

ArAIEval Shared Task: Propagandistic Techniques Detection in Unimodal and Multimodal Arabic Content

Maram Hasanain^{1*}, Md. Arid Hasan^{2*}, Fatema Ahmed¹, Reem Suwaileh³,
Md. Rafiul Biswas³, Wajdi Zaghrouani³, Firoj Alam¹

¹Qatar Computing Research Institute, HBKU, Doha, Qatar

²University of New Brunswick, Fredericton, NB, Canada

³Hamad Bin Khalifa University, HBKU, Doha, Qatar

{mhasanain, rsuwaileh, faktor, mdbi30331, wzaghrouani, fialam}@hbku.edu.qa,
arid.hasan@unb.ca

<https://araieval.gitlab.io/>

Abstract

We present an overview of the second edition of the ArAIEval shared task, organized as part of the ArabicNLP 2024 conference co-located with ACL 2024. In this edition, ArAIEval offers two tasks: (i) detection of propagandistic textual spans with persuasion techniques identification in tweets and news articles, and (ii) distinguishing between propagandistic and non-propagandistic memes. A total of 14 teams participated in the final evaluation phase, with 6 and 9 teams participating in Tasks 1 and 2, respectively. Finally, 11 teams submitted system description papers. Across both tasks, we observed that fine-tuning transformer models such as AraBERT was at the core of the majority of the participating systems. We provide a description of the task setup, including a description of the dataset construction and the evaluation setup. We further provide a brief overview of the participating systems. All datasets and evaluation scripts are released to the research community.¹ We hope this will enable further research on these important tasks in Arabic.

1 Introduction

Online media has become a primary channel for information dissemination and consumption, with numerous individuals considering it their main source of news (Perrin, 2015). While online media (including news and social media platforms) offers a plethora of benefits, it is periodically exploited by malicious actors aiming to manipulate and mislead a broad audience. They often engage in sharing inappropriate content, misinformation, and disinformation (Alam et al., 2022a; Sharma et al., 2022). Among the various forms of misleading and harmful content, propaganda is another communication tool designed to influence opinions and actions to

achieve a specific objective. Propaganda is defined as a form of communication that is aimed at influencing the attitude of a community toward some cause or position by presenting only one side of an argument, which is achieved by means of well-defined rhetorical and psychological devices (IPA, 1938). It is often biased or misleading in nature and is used to promote a particular political cause or point of view.

News reporting in the mainstream media channels often use persuasion techniques to promote a particular editorial agenda. In different communication channels, propaganda is conveyed through the use of diverse persuasion techniques (Miller, 1939), which range from leveraging the emotions of the audience, such as using *emotional technique* or logical fallacies such as *straw man* (misrepresenting someone’s opinion), hidden *ad-hominem fallacies*, and *red herring* (presenting irrelevant data).

In the past years, propaganda has been widely used on social media to influence and/or mislead the audience, which became a major concern for different stakeholders including social media platforms and government agencies. To combat the proliferation of propaganda on online platforms, there has been a significant surge in research in recent years. The aim is to automatically identify propagandistic content in textual, visual, and multimodal modalities, such as memes (Chen et al., 2023; Dimitrov et al., 2021; Da San Martino et al., 2019). These efforts mainly focused on English, but with increased interest in other languages too.

We enrich the Arabic AI research on that problem, by organizing a shared task on fine-grained propaganda techniques detection for Arabic, which attracted many participants (Alam et al., 2022b). We followed that by organizing the **Arabic AI Evaluation (ArAIEval)** shared task at ArabicNLP 2023, targeting binary and multilabel propaganda

* Equal contribution.

¹<https://araieval.gitlab.io/>

detection in Arabic text, with 14 and 8 teams participating in these subtasks, respectively (Hasanain et al., 2023b).

Following the success of our previous shared tasks, and given the interest from the community in further pushing research in this domain, we organize the second edition of the ArAIEval shared task covering the following two tasks: (i) detection of propagandistic textual spans with persuasion techniques identification (unimodal), and (ii) distinguishing between propagandistic and non-propagandistic memes (multimodal).

This edition of the shared task attracted good participation with a total of 45 registrations. Fourteen teams submitted their systems, and 11 teams submitted system description papers. For all tasks, majority of participating systems utilized transformer-based models, and several of them applied data augmentation techniques. All systems outperformed the respective random baselines.

In the remainder of this paper, we define each of the two tasks, describe the manually constructed Arabic evaluation datasets, and provide an overview of the participating systems and their official scores.

2 Related Work

2.1 Persuasion Techniques

The detection of persuasion techniques has brought significant attention across multiple research domains, including natural language processing (NLP), computational linguistics, psychology, and communication studies. Researchers have developed various approaches to automatically identify persuasive elements in digital content, leveraging advancements in NLP and machine learning.

