

HateGPT: Unleashing GPT-3.5 Turbo to Combat Hate Speech on X

Aniket Deroy^{1,*†}, Subhankar Maity¹

¹IIT Kharagpur, Kharagpur, India

Abstract

The widespread use of social media platforms like Twitter and Facebook has enabled people of all ages to share their thoughts and experiences, leading to an immense accumulation of user-generated content. However, alongside the benefits, these platforms also face the challenge of managing hate speech and offensive content, which can undermine rational discourse and threaten democratic values. As a result, there is a growing need for automated methods to detect and mitigate such content, especially given the complexity of conversations that may require contextual analysis across multiple languages, including code-mixed languages. We participated in the English task where we have to classify English tweets into two categories namely Hate and Offensive and Non Hate-Offensive. In this work, we experiment with state-of-the-art large language models like GPT-3.5 Turbo via prompting to classify tweets into Hate and Offensive or Non Hate-Offensive. We modified the temperature as an experimental parameter. In this study, we evaluate the performance of a classification model using Macro-F1 scores across three distinct runs. The Macro-F1 score, which balances precision and recall across all classes, is used as the primary metric for model evaluation. The scores obtained are 0.756 for run 1, 0.751 for run 2, and 0.754 for run 3, indicating a high level of performance with minimal variance among the runs. The results suggest that the model consistently performs well in terms of precision and recall, with run 1 showing the highest performance. These findings highlight the robustness and reliability of the model across different runs.

Keywords

GPT, Hate Speech, Classification, English, Prompt Engineering

1. Introduction

The advent of social media platforms such as Twitter (currently known as X) and Facebook has revolutionized the way individuals communicate, enabling people from diverse backgrounds to share their thoughts, experiences, and opinions freely [1]. This democratization of content creation has led to an exponential increase in user-generated data. While these platforms have facilitated global connectivity and discourse, they have also become hotbeds for hate speech [2, 3, 4, 5] and offensive content. Such content not only disrupts meaningful communication but also poses significant threats to social cohesion and democratic values.

Addressing the proliferation of hate speech [6, 7, 8] on social media is a complex challenge. The nature of online communication, where context and nuance often play a crucial role, makes it difficult to detect offensive language accurately [9]. This challenge is further compounded by the multilingual nature of online communities, where users frequently employ code-mixed languages such as Hinglish (a mix of Hindi and English), German-English, and Bangla, among others [10]. As these languages blend cultural and linguistic elements, the task of identifying hate speech becomes even more intricate.

In response to this growing concern, technology companies and social media platforms have begun to invest in automated methods to detect and manage offensive content [2, 11, 12]. The goal is to strike a balance between preserving open and free dialogue while preventing the spread of harmful speech. In this work, we focus on the classification of English tweets into two categories: Hate and Offensive and Non Hate-Offensive. By leveraging state-of-the-art large language models such as GPT-3.5 Turbo, we experiment with prompting techniques to classify tweets accurately.

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)

🆔 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).

Run 1 achieved the highest Macro-F1 score at 0.756, indicating it balanced precision and recall across different classes better than the other runs. Run 2 had a slightly lower score of 0.751, suggesting a small decline in performance, either in precision, recall, or both. Run 3 scored 0.754, which was slightly lower than Run 1 but higher than Run 2, indicating its performance was similar to Run 2's with a minor improvement.

2. Related Work

The detection of hate speech and offensive content on social media has garnered significant attention in recent years, driven by the growing need to maintain safe and constructive online environments [13]. Researchers have explored a variety of approaches to address this issue, ranging from traditional machine learning techniques [14, 15] to the application of advanced deep learning models [16, 2].

Early approaches to hate speech detection [14] used simple machine learning algorithms. These models [15] used manually crafted features, including word n-grams, part-of-speech tags, and sentiment scores, to classify text. For instance, Badjatiya et al. [17], Chiu et al. [18] employed a logistic regression model with n-grams and part-of-speech features to classify tweets into hate speech, offensive language, and neither. However, the performance of these models was often limited by their reliance on surface-level features, which could not fully capture the complexities of language and context.

Zhang and Luo [6] experimented with LSTM networks and Gradient Boosted Decision Trees (GBDT) to classify hate speech on Twitter, demonstrating improvements over traditional machine learning methods. Similarly, Liu and Avci [7] utilized a CNN-LSTM architecture to detect offensive language, showing that deep learning models could capture both local and sequential patterns in text.

