wangkongqiang@EEUCA 2026: Multimodal Identification of Vaccine Critical Content on Social Media

Kongqiang Wang
School of Information Science and Engineering, Yunnan University,
Yunnan Baiyao Street, 650500,
Kunming, Yunnan, China.
wangkongqiang60@gmail.com

Peng Zhang
School of Information Science and Engineering, Yunnan University,
Yunnan Baiyao Street, 650500,
Kunming, Yunnan, China.
zpp1219@gmail.com

Qingli Tan
College of Ecology and Environment,
Yunnan University,
Yunnan Baiyao Street, 650500,
Kunming, Yunnan, China.
tanqingli@stu.ynu.edu.cn

Abstract

Our team was interested in content classification and labeling from multimodal meme detection of vaccine critical content on social media. We joined the shared task on Multimodal Identification of Vaccine Critical Content on Social Media@EEUCA with ACL 2026. In this task, our goal is to assign a content classification label to vaccine-related discourse (e.g., Vaccine critical, Neutral, Pro-vaccine). The objective is to develop systems that can classify the intent of a vaccine-related meme. The dataset for this task will have three labels: Vaccine critical (0), Neutral (1), and Pro-vaccine (2). The performance will be ranked by F1-score (Macro). This shared task is based on the *VaxMeme* dataset, a collection of over 10,000 manually annotated vaccination-related memes, designed to support multimodal vaccine-critical meme detection. Our group used a supervised learning method on finetuning pre-trained models and Large Language Model (LLM), including Qwen2 LLMs and Llama series LLMs based on Llama-Factory. The best result on the test set for shared task were Macro F1 score of 0.8153, Accuracy 0.8185, Precision (Macro) 0.8151, and Recall (Macro) 0.8159 from finetuning qwen2_1.5B LLM method, ranking 12th among all teams. The complete code of this entire project can be found at our GitHub address[1].

1 Introduction

First of all, let's introduce the overview of shared task on Multimodal Identification of Vaccine Critical Content on Social Media@EEUCA with ACL 2026 (Thapa et al., 2026b). Memes have become a powerful and fast-spreading medium for sharing information online, especially around high-impact public health issues such as COVID-19 vaccination. While memes can be used to promote awareness and positive behavior, they are also frequently used to spread misinformation (Thapa et al., 2026a), skepticism (Thapa et al., 2024b), and vaccine-critical narratives, often through sarcasm and implicit context that make automated analysis challenging. In this context, the distinction between vaccine critical and pro-vaccine becomes blurred, as vaccination-related images straddle the line between satire and offense, challenging researchers and platforms alike to navigate the complexities of memes content moderation. As one label generally fails to encompass multiple aspects of linguistics, this shared task classifies memes on three aspects: Vaccine critical (0), Neutral (1), and Pro-vaccine (2).

This shared task is based on the *VaxMeme* dataset (Naseem et al., 2023), a collection of over 10,000 manually annotated vaccination-related memes, designed to support multimodal vaccine-critical meme detection. The task invites participants to develop models that jointly leverage both visual and textual representations to capture the global and local contextual cues embedded in memes. By focusing on fine-grained multimodal understanding, this challenge aims to advance more reliable systems for monitoring vaccine-related discourse, supporting myth debunking efforts, and informing the design of effective public health communication strategies on social media platforms.

2 Background

2.1 Content Detection for Vaccine Critical Posts

Majority of research (Wang et al., 2020) on identifying vaccine critical posts on social media has mainly focused on textual content. (Zhang et al., 2020) presented three models for analysing public sentiment on HPV vaccines on Twitter using

[1] https://github.com/WangKongQiang/EEUCA2026_Multimodal_Identification_of_Vaccine_Critical_Content_on_Social_Media

Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications, pages 112–121
July 3, 2026 ©2026 Association for Computational Linguistics

transfer learning. They fine-tuned bidirectional encoder representations from Transformers (BERT) (Devlin et al., 2019), and their results demonstrated the effectiveness of the proposed framework, which also aided in the discovery of vaccine uptake strategies. Recently, (Naseem et al., 2021) categorised vaccine-related tweeter posts using word representation from the domain-specific context with common knowledge and sentiment data. Their proposed method outperformed several traditional and recent transformer-based pre-trained language models. Previously published architectures, however, only focus on local semantic word representations using a sliding window for textual content. However, long-range and non-consecutive semantic links among feature representation words are required to capture global characteristics. We address this limitation by using a graph-based method to capture both local and global features of textual content.