One prominent approach focuses on the linguistic analysis of persuasive texts. For instance, Hidey et al. (2017) proposed a framework for identifying persuasion techniques in online discussions by analyzing rhetorical strategies such as appeals to emotion, authority, and logic. Their study demonstrated the potential of NLP techniques to classify and detect persuasive elements with high accuracy.

Building on this foundation, researchers have explored the use of machine learning models to enhance the detection of persuasion techniques. Habernal and Gurevych (2015) introduced the ArgumenText system, which applies supervised learning to classify argument components and identify persuasion techniques in user-generated content.

This system showed promising results in distinguishing between different types of persuasive appeals.

In the domain of social media, the detection of persuasion techniques has been particularly challenging due to the diverse and dynamic nature of user interactions. One notable study by Morstatter et al. (2018) utilized deep learning models to detect persuasion techniques in tweets related to political campaigns. By analyzing both textual features and user metadata, the researchers were able to identify persuasive content employed by political actors and their impact on public opinion.

The advent of a large language model has paved a new direction in persuasion techniques detection. Hasanain et al. (2024) developed *ArPro* containing 8K paragraphs labeled at the text span level for 23 propagandistic techniques and used it to evaluate the performance of GPT-4 in fine-grained propaganda detection. GPT-4 model has also shown improved performance when using a large-scale dataset with annotations from human annotators of varying expertise and providing more information to the model as prompts (Hasanain et al., 2023a).

2.2 Multimodal Detection

Multilingual and multimodal detection of persuasion techniques has also gained attention in recent years. Park et al. (2016) created a multimodal corpus including acoustic, verbal, and visual features of 1,000 movie review videos to predict persuasiveness. Fang et al. (2022a) embedded text and images using different networks and fused the multi-modal embeddings. A split-and-share module with multi-level representations then processes these fused features to improve the detection performance of persuasive techniques.

Nojavanasghari et al. (2016) introduced a deep neural network architecture for predicting persuasiveness by combining three modalities: visual, acoustic, and text modalities. The multimodal fusion model outperforms unimodal models and can handle limited labeled data effectively. Another significant contribution to the field is the work by Chen et al. (2017) on multimodal sentiment analysis, which is closely related to the detection of persuasive content.

2.3 Related Shared Tasks

Several shared tasks have been organized to engage the research community in persuasion techniques detection. The shared task organized by

Barrón-Cedeno et al. at the CLEF 2022 conference focused on detecting propaganda and persuasion techniques in news articles across multimodal data. Similarly, SemEval-2021 Task 6 (Dimitrov et al., 2021) focused on the detection of persuasion techniques in multimodal content, such as memes and videos, combining textual and visual analysis. The shared task in SemEval-2023 Task 3 (Piskorski et al., 2023) consists of 23 persuasion techniques detection and has multiclass and multilabel classification setups focusing on the paragraph level.

Along with such initiatives, we have primarily focused on Arabic content. The propaganda detection shared task, co-located with WANLP 2022, was mainly focused on tweets in both binary and multilabel settings (Alam et al., 2022b). Following the success of that task, we the ArAIEval shared task at Arabic NLP 2023, aimed to identify persuasive elements in Arabic tweets and news articles (Hasanain et al., 2023b) in a multilabel setting. This year, we have expanded our dataset to include multimodal content in form of memes. We also focus on fine-grained propaganda detection at the span-level. This can motivate developing more robust models for detecting persuasion techniques in various Arabic media contexts.

3 Task 1: Unimodal (Text) Propagandistic Technique Detection

3.1 Task Definition

The objective of this task is to develop and evaluate systems capable of detecting specific propagandistic techniques within textual content. The task is defined as follows. Given a multigenre text snippet (a news paragraph or a tweet), the task is to detect the propaganda techniques used in the text together with the exact span(s) in which each propaganda technique appears. This is a multilabel sequence tagging task.

3.2 Data

The dataset for this task covers two genres, tweets and paragraphs extracted from news articles. We collect raw data to annotate as follows.

Tweets: We start from the same tweets dataset collected from Twitter accounts of Arabic news sources, as described in the previous editions of the shared task (Alam et al., 2022b; Hasanain et al., 2023b). In addition to that, we collect a fresh set of tweets targeting the current events in Gaza. The tweet set was collected by searching through the



Figure 1: An example of a news paragraph annotated with propagandistic techniques.

Twitter search API using manually developed Arabic keywords.