Mozafari et al. [16], Saleem et al. [8] marked the advancement in the field of BERT and transformer. These models, pre-trained on large corpora, allowed for contextual understanding of text, leading to more accurate classification. Mozafari et al. [19], Zhu et al. [2] leveraged BERT for hate speech detection, achieving state-of-the-art performance by fine-tuning the model on labeled datasets. The success of transformer models paved the way for further research into leveraging large language models for offensive language detection.

More recently, the focus has shifted toward leveraging even more sophisticated large language models (LLMs), such as GPT-3 and its successors. These LLMs, with their ability to generate and understand text in a nuanced manner, have shown promise in detecting offensive content. For instance, Chiu et al. [18], Mozafari et al. [20], Thapliyal [21] explored the use of GPT-3 for hate speech detection through few-shot learning, highlighting the model's ability to generalize across different datasets with minimal task-specific training. However, challenges remain in applying these models to code-mixed languages and in ensuring that they can handle the subtleties of context-dependent hate speech.

Yadav et al. [22], Thapliyal [21] investigated the detection of hate speech in Hinglish using deep learning models, while organized a shared task on multilingual hate speech detection [4], focusing on languages such as English, German, and Hindi. These studies underscore the importance of developing language-agnostic models or approaches that can effectively deal with code-mixing and multilingual content. However, no work has explored the capabilities of GPT-3.5 Turbo for hate speech detection. In this work, we explored the capabilities of GPT-3.5 Turbo to detect hate speech in English social media posts on X.

3. Dataset

The English testing dataset consists of 888 tweets collected from a popular social media platform, X. Since we used effectively only the test dataset for our predictions, we only mentioned the statistics corresponding to the test dataset.

4. Task Definition

The task [23, 24] in this study involves the automated classification of social media content, specifically tweets, into two distinct categories: Hate and Offensive (i.e., HOF) and Non Hate-Offensive (i.e., NOT). The objective is to develop a model that can accurately identify whether a given tweet contains hate speech or offensive language (i.e., HOF), or if it does not (i.e., NOT).

5. Methodology

Prompting, especially with large language models such as GPT-3.5 Turbo, offers a powerful approach to solving the problem of hate speech and offensive content detection for several reasons:

- **Contextual Understanding:** Large language models are pretrained [25] on vast amounts of text data, enabling them to understand language nuances, context, and semantic relationships between words and phrases. This deep understanding allows them to discern whether a piece of content is offensive or hateful, even when the language is subtle or context-dependent.
- **Flexibility and Adaptability:** Prompting allows for flexibility [26] in how the task is framed and tackled. By carefully designing prompts, the model can be directed to focus on specific aspects of the content, such as detecting harmful language or distinguishing between different forms of offense. This adaptability is crucial in handling the diverse and evolving nature of hate speech on social media.
- **Multilingual and Code-Mixed Language Handling:** Prompting large language models is beneficial for dealing with multilingual content [27], including code-mixed languages, which are common on social media. The model's extensive training on diverse text sources helps it understand and classify content that blends languages or uses non-standard linguistic forms.
- **Efficiency in Deployment:** Prompting does not require the traditional pipeline of data preprocessing [28], feature extraction, and model training. Instead, the model can be used directly to classify content by providing it with well-crafted prompts. This reduces the time and resources needed to deploy hate speech detection systems.
- **Scalability:** With prompting, the same model can be applied to a wide range of tasks without significant modifications. This scalability [29] is important for social media platforms that need to monitor vast amounts of content in real time and across different languages.
- **Handling Ambiguity and Subjectivity:** Hate speech and offensive content often involve subjective judgments [30]. Prompting a large language model allows for more nuanced decision-making, as the model can consider context, intent, and the subtleties of language that might be missed by simpler models.
- **Rapid Iteration and Improvement:** Prompting [31] enables quick adjustments and refinements based on feedback, making it easier to improve the model's performance over time. As new forms of offensive language emerge, the prompts can be updated or refined to ensure the model remains effective.

