

Classify and Label the Content from the Cross-Cultural Misogynistic Meme Detection Using Zero-shot Prompts

Kongqiang Wang^{1*}, Qingli Tan², and Peng Zhang¹

¹School of Information Science and Engineering, Yunnan University, 650500, Kunming, Yunnan, China

²College of Ecology and Environment, Yunnan University, 650500, Kunming, Yunnan, China

*Corresponding author: wangkongqiang60@gmail.com

Abstract

Online misogyny increasingly appears in multimodal formats such as memes. Memes combine text and images to convey humor, sarcasm, and ideology. Misogynistic meaning is often implicit and culturally grounded. A meme interpreted as harmful in one cultural setting may be perceived differently in another. Based on the complex scenario of cross-cultural misogynistic meme detection, we proposed to conduct predictions using large language models (LLMs) through zero-shot prompts in the absence of training datasets. This enables our method to effectively address the key pain points such as the difficulty in annotating the training dataset, the high cost and long time consumption of manually labeling the training dataset labels. We applied this zero-shot prompts method to the Tamil (India) MDMD dataset. The best model results achieved Macro F1-score and Accuracy of 0.628995 and 0.696629 respectively for Task A: Single-culture prediction. Overall contexts including Indian context, Chinese context, and Western context, the best model results achieved Macro F1-score and Accuracy of 0.601143 and 0.695693 respectively for Task B: Cross-cultural prediction. This approach of deliberately not using the training dataset for fine-tuning and training large language models (LLMs) or pre-trained models enables our method to have greater flexibility in real-world scenarios. For further information, the project code and user's guide is available from https://github.com/WangKongQiang/CC-MMD_Grand_Challenge_ICMI_2026.

Index Terms— Multimodal Learning, Zero-shot Prompts, Large Language Models (LLMs), Cross-Cultural Detection, Misogynistic Classification, Meme Detection.

1 Introduction

The Cross-Cultural Misogynistic Meme Detection (CC-MMD) challenge is hosted as an ICMI 2026 Grand Challenge. It benchmarks culturally robust multimodal systems for binary misogyny classification in memes across Indian, Chinese, and Western (English) contexts [1]. Online misogyny increasingly appears in multimodal formats such as memes. Memes combine text and images to convey humor, sarcasm, and ideology [2]. Misogynistic meaning is often implicit and culturally

grounded [3]. A meme interpreted as harmful in one cultural setting may be perceived differently in another.

Current multimodal models achieve strong performance in high-resource settings but often lack cultural grounding. They struggle under cross-lingual and cross-cultural distribution shifts. Prior findings show that cultural background significantly influences misogyny annotation and interpretation. CC-MMD dataset addresses this challenge through multilingual data and systematic cross-cultural annotation. The benchmark explicitly evaluates robustness across cultural partitions rather than a single pooled test set.

This grand challenge invites participants to build culture-aware multimodal systems that generalize reliably under cultural shifts. Based on the setup of this task, we made some special modifications to enhance the performance of our results in zero-shot prompt-based large language models (LLMs). Because the task provided the training set and development set content, we deliberately did not use them for fine-tuning the large language model (LLM) or pre-training model, in order to evaluate cross-cultural misogynistic meme detection in the zero-sample setting. Content moderation systems are often trained in dominant cultural contexts [4]. This risks biased predictions and inconsistent decisions across regions. Misogyny detection requires sensitivity to implicit language, humor, symbolism, and local norms. By explicitly evaluating cultural robustness, this challenge advances research toward more inclusive, reliable, and globally deployable multimodal systems. However, the cost of manually annotating datasets and collecting valuable samples that meet the requirements of a specific cultural background is extremely high and takes a long time [5]. Therefore, zero-sample prediction for cross-cultural misogynistic meme detection is more in line with real-world scenarios and practical needs [6]. We set up the environment based on this experiment, and the main contributions are as follows:

- In the scenario without manually annotated training data, conduct zero-shot prediction of the large language model (LLM) for cross-cultural misogynistic meme detection, making the task more in line with real-world scenarios and practical needs [7].
- This paper evaluated the performance of various mainstream large language models (LLMs) from both Chinese

and the rest of the world under both transcription text settings and multimodal meme settings. Pioneering the latest approach for zero-sample cross-cultural misogynistic meme detection.