News paragraphs: The news paragraphs were sourced and annotated following the same approach followed to develop *ArPro* dataset (Hasanain et al., 2024). News paragraphs are extracted from two collections of news articles: (i) AraFacts, and (ii) a large-scale in-house collection. The AraFacts dataset (Ali et al., 2021) contains true and false Arabic claims verified by fact-checking websites, and each claim is associated with online sources propagating or negating the claim. We only keep Web pages that are from news media in the set (e.g., www.alquds.co.uk). We automatically parsed the news articles and selected paragraphs for the annotation. As for the **in-house collection**, it consists of 600K news articles from over 400 news media outlets. The dataset is very diverse, covering 14 different broad topics and a wide coverage of Arabic news media.

Data annotation: The annotation process includes two phases as detailed in previous studies (Hasanain et al., 2024, 2023a). The two annotation phases consist of: (i) three annotators independently annotate the same text snippet by persuasion techniques on the span level, through an annotation interface designed for the task and (ii) two consolidators (expert annotators) discuss each instance and finalize the gold annotations. Since annotations are at the fragment level, it might happen that an an-

notation is spotted by only one annotator. In order to train the annotators, we provide clear annotation instructions with examples and ask them to annotate a sample of text (tweets, or paragraphs from news articles). Then, we revise their annotations and provide feedback. In Figure 1, we provide an example of span annotation along with its English translation, which shows different text spans annotated with propagandistic techniques.

The annotation taxonomy consists of a set of 23 persuasion techniques that is adopted from existing research (Piskorski et al., 2023). We should note here that multiple techniques can be found in the same text snippet even with an overlap. Below we give an example subset of the persuasion techniques, and briefly summarize them:

1. **Name calling or labeling:** labeling the object of the propaganda campaign as something that the target audience fears, hates, finds undesirable, or loves, praises.
2. **Appeal to Fear, Prejudice:** building support or rejection for an idea by instilling fear or repulsion towards it, or to an alternative idea.
3. **Strawman:** giving the impression that an argument is refuted, whereas the real subject of the argument was not addressed or refuted, but instead was replaced with a different one.

Data splits: The full set of annotated paragraphs is divided into three subsets: train, development, and test, using a stratified splitting approach to ensure that the distribution of persuasion techniques is consistent across the splits. For the tweets set, we split the full annotated tweet set from the previous edition of the lab (Alam et al., 2022b) into train and development subsets, while the test set is annotated for this shared task. Finally, we construct the multi-genre subsets for the task by merging the sets of paragraphs and tweets. The full dataset is $\sim 9\text{K}$ snippets, randomly split into $\sim 78\%$, $\sim 11\%$ and $\sim 11\%$ training, development and test splits, respectively.

Statistics: In Tables 1 we show the distribution of labels across splits for Task 1. It reports that the techniques *loaded language*, *name calling*, and *exaggeration-minimisation* are the most frequent among all techniques. Overall, the label space is highly imbalanced, which is typical for these types of tasks, as seen in previous studies (Da San Martino et al., 2019; Dimitrov et al., 2021).

Technique	Train	Dev	Test
Appeal to Authority	218	28	13
Appeal to Fear-Prejudice	180	32	7
Appeal to Hypocrisy	108	16	7
Appeal to Popularity	46	4	3
Appeal to Time	54	6	8
Appeal to Values	116	24	5
Causal Oversimplification	310	38	21
Consequential Oversimplification	83	11	6
Conversation Killer	69	10	6
Doubt	293	44	26
Exaggeration-Minimisation	1077	141	133
False Dilemma-No Choice	74	9	5
Flag Waving	218	34	17
Guilt by Association	23	2	4
Loaded Language	8779	1073	1429
Name Calling-Labeling	2243	330	432
Obfuscation-Vagueness-Confusion	578	67	55
Questioning the Reputation	872	128	106
Red Herring	41	5	3
Repetition	144	17	17
Slogans	190	37	7
Straw Man	23	3	2
Whataboutism	26	5	1

Table 1: Distribution of the persuasion techniques in Task 1 dataset across different splits.

3.3 Evaluation Setup

The task was organized into two phases:

- **Development phase:** we released the train and development subsets, and participants submitted runs on the development set through a competition on Codalab.²
- **Test phase:** we released the official test subset, and the participants were given a few days to submit their final predictions through the same Codalab competition. Only the latest submission from each team was considered official and was used for the final team ranking.

Measures: Task 1 is a multilabel sequence tagging task. We measure the performance of the participating systems using a modified F_1 measure that accounts for partial matching between the spans across the gold labels and the predictions (Alam et al., 2022b).