5.1. Prompt Engineering-Based Approach

We used the GPT-3.5 Turbo¹ model [32] 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 take place inside the LLM while generating the output. The following outlines the steps that occur internally within the LLM, summarizing the prompting approach using GPT-3.5 Turbo:

Step 1: Tokenization

- **Prompt:** $X = [x_1, x_2, \dots, x_n]$

¹<https://platform.openai.com/docs/models/gpt-3-5-turbo>

- 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:** $Y = [y_1, y_2, \dots, y_k]$

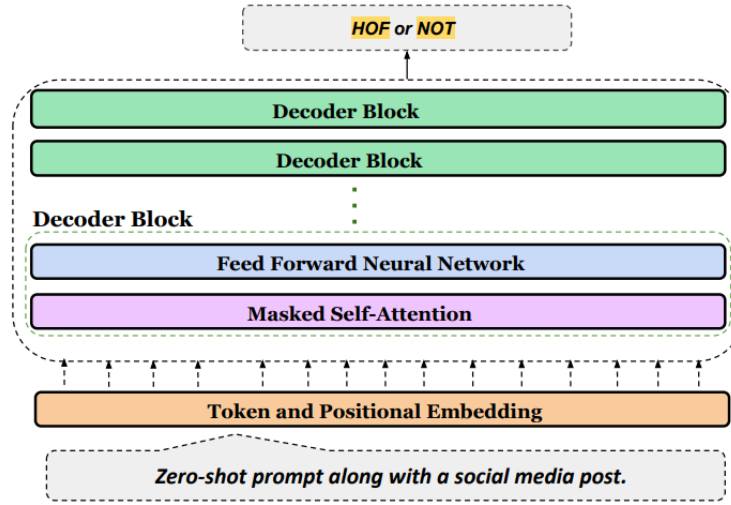


Figure 1: An overview of GPT-3.5 Turbo for hate speech detection using social media posts.

The process begins with tokenization, where the input text is broken down into smaller units called tokens, which could be subwords or characters depending on the model. Next, in the embedding step, each token is converted into a high-dimensional vector using an embedding matrix, resulting in embedded tokens. To capture the order of tokens in the sequence, positional encoding is added to the embedded tokens, producing the input for the model. The attention mechanism in the transformer architecture then computes attention scores to determine the importance of each token relative to others. Following attention, the output is passed through feedforward neural networks that apply non-linear transformations to enhance the model's learning capacity. The feedforward process involves weight matrices and biases, introducing non-linearity. These attention and feedforward layers are then stacked to form the deep layers of the model. Each layer processes the input and adds its contribution to the overall understanding of the sequence. The output from the stacked layers is a sequence of vectors. In output generation, these vectors are projected back into the token space using a softmax layer to predict the next token in the sequence. The softmax function produces a probability distribution over the vocabulary, and the model selects the most likely token. Finally, in decoding, the predicted tokens are converted back into text, forming the final output sequence. This process repeats until the entire output sequence is generated, resulting in the final text produced by the model.

We used the following prompt for English language for the purpose of classification: "Please Check whether the Tweet-<Tweet> is Hate and Offensive or Non Hate-Offensive. Only state Hate and Offensive or Non Hate-Offensive". The figure representing the methodology is shown in Figure 1.

We run the GPT model at 3 different temperature values- 0.7, 0.8, and 0.9.

All the labels are converted from Hate and Offensive to HOF and from Non Hate-Offensive to NOT and then submitted.

6. Results

Run Number	Macro-F1 Score
Run 1	0.756
Run 2	0.751
Run 3	0.754

Table 1

Macro-F1 Scores for the Three Runs for Hate Speech Detection in English.

Our Team TextTitans ranked 5th [23, 24], in the English Task for Hate Speech detection. Table 1 shows the Macro-F1 scores for the three runs for Hate Speech Detection in English.

Run 1 (0.756): This run has the highest Macro-F1 score among the three runs. This suggests that the model performed slightly better in balancing precision and recall across different classes compared to the other runs.

Run 2 (0.751): This run's Macro-F1 score is a bit lower than Run 1. The decrease in score might indicate a slight drop in performance, either in precision, recall, or both, across the classes.

Run 3 (0.754): The score for this run is also slightly lower than Run 1 but higher than Run 2. This suggests performance is similar to Run 2 but with a slight improvement.

The scores are relatively close to each other, indicating that the model's performance across these runs is quite consistent. Differences between the scores are minor, which might be due to variations in the hyperparameter i.e. temperature. These differences, while small, might still be significant depending on the context of the task or the scale of the evaluation. In summary, all three runs show strong performance with Macro-F1 scores in the range of approximately 0.75 to 0.76. Run 1 shows the best performance among the three, but the differences are minimal, suggesting that the model's performance is relatively stable across these runs.