2 Related Works

2.1 Cross-Cultural Misogynistic Meme Detection

Misogynistic content increasingly appears in multimodal memes combining images and text, where harmful meanings are often implicit and culturally dependent [8]. Existing misogyny detection methods mainly focus on English and single-culture settings, limiting their robustness under cross-cultural distribution shifts. Recent multimodal models such as CLIP and multimodal large language models (MLLMs) improve meme understanding by jointly modeling visual and textual information, but they still struggle with culturally grounded interpretations. The CC-MMD challenge addresses this issue by evaluating misogyny detection across Indian, Chinese, and Western contexts. In this work, we explore zero-shot prompt-based large language models (LLMs) for cross-cultural misogynistic meme detection without task-specific fine-tuning [9].

2.2 Zero-shot Prompting Engineering

Zero-shot prompting enables large language models (LLMs) to perform downstream tasks without task-specific fine-tuning by leveraging natural language instructions. Recent studies demonstrate that carefully designed prompts can significantly improve reasoning, classification, and multimodal understanding capabilities. In multimodal hate speech and meme detection, prompt engineering helps models better interpret implicit semantics, sarcasm, and contextual information contained in images and text [10]. Compared with supervised learning approaches, zero-shot prompting reduces annotation costs and improves adaptability to unseen domains. Motivated by these advantages, this work employs prompt-based zero-shot inference to evaluate the robustness of large language models (LLMs) [11] in cross-cultural misogynistic meme detection scenarios.

3 Task Information

This task can be divided into two consecutive subtasks: **Task A: Single-culture prediction** and **Task B: Cross-cultural prediction**.

Task A evaluates performance on a single target culture. Authors may choose and submit predictions for any one target culture they want to predict. Target cultures correspond to the culture-specific test partitions in CC-MMD (India, China, Western/Ireland). Input content are meme image (referenced by `image_id`), transcription field, and test partition from one

selected culture. Output content is one binary prediction per meme ($I = \text{misogyny}$, $0 = \text{not-misogyny}$).

Task B evaluates cross-cultural robustness. Authors must submit predictions for all provided target cultures. The official ranking uses the culture-balanced Macro F1 averaged across the three culture partitions. Input content are meme image (referenced by `image_id`), transcription field, and all cultural test partitions. Output content is one prediction per meme for each culture ($I = \text{misogyny}$, $0 = \text{not-misogyny}$).

4 Evaluation Metrics

First, let me clarify the task-specific requirements. For Task A: Single Culture Evaluation, the objective is to predict the label from one specific cultural perspective [12]. The requirements for output is your file should contain the prediction for your chosen perspective. For Task B: Cross-Culture Evaluation, the objective is models must capture multi-perspective cultural nuances. The requirements for output is the model must return all cultural perspective labels relevant to the chosen dataset context within the same file. Note that depending on the dataset, this might include the original culture plus one or two target cultures.

Evaluation is conducted separately for each subtask. All metrics are computed on the official hidden test sets. For Task A: Single-Culture Evaluation. Participants select one cultural partition (India, China, or Western). Evaluation is performed only on the selected partition. **Primary Metric.** Macro F1-score on the selected cultural test set. Macro F1 balances performance across both classes (*misogyny* and *not-misogyny*). **Secondary Metric.** Accuracy on the selected cultural test set. Leaderboard ranking for Task A is based on Macro F1.

For Task B: Cross-Cultural Evaluation. Participants must submit predictions for all cultural partitions. Evaluation measures cross-cultural robustness. **Primary Metric.** Macro F1 is computed independently for each cultural partition: $F1_{India}$, $F1_{China}$, $F1_{Western}$. The final Task B score is the unweighted average: $FinalScore = (F1_{India} + F1_{China} + F1_{Western})/3$. This culture-balanced score ensures systems perform consistently across all cultural contexts. **Secondary Metric.** Accuracy reported per culture. Mean accuracy across all cultures. **Tie-Breaking Rules (Task B).** Higher worst-culture Macro F1 (minimum among the three cultures). If still tied, higher mean accuracy. Important Notes: Test labels are hidden. Scores are computed using the official evaluation script. All submissions must follow the required CSV format.

