

YouTube Comments Decoded: Leveraging LLMs for Low Resource Language Classification

Aniket Deroy^{1,*}, Subhankar Maity¹

¹IIT Kharagpur, Kharagpur, India

Abstract

Sarcasm detection is a significant challenge in sentiment analysis, particularly due to its nature of conveying opinions where the intended meaning deviates from the literal expression. This challenge is heightened in social media contexts where code-mixing, especially in Dravidian languages, is prevalent. Code-mixing involves the blending of multiple languages within a single utterance, often with non-native scripts, complicating the task for systems trained on monolingual data. This shared task introduces a novel gold standard corpus designed for sarcasm and sentiment detection within code-mixed texts, specifically in Tamil-English and Malayalam-English languages. The primary objective of this task is to identify sarcasm and sentiment polarity within a code-mixed dataset of Tamil-English and Malayalam-English comments and posts collected from social media platforms. Each comment or post is annotated at the message level for sentiment polarity, with particular attention to the challenges posed by class imbalance, reflecting real-world scenarios. In this work, we experiment with state-of-the-art large language models like GPT-3.5 Turbo via prompting to classify comments into sarcastic or non-sarcastic categories. We obtained a macro-F1 score of 0.61 for Tamil language. We obtained a macro-F1 score of 0.50 for Malayalam language.

Keywords

GPT, Sarcasm, Classification, Low-resource Languages, Prompt Engineering

1. Introduction

Sarcasm, a form of expression where the intended meaning sharply diverges from the literal meaning, poses a significant challenge for sentiment analysis systems [1]. This complexity arises because sarcasm often involves subtle linguistic cues and context-dependent interpretation, making it difficult for automated systems to accurately detect and classify [2, 3].

Sarcasm is a sharp, often humorous form of verbal irony where someone says the opposite of what they truly mean, typically to convey disdain or mockery. It relies heavily on tone and context, as the delivery can transform a seemingly straightforward statement into a biting comment [2, 3]. While it can be a playful way to highlight absurdities or critique behavior, sarcasm can also lead to misunderstandings, especially in written communication where vocal inflections are absent. Ultimately, it's a double-edged sword that can entertain or offend, depending on the audience and the intent behind it.

Overall, the study of code-mixing in NLP is an evolving field with ongoing research aimed at improving the performance of various NLP tasks in multilingual contexts. As multilingualism continues to grow globally, the development of models that can effectively handle code-mixed data is becoming increasingly important.

Code-mixing [4] in NLP refers to the blending of two or more languages within a single sentence or conversation, commonly seen in multilingual societies. It poses unique challenges for tasks like language identification, part-of-speech tagging, and sentiment analysis. Traditional NLP models often struggle with code-mixed data due to the complex interaction between languages. To address this, researchers have developed advanced neural network models, such as RNNs and CNNs, that better capture linguistic context. Code-mixed text also complicates machine translation and speech recognition, requiring specialized approaches to handle language switching effectively. As multilingualism becomes

Forum for Information Retrieval Evaluation, December 12-15, 2024, India

*Corresponding author.

✉ roydanik18@kgpian.iitkgp.ac.in (A. Deroy); subhankar.ai@kgpian.iitkgp.ac.in (S. Maity)

id 0000-0001-7190-5040 (A. Deroy); 0009-0001-1358-9534 (S. Maity)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

more prevalent, the importance of developing robust NLP systems capable of processing code-mixed data continues to grow.

Code-mixing is a widespread phenomenon in multilingual societies, particularly in Dravidian languages such as Tamil and Malayalam, where users frequently switch between native languages and English [4]. These code-mixed texts often appear in non-native scripts, such as Roman, due to the convenience of typing on digital platforms [4]. This creates a unique set of challenges for natural language processing (NLP) systems, which are typically trained on monolingual data and thus struggle to handle the intricacies of code-switching across different linguistic levels, such as syntax, semantics, and phonetics [5].

The Dravidian languages, with their rich cultural and linguistic history, present a distinct case for code-mixed text analysis [6]. Tamil, with its official status in India, Sri Lanka, and Singapore, and a significant diaspora community, has a script that evolved from ancient scripts like Tamil and the Chola-Pallava script [7]. Similarly, Malayalam, predominantly spoken in Kerala, India, employs an alpha-syllabic script that blends alphabetic and syllable-based writing systems [8]. However, on social media, these languages are often typed using the Roman script, contributing to a growing corpus of code-mixed data.

To address these challenges, this paper introduces a prompt-based method [9] using GPT-3.5 Turbo [10] for sarcasm and sentiment detection in code-mixed texts, specifically focusing on Tamil-English and Malayalam-English language pairs. This shared task [11, 12, 13, 14, 15, 16] aims to advance research in sentiment analysis by providing a dataset that reflects real-world scenarios where code-mixing is prevalent and by encouraging the development of systems that can accurately classify sarcasm and sentiment polarity in these under-resourced languages.