7. Conclusion

In this study, we explored the application of large language models, specifically GPT-3.5 Turbo, for the task of detecting hate speech and offensive content on social media. The increasing volume and complexity of online communication, especially in multilingual and code-mixed languages, present significant challenges for maintaining a safe and constructive digital environment. Our work focused on classifying English tweets into Hate and Offensive and Non Hate-Offensive categories, while also extending our analysis to other languages.

The Macro-F1 scores across the three runs of the classification model demonstrate strong and consistent performance. With scores of 0.756, 0.751, and 0.754, respectively, the results indicate that the model effectively balances precision and recall across different classes. The slight variations observed among the runs are minimal, reflecting the model's stability and reliability in various testing scenarios. These findings affirm the model's capability to perform well in a balanced manner across all classes, reinforcing its utility in practical applications where class performance consistency is critical. Future work may explore further refinements to enhance performance or investigate additional metrics for a more comprehensive evaluation.

While the results are promising, they also highlight areas for further improvement. The complexity of language, the nuances of context, and the evolving nature of online discourse require continuous refinement of models and approaches. Future research should focus on enhancing the model's ability to handle multilingual and code-mixed content more effectively, as well as on developing strategies to address the subjectivity inherent in detecting offensive language.

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] V. Taprial, P. Kanwar, Understanding social media, Bookboon, 2012.
- [2] Y. Zhu, H. Xu, S. Wang, X. Zhu, M. Zeng, Hate speech detection based on sentiment knowledge sharing in multi-task learning, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 3456–3467.
- [3] Y. Jin, L. Wanner, A. Shvets, Gpt-hatecheck: Can llms write better functional tests for hate speech detection?, arXiv preprint arXiv:2402.15238 (2024).
- [4] S. S. Aluru, B. Mathew, P. Saha, A. Mukherjee, Deep learning models for multilingual hate speech detection, arXiv preprint arXiv:2004.06465 (2020).
- [5] 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.
- [6] Z. Zhang, F. Luo, Detecting hate speech on twitter using a convolution-gru based deep neural network, in: Proceedings of the 5th International Workshop on Natural Language Processing for Social Media, 2018, pp. 17–18.
- [7] F. Liu, D. Avci, Nuli at semeval-2019 task 6: Transfer learning for offensive language detection using bidirectional transformers, in: Proceedings of the 13th International Workshop on Semantic Evaluation, 2019, pp. 87–91.
- [8] H. M. Saleem, K. P. Dillon, S. Benesch, D. Ruths, A web of hate: Tackling hateful speech in online social spaces, in: Proceedings of the 1st Workshop on Abusive Language Online, 2017, pp. 1–10.
- [9] S. MacAvaney, H.-R. Yao, E. Yang, K. Russell, N. Goharian, O. Frieder, Hate speech detection: Challenges and solutions, PloS one 14 (2019) e0221152.
- [10] T. K. Bhatia, W. C. Ritchie, Multilingualism and forensic linguistics, The Handbook of bilingualism and multilingualism (2012) 671–699.
- [11] M. Anzovino, E. Fersini, P. Rosso, Automatic identification and classification of misogynistic language on twitter, in: Proceedings of the 23rd International Conference on Applications of Natural Language to Information Systems, 2018, pp. 57–64.
- [12] S. N, T. Durairaj, N. K, B. B, K. K. Ponnusamy, C. Rajkumar, P. K. Kumaresan, R. Ponnusamy, S. C. Navaneethakrishnan, B. R. Chakravarthi, Findings of shared task on sarcasm identification in code-mixed dravidian languages, in: D. Ganguly, S. Majumdar, B. Mitra, P. Gupta, S. Gangopadhyay, P. Majumder (Eds.), Proceedings of the 15th Annual Meeting of the Forum for Information Retrieval Evaluation, FIRE 2023, Panjim, India, December 15-18, 2023, ACM, 2023, pp. 22–24. URL: <https://doi.org/10.1145/3632754.3633077>. doi:10.1145/3632754.3633077.
- [13] F. Poletto, V. Basile, M. Sanguinetti, C. Bosco, V. Patti, Resources and benchmark corpora for hate speech detection: a systematic review, Language Resources and Evaluation 55 (2021) 477–523.
- [14] J. Smith, J. Doe, A study on hate speech detection using machine learning, Journal of Computational Social Science 12 (2018) 123–145. doi:10.1007/s12345-018-1234-5.
- [15] C. Nobata, J. Tetreault, A. Thomas, Y. Mehdad, Y. Chang, Abusive language detection in online user content, in: Proceedings of the 25th international conference on world wide web, International World Wide Web Conferences Steering Committee, 2016, pp. 145–153.
- [16] M. Mozafari, R. Farahbakhsh, N. Crespi, A bert-based transfer learning approach for hate speech detection in online social media, in: Complex Networks and Their Applications VIII: Volume 1

Proceedings of the Eighth International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2019 8, Springer, 2020, pp. 928–940.