5 CC-MMD Datasets Analysis

Official organizers present the Cross-Cultural Misogyny Meme Dataset (CC-MMD), the benchmark dataset for this grand challenge containing multimodal memes annotated for misogyny across different cultural contexts. It is multi-lingual, multimodal, and cross-culturally annotated. Each meme contains an image and its associated OCR-extracted

text [13]. CC-MMD integrates memes organized into three culture-specific partitions: **Indian Context**. Tamil and Malayalam memes (MDMD dataset). **Chinese Context**. Chinese memes (CMMD dataset [14]). **Western Context**. English/Ireland memes (MAMI dataset). Each meme is annotated from multiple cultural perspectives, enabling cross-cultural analysis. These partitions are used in Task A and Task B to evaluate single-culture and cross-cultural robustness. The MAMI dataset [15] (Western Context) is focused exclusively on cross-cultural labels and only contains annotations for `target_culture1` and `target_culture2`. The data is provided in CSV format alongside the image files. An abstract example situation is shown in the Table 1 below.

Table 1: Special characters of data format & fields.

image_id	transcription	original_culture	target_culture1	target_culture2
meme_001.jpg	"Example OCR extracted text here"	1	0	1

Note: image_id – Unique meme identifier. *transcription* – OCR-extracted meme text. *original_culture* – Label from source culture annotators. *target_culture1* – Cross-cultural annotation. *target_culture2* – Cross-cultural annotation. Labels are binary: 1 = misogyny and 0 = not-misogyny.

The dataset is divided into training, development, and test sets. **Training Set**. Provided with labels (and images) for model development. **Development Set**. Used for validation and tuning. **Test Set**. Labels withheld for official evaluation. Participants must use the provided splits and may not modify them. The number of images and transcriptions in the Train, Dev, and Test sets of the MDMD - Tamil (India) Dataset [16], which was mainly used for our experiment, is shown in Table 2.

Table 2: The number of images and transcriptions splits in the Train, Dev, and Test sets of the MDMD - Tamil (India) Dataset.

Dataset - Language (Cultural Context)	Data Splits	Transcriptions/Text Length	Original_Labels	Irish_Labels	Chinese_Labels
MDMD - Tamil (India)	Train Set	145,41,921-11	273,863	268,868	242,894
	Dev Set	137,65,996-17	75,209	69,215	70,214
	Test Set	145,09,822-34	100,256	88,268	85,271

Note: The transcription text length is indicated in the format of average length-maximum length-minimum length. The original_labels, irish_labels, and chinese_labels indicate the format as misogyny(1);not-misogyny(0).

6 Methodology

6.1 Textual Prompt Engineering for Task A: Single-culture prediction

To investigate the effectiveness of text-only misogyny detection [17], we design a textual prompt engineering framework for Task A: Single-culture Prediction. Unlike the multimodal setting, this task focuses exclusively on OCR-extracted meme transcriptions without visual information [18]. The objective is to determine whether a meme transcription contains misogynistic content within a specific cultural context. The main test we conducted was on the tamil language in the indian country context, focusing on its single-cultural misogynistic meme detection aspects. Given that online meme texts often contain sarcasm, slang, transliterated expressions, implicit

hate speech, and culturally dependent semantics, carefully designed prompts are essential for guiding large language models (LLMs) toward reliable classification behavior.

6.1.1 Prompt Design

The proposed prompt explicitly instructs the large language model (LLM) to: Analyze the meme transcription carefully; Consider the linguistic and socio-cultural context; Detect implicit or explicit misogynistic meaning; Produce structured JSON outputs for downstream processing. The prompt template is defined as follows shown in Figure 1. The transcription text is dynamically appended at the end of the prompt. This prompting strategy encourages the model to perform contextual reasoning rather than relying on superficial keyword matching [19].




Figure 1: The overall zero-shot prompts are illustrated conceptually and the model performs textual modality binary classification for Task A: Single-culture prediction.