Previously research has examined the use of multimodal content for detecting hateful memes (Lee et al., 2021), misleading information (Volkova et al., 2019), antisemitism (Chandra et al., 2021), and fake news detection (Wang et al., 2018). Experiments conducted using unimodal and multimodal in previous studies showed that understanding both modalities is essential for detection. Limited research has explored multimodal data to identify vaccine critical memes on social media. Recently, (Wang et al., 2020) created a multimodal dataset from Instagram posts and presented a multimodal framework with semantic and task-level attention to identifying vaccine critical information on social media. In contrast, our work jointly learns global and local representations of the textual and visual content of memes, which provide complementary information to improve the identification of vaccine critical memes on Twitter. We suggest that releasing a robustly annotated dataset to the community will support further advances and benchmarking of methods in this space.

2.2 Vaccine and Multimodal Datasets

Social media is a valuable source of information and has been widely used for various tasks like health mention classification (Naseem et al., 2022c), identifying suicide (Naseem et al., 2022b) and depression (Naseem et al., 2022a) and others. Systematic reviews show the wide range of applications for classifying user-generated content for vaccine hesitancy on social media, such as infectious diseases and outbreaks such as human papillomavirus, measles Influenza, mining misinformation mining.

Only two multimodal datasets are used in the previous studies to identify vaccine critical information on social media. The first of them was presented by (Wang et al., 2020), where authors used Instagram posts with text and visual content collected from January 2016 to October 2019 to identify vaccine critical information on Instagram posts. MMCoVaR (Chen et al., 2021), a multimodal COVID-19 vaccine focused data repository is the second dataset. MMCoVaR comprises 2,593 articles and 24,184 tweets from February 2020 to March 2021 and is limited to only COVID vaccine related posts. Both mentioned datasets are not publicly available, whereas we make our dataset publicly available for further research.

3 Dataset

In this section, we describe various aspects of task dataset including data collection, meme annotation, and dataset statistics. Task dataset comprises 10,244 memes that encompass different intent content relevant to the vaccination.

3.1 Utterance Annotation

Each meme was labeled to one of 3 labels: Vaccine critical (0), Neutral (1), and Pro-vaccine (3). Neutral if meme was not concept of vaccination. meme annotation for each label are mentioned below.

Vaccine critical: A meme (text or image or both) criticises vaccines, contains vaccine misinformation about vaccine side effects, vaccine conspiracy theories, and cases or statistical conclusions against vaccines.

Neutral: A meme (text or image or both) reports the events or others' opinions objectively related to vaccines, such as talking about rights of people related to vaccines, or news or statistical charts about vaccines showing no content in favor or against vaccines.

Pro-vaccine: A meme (text or image or both) contains a content in favor of vaccines, advising people to get vaccinated, a content about any event or place that is open only for vaccinated people or promoting and selling products with slogans in favor of vaccines.

Table 1: Dataset statistics for *VaxMeme*. The data consists of 10,244 samples for the multimodal identification task of vaccine critical memes on Twitter.

Data	Number of Pro-Vaccine	Number of Vaccine critical	Number of Neutral	Total
Full	3983	3441	2820	10244
Timeline	**Number of Pro-Vaccine**	**Number of Vaccine critical**	**Number of Neutral**	**Total**
T1	452	1679	1027	3158
T2	1040	747	1062	2849
T3	2491	1015	731	4237

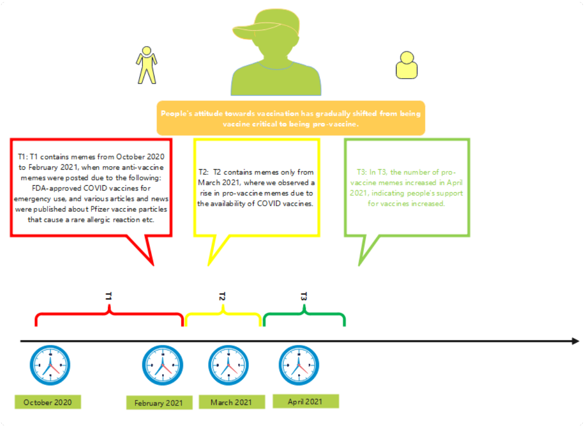

Figure 1: Timelines of the general vaccine critical memes gradually shifted in intention label.