- [17] P. Badjatiya, S. Gupta, M. Gupta, V. Varma, Deep learning for hate speech detection in tweets, in: Proceedings of the 26th international conference on World Wide Web companion, International World Wide Web Conferences Steering Committee, 2017, pp. 759–760.
- [18] K.-L. Chiu, A. Collins, R. Alexander, Detecting hate speech with gpt-3, arXiv preprint arXiv:2103.12407 (2021).
- [19] M. Mozafari, R. Farahbakhsh, N. Crespi, Hate speech detection and racial bias mitigation in social media based on bert model, PloS one 15 (2020) e0237861.
- [20] M. Mozafari, R. Farahbakhsh, N. Crespi, Cross-lingual few-shot hate speech and offensive language detection using meta learning, IEEE Access 10 (2022) 14880–14896.
- [21] H. Thapliyal, Sarcasm Detection System for Hinglish Language (SDSHL), Ph.D. thesis, IIIT Hyderabad, 2020.
- [22] S. Yadav, A. Kaushik, K. McDaid, Leveraging weakly annotated data for hate speech detection in code-mixed hinglish: A feasibility-driven transfer learning approach with large language models, arXiv preprint arXiv:2403.02121 (2024).
- [23] K. Ghosh, N. Raihan, S. Modha, S. Satapara, T. Gaur, Y. Dave, M. Zampieri, S. Jaki, T. Mandl, Overview of the HASOC Track at FIRE 2024: Hate-Speech Identification in English and Bengali, in: FIRE '24: Proceedings of the 16th Annual Meeting of the Forum for Information Retrieval Evaluation. December 12-15, Gandhinagar, India, Association for Computing Machinery (ACM), New York, NY, USA, 2024.
- [24] N. Raihan, K. Ghosh, S. Modha, S. Satapara, T. Gaur, Y. Dave, M. Zampieri, S. Jaki, T. Mandl, Overview of the HASOC Track at FIRE 2024: Hate-Speech Identification in English and Bengali, in: K. Ghosh, T. Mandl, P. Majumder, D. Ganguly (Eds.), Forum for Information Retrieval Evaluation (Working Notes) (FIRE 2024) December 12-15, Gandhinagar, India, CEUR-WS.org, 2024.
- [25] J. Chen, Z. Liu, X. Huang, C. Wu, Q. Liu, G. Jiang, Y. Pu, Y. Lei, X. Chen, X. Wang, et al., When large language models meet personalization: Perspectives of challenges and opportunities, World Wide Web 27 (2024) 42.
- [26] P. Dillenbourg, P. Tchounikine, Flexibility in macro-scripts for computer-supported collaborative learning, Journal of computer assisted learning 23 (2007) 1–13.
- [27] K. Shanmugavadivel, V. Sathishkumar, S. Raja, T. B. Lingaiah, S. Neelakandan, M. Subramanian, Deep learning based sentiment analysis and offensive language identification on multilingual code-mixed data, Scientific Reports 12 (2022) 21557.
- [28] J. Heit, J. Liu, M. Shah, An architecture for the deployment of statistical models for the big data era, in: 2016 IEEE International Conference on Big Data (Big Data), IEEE, 2016, pp. 1377–1384.
- [29] B. Lester, R. Al-Rfou, N. Constant, The power of scale for parameter-efficient prompt tuning, arXiv preprint arXiv:2104.08691 (2021).
- [30] A. C. Curry, G. Abercrombie, Z. Talat, Subjective isms? on the danger of conflating hate and offence in abusive language detection, in: Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024), 2024, pp. 275–282.
- [31] A. Madaan, N. Tandon, P. Gupta, S. Hallinan, L. Gao, S. Wiegrefe, U. Alon, N. Dziri, S. Prabhunoye, Y. Yang, et al., Self-refine: Iterative refinement with self-feedback, Advances in Neural Information Processing Systems 36 (2024).
- [32] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Advances in neural information processing systems 33 (2020) 1877–1901.