6.1.2 Culture-aware Reasoning

To improve cross-cultural understanding, the prompt incorporates both language and country information. Since misogynistic expressions can vary significantly across regions and communities, the model is guided to interpret the transcription within the target cultural environment [20]. For example: Tamil memes are interpreted under the cultural context of India; Chinese memes are interpreted under the cultural context of China; English memes may be interpreted under the cultural context of Ireland. Our main focus was on testing Tamil memes within the cultural framework of India. This culture-aware prompting mechanism enables the model to capture implicit misogynistic semantics, local stereotypes, and culturally dependent offensive expressions [21].

6.1.3 Large Language Model Inference

The textual classification framework utilizes the different large language models (LLMs) through the Ollama inference API. The model performs zero-shot classification using prompt-based instruction following. The inference configuration is defined as shown in Table 3. A deterministic decoding strategy with low temperature is adopted to ensure stable and reproducible predictions.

Table 3: Some typical commands hyper-parameter for textual modality large language models (LLMs)

Model	Options
deepseek-r1-1.5b	{"temperature": 0.3, "num_predict": 1024}
deepseek-r1-7b	{"temperature": 0.0, "num_predict": 512}
deepseek-r1-8b	{"temperature": 0.0, "num_predict": 512}
gemma-2b	{"temperature": 0.3, "num_predict": 1024}
gemma-7b	{"temperature": 0.0, "num_predict": 512}
glm4-9b	{"temperature": 0.0, "num_predict": 512}
llama3-8b	{"temperature": 0.0, "num_predict": 512}
mistral-instruct	{"temperature": 0.3, "num_predict": 1024}
mistral-latest	{"temperature": 0.3, "num_predict": 1024}
qwen2.5-7b	{"temperature": 0.0, "num_predict": 512}

6.1.4 Advantages of Textual Prompt Engineering

The proposed textual prompt engineering strategy offers several advantages: Culture-aware text understanding [22]; Zero-shot misogyny detection without task-specific training [23]; Structured and interpretable outputs; Robust handling of multilingual meme transcriptions; Flexible integration with different large language models (LLMs). By explicitly incorporating cultural and linguistic context into the prompt design, the framework improves the detection of implicit misogynistic expressions in multilingual online meme environments [24].

6.2 Multimodal Prompt Engineering for Task B: Cross-cultural prediction

This study proposes a multimodal misogynistic meme detection framework based on the vision-language model (VLM) Alibaba Cloud's Qwen2.5-VL 3B and 7B. The framework jointly analyzes meme images and OCR-extracted textual transcriptions to determine whether a meme contains misogynistic content under different cultural contexts. At the same time, we also referred to the models Gemma-3 1B and 4B, which only contain OCR-extracted textual transcriptions for determination. The proposed pipeline consists of four major stages: Data Loading and Preprocessing; Prompt-based Multimodal Reasoning; Vision-Language Inference; Prediction Parsing and Result Generation. The overall workflow is illustrated conceptually as shown in Figure 2 (left).




Figure 2: The overall workflow is illustrated conceptually and the model performs binary classification for Task B: Cross-cultural prediction (left). The prompt template is formulated for Qwen2.5-VL 3B/7B and Gemma-3 1B/4B (right).

Internet memes typically communicate meaning through the interaction between visual and textual elements. Therefore, re-

Table 4: Test dataset results for textual modality large language models (LLMs) in the MDMD - Tamil (India) Dataset. Bold font represents the best model performance for this Task A.

Model	Language	Cultural Context	Macro-F1	Accuracy
deepseek-r1-1.5b	Tamil	India	0.513717	0.595506
deepseek-r1-7b	Tamil	India	0.52647	0.648876
deepseek-r1-8b	Tamil	India	0.571072	0.654494
gemma-2b	Tamil	India	0.279834	0.311798
gemma-7b	Tamil	India	0.27615	0.317416
glm4-9b	Tamil	India	0.628995	0.696629
llama3-8b	Tamil	India	0.458194	0.727528
mistral-instruct	Tamil	India	0.558774	0.573034
mistral-latest	Tamil	India	0.4135	0.418539
qwen2.5-7b	Tamil	India	0.615254	0.685393

lying solely on text or image information may lead to incomplete semantic understanding. To address this issue, the proposed framework adopts a multimodal reasoning strategy that integrates: Visual content from meme images; OCR-extracted textual transcription; Cultural context information. For each meme sample, the input consists of $X = \{I, T, C, L\}$ where: I denotes the meme image, T represents the OCR transcription, C denotes the target cultural context (e.g., India, China, Ireland), L represents the meme language (e.g., Tamil, Malayalam, Chinese, English).