3.2 Dataset Statistics

Table 1 (*Upper*) provides the class distribution of intent across the 10,244 English memes, and Table 1 (*Down*) provides the different timelines dividing distribution across all memes. Most memes are Pro-Vaccine or Vaccine critical in nature and a approximately data balance is present. However, this is in line with real world data distributions, where different people have different views on getting vaccinated. Figure 1 illustrates the timelines for vaccine critical memes.

4 System Overview

4.1 Fine-tuning Pre-trained Models

Introduction. In recent years, with the rapid development of deep learning technology, large-scale pre-trained models have achieved remarkable results in fields such as natural language processing, computer vision, and multimodal learning. Compared with traditional models trained from scratch, pre-trained models can learn rich semantic representations and general knowledge by pre-training on large-scale general corpora, thereby significantly improving the performance and training efficiency of downstream tasks. However, the knowledge learned by pre-trained models on general corpora often has strong generalization, while specific tasks usually have obvious domain characteristics. Therefore, directly applying pre-trained models to downstream tasks often fails to achieve the best results. To solve this problem, researchers usually adopt the fine-tuning strategy, that is, further optimize the model parameters using the data of specific tasks on the basis of the pre-trained model, so that it can better adapt to the target task. In this study, to enhance the model's performance in the task of vaccine critical content analysis on social media, a pre-trained language model was adopted as the base model and fine-tuned in combination with specific task data (*VaxMeme* dataset), enabling the model to effectively learn the semantic relationship between vaccination-related meme expressions and their underlying intents.

The classifier in the pre-trained model uses a transformer based classifier. The specific pre-trained models of the classifier are shown in the Table 2.

Table 2: pre-trained model classifier structure for vaccine critical content classification.

Model	Batch Size	Num Epochs	Learning Rate
alber/albert-base-v2	4	5	2e-5
google-bert/bert-base-uncased	4	5	2e-5
nghuyong/ernie-2.0-large-en	4	5	1e-5
nghuyong/ernie-1.0-base-zh	4	5	1e-5
FacebookAI/roberta-large-mnli	4	5	2e-5
cambridgeltl/trans-encoder-bi-simcse-roberta-large	4	5	1e-5

The Principle of Fine-tuning Pre-trained Models. Pre-trained models typically use large-scale corpora for self-supervised learning, such as tasks like language model prediction, masked language modeling, or autoregressive modeling, thereby learning common language representations. After pre-training, the model parameters already contain a large amount of language knowledge and semantic information. Therefore, in downstream tasks, only a small amount of labeled data is needed to achieve good performance.

The basic idea of fine-tuning is to introduce supervisory signals from downstream tasks on the basis of the parameters of the pre-trained model and further optimize the model parameters through the gradient descent algorithm. Let the parameters of the pre-trained model be θ, Given the downstream task training dataset $\mathcal{D} = (x_i, y_i)_{i=1}^{N}$, Where x_i represents the input text and y_i represents the corresponding label, then the model training objective can be expressed as:

$$\mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \ell\left(f(x_i;\theta), y_i\right) \tag{1}$$

Here, $f(\cdot)$ represents the model prediction function, and $\ell(\cdot)$ is the loss function (such as cross-entropy loss). By minimizing this loss function, the model can gradually adapt to the data distribution of a specific task, thereby enhancing the prediction performance.

In practical applications, fine-tuning typically includes the following two methods:

- Full Fine-tuning: Update all the parameters of the pre-trained model to enable it to fully adapt to the target task.
- Parameter-efficient Fine-tuning: Only update some parameters or introduce additional lightweight modules to reduce training costs, such as Adapter, LoRA and other methods.

In this study, based on the characteristics of the task and the availability of computing resources, the pre-trained model was trained using a parameter-efficient fine-tuning strategy.

Input Data Construction. First of all, the original data needs to be converted into an input format that the model can handle. For text tasks, the following steps are usually required:

- Text Cleaning and Preprocessing: Remove irrelevant symbols or abnormal characters;
- Word Segmentation and Encoding: Use the tokenizer corresponding to the pre-trained model to convert the text into a token sequence;
- Input Sequence Construction: Ultimately, the input text will be represented as a sequence of token ids and input into the pre-trained model for feature encoding.