Participants in this task were provided development, training, and test datasets composed of comments and posts collected from social media platforms such as YouTube.

In this work, we explore the capabilities of state-of-the-art large language models, specifically GPT-3.5 Turbo [10], by leveraging prompt-based techniques to classify social media comments into sarcastic or non-sarcastic categories. GPT-3.5 Turbo [10], a variant of OpenAI’s GPT-3.5 model, is known for its advanced natural language understanding and generation capabilities, making it a promising candidate for handling the nuanced task of sarcasm detection, particularly in code-mixed texts.

Our team, TextTitans, participated in a competition aimed at advancing sarcasm detection in Tamil and Malayalam. The results of our systems’ performance provide valuable insights into the current state of NLP for these languages. For Tamil, our system achieved a macro-F1 score of 0.61, placing us 9th among competing teams. The macro-F1 score, which balances precision and recall across multiple classes, reflects our system’s ability to effectively handle the nuances of Tamil sarcasm. However, despite the relatively strong performance indicated by the F1 score, we achieved 9th rank.

In contrast, our system’s performance for Malayalam resulted in a macro-F1 score of 0.50, securing the 13th position in the rankings. This score, while moderate, highlights the challenges inherent in processing Malayalam text, particularly in capturing the subtleties of sarcastic expressions. The lower rank compared to Tamil indicates that our system faced greater difficulties with Malayalam, and that several other systems were more adept at handling these challenges.

These results underscore the varying levels of difficulty and competition present in sarcasm detection for Tamil and Malayalam. The higher performance in Tamil suggests that our system is more attuned to the linguistic characteristics of Tamil sarcasm, while the more modest performance in Malayalam points to areas where further research and refinement are necessary. This study not only contributes to the growing body of work in multilingual sarcasm detection but also highlights the importance of continuing to develop robust NLP systems capable of addressing the complexities of underrepresented languages.

2. Related Work

Sarcasm detection has been a topic of significant interest in the field of natural language processing (NLP) due to its complex and context-dependent nature [17]. Traditional approaches to sarcasm detection primarily relied on feature-based methods, where linguistic cues such as punctuation, lexical patterns, and syntactic structures were used to identify sarcasm [17]. For instance, works by [18] and [19] utilized syntactic patterns and hashtag-based supervision to improve sarcasm detection in English texts. These approaches, however, often struggled with generalization, particularly in scenarios involving nuanced and culturally specific forms of sarcasm.

[20] introduced an approach using deep neural networks that leveraged context to improve sarcasm detection. Similarly, [21] explored attention-based models that better captured the dependencies within a sentence that could indicate sarcasm. These deep learning models significantly advanced the state-of-the-art in sarcasm detection, particularly in monolingual contexts.

One of the primary challenges in dealing with code-mixed text is language identification [22], which involves determining the language of each word or phrase within a mixed-language sentence. Traditional methods for this task often relied on n-gram models, dictionary-based approaches, and supervised learning techniques. However, these methods sometimes struggle with the complex linguistic patterns present in code-mixed data. These models are particularly effective at capturing the contextual information necessary to accurately identify languages in code-mixed text.

Another critical area of research is part-of-speech tagging [23] for code-mixed sentences. This task becomes more complicated due to the interaction between different languages' grammatical structures. Researchers have employed techniques like Conditional Random Fields (CRFs) and sequence-to-sequence models, adapting them to handle the intricacies of mixed-language data. The use of multilingual embeddings, where words from different languages are mapped into a shared vector space, has also been shown to improve the accuracy of part-of-speech tagging in code-mixed scenarios.

In the context of sentiment analysis [24], code-mixing introduces additional complexity due to the potential for differing sentiment expressions across languages. Traditional sentiment analysis tools often fail when applied directly to code-mixed text. To overcome this, researchers have explored methods such as using bilingual lexicons, translating the text into a single language before analysis, and employing deep learning models that can process code-mixed input directly.

Machine translation [25] for code-mixed text is another challenging area. Standard machine translation systems are usually trained on monolingual data, making them less effective at handling code-mixed sentences. Researchers have proposed various strategies to address this, including the use of synthetic code-mixed data for training, as well as incorporating language tags and contextual embeddings into translation models. These approaches aim to create more robust translation systems capable of handling the nuances of code-mixing.

Additionally, speech recognition [26] in code-mixed environments requires models that can accurately transcribe spoken language that switches between different languages. This task is particularly challenging because traditional speech recognition systems are typically designed for single-language input. Recent advances involve the use of end-to-end models, such as attention-based encoder-decoder architectures, which have shown promise in recognizing and transcribing code-mixed speech.