The OCR transcription is appended to the textual prompt to help the model reason jointly over textual and visual semantics. At the same time, we also referred to the models Gemma-3 1B and 4B, which only contain OCR-extracted textual transcriptions for determination. Therefore, the input of these two models ignores the visual content from meme images. However, the results showed that these two textual models actually achieved better overall Macro F1-score and Accuracy. The possible reason is that the visual language model (VLM) is still unable to fully understand the content information of meme images, which instead leads to excessive redundancy of the data, causing interference for the large language model (LLM) when performing cross-cultural misogynistic meme detection.

A prompt-based classification strategy is employed to guide the model toward culturally-aware misogyny detection [25]. The prompt explicitly instructs the model to: Analyze both image and text simultaneously; Consider the specified cultural context; Detect implicit misogynistic meaning; Output predictions in structured JSON format. The prompt template is formulated as shown in Figure 2 (right). The meme image and OCR transcription is appended dynamically. This design encourages the model to perform cross-modal semantic reasoning rather than relying on isolated modalities.

7 Results and Analysis

For Task A: Single-culture Prediction, the main dataset we use is the MDMD - Tamil (India) Dataset, which predicts the situation of misogynistic meme detection based on the Tamil language within the Indian cultural context. The results of its test set are shown in Table 4.

For Task Task B: Cross-cultural prediction, the main dataset we use is the MDMD - Tamil (India) Dataset, which predicts the situation of misogynistic meme detection based on the Tamil language within the Indian/Irish/Chinese cultural con-

Table 5: Test dataset results for multimodal or textual modality large language models (LLMs) in the MDMD - Tamil (India) Dataset. Bold font represents the best model performance for this Task B.

Model	Language	Cultural Context	Macro-F1	Accuracy
gemma-3-1b-it-abiliterated	Tamil	India	0.461018	0.716292
	Tamil	Irish	0.476471	0.747191
	Tamil	Chinese	0.462167	0.735955
	Tamil	Overall	0.466552	0.731446
	Tamil	India	0.653783	0.708056
gemma-3-4b-it-abiliterated	Tamil	Irish	0.555935	0.679775
	Tamil	Chinese	0.593712	0.702247
	Tamil	Overall	0.601143	0.695693
	Tamil	India	0.217582	0.278090
	Tamil	Irish	0.196388	0.244382
qwen2.5-vl-3b-instruct	Tamil	Chinese	0.192744	0.238764
	Tamil	Overall	0.202238	0.253745
	Tamil	India	0.301822	0.337079
	Tamil	Irish	0.377359	0.384831
	Tamil	Chinese	0.443188	0.443820
qwen2.5-vl-7b-instruct	Tamil	Overall	0.374123	0.388577

text. The results of its test set are shown in Table 5.

8 Discussion and Limitation

We merely discussed the application effect of zero-shot prompt engineering for Cross-Cultural Misogynistic Meme Detection (CC-MMD). In cases where manual annotation of the dataset is costly, considering transfer learning or introducing existing official datasets might further enhance the effect [26]. This is an important part of our future work.

9 Conclusion

For Task A: Single-culture prediction and Task B: Cross-cultural prediction, we discussed the effectiveness of zero-shot prompt engineering in single-culture prediction and multi-culture prediction. Interestingly, the introduction of meme image did not improve the model's prediction performance. Instead, it misled the model's prediction results. Simply using the meme transcription could achieve good results. This was a great attempt in the application scenarios where the training dataset is scarce [27] and it is more in line with real-life situations [28].

10 Acknowledgments

This material is based upon work supported by the organizers of CC-MMD Grand Challenge under ICMI 2026. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the organizers of CC-MMD Grand Challenge.