Task Structure Design. During the fine-tuning process, it is necessary to design the corresponding prediction structure based on the specific task. For instance, in the task of vaccine critical memes intent analysis on social media, the model needs to simultaneously identify vaccination-related discourse intent categories as well as the corresponding memes content information. Therefore, a model is usually composed of the following parts:

- Pre-trained Encoding Layer: Used for extracting semantic representations of text;
- Task-specific Layer: For example, the classification layer or the sequence labeling layer;

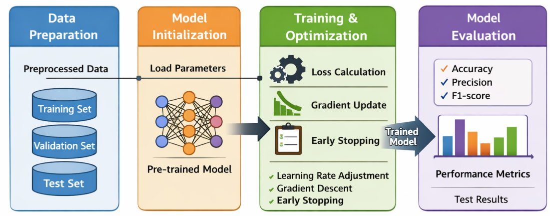

Figure 2: The framework diagram of the fine-tuning pre-trained classification model.

- Output Layer: Generate the final prediction result.

By adding task-related structures at the top of the pre-trained model, the model's adaptability to specific tasks can be effectively enhanced.

The overall architecture diagram of the fine-tuning pre-trained model is shown in the Figure 2.

Design of Loss Function. During the training process, it is necessary to select an appropriate loss function based on the type of task. For classification tasks, the cross-entropy loss function is usually adopted.

$$\mathcal{L} = -\sum_{i=1}^{C} y_i \log(p_i) \tag{2}$$

Here, C represents the number of categories, y_i is the true label, and p_i is the model prediction probability. By minimizing the loss function, the consistency between the model's prediction results and the true labels can be gradually improved.

4.2 Hard Voting Mechanism

The hard voting mechanism is a common model fusion strategy in ensemble learning, mainly used for classification tasks. Its basic idea is: multiple base learners make predictions on the same sample respectively, and then determine the final prediction category through majority voting.

The Principle of Hard Voting Mechanism. Assume that the ensemble model consists of M base classifiers: $h_1(x), h_2(x), \ldots, h_M(x)$, where $h_i(x)$ denotes the prediction of the i-th classifier for input sample x. The final prediction of the hard voting ensemble is determined by majority voting:

$$\hat{y} = \arg\max_{c \in C} \sum_{i=1}^{M} I(h_i(x) = c) \tag{3}$$

where C represents the set of all possible classes and $I(\cdot)$ is an indicator function defined as:

$$I(h_i(x) = c) = \begin{cases} 1, & \text{if } h_i(x) = c \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

In a weighted hard voting scheme, each classifier is assigned a weight w_i, and the final prediction can be written as:

$$\hat{y} = \arg\max_{c \in C} \sum_{i=1}^{M} w_i \cdot I(h_i(x) = c) \quad (5)$$

In our experiment, the weights of each classifier were the same.

4.3 Fine-tuning of the Qwen2 and Llama series Large Language Model (LLM)

Qwen2 is an open-source large language model (LLM) developed by the Tongyi Qianwen team and created by Alibaba Cloud's Tongyi Lab. Using Qwen2 as the base large language model (LLM) and achieving high-accuracy text classification through instruction fine-tuning is an introductory task for learning the fine-tuning of large language models (LLMs). The Llama-2 7B (Touvron et al., 2023b) model approximately 7 billion parameters in an open-source large language model (LLM) released by Meta in 2023, belongs to a typical Transformer decoder architecture. Parameter scale is approximately 7B; Context length is approximately 4K tokens; Architecture is Standard Transformer + Multi-Head Attention (MHA). Features: Mature structure and stable reasoning; Suitable for deployment in resource-constrained environments; It is mostly used for basic NLP tasks and lightweight applications. The Llama-3 8B (Grattafiori et al., 2024) model of is a new generation version released in 2024, featuring significant upgrades in both architecture and training data. Parameter scale is approximately 8B; Context length is approximately 8K tokens (longer context); Architecture improvement based on enhance efficiency by using Grouped Query Attention (GQA); Tokenizer based on word list from 32K to 128K (stronger expressive power); The training data scale has significantly expanded (about 7 times).