Our contribution is situated at the intersection of sarcasm detection, code-mixed text processing, and the application of large language models, addressing a critical gap in the literature and advancing the state-of-the-art in this complex domain.

3. Dataset

The Tamil testing dataset is 6338 comments. The Malayalam testing dataset is 2826 comments. Since we used effectively only the test dataset for our predictions, we only mention the statistics corresponding to the test dataset.

Given that our evaluation focused entirely on the test datasets, the statistics mentioned correspond

to these specific datasets. The decision to use only the test datasets for prediction ensures that the performance metrics, such as the macro-F1 scores, accurately reflect the model's generalization capabilities on unseen data.

The Tamil testing dataset, with its larger size of 6,338 comments, provides a broad spectrum of inputs for the model to process, potentially leading to more robust insights into the model's strengths and weaknesses. On the other hand, the smaller Malayalam testing dataset, with 2,826 comments, presents its own set of challenges, particularly in a low-resource language context where data is often scarcer.

By focusing on these testing datasets, our analysis remains rooted in real-world applicability, offering a clear view of how well the models perform in practical scenarios without additional interventions or training phases.

4. Task Definition

The task is to accurately classify YouTube comments that are written in Tamil-English and Malayalam-English (code-mixed) into two distinct categories: Sarcastic and Non-sarcastic, using a AI model.

5. Methodology

Prompting [9] is an effective approach for solving the problem of sarcasm detection in code-mixed texts, particularly in under-resourced languages like Tamil-English and Malayalam-English, for several reasons:

- **Natural Language Understanding:** Prompting allows the model to interpret and execute tasks based on human-like instructions. By framing the classification task as a natural language prompt [27], the model can utilize its language understanding capabilities to classify text according to the specified categories. The prompt provides context to the model, guiding it to focus on specific aspects of the input (e.g., determining sarcasm). This ensures that the model's response is aligned with the task at hand.
- **Leverage of Pretrained Knowledge:** Large language models like GPT-3.5 Turbo have been trained on vast amounts of text data, encompassing diverse linguistic patterns, cultural contexts, and language nuances [28]. Prompting allows us to tap into this rich, pre-trained knowledge base, enabling the model to interpret sarcasm even in challenging code-mixed scenarios [9]. By carefully crafting prompts, we can guide the model to focus on specific aspects of the input, such as tone, context, and language switching, which are critical for sarcasm detection.
- **Reduced Need for Fine-Tuning:** With the right prompt, large pre-trained models (like GPT) [29] can perform various classification tasks without the need for extensive fine-tuning on task-specific data. This is particularly useful when data or computational resources for fine-tuning are limited. Prompts leverage the model's existing knowledge, gained from vast pre-training on diverse text corpora, enabling it to generalize to new tasks with minimal additional training.
- **Adaptability to Multilingual and Code-Mixed Texts:** Traditional machine learning models often require large, labeled datasets for each specific task and language, which are scarce for many code-mixed languages [30]. In contrast, prompting large language models provides a flexible way to adapt to different languages and dialects, even when training data is limited [9]. The model's inherent ability to understand and generate text in multiple languages, including code-mixed varieties, makes prompting an attractive approach for this task.
- **Improved Interpretability:** The explicit nature of prompts makes it easier to understand how the model is interpreting the task [31], as the instructions are directly embedded in the input. This transparency can help in analyzing and improving the model's performance. By examining the model's responses to different prompts, researchers can gain insights into where and why the model might be making classification errors, leading to more targeted improvements.

- **Reduced Need for Task-Specific Data:** Building a model from scratch or fine-tuning a model for sarcasm detection in code-mixed texts would typically require a substantial amount of labeled data, which is often difficult to obtain for under-resourced languages [32]. Prompting, however, minimizes the need for such extensive labeled datasets, as it allows the model to utilize its general language understanding capabilities with minimal task-specific adjustments [9]. This makes it a more feasible and efficient solution for resource-constrained settings.
- **Flexibility and Rapid Experimentation:** Prompting enables rapid experimentation with different strategies to optimize model performance [9]. Researchers can quickly test various prompt formulations to see which one elicits the most accurate and context-aware responses from the model. This flexibility is particularly beneficial for complex tasks like sarcasm detection, where the model's interpretation can vary significantly depending on how the problem is framed.
- **Scalability Across Different Scenarios:** Once an effective prompting strategy is developed, it can be easily adapted or scaled to other similar tasks or languages without the need for extensive retraining [33, 9]. This is particularly useful in multilingual environments where a single model might need to handle multiple languages or dialects. Prompting allows for the reuse of the same model across different linguistic contexts with minimal modifications.

