi. 2023. [Text-based fine-grained emotion prediction](#). *IEEE Transactions on Affective Computing*, pages 12–12.
- Paramansh Singh, Siraj Sandhu, Subham Kumar, and Ashutosh Modi. 2020. [newsSweeper at SemEval-2020 task 11: Context-aware rich feature representations for propaganda classification](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1764–1770, Barcelona (online). International Committee for Computational Linguistics.
- Gabriella Szabo. 2020. [Emotional communication and participation in politics](#). *Intersections*, 6:5–21.
- Rúben Teimas and José Saias. 2023. Detecting persuasion attempts on social networks: Unearthing the potential of loss functions and text pre-processing in imbalanced data settings. *Electronics*, 12(21):4447.
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT – building open translation services for the world](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Shehel Yoosuf and Yin Yang. 2019. Fine-grained propaganda detection with fine-tuned bert. In *Proceedings of the second workshop on natural language processing for internet freedom: censorship, disinformation, and propaganda*, pages 87–91.