3.4 Results and Overview of the Systems

A total of 6 teams submitted runs for the evaluation phase of the task. In Table 2, we provide an overview of the participating systems for which a description paper was submitted. As shown in

²<https://codalab.lisn.upsaclay.fr/competitions/18111>

Team	Transformer	DL	ML	Misc.
	AraBERTv2 MARBERT Arabic-BERT mBERT CAMELBERT	BiLSTM/CNN/LSTM	RF/SVM/MNB	Data augm. Preprocessing Hyperparameter tuning Feature Engineering Model Layer
CUET_sstm (Labib et al., 2024)	☑	☑	☑	☑
Mela (Riyadh and Nabhani, 2024)	☑			☑
MemeMind (Biswas et al., 2024)	☑			☑
Nullpointer (Abrar Abir, 2024)	☑			☑
SussexAI (Fouad and Weeds, 2024)	☑			☑
SemanticCuetSync (Shohan et al., 2024)	☑	☑	☑	☑

Table 2: **Task 1** Overview of the approaches. DL: Deep Learning, ML: Classic Machine Learning.

Team	Rank	Micro F1
CUET_sstm	1	0.2995
Mela	2	0.2833
MemeMind	3	0.2774
Nullpointer	4	0.2541
SussexAI	5	0.1228
SemanticCuetSync	6	0.0783
Baseline		0.0151

Table 3: Official results for **Task 1**. Runs ranked by the official measure: modified micro F1.

Table 2, fine-tuning pre-trained Arabic models, specifically AraBERT (Antoun et al., 2020), multilingual BERT (mBERT) (Devlin et al., 2019), Arabic-BERT (Safaya et al., 2020), and CAMEL-BERT (Inoue et al., 2021) is the most common system architecture.

In Table 3, we report the results of the participants’ systems including a random baseline. Results are competitive among the top four teams, performance differences are relatively lower among them. All systems surpass the random baseline.

Team **CUET_sstm** (Labib et al., 2024) used data augmentation with the synonym replacement approach from the nlpaug library. The data was preprocessed and encoded with the BIO technique to transform it into a token classification problem for span detection. They experimented with fine-tuning different models, including transformer-based models such as MAREFANER and AraBERT; however, Arabic-BERT outperformed the other models.

Team **Mela** (Riyadh and Nabhani, 2024) fine-tuned mBERT on the training set and experimented with different hidden layers of the model. Their experiments showed that the 10th hidden layer yielded the best F1 score on the development set.

Team **SemanticCuetSync** (Shohan et al., 2024) experimented with different types of models. Firstly, they trained typical machine learning models, including: Logistic Regression, Support Vector Machine, and Multinomial Naive Bayes. Secondly, they also trained deep learning-based models such as CNN, CNN+LSTM, and CNN+BiLSTM. Lastly, several transformer-based models were fine-tuned: AraBERTv2, CAMELBERT, and Arabic-BERT. Among all the models, CAMELBERT resulted in the highest micro-F1 score.

Team **SussexAI** (Fouad and Weeds, 2024) performed data augmentation on least represented labels to reduce label imbalance. They also generated synthetic data using random masking. They truncated ‘non-propaganda’ tokens outside ‘propaganda’ spans and reduced ‘non-propaganda’ tokens. Finally, they fine-tuned the AraBERT transformer model with three dataset setups: augmented, non-augmented, and random truncation.

Team **MemeMind** (Biswas et al., 2024) encoded each label and created BIO tags, modeling the problem as a token classification task. Then, they fine-tuned AraBERTv2 in two steps: trained only the classification layer for a few epochs and then fine-tuned the full model by updating all parameters. They tried different transformer models. Among them, AraBERTv2 showed the best performance.

Team **Nullpointer** (Abrar Abir, 2024) applied dif-

Class labels	Train	Dev	Test	Total
Not propaganda	1,540	224	436	2,200
Propaganda	603	88	171	862
Total	2,143	312	607	3,062

Table 4: Distribution of Task 2 dataset.

ferent preprocessing techniques, such as filtering out Unicode characters. Then, they fine-tuned AraBERTv2 following two setups to model the task: (i) model the problem as a token classification task, with different label aggregation applied over tokens at prediction time to find per-word label, and (ii) model the problem as a word classification task, representing each word using max-pooling of embeddings of its constituent tokens.

4 Task 2: Multimodal Propagandistic Memes Classification

4.1 Task Definition

The goal of this task is to develop and evaluate systems capable of classifying multimodal propagandistic memes. Memes typically consist of a background image, and a layer of text that adds context, humour, or commentary to the image. The combination of the image and the text creates a specific message, joke, or commentary that is meant to be easily understood, relatable, and shareable.

The aim of this task is to foster research in the intersection of multimodal learning and propaganda detection, encouraging innovative solutions that effectively combine visual and textual information to identify and understand propagandistic contents. To account for the different modalities covered by this task, we model the task into three binary classification subtasks.