In summary, prompting is a powerful and efficient approach to tackle the problem of sarcasm detection in code-mixed texts, offering the benefits of leveraging large-scale pretrained models, adapting to multilingual and code-mixed scenarios, reducing dependency on large annotated datasets, enabling rapid experimentation, and providing scalability across different linguistic contexts.

5.1. Prompt Engineering-Based Approach

We used the GPT-3.5 Turbo model via prompting to solve the classification task. We used GPT-3.5 Turbo in zero-Shot mode via prompting. After the prompt is provided to the LLM, the following steps happen internal to the LLM while generating the output. The following outlines the steps [10, 34, 35] that occur internally within the LLM, summarizing the prompting approach using GPT-3.5 Turbo:

Step 1: Tokenization

- **Prompt:** $X = [x_1, x_2, \dots, x_n]$
- The input text (prompt) is first tokenized into smaller units called tokens. These tokens are often subwords or characters, depending on the model's design.
- **Tokenized Input:** $T = [t_1, t_2, \dots, t_m]$

Step 2: Embedding

- Each token is converted into a high-dimensional vector (embedding) using an embedding matrix E .
- **Embedding Matrix:** $E \in \mathbb{R}^{|V| \times d}$, where $|V|$ is the size of the vocabulary and d is the embedding dimension.
- **Embedded Tokens:** $T_{\text{emb}} = [E(t_1), E(t_2), \dots, E(t_m)]$

Step 3: Positional Encoding

- Since the model processes sequences, it adds positional information to the embeddings to capture the order of tokens.
- **Positional Encoding:** $P(t_i)$
- **Input to the Model:** $Z = T_{\text{emb}} + P$

Step 4: Attention Mechanism (Transformer Architecture)

- **Attention Score Calculation:** The model computes attention scores to determine the importance of each token relative to others in the sequence.

- **Attention Formula:**

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (1)$$

- where Q (query), K (key), and V (value) are linear transformations of the input Z .
- This attention mechanism is applied multiple times through multi-head attention, allowing the model to focus on different parts of the sequence simultaneously.

Step 5: Feedforward Neural Networks

- The output of the attention mechanism is passed through feedforward neural networks, which apply non-linear transformations.

- **Feedforward Layer:**

$$\text{FFN}(x) = \max(0, xW_1 + b_1)W_2 + b_2 \quad (2)$$

- where W_1, W_2 are weight matrices and b_1, b_2 are biases.

Step 6: Stacking Layers

- Multiple layers of attention and feedforward networks are stacked, each with its own set of parameters. This forms the "deep" in deep learning.

- **Layer Output:**

$$H^{(l)} = \text{LayerNorm}(Z^{(l)} + \text{Attention}(Q^{(l)}, K^{(l)}, V^{(l)})) \quad (3)$$

$$Z^{(l+1)} = \text{LayerNorm}(H^{(l)} + \text{FFN}(H^{(l)})) \quad (4)$$

Step 7: Output Generation

- The final output of the stacked layers is a sequence of vectors.
- These vectors are projected back into the token space using a softmax layer to predict the next token or word in the sequence.

- **Softmax Function:**

$$P(y_i|X) = \frac{\exp(Z_i)}{\sum_{j=1}^{|V|} \exp(Z_j)} \quad (5)$$

- where Z_i is the logit corresponding to token i in the vocabulary.
- The model generates the next token in the sequence based on the probability distribution, and the process repeats until the end of the output sequence is reached.

Step 8: Decoding

- The predicted tokens are then decoded back into text, forming the final output.
- **Output Text:** $Y = [y_1, y_2, \dots, y_k]$

We used the following prompt for Tamil language for the purpose of classification: *"Please Check whether the comment-<Text> is Sarcastic or Non-sarcastic. Only state Sarcastic or Non-sarcastic"*. The figure representing the methodology is shown in Figure 1.

We used the following prompt for Malayalam language for the purpose of classification: *"Please Check whether the comment-<Text> is Sarcastic or Non-sarcastic. Only state Sarcastic or Non-sarcastic"*. The figure representing the methodology is shown in Figure 2. We run the GPT model at temperatures of 0.7, 0.8, and 0.9.

We have used two different figures (Figure 1 and Figure 2) for the two sarcasm detection models in Tamil and Malayalam.