References

[1] Aashish Bhandari, Siddhant B Shah, Surendrabikram Thapa, Usman Naseem, Mehwish Nasim, "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, 1994–2003, 2023.

- [2] Koyel Ghosh, Mithun Das, Sumukh Patel, Nilotpal Bhandary, Alloy Das, Animesh Mukherjee, Sandip Modha, Debasis Ganguly, Utpal Garain, Sylvia Jaki, "Overview of the hasoc track at fire 2025: Abusive meme identification—shadows behind the laughter," In Proceedings of the 17th annual meeting of the Forum for Information Retrieval Evaluation, 28–31, 2025.
- [3] Kongqiang Wang, Qingli Tan, "Hate Speech and Offensive Content Identification in Memes in Bangla, Hindi and Gujarati using Auxiliary Text Supervised Learning," 2025.
- [4] Kongqiang Wang, Xiaobing Zhou, "Detecting Hate Speech and Offensive Content in English: A Model Using DistilBERT and A Model Based on FastText Strategies," 2024.
- [5] Wang Kongqiang, Zhang Peng, "wangkongqiang@ CASE 2025: Detection and Classifying Language and Targets of Hate Speech using Auxiliary Text Supervised Learning," In Proceedings of the 8th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Texts, 62–70, 2025.
- [6] Koyel Ghosh, Saptarshi Saha, Thomas Mandl, Sandip Modha, "Findings from shared tasks on hate speech detection: Performance patterns for low-resource languages," Pattern Recognition Letters, 2025.
- [7] Wiktor Soral, Michał Bilewicz, Mikołaj Winiewski. "Exposure to hate speech increases prejudice through desensitization," Aggressive behavior, 44, 2:136–146, 2018.
- [8] Koyel Ghosh, Neeraj Kumar Singh, Joy Mahapatra, Saptarshi Saha, Apurbalal Senapati, Utpal Garain, "Safe-Speech: a three-module pipeline for hate intensity mitigation of social media texts in Indic languages," Social Network Analysis and Mining 14, 1, 245, 2024.
- [9] Kongqiang Wang, "wangkongqiang at Mental-RiskES@ IberLEF 2025: Early Detection of Mental Disorders Risk in Spanish," In IberLEF (Working Notes), CEUR Workshop Proceedings, 2025.
- [10] Wang Kongqiang, "wangkongqiang at SemEval-2025 Task 11: Bridging the Gap in Text-Based Emotion Detection," In Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025), 2025.
- [11] Kongqiang Wang, Qingli Tan, Peng Zhang. "Emotional Cause Analysis Based on In-Context Learning of Large Language Models," INNO-PRESS: Journal of Emerging Applied AI 2, 1, 2026.
- [12] Nauros Romim, Mosahed Ahmed, Hriteshwar Talukder, Md Saiful Islam, "Hate speech detection in the bengali language: A dataset and its baseline evaluation," In Proceedings of International Joint Conference on Advances

