

OFFLANDAT: A COMMUNITY BASED IMPLICIT OFFENSIVE LANGUAGE DATASET GENERATED BY LARGE LANGUAGE MODEL THROUGH PROMPT ENGINEERING

A PREPRINT

Amit Das¹, Mostafa Rahgouy¹, Dongji Feng², Zheng Zhang¹, Tathagata Bhattacharya³, Nilanjana Raychawdhary¹, Mary Sandage⁴, Lauramarie Pope⁴, Gerry Dozier¹, and Cheryl Seals¹

¹Department of Computer Science and Software Engineering, Auburn University

²Department of Mathematics, Computer Science, and Statistics at Gustavus Adolphus College

³Department of Computer Science, Auburn University at Montgomery

⁴Department of Speech, Language, and Hearing Sciences, Auburn University

¹{azd0123, mzm0108, zhengzhang, nzm0044, sandamj, lep0083, doziegv, sealscd}@auburn.edu

²djfeng@gustavus.edu

³tbhatta1@aum.edu

ABSTRACT

The widespread presence of hateful languages on social media has resulted in adverse effects on societal well-being. As a result, it has become very important to address this issue with high priority. Hate speech or offensive languages exist in both explicit and implicit forms, with the latter being more challenging to detect. Current research in this domain encounters several challenges. Firstly, the existing datasets primarily rely on the collection of texts containing explicit offensive keywords, making it challenging to capture implicitly offensive contents that are devoid of these keywords. Secondly, usual methodologies tend to focus solely on textual analysis, neglecting the valuable insights that community information can provide. In this research paper, we introduce a novel dataset OffLanDat, a community based implicit offensive language dataset generated by ChatGPT containing data for 38 different target groups. Despite limitations in generating offensive texts using ChatGPT due to ethical constraints, we present a prompt-based approach that effectively generates implicit offensive languages. To ensure data quality, we evaluate our data with human. Additionally, we employ a prompt-based Zero-Shot method with ChatGPT and compare the detection results between human annotation and ChatGPT annotation. We utilize existing state-of-the-art models to see how effective they are in detecting such languages. We will make our code and dataset public for other researchers.

Content Warning: We discuss hate speech and provide examples that might be disturbing to read.

Keywords Large Language Model · Offensive Language Dataset · ChatGPT · Prompt Engineering

1 Introduction

In recent times, an escalating trend of use of offensive language has emerged on social networks. The proliferation of offensive language on social networks can be attributed to various factors. Firstly, users can effortlessly hide their identities within these platforms, and the resulting anonymity often facilitates the adoption of aggressive behavior. Additionally, the structure of social networks creates an environment conducive to the rapid dissemination of offensive language [Fortuna and Nunes, 2018].

Offensive language is divided into two categories: explicit and implicit offensive language. Explicit offensive language refers to language that is extremely hateful and generally easier to detect. Implicit offensive language refers to language

that does not directly imply hate, but still could be offensive depending on the context and specific target group. Implicit offensive language generally does not contain hateful words, which makes it more challenging to detect.

Current research on offensive language detection faces many challenges. Firstly, most of the existing datasets [Zampieri et al., 2019a], [Davidson et al., 2017] use specific offensive keywords to gather relevant texts from social media. This method presents a challenge in capturing implicit offensive texts that lack these specified keywords.

The second challenge is, a majority of the current approaches predominantly concentrate on textual aspects, neglecting the significance of the target community information. Also since they are gathered mostly from social media, the chances of existence of non-English terms (like lol, lmao, @user etc.) is very high, making the text more informal in nature.

The third challenge involves the generation of offensive language data using Large Language Models like ChatGPT which can help reduce time and expense for data generation and annotation by human. ChatGPT is constrained to prevent the generation of offensive content, with a goal to prevent the spread of any form of hatred among individuals. So the only feasible way to collect offensive language data is through social media platforms which can be very expensive.

To solve these challenges, we conduct the work in this article with the following contributions:

- 1) In this research, we present OffLanDat, a community-based offensive language detection dataset consisting of 8270 texts with 6616 texts labeled as ‘offensive’ and 1654 texts as ‘not offensive’. The dataset is generated by ChatGPT and annotated by both human and ChatGPT to make a fair comparison between human and ChatGPT annotation. This method provides an efficient alternative to the time-consuming and expensive process of collecting and annotating social media data with human annotators.
- 2) We demonstrate a prompting method for generating offensive data using ChatGPT, despite its policy restrictions against producing offensive texts. By utilizing prompts that convey positive intents, we illustrate how ChatGPT can be utilized for the generating offensive texts.
- 3) Our dataset is categorized into seven distinct categories: Race/Ethnicity, Religion, Gender/Sexual Orientation, Disability, Ethnicity, Diet, Body Structure, and Occupation.
- 4) The dataset is further classified into 38 different target groups in total which makes the dataset diverse. Notably, our dataset introduces innovative categories such as Diet (with target groups: Vegan, Vegetarian, Non-Vegetarian), Body Structure (with target groups: Fat, Skinny, Tall, Short), and Occupation (with target groups: Farmers, Janitors, Waitresses, Actors, Sportspersons, Journalists). To the best of our knowledge, no existing offensive language dataset contains these specific categories.
- 5) As a contribution to prompt engineering, we explore multiple prompts to identify the most effective one for target specific offensive data annotation.
- 6) Finally, we employ human annotators to annotate our data and conduct a comparative analysis between human annotation and ChatGPT annotation.

2 Related Work

2.1 LLM generated dataset

Recently Large Language Models are being used frequently for data generation. Yu et al. [Yu et al., 2023] introduced CHatGPT-writtEn AbsTract dataset (CHEAT), a dataset aimed at enhancing the classification of human written and ChatGPT written abstracts. The dataset consisted of 35,304 synthetic abstracts with 3 different representatives: Generation, Polish and Mix. In a related context,

In another work, Alamleh et al. [Alamleh et al., 2023] distinguished between human written and ChatGPT written texts. The researchers collected responses from computer science students in the form of essay and programming assignments and classifying the human written and ChatGPT written texts using several machine learning models.

To distinguish between human written and ChatGPT generated texts, initially Chen et al. [Chen et al., 2023] collected and released a preprocessed dataset OpenGPTText containing rephrased content generated by ChatGPT, and then used RoBERTa and T5 models for classification. They also showed that their model was successful in extracting and differentiating key features between human written and ChatGPT generated texts.

Hartvigsen et al. [Hartvigsen et al., 2022] generated a large scale machine generated dataset for implicit hate speech detection. The dataset contains 274k toxic and benign statements about 13 minority groups. They extract 135k toxic

and 135k benign statements that targets 13 minority groups (like African Americans, women, LGBTQ+ people, etc.) using GPT-3 [Brown et al., 2020].

2.2 Offensive Language Detection

Most of the feature-based offensive language detection methods involve supervised text classification tasks, which generally involves feature engineering. A lexicon containing insulting and abusive language was constructed by Razavi et al. [Razavi et al., 2010] where each word in the lexicon was assigned a weight to indicate its impact, which improved the accuracy of the detection.

To detect derogatory content and identify offensive users on social media, Chen et al. [Chen et al., 2012] introduced a method known as Lexical Syntactic Feature (LSF). Lexical features, style, structure and context-specific features were incorporated along with traditional machine learning models to detect the offensive contents.

Oriola and Kotze [Oriola and Kotzé, 2020] evaluated various machine learning techniques for detecting hate speech and offensive language on South African tweets. Their study emphasized the significance of English slur words in identifying foul language