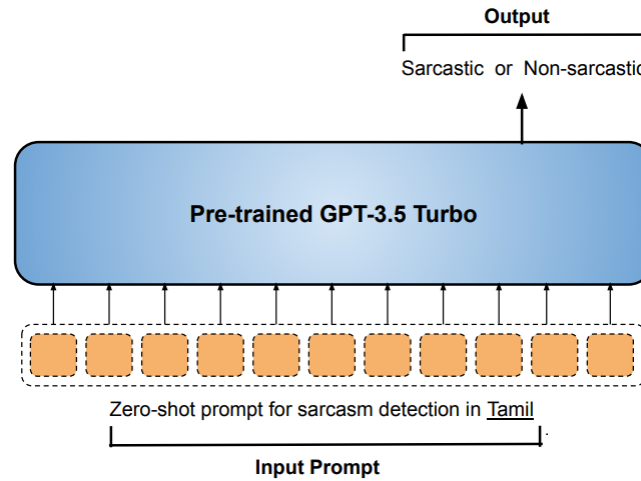


Figure 1: Architechture diagram for sarcasm detection in Tamil.

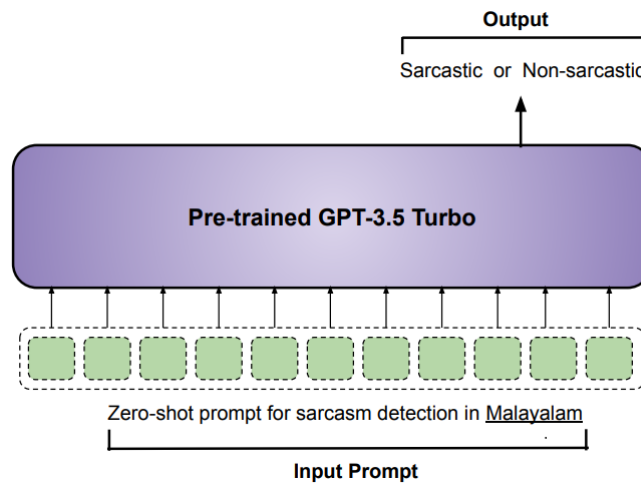


Figure 2: Architechture diagram for sarcasm detection in Malayalam.

The methodology employed for the classification task involved using language-specific prompts to determine whether a comment is sarcastic or non-sarcastic, with this approach applied to both Tamil and Malayalam languages. For Tamil, the prompt used was: "Please Check whether the comment-<Text> is Sarcastic or Non-sarcastic. Only state Sarcastic or Non-sarcastic". This prompt was carefully crafted to direct the model's focus solely on the content of the comment, ensuring that the output would be either "Sarcastic" or "Non-sarcastic" without any additional context or information.

Similarly, for Malayalam, the prompt used was: "Please Check whether the comment-<Text> is Sarcastic or Non-sarcastic. Only state Sarcastic or Non-sarcastic". This prompt, is similar in structure to the Tamil prompt, was designed to maintain consistency in the classification task across both languages.

The process began with the preparation of input, where comments in Tamil and Malayalam were fed into the model using their respective prompts. The model then processed each comment according to the instructions provided in the prompt, returning a binary classification of either "Sarcastic" or "Non-sarcastic" for each comment.

In conclusion, this methodology ensured a consistent classification task across different languages while accommodating the unique characteristics of Tamil and Malayalam. The use of clear and concise prompts enabled the model to focus on the core task of sarcasm detection, leading to reliable and interpretable results.

6. Results

Our team named **TextTitans** [36], for the Tamil language the macro-F1 score is 0.61 with a rank of 9th. For Malayalam language the macro-F1 score is 0.50 with a rank of 13th.

For Malayalam, the system has achieved a macro-F1 score of 0.50, which represents the harmonic mean of precision and recall across multiple classes, treating each class equally. This score places the system in the 13th position among other competing models or systems. A rank of 13th suggests that while the performance is moderate, there are several other systems that have performed better.

For Tamil, the system has a macro-F1 score of 0.61, indicating better performance than the Malayalam system in terms of the balance between precision and recall. However, despite the higher F1 score, this system is ranked 9th. This lower rank, despite a better score, suggests that the competition for Tamil language tasks was fiercer, with more systems achieving higher or comparable performance levels.

| | Precision | Recall | F1-Score | Support |
|---------------|-----------|--------|----------|---------|
| Non-sarcastic | 0.82 | 0.73 | 0.77 | 2314 |
| Sarcastic | 0.18 | 0.27 | 0.22 | 512 |
| Micro avg | 0.65 | 0.65 | 0.65 | 2826 |
| Macro avg | 0.50 | 0.50 | 0.50 | 2826 |
| Weighted avg | 0.70 | 0.65 | 0.67 | 2826 |

Table 1

Classification Report for Sarcasm Detection for Malayalam language.

| | Precision | Recall | F1-Score | Support |
|---------------|-----------|--------|----------|---------|
| Non-sarcastic | 0.79 | 0.79 | 0.79 | 4621 |
| Sarcastic | 0.43 | 0.43 | 0.43 | 1717 |
| Micro avg | 0.69 | 0.69 | 0.69 | 6338 |
| Macro avg | 0.61 | 0.61 | 0.61 | 6338 |
| Weighted avg | 0.69 | 0.69 | 0.69 | 6338 |

Table 2

Classification Report for Sarcasm Detection for Tamil language.