Instruction fine-tuning is a process of further training an LLMs on a dataset composed of (instruction, input, output) combine pairs. Among them, the instructions represent the human instructions of the model, the input represent the raw data content from specific dataset, and the output represents the expected output that follows the instructions. This process helps bridge the gap between the next word prediction target of LLMs and the goal of users to have LLMs follow human instructions.

In this vaccine-related behavioral intent classification task on social media, we will use the Qwen2-1.5B-Instruct [2] and Qwen2-7B-Instruct [3] model to perform instruction fine-tuning tasks on the dataset, while using SwanLab [4] for monitoring and visualization. Llama3-8B, while maintaining a lightweight scale, has achieved a significant performance improvement over Llama2-7B through a larger vocabulary, longer context window, and GQA attention mechanism, making it a strong competitor among current open-source small-parameter models. The fine-tuning of the Llama series models (Touvron et al., 2023a) is mainly carried out using the Llama-Factory tool [5]. The following presents three demonstration formatted data samples for fine-tuning LLM data in the train dataset. Our training task is to ensure that the fine-tuned large language model (LLM) can predict the correct output based on the prompt words composed of post_text, image_text and selectable types.

The complete process of fine-tuning the Qwen2 and Llama series large language model (LLM) using the train dataset and conducting model inference on validation or test dataset with the fine-tuned large language model (LLM), as shown in Figure 3.

5 Results and Analysis

For results obtained by our fine-tuning pre-trained models and fine-tuning Qwen2 or Llama series LLM methods on the train dataset, and implement model inference on validation dataset and test dataset are shown in Table 3 and Table 4 respectively. roberta-RNN indicates the addition of a layer of recurrent neural network (RNN) after *FacebookAI/roberta-large-mnli* model, which is the LSTM layer. roberta-cnn indicates adding a convolutional neural network (CNN) layer after *FacebookAI/roberta-large-mnli*, which is the Conv2d layer. bagging refers to the model ensemble decision-making method that employs hard vot-

[2] https://huggingface.co/Qwen/Qwen2-1.5B-Instruct
[3] https://huggingface.co/Qwen/Qwen2-7B-Instruct
[4] https://swanlab.cn
[5] GitHub: https://github.com/hiyouga/LlamaFactory

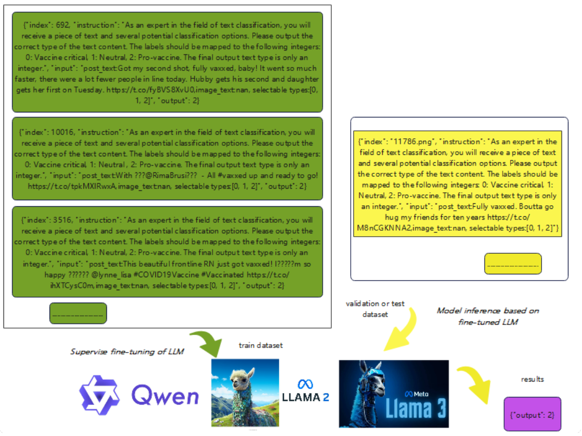

Figure 3: The framework diagram of the fine-tuning qwen2 and Llama series classification LLM and model inference.

ing mechanism in the four models: roberta, roberta-RNN, simcse-roberta-lstm, roberta-lstm-gru.

Table 3: The results obtained by our fine-tuning pre-trained models and fine-tuning Qwen2 or Llama series LLM methods for vaccine-related discourse intent classification task on the validation dataset.

Pre-trained Model	**Recall (Macro)**	**Precision (Macro)**	**F1 (Macro)**	**Accuracy**
albert/albert-base-v2	-	-	-	-
google-bert/bert-base-uncased	-	-	-	-
nghuyong/ernie-2.0-large-en	-	-	-	-
nghuyong/ernie-1.0-base-zh	-	-	-	-
FacebookAI/roberta-large-mnli	-	-	-	-
cambridgeltl/trans-encoder-bi-simcse-roberta-large	-	-	-	-
roberta-RNN	-	-	-	-
roberta-cnn	-	-	-	-
roberta-lstm-gru	-	-	-	-
simcse-roberta-lstm	-	-	-	-
bagging	-	-	-	-
Large Language Model	**Recall (Macro)**	**Precision (Macro)**	**F1 (Macro)**	**Accuracy**
qwen2_1.5B	0.7945	0.7948	0.7943	0.7969
qwen2_7B	-	-	-	-
Llama2_7B	-	-	-	-
Llama3_8B	-	-	-	-

Table 4: The results obtained by our fine-tuning pre-trained models and fine-tuning Qwen2 or Llama series LLM methods for vaccine-related discourse intent classification task on the test set.