- **Subtask 2A:** Given a text extracted from a meme, categorize whether it is propagandistic or not.
- **Subtask 2B:** Given a meme (text overlaid image), detect whether the content is propagandistic.
- **Subtask 2C:** Given multimodal content (text extracted from meme and the meme itself), detect whether the content is propagandistic.

4.2 Dataset

The dataset consists of ~ 3 K memes annotated as propagandistic vs not-propagandistic, which were collected from different social media (e.g., Face-

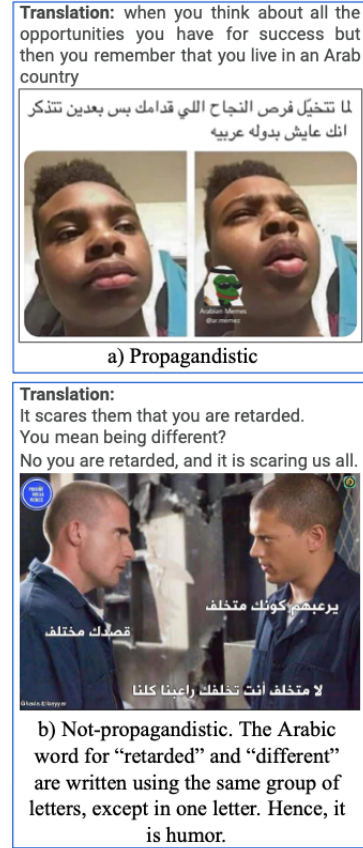


Figure 2: An example of memes with propagandistic and not-propagandistic categories.

book, Twitter, Instagram and Pinterest). Each of them were annotated by three annotators and majority decision is considered as the final label. Texts from the memes were extracted using an off-the-shelf OCR,³ followed by manual editing of the propagandistic memes. An example of such memes is provided in Figure 2. The distribution of propagandistic and not-propagandistic labels are 40% and 60%, respectively. More details of this dataset can be found in (Alam et al., 2024).

Data splits: The dataset is split into 70%, 10% and 20% for training, development, and test, respectively.

Statistics: Table 4 reports the distribution of labels for Task 2. The proportion of propagandistic memes is low, with a total of 28%, which makes the task more challenging.

4.3 Evaluation Setup

Similar to Task 1, we also conducted this task in two phases (i.e., development and test) as dis-

³<https://github.com/JaidedAI/EasyOCR>

Team	Subtask	Text Models	Image Models	Misc.
	2A 2B 2C	AraBERTv2 MARBERT ARABERT QARiB CAMELBERT GigaBERT BLOOMZ	CLIP EfficientFormer ConvNeXt-tiny CAFormer EVA MaxViT SwinV2 ResNet DenseNet LSTM	Data aug. Preprocessing
AlexUNLP-MZ (Zaytoon et al., 2024)	✓	✓	✓	✓
ASOS (Alhabashi et al., 2024)	✓	✓	✓	✓
CLTL (Wang and Markov, 2024)	✓	✓	✓	✓
MemeMind (Shah et al., 2024)	✓	✓	✓	✓
MODOS (Haouhat et al., 2024)	✓	✓	✓	✓

Table 5: **Task 2** Overview of the approaches. DL: Deep Learning, Data aug: Data augmentation

Team	Rank	Macro F1
Subtask 2A		
AlexUNLP-MZ	1	0.787
CLTL	2	0.779
MemeMind	3	0.746
DLRG	4	0.739
One_by_zero	5	0.674
Z-Index	6	0.633
Baseline		0.453
Subtask 2B		
CLTL	1	0.711
MemeMind	2	0.664
AlexUNLP-MZ	3	0.659
Baseline		0.475
Subtask 2C		
AlexUNLP-MZ	1	0.805
ASOS	2	0.798
CLTL	3	0.798
MemeMind	4	0.797
Team Engima	5	0.753
MODOS	6	0.729
Z-Index	7	0.712
Baseline		0.493

Table 6: Official results for **Subtask 2A, 2B and 2C**. Runs ranked by the official measure: Macro F1.

cussed in Section 3.3. Systems were evaluated using macro F1 as the official measure.

4.4 Results Overview of the Systems

A total of 6, 3 and 7 teams submitted runs for Subtask 2A, 2B and 2C, respectively with nine unique teams. Table 5 gives an overview of the participating systems for which a description paper was submitted. Three out of five teams participated in all subtasks. Among the teams, fine-tuning trans-

former models such as MARBERT (Abdul-Mageed et al., 2021) and CAMELBERT (Inoue et al., 2021) is the most popular architecture. As for the vision models ResNet was the most popular choice. In Table 6, we report the results and rankings for *all* systems across all subtasks. All systems outperformed the random baselines by a large margin. Among the three modalities, performances in the text and multimodal modalities (Subtasks 2A and 2C) were relatively higher compared to those in the image-only modality.