Table 1 shows the classification report for Sarcasm Detection for Malayalam language. The model performs well with a precision of 0.82, indicating that most of the time when it predicts a comment as non-sarcastic, it is correct. However, with a recall of 0.73, it misses some non-sarcastic comments, leading to a moderately high F1-score of 0.77. The model struggles with sarcasm detection, evident by a low precision (0.18) and recall (0.27). This indicates the model has difficulty both identifying sarcastic comments and correctly predicting them, resulting in a low F1-score of 0.22. The overall performance across both classes is quite balanced (0.50 for precision, recall, and F1-score), showing the disparity in performance between the two classes. Weighted by the number of instances in each class, this metric gives a better sense of the model's performance on the dataset as a whole, showing the skew towards the Non-sarcastic class due to its larger support.

Table 2 shows the classification report for Sarcasm Detection for Tamil language. The model shows strong performance with a precision, recall, and F1-score all at 0.79, indicating balanced detection and prediction of non-sarcastic comments. While still less effective at detecting sarcasm, the model has improved significantly from the first table, with a precision, recall, and F1-score of 0.43. This shows better balance in detecting sarcasm compared to the first report. The macro average has improved to 0.61 for all three metrics, indicating better overall performance across both classes, even though the model is still biased towards the Non-sarcastic class. Consistent across precision, recall, and F1-score at 0.69, showing that the model's performance is more balanced in this report compared to the first one, likely due to better handling of the Sarcastic class.

Error Analysis: According to Table 1, the sarcasm detection results for the Malayalam language using GPT-3.5 reveal significant performance gaps, particularly in identifying sarcastic remarks. The model's low precision of 0.18 for the sarcastic class indicates a high rate of false positives, suggesting that non-sarcastic comments are often misclassified as sarcastic. Additionally, a recall of 0.27 shows that many actual sarcastic instances go undetected, resulting in a low F1-score of 0.22. In contrast, the non-sarcastic class performs better, with a precision of 0.82 and an F1-score of 0.77, reflecting the

model’s stronger ability to identify straightforward statements. This imbalance highlights the model’s struggle with the nuances of sarcasm, indicating the need for further refinement in training strategies or additional data to enhance performance in Malayalam.

Similarly, the sarcasm detection results for the Tamil language (See Table 2) indicate critical areas for improvement. The model’s precision for sarcastic comments is only 0.43, which signifies a considerable number of false positives, where non-sarcastic statements are misclassified as sarcastic. The recall for the sarcastic class is also low at 0.43, suggesting that the model fails to identify many genuine sarcastic instances, leading to an F1-score of 0.43. Conversely, the non-sarcastic class shows stronger performance, achieving a precision, recall, and F1-score of 0.79. This indicates effective detection of straightforward statements. The overall balanced performance, with a macro average F1-score of 0.61, highlights the inconsistency in sarcasm detection. These findings underscore the necessity for improved training methods or more diverse datasets to bolster the model’s capability to recognize sarcasm in Tamil more effectively.

7. Conclusion

The complexity of code-switching, coupled with the nuanced nature of sarcasm, presents significant hurdles for traditional sentiment analysis systems, particularly in under-resourced languages. To tackle these challenges, we experimented with state-of-the-art large language models like GPT-3.5 Turbo, leveraging prompt-based techniques to classify comments and posts as sarcastic or non-sarcastic. Our approach demonstrated the potential of prompting as an effective method for sarcasm detection in multilingual and code-mixed environments, allowing us to utilize the extensive pretrained knowledge of the model with minimal task-specific data. The presence of a new gold standard corpus for Tamil-English and Malayalam-English code-mixed text contributes to the advancement of research in this area, providing a valuable resource for the development and evaluation of sarcasm detection systems. Overall, our work underscores the viability of prompt-based techniques for addressing the unique challenges of sarcasm detection in code-mixed languages and paves the way for future studies aimed at improving systems in multilingual social media contexts. We hope that this research will inspire continued exploration of advanced language models in tackling the diverse and evolving landscape of digital communication.

The evaluation of our sarcasm detection systems for Malayalam and Tamil languages provides valuable insights into the performance and challenges associated with handling these languages. For Malayalam, the system achieved a macro-F1 score of 0.50, which, while representing a moderate balance between precision and recall across classes, placed the system at the 13th position among competing models. This rank indicates that while our system performs adequately, there are several other models that have outperformed it in the specific context of Malayalam sarcasm detection.