Model	**Recall (Macro)**	**Precision (Macro)**	**F1 (Macro)**	**Accuracy**
albert/albert-base-v2	0.7618	0.7916	0.7611	0.76
google-bert/bert-base-uncased	0.763	0.7734	0.7649	0.7727
nghuyong/ernie-2.0-large-en	0.7918	0.8016	0.7915	0.7941
nghuyong/ernie-1.0-base-zh	0.756	0.7652	0.7552	0.76
FacebookAI/roberta-large-mnli	0.8029	0.808	0.8025	0.8049
cambridgeltl/trans-encoder-bi-simcse-roberta-large	0.7933	0.7986	0.7935	0.7971
roberta-RNN	0.7985	0.8019	0.7971	0.799
roberta-cnn	0.7982	0.8	0.7988	0.8049
roberta-lstm-gru	0.7945	0.8076	0.7928	0.7922
simcse-roberta-lstm	0.804	0.8047	0.8036	0.8078
bagging	0.804	0.809	0.8025	0.8039
Large Language Model	**Recall (Macro)**	**Precision (Macro)**	**F1 (Macro)**	**Accuracy**
qwen2_1.5B	0.8159	0.8151	0.8153	0.8185
qwen2_7B	0.8135	0.815	0.8134	0.8166
Llama2_7B	0.8006	0.8026	0.7998	0.8029
Llama3_8B	0.8051	0.8044	0.8046	0.8098

6 Discussion

It is highly unusual that Qwen2-1.5B outperformed Llama-3-8B. A plausible explanation for this anomaly is that the 8B model may have been more sensitive to hyperparameter settings, such as learning rate or regularization, leading to suboptimal fine-tuning or mild overfitting on the training data. In contrast, the smaller 1.5B model could have benefited from better generalization under the same setup. Additionally, differences in pretraining data quality and alignment strategies between Qwen and Llama models may also have contributed to the performance gap. For shared task on Multimodal Identification of Vaccine Critical Content on Social Media@EEUCA with ACL 2026, we referred to the relevant tasks of CASE 2025 (Hurriyetoglu et al., 2025), CASE 2024 (Thapa et al., 2024a) and CASE 2023 (Thapa et al., 2023) shared tasks on multimodal hate speech detection and derived our own method. Although the effect of the experiment needs to be strengthened. However, these contents and ideas have given us a lot of inspiration. Vaccine-related discourse intent content analysis is a longstanding tradition of the EEUCA workshop series. We believe that with our further research and more detailed optimization on training of the model, we will achieve even greater success in future competitions.

7 Conclusion

We employed multiple methods in detection of vaccine-related discourse behavioral intent on social media, which respectively involved the transformer pre-trained models and Qwen2 or Llama series LLM. Our final leaderboard is shown in the Table 5. The best result of this task was achieved by fine-tuning the Qwen2_1.5B large language model (LLM) and conducting inference on the test set.

Table 5: The final leaderboard of shared task on Multimodal Identification of Vaccine Critical Content on Social Media@EEUCA with ACL 2026.