Team **AlexUNLP-MZ** (Zaytoon et al., 2024) used a Large Language Model (LLM) to extract features from the dataset. They considered ‘weighted loss’ and ‘contrastive loss’ while training the model. For text classification, they used the BLOOMZ model. For image classification, they tried using CNN-based ResNet and DenseNet architectures. For the multimodal data, they used a fusion of two architectures, such as BLOOMZ-1b1 and ResNet101. Team **ASOS** (Alhabashi et al., 2024) used the MARBERT model for text classification and ResNet50 for image classification, and finally fused these two models for multimodality.

Team **CLTL** (Wang and Markov, 2024) applied MARBERT, CAMELBERT and GigaBERT (Lan et al., 2020) for text classification. CAMELBERT showed superior performance compared to the other models. For the image modality, they examined two models: EVA (Fang et al., 2022b) and CAFormer (Yu et al., 2023). For multimodality, they merged the embeddings of text and image using the multilayer perceptron technique.

Team **MODOS** (Haouhat et al., 2024) preprocessed raw data before feeding it into the model. Firstly, they used the Segment Anything Model⁴ for image segmentation of the meme images. Then, they em-

⁴<https://segment-anything.com/>

ployed the state-of-the-art image encoder CLIP⁵ to extract the image embeddings. Finally, they used LSTM for multimodal classification.

Team **MemeMind** (Shah et al., 2024) used synthetic text generated by GPT-4 (OpenAI, 2023) for text data augmentation. For image modality, they augmented data using DALL-E-2 and fine-tuned ResNet50, EfficientFormer (v2), and ConvNeXt-tiny architectures. For the multimodality, they fused the ConvNeXt-tiny and BERT.

5 Conclusion and Future Work

We presented an overview of the ArAIEval shared task, which consists of two tasks: (i) detection of propagandistic textual spans with persuasion techniques identification in tweets and news articles, and (ii) distinguishing between propagandistic and non-propagandistic memes. The task attracted the attention of many teams: a total of 45 teams registered to participate during the evaluation phase, with 6 and 9 teams eventually making an official submission on the test set for tasks 1 and 2, respectively. Finally, 11 teams submitted task description papers. Task 1 aimed to identify the propaganda techniques used in multi-genre text snippets, including tweets and news articles. The task was to detect textual spans with propagandistic techniques. On the other hand, Task 2 aimed to detect propaganda in memes (multimodal content) in binary classification settings for different modalities. For both tasks, majority of the systems fine-tuned pre-trained transformer models.

Future editions of the ArAIEval shared task will explore more complex tasks, such as incorporating Arabic dialects. Increasing the dataset size and focusing on low-prevalence class labels will be crucial to improve model robustness. For task 2, future work could involve labeling the dataset with specific propagandistic techniques.

Limitations

The datasets for both tasks are skewed in label distribution, making system development more challenging. To address this problem, increasing the data size with a focus on low-prevalence class labels could be beneficial. As for Task 2, we observe that the systems achieved relatively better performance, likely because the task was relatively simple in nature – a binary classification setting.

Ethical Considerations

The ArAIEval shared task involves the development of models to detect propagandistic content in Arabic text and memes. While this research aims to combat the spread of misleading and manipulative information, it is essential to consider the potential ethical implications.

The datasets used in the shared task may contain biases, as propaganda is often subjective and can be influenced by cultural, political, and social factors. It is crucial to acknowledge these biases and strive for diverse and representative datasets to ensure the models developed are as unbiased as possible.

The models developed in this shared task could potentially be misused by malicious actors to create more sophisticated propaganda or to target specific individuals or groups. Therefore, we ask developers to be aware of this issue while deploying models.

Acknowledgments

The work of M. Hasanain, F. Ahmed, R. Suwaileh, W. Zaghouani and F. Alam is partially supported by NPRP 14C-0916-210015 from the Qatar National Research Fund, which is a part of Qatar Research Development and Innovation Council (QRDI). The findings achieved herein are solely the responsibility of the authors.