In contrast, the Tamil sarcasm detection system achieved a higher macro-F1 score of 0.61, reflecting a better overall balance between precision and recall compared to the Malayalam system. Despite this improved score, the Tamil system was ranked 9th among competitors. This lower ranking, even with a higher F1 score, highlights the more competitive landscape for Tamil language tasks, where several systems have achieved similar or better performance levels.

These results underscore the varying degrees of difficulty in sarcasm detection across different languages and the importance of context in interpreting performance metrics like the macro-F1 score. The differences in ranking also suggest that the complexity of the language, the nature of the testing data, and the sophistication of competing models play significant roles in determining a system’s success.

In future work, it will be crucial to explore ways to enhance the performance of both systems, particularly by focusing on the nuances of sarcasm in these languages, improving the quality and diversity of the training data, and leveraging more advanced modeling techniques. Additionally, a more in-depth analysis of the specific errors made by the systems could provide further insights into areas that require targeted improvements. Ultimately, the goal is to develop robust and competitive systems

that can more effectively handle the challenges of sarcasm detection in both Malayalam and Tamil, contributing to the broader field of natural language processing for low-resource languages.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT in order to: Drafting content, Grammar and spelling check, etc. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] H. T. Uchiyama, D. N. Saito, H. C. Tanabe, T. Harada, A. Seki, K. Ohno, T. Koeda, N. Sadato, Distinction between the literal and intended meanings of sentences: A functional magnetic resonance imaging study of metaphor and sarcasm, *cortex* 48 (2012) 563–583.
- [2] C. Hong, Sarcasm Detection Incorporating Context & World Knowledge, Ph.D. thesis, COOPER UNION, 1914.
- [3] M. S. Bin Razali, Figurative language detection using deep learning and contextual features (2023).
- [4] N. Jose, B. R. Chakravarthi, S. Suryawanshi, E. Sherly, J. P. McCrae, A survey of current datasets for code-switching research, in: 2020 6th international conference on advanced computing and communication systems (ICACCS), IEEE, 2020, pp. 136–141.
- [5] L. Nguyen, Z. Yuan, G. Seed, Building educational technologies for code-switching: Current practices, difficulties and future directions, *Languages* 7 (2022) 220.
- [6] B. R. Chakravarthi, R. Priyadharshini, V. Muralidaran, N. Jose, S. Suryawanshi, E. Sherly, J. P. McCrae, Dravidiancodemix: Sentiment analysis and offensive language identification dataset for dravidian languages in code-mixed text, *Language Resources and Evaluation* 56 (2022) 765–806.
- [7] V. Ramaswamy, Historical dictionary of the Tamils, Rowman & Littlefield, 2017.
- [8] N. B. FalDessai, Development of a Text to Speech System for Devanagari Konkani, Ph.D. thesis, Goa University, 2017.
- [9] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, G. Neubig, Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing, *ACM Computing Surveys* 55 (2023) 1–35.
- [10] T. B. Brown, Language models are few-shot learners, *arXiv preprint ArXiv:2005.14165* (2020).
- [11] B. R. Chakravarthi, R. Priyadharshini, J. Jose, P. Kumaresan, S. Muralidaran, Findings of the shared task on sentiment analysis for dravidian languages in code-mixed text, in: Proceedings of the first workshop on speech and language technologies for Dravidian languages, 2021, pp. 133–139.
- [12] B. R. Chakravarthi, A. Hande, R. Ponnusamy, P. K. Kumaresan, R. Priyadharshini, How can we detect homophobia and transphobia? experiments in a multilingual code-mixed setting for social media governance, *International Journal of Information Management Data Insights* 2 (2022) 100119.
- [13] B. R. Chakravarthi, Hope speech detection in youtube comments, *Social Network Analysis and Mining* 12 (2022) 75.
- [14] B. R. Chakravarthi, N. Sripriya, B. Bharathi, K. Nandhini, S. C. Navaneethakrishnan, T. Durairaj, R. Ponnusamy, P. K. Kumaresan, K. K. Ponnusamy, C. Rajkumar, Overview of the shared task on sarcasm identification of dravidian languages (malayalam and tamil) in dravidiancodemix, in: Forum of Information Retrieval and Evaluation FIRE-2023, 2023.
- [15] B. R. Chakravarthi, N. Sripriya, B. Bharathi, K. Nandhini, S. Chinnaudayar Navaneethakrishnan, T. Durairaj, R. Ponnusamy, P. K. Kumaresan, K. K. Ponnusamy, C. Rajkumar, Overview of the shared task on sarcasm identification of Dravidian languages (Malayalam and Tamil) in DravidianCodeMix, in: Forum of Information Retrieval and Evaluation FIRE - 2023, 2023.
- [16] N. Sripriya, T. Durairaj, K. Nandhini, B. Bharathi, K. K. Ponnusamy, C. Rajkumar, P. K. Kumaresan,