#	Username	Recall (Macro)	Precision (Macro)	F1 (Macro)	Accuracy
1	lili12-637947	0.8517	0.8494	0.8494	0.8517
2	wangxiuxian-637268	0.8409	0.8386	0.8389	0.842
3	rishta_19-611897	0.8359	0.8383	0.8357	0.839
4	_alexcristea-636983	0.8351	0.8338	0.834	0.838
5	sumaiya_110-594217	0.834	0.8345	0.8332	0.8361
6	anchy-637928	0.8309	0.8309	0.8308	0.8341
7	myname-637930	0.8309	0.8309	0.8308	0.8341
8	quasar-637336	0.8324	0.8331	0.8306	0.8322
9	wenbin-634065	0.8218	0.8205	0.8205	0.8244
10	naturia_beast-636958	0.8209	0.8212	0.8201	0.8244
11	vinaybabu-637935	0.819	0.8216	0.8184	0.8215
12	**wangkongqiang-495416**	**0.8159**	**0.8151**	**0.8153**	**0.8185**
13	ratpier-637076	0.8161	0.817	0.815	0.8176
14	yjwong1999-494691	0.8141	0.8189	0.8122	0.8137
15	linus-637363	0.8123	0.8106	0.8105	0.8137
16	havnis-636808	0.8083	0.808	0.8067	0.8117
17	alishba-wazir-604227	0.8071	0.8132	0.8067	0.8088
18	zmin123-553584	0.8013	0.8005	0.7997	0.8039
19	lin123-637530	0.8007	0.7992	0.7994	0.8039
20	barkion-636765	0.7986	0.808	0.7976	0.799
21	merrli-636903	0.7982	0.8058 (19)	0.7972	0.799
22	exterio-636705	0.7846	0.7964	0.7861	0.7912
23	abs123-504332	0.7864	0.7868	0.7846	0.7912
24	thatgrass-519137	0.7802	0.7858	0.7754	0.7844
25	kannanrrk-615633	0.7437	0.7435	0.7436	0.7502

8 Limitations of the Work

we are interested in learning about LLMs in computational social science (Thapa et al., 2025), our paper mainly focuses on making discussions on vaccine-related discourse for this meme behavioral intent classification task. This is because we are quite interested in and good at identifying hate (Bhandari et al., 2023) and offense categories in the text (Parihar et al., 2021). Due to our lack of utilization of context features, we are unable to make good use of the image content in the train dataset of this sharing task. we have chosen the 7B version of Qwen2 due to the limited computing resources. If we could use a larger language model with more parameters, we would achieve better prediction results. These are all our future tasks.

Acknowledgments

We are very grateful to the organizers of the Shared Task on Multimodal Identification of Vaccine Critical Content on Social Media@EEUCA with ACL 2026 (Hürriyetoğlu et al., 2026) and the School of Information Science and Engineering of Yunnan University for providing the experimental environment and equipment.

References

Aashish Bhandari, Siddhant B Shah, Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023. Crisishatemm: Multimodal analysis of directed and undirected hate speech in text-embedded images from russia-ukraine conflict. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1994–2003.

Mohit Chandra, Dheeraj Pailla, Himanshu Bhatia, Aadilmehdi Sanchawala, Manish Gupta, Manish Shrivastava, and Ponnurangam Kumaraguru. 2021. "subverting the jewtocracy": Online antisemitism detection using multimodal deep learning. In *13th ACM Web Science Conference*, page 148–157.

Mingxuan Chen, Xinqiao Chu, and KP Subbalakshmi. 2021. Mmcovar: multi-modal covid-19 vaccine focused data repository for fake news detection anda baseline architecture for classification. In *Proceedings of the 2021 IEEE/ACMInternational Conference on Advances in Social Networks Analysis and Mining*, pages 31–38.

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 Associationfor Computational Linguistics: Human Language Technologies*, volume 1 (Long andShort Papers), pages 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias

Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. The llama 3 herd of models.

Ali Hurriyetoglu, Surendrabikram Thapa, Hristo Tanev, and Surabhi Adhikari. 2025. Findings and insights from the 8th workshop on challenges and applications of automated extraction of socio-political events from text. In *Proceedings of the 8th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Texts*, pages 1–5, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Ali Hürriyetoğlu, Surendrabikram Thapa, Hristo Tanev, Laxmi Thapa, and Surabhi Adhikari. 2026. Overview of the workshop on event extraction and understanding: Challenges and applications. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Roy Ka-Wei Lee, Rui Cao, Ziqing Fan, Jing Jiang, and Wen-Haw Chong. 2021. Disentangling hate in online memes. In *Proceedings of the 29th ACM International Conference on Multimedia*, page 5138–5147.

Usman Naseem, Adam G Dunn, Jinman Kim, and Matloob Khushi. 2022a. Early identification of depression severity levels on reddit using ordinal classification. In *Proceedings of the ACM Web Conference*, pages 2563–2572.

Usman Naseem, Matloob Khushi, Jinman Kim, and Adam G. Dunn. 2021. Classifying vaccine sentiment tweets by modelling domain-specific representation and commonsense knowledge into context-aware attentive gru.