References

- Muhammad Abdul-Mageed, AbdelRahim Elmadany, and El Moatez Billah Nagoudi. 2021. **ARBERT & MARBERT: Deep bidirectional transformers for Arabic**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7088–7105, Online. Association for Computational Linguistics.
- Kemal Ofrazier Abrar Abir. 2024. Nullpointer at ArAIEval Shared Task: Arabic propagandist technique detection with token-to-word mapping in sequence tagging. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Firoj Alam, Stefano Cresci, Tanmoy Chakraborty, Fabrizio Silvestri, Dimiter Dimitrov, Giovanni Da San Martino, Shaden Shaar, Hamed Firooz, and Preslav Nakov. 2022a. **A survey on multimodal disinformation detection**. In *Proceedings of the 29th International Conference on Computational Linguistics*,

⁵<https://github.com/openai/CLIP>

- pages 6625–6643, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Firoj Alam, Abul Hasnat, Fatema Ahmed, Md Arid Hasan, and Maram Hasanain. 2024. ArMeme: Propagandistic content in arabic memes. *arXiv:2406.03916*.
- Firoj Alam, Hamdy Mubarak, Wajdi Zaghouani, Preslav Nakov, and Giovanni Da San Martino. 2022b. Overview of the WANLP 2022 shared task on propaganda detection in Arabic. In *Proceedings of the Seventh Arabic Natural Language Processing Workshop, WANLP '22*, Abu Dhabi, UAE.
- Yasser Alhabashi, Serry Sibae, Abdullah Alharbi, Omar Najar, and Samar Ahmed. 2024. ASOS at ArAIEval Shared Task: persuasion techniques detection: an interdisciplinary approach to identifying and counteracting manipulative strategies. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Zien Sheikh Ali, Watheq Mansour, Tamer Elsayed, and Abdulaziz Al-Ali. 2021. AraFacts: the first large arabic dataset of naturally occurring claims. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 231–236.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. AraBERT: Transformer-based model for arabic language understanding. In *LREC 2020 Workshop Language Resources and Evaluation Conference 11–16 May 2020*, page 9.
- A Barrón-Cedeno et al. Overview of the CLEF–2022 CheckThat! Lab on fighting the covid-19 infodemic and fake news detection. *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, page 495.
- Md. Rafiul Biswas, Zubair Shah, and Wajdi Zaghouani. 2024. MemeMind at ArAIEval Shared Task: spotting persuasive spans in arabic text with persuasion techniques identification. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Minghai Chen, Sen Wang, Paul Pu Liang, Tadas Baltrušaitis, Amir Zadeh, and Louis-Philippe Morency. 2017. Multimodal sentiment analysis with word-level fusion and reinforcement learning. In *Proceedings of the 19th ACM international conference on multimodal interaction*, pages 163–171.
- Pengyuan Chen, Lei Zhao, Yangheran Piao, Hongwei Ding, and Xiaohui Cui. 2023. Multimodal visual-textual object graph attention network for propaganda detection in memes. *Multimedia Tools and Applications*, pages 1–16.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news articles. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, EMNLP-IJCNLP 2019, Hong Kong, China.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT '19*, pages 4171–4186, Minneapolis, MN, USA. Association for Computational Linguistics.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. SemEval-2021 Task 6: Detection of persuasion techniques in texts and images. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 70–98, Online. Association for Computational Linguistics.
- Taosong Fang, Bo Xu, and Linlin Zong. 2022a. Emotion-enhanced multi-modal persuasive techniques detection using split features. In *Fuzzy Systems and Data Mining VIII*, pages 172–180. IOS Press.
- Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. 2022b. EVA: Exploring the limits of masked visual representation learning at scale. *arXiv preprint arXiv:2211.07636*.
- Mary Fouad and Julie Weeds. 2024. SussexAI at ArAIEval Shared Task: mitigating class imbalance in arabic propaganda detection. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Ivan Habernal and Iryna Gurevych. 2015. Exploiting debate portals for semi-supervised argumentation mining in user-generated web discourse. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2127–2137, Lisbon, Portugal. Association for Computational Linguistics.
- Abdelhamid Haouhat, Hadda Cherroun, Slimane Bellaouar, and Attia Nehar. 2024. MODOS at ArAIEval Shared Task: Multimodal propagandistic memes classification using weighted sam, clip and arabiangpt. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Maram Hasanain, Fatema Ahmed, and Firoj Alam. 2023a. Large language models for propaganda span annotation. *arXiv preprint arXiv:2311.09812*.