- R. Ponnusamy, C. Subalalitha, B. R. Chakravarthi, Findings of shared task on sarcasm identification in code-mixed dravidian languages, FIRE 2023 16 (2023) 22.
- [17] A. D. Dave, N. P. Desai, A comprehensive study of classification techniques for sarcasm detection on textual data, in: 2016 International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT), IEEE, 2016, pp. 1985–1991.
 - [18] D. Davidov, O. Tsur, A. Rappoport, Enhanced sentiment learning using twitter hashtags and smileys, in: Coling 2010: Posters, 2010, pp. 241–249.
 - [19] E. Riloff, A. Qadir, P. Surve, L. De Silva, N. Gilbert, R. Huang, Sarcasm as contrast between a positive sentiment and negative situation, in: Proceedings of the 2013 conference on empirical methods in natural language processing, 2013, pp. 704–714.
 - [20] A. Ghosh, T. Veale, Fracking sarcasm using neural network, in: Proceedings of the 7th workshop on computational approaches to subjectivity, sentiment and social media analysis, 2016, pp. 161–169.
 - [21] Y. Tay, L. A. Tuan, S. C. Hui, J. Su, Reasoning with sarcasm by reading in-between, arXiv preprint arXiv:1805.02856 (2018).
 - [22] A. F. Hidayatullah, A. Qazi, D. T. C. Lai, R. A. Apong, A systematic review on language identification of code-mixed text: techniques, data availability, challenges, and framework development, IEEE access 10 (2022) 122812–122831.
 - [23] R. Sequiera, M. Choudhury, K. Bali, Pos tagging of hindi-english code mixed text from social media: Some machine learning experiments, in: Proceedings of the 12th international conference on natural language processing, 2015, pp. 237–246.
 - [24] B. R. Chakravarthi, R. Priyadharshini, V. Muralidaran, S. Suryawanshi, N. Jose, E. Sherly, J. P. McCrae, Overview of the track on sentiment analysis for dravidian languages in code-mixed text, in: Proceedings of the 12th Annual Meeting of the Forum for Information Retrieval Evaluation, 2020, pp. 21–24.
 - [25] G. Jawahar, E. M. B. Nagoudi, M. Abdul-Mageed, L. V. Lakshmanan, Exploring text-to-text transformers for english to hinglish machine translation with synthetic code-mixing, arXiv preprint arXiv:2105.08807 (2021).
 - [26] Y. Long, Y. Li, Q. Zhang, S. Wei, H. Ye, J. Yang, Acoustic data augmentation for mandarin-english code-switching speech recognition, Applied Acoustics 161 (2020) 107175.
 - [27] J. Allen, Natural language understanding, Benjamin-Cummings Publishing Co., Inc., 1988.
 - [28] M. U. Hadi, R. Qureshi, A. Shah, M. Irfan, A. Zafar, M. B. Shaikh, N. Akhtar, J. Wu, S. Mirjalili, et al., Large language models: a comprehensive survey of its applications, challenges, limitations, and future prospects, Authorea Preprints (2023).
 - [29] P. Colon-Hernandez, C. Havasi, J. Alonso, M. Huggins, C. Breazeal, Combining pre-trained language models and structured knowledge, arXiv preprint arXiv:2101.12294 (2021).
 - [30] A. F. Hidayatullah, A. Qazi, D. T. C. Lai, R. A. Apong, A systematic review on language identification of code-mixed text: techniques, data availability, challenges, and framework development, IEEE access 10 (2022) 122812–122831.
 - [31] Y. W. Jie, R. Satapathy, G. S. Mong, E. Cambria, et al., How interpretable are reasoning explanations from prompting large language models?, arXiv preprint arXiv:2402.11863 (2024).
 - [32] K. Goswami, Unsupervised deep representation learning for low-resourced, LASER 77 (2012) 79–7.
 - [33] M. G. Z. A. Khan, M. F. Naeem, L. Van Gool, D. Stricker, F. Tombari, M. Z. Afzal, Introducing language guidance in prompt-based continual learning, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 11463–11473.
 - [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, Advances in neural information processing systems 30 (2017) 5998–6008.
 - [35] A. Radford, Improving language understanding by generative pre-training (2018).
 - [36] B. R. Chakravarthi, S. N, B. B, N. K, T. Durairaj, R. Ponnusamy, P. K. Kumaresan, K. K. Ponnusamy, C. Rajkumar, Overview of sarcasm identification of dravidian languages in dravidiancodemix@fire-2024, in: Forum of Information Retrieval and Evaluation FIRE - 2024, DAICT, Gandhinagar, 2024.