- in Computational Intelligence: IJCACI 2020, Springer, 457–468, 2021.
- [13] Kongqiang Wang, Peng Zhang, Qingli Tan, "Misinformation Detection in Social Media Texts and LLM Generated Text using Auxiliary Text Supervised Learning," Working Notes of FIRE (2025), 17–20, 2025.
- [14] Bharathi Raja Chakravarthi, Rahul Ponnusamy, Ping Du, Xiaojian Zhuang, Saranya Rajiakodi, Paul Buitelaar, Premjith B, Bhuvaneswari Sivagnanam, Anshid Kizhakkeparambil, Lavanya S.K., "An Overview of the Misogyny Meme Detection Shared Task for Chinese Social Media," In Proceedings of the 5th Conference on Language, Data and Knowledge: Fifth Workshop on Language Technology for Equality, Diversity, Inclusion. Unior Press, Naples, Italy, 200–208, 2025.
- [15] Elisabetta Fersini, Francesca Gasparini, Giulia Rizzi, Aurora Saibene, Berta Chulvi, Paolo Rosso, Alyssa Lees, Jeffrey Sorensen, "SemEval-2022 Task 5: Multimedia Automatic Misogyny Identification," In Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022), Guy Emerson, Natalie Schluter, Gabriel Stanovsky, Ritesh Kumar, Alexis Palmer, Nathan Schneider, Siddharth Singh, and Shyam Ratan (Eds.). Association for Computational Linguistics, Seattle, United States, 533–549. 2022.
- [16] Rahul Ponnusamy, Kathiravan Pannerselvam, Saranya Rajiakodi, Prasanna Kumar Kumaresan, Sajeetha Thavaresan, Bhuvaneswari Sivagnanam, Anshid K.A, Susminu S Kumar, Paul Buitelaar, Bharathi Raja Chakravarthi, "From Laughter to Inequality: Annotated Dataset for Misogyny Detection in Tamil and Malayalam Memes," In Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024). ELRA and ICCL, Torino, Italia, 7480–7488, 2024.
- [17] Alexandra Olteanu, Carlos Castillo, Jeremy Boy, Kush Varshney, "The effect of extremist violence on hateful speech online," In Proceedings of the international AAAI conference on web and social media, Vol. 12. 2018.
- [18] Georgios K Pitsilis, Heri Ramampiaro, Helge Langseth, "Detecting offensive language in tweets using deep learning," arXiv preprint arXiv:1801.04433, 2018.
- [19] Ying Chen, Yilu Zhou, Sencun Zhu, Heng Xu, "Detecting offensive language in social media to protect adolescent online safety," In 2012 international conference on privacy, security, risk and trust and 2012 international conference on social computing. Ieee, 71–80, 2012.
- [20] Ziqi Zhang, David Robinson, Jonathan Tepper "Detecting hate speech on twitter using a convolution-gru based deep neural network," In European semantic web conference. Springer, 745–760, 2018.
- [21] Somnath Banerjee, Maulindu Sarkar, Nancy Agrawal, Punyajoy Saha, Mithun Das, "Exploring transformer based models to identify hate speech and offensive content in english and indo-aryan languages," arXiv preprint arXiv:2111.13974, 2021.
- [22] Bharathi Raja Chakravarthi, Ruba Priyadharshini, Navya Jose, Thomas Mandl, Prasanna Kumar Kumaresan, Rahul Ponnusamy, John Philip McCrae, Elizabeth Sherly, "Findings of the shared task on offensive language identification in Tamil, Malayalam, and Kannada," In Proceedings of the first workshop on speech and language technologies for Dravidian languages. 133–145, 2021.
- [23] Sandip Modha, Thomas Mandl, Gautam Kishore Shahi, Hiren Madhu, Shrey Satapara, Tharindu Ranasinghe, Marcos Zampieri, "Overview of the hasoc subtrack at fire 2021: Hate speech and offensive content identification in english and indo-aryan languages and conversational hate speech," In Proceedings of the 13th annual meeting of the forum for information retrieval evaluation. 1–3, 2021.
- [24] Mithun Das, Animesh Mukherjee, "Banglaabusememe: A dataset for bengali abusive meme classification," In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. 15498–15512, 2023.
- [25] Armand Joulin, Edouard Grave, Piotr Bojanowski, Tomáš Mikolov, "Bag of tricks for efficient text classification," In Proceedings of the 15th conference of the European chapter of the association for computational linguistics: volume 2, short papers. 427–431, 2017.
- [26] Zeerak Talat, Dirk Hovy, "Hateful symbols or hateful people? predictive features for hate speech detection on twitter," In Proceedings of the NAACL student research workshop. 88–93, 2016.
- [27] Mithun Das, Somnath Banerjee, Animesh Mukherjee, "Data bootstrapping approaches to improve low resource abusive language detection for indic languages," In Proceedings of the 33rd ACM conference on hypertext and social media. 32–42, 2022.
- [28] Mithun Das, Punyajoy Saha, Ritam Dutt, Pawan Goyal, Animesh Mukherjee, Binny Mathew, "You too brutus! trapping hateful users in social media: Challenges, solutions & insights," In Proceedings of the 32nd ACM Conference on Hypertext and Social Media. 79–89, 2021.