- Maram Hasanain, Fatema Ahmed, and Firoj Alam. 2024. Can GPT-4 identify propaganda? annotation and detection of propaganda spans in news articles. In *Proceedings of the 2024 Joint International Conference On Computational Linguistics, Language Resources And Evaluation, LREC-COLING 2024*.
- Maram Hasanain, Firoj Alam, Hamdy Mubarak, Samir Abdaljalil, Wajdi Zaghouni, Preslav Nakov, Giovanni Da San Martino, and Abed Freihat. 2023b. [ArAIEval Shared Task: Persuasion techniques and disinformation detection in Arabic text](#). In *Proceedings of ArabicNLP 2023*, pages 483–493, Singapore (Hybrid). Association for Computational Linguistics.
- Christopher Hidey, Elena Musi, Alyssa Hwang, Smaranda Muresan, and Kathy McKeown. 2017. Analyzing the semantic types of claims and premises in an online persuasive forum. In *Proceedings of the 4th Workshop on Argument Mining*. Columbia Univ., New York, NY (United States).
- Go Inoue, Bashar Alhafni, Nurpeiis Baimukan, Houda Bouamor, and Nizar Habash. 2021. The interplay of variant, size, and task type in Arabic pre-trained language models. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 92–104.
- IPA. 1938. In Institute for Propaganda Analysis, editor, *Propaganda Analysis. Volume I of the Publications of the Institute for Propaganda Analysis*. New York, NY, USA.
- Momtazul Labib, Samia Rahman, Hasan Murad, and Uday Das. 2024. CUET_sstm at ArAIEval Shared Task: unimodal (text) propagandistic technique detection using transformer-based model. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Wuwei Lan, Yang Chen, Wei Xu, and Alan Ritter. 2020. Gigabert: Zero-shot transfer learning from english to arabic. In *Proceedings of The 2020 Conference on Empirical Methods on Natural Language Processing (EMNLP)*.
- Clyde R. Miller. 1939. The Techniques of Propaganda. From “How to Detect and Analyze Propaganda,” an address given at Town Hall. The Center for learning.
- Fred Morstatter, Yunqiu Shao, Aram Galstyan, and Shanika Karunasekera. 2018. From alt-right to alt-rechts: Twitter analysis of the 2017 german federal election. In *Companion Proceedings of the The Web Conference 2018*, pages 621–628.
- Behnaz Nojavanasghari, Deepak Gopinath, Jayanth Koushik, Tadas Baltrušaitis, and Louis-Philippe Morency. 2016. Deep multimodal fusion for persuasiveness prediction. In *Proceedings of the 18th ACM international conference on multimodal interaction*, pages 284–288.
- OpenAI. 2023. [GPT-4 technical report](#). Technical report, OpenAI.
- Sunghyun Park, Han Suk Shim, Moitreyia Chatterjee, Kenji Sagae, and Louis-Philippe Morency. 2016. Multimodal analysis and prediction of persuasiveness in online social multimedia. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 6(3):1–25.
- Andrew Perrin. 2015. Social media usage. *Pew research center*, pages 52–68.
- 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 multilingual setup](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361, Toronto, Canada. Association for Computational Linguistics.
- Md Abdur Razzaq Riyadh and Sara Nabhani. 2024. Mela at ArAIEval Shared Task: propagandistic techniques detection in arabic with a multilingual approach. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Ali Safaya, Moutasem Abdullatif, and Deniz Yuret. 2020. [KUISAIL at SemEval-2020 task 12: BERT-CNN for offensive speech identification in social media](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 2054–2059, Barcelona (online). International Committee for Computational Linguistics.
- Uzair Shah, Md. Rafiul Biswas, Marco Agus, Mowafa Househ, and Wajdi Zaghouni. 2024. MemeMind at ArAIEval Shared Task: generative augmentation and feature fusion for multimodal propaganda detection in arabic memes through advanced language and vision models. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.
- Shivam Sharma, Firoj Alam, Md. Shad Akhtar, Dimitar Dimitrov, Giovanni Da San Martino, Hamed Firooz, Alon Halevy, Fabrizio Silvestri, Preslav Nakov, and Tanmoy Chakraborty. 2022. [Detecting and understanding harmful memes: A survey](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI ’22*, pages 5597–5606, Vienna, Austria. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Symom Hossain Shohan, Md. Sajjad Hossain, Ashraf Islam Paran, Shawly Ahsan, Jawad Hossain, and Mohammed Moshikul Hoque. 2024. SemanticCuet-Sync at ArAIEval Shared Task: finetuning transformer models for propagandistic technique detection and span prediction. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.

Yeshan Wang and Ilia Markov. 2024. CLTL at ArAIEval Shared Task: multimodal propagandistic memes classification using transformer models. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.

Weihao Yu, Chenyang Si, Pan Zhou, Mi Luo, Yichen Zhou, Jiashi Feng, Shuicheng Yan, and Xinchao Wang. 2023. Metaformer baselines for vision. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (01):1–17.

Mohamed Zaytoon, Nagwa ElMakky, and Marwan Torki. 2024. AlexUNLP-MZ at ArAIEval Shared Task: contrastive learning, llm features extraction and multi-objective optimization for arabic multi-modal meme propaganda detection. In *Proceedings of The Second Arabic Natural Language Processing Conference (ArabicNLP 2024)*, Bangkok. Association for Computational Linguistics.