Usman Naseem, Matloob Khushi, Jinman Kim, and Adam G Dunn. 2022b. Hybrid text representation for explainable suicide risk identification on social media. *IEEE transactions on computational social systems*.

Usman Naseem, Jinman Kim, Matloob Khushi, and Adam G Dunn. 2022c. Identification of disease or symptom terms in reddit to improve health mention classification. In *Proceedings of the ACM Web Conference*, pages 2573–2581.

Usman Naseem, Jinman Kim, Matloob Khushi, and Adam G Dunn. 2023. A multimodal framework for the identification of vaccine critical memes on twitter. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 706–714.

Anil Singh Parihar, Surendrabikram Thapa, and Sushruti Mishra. 2021. Hate speech detection using natural language processing: Applications and challenges. In *2021 5th International Conference on Trends in Electronics and Informatics (ICOEI)*, pages 1302–1308. IEEE.

Laxmi Thapa, Aryaman Jain, Lakshmojee Koduru, Surabhi Adhikari, Junaid Rashid, Jungeun Kim, Surendrabikram Thapa, and Usman Naseem. 2026a. Concept-grounded detection of vaccine misinformation in multimodal content using interpretable vision-language models. In *Companion Proceedings of the ACM on Web Conference 2026*.

Surendrabikram Thapa, Farhan Ahmad Jafri, Ali Hürriyetoğlu, Francielle Vargas, Roy Ka Wei Lee, and Usman Naseem. 2023. Multimodal hate speech event detection-shared task 4. In *CASE 2023-Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-Political Events from Text, associated with 14th International Conference on Recent Advances in Natural Language Processing, RANLP 2023*, pages 151–159. Association for Computational Linguistics.

Surendrabikram Thapa, Kritesh Rauniyar, Farhan Ahmad Jafri, Hariram Veeramani, Raghav Jain, Sandesh Jain, Francielle Vargas, Ali Hürriyetoğlu, and Usman Naseem. 2024a. Extended multimodal hate speech event detection during russia-ukraine crisis-shared task at case 2024. In *7th Workshop on Challenges and Applications of Automated Extraction of Socio-Political Events from Text, CASE 2024*, pages 221–228. Association for Computational Linguistics.

Surendrabikram Thapa, Kritesh Rauniyar, Hariram Veeramani, Aditya Shah, Imran Razzak, and Usman Naseem. 2024b. Did you tell a deadly lie? evaluating large language models for health misinformation identification. In *International Conference on Web Information Systems Engineering*, pages 391–405. Springer.

Surendrabikram Thapa, Shuvam Shiwakoti, Sidhant Bikram Shah, Surabhi Adhikari, Hariram Veeramani, Mehwish Nasim, and Usman Naseem. 2025. Large language models (llm) in computational social science: prospects, current state, and challenges. *Social Network Analysis and Mining*, 15(1):1–30.

Surendrabikram Thapa, Shuvam Shiwakoti, Sidhant Bikram Shah, Kritesh Rauniyar, Laxmi Thapa, Surabhi Adhikari, Kristina T. Johnson, Ali Hürriyetoğlu, Hristo Tanev, and Usman Naseem. 2026b. Multimodal identification of vaccine content stance on social media. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. Llama: Open and efficient foundation language models.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. Llama 2: Open foundation and fine-tuned chat models.

Svitlana Volkova, Ellyn Ayton, Dustin L Arendt, Zhuanyi Huang, and Brian Hutchinson. 2019. Explaining multimodal deceptive news prediction model. In *Proceedings of the international AAAI conference on web and social media*, volume 13, pages 659–662.

Yaqing Wang, Fenglong Ma, Zhiwei Jin, Ye Yuan, Guangxu Xun, Kishlay Jha, Lu Su, and Jing Gao. 2018. Eann: Event adversarial neural networks for multi-modal fake news detection. In *Proceedings of the 24th acm sigkdd international conferenceon knowledge discovery data mining*, pages 849–857.

Zuhui Wang, Zhaozheng Yin, and Young Anna Argyris. 2020. Detecting medical misinformation on social media using multimodal deep learning. *IEEE Journalof Biomedical and Health Informatics*, 25(6):2193–2203.

Li Zhang, Haimeng Fan, Chengxia Peng, Guozheng Rao, and Qing Cong. 2020. Sentiment analysis methods for hpv vaccines tweets based on transfer learning. *Healthcare*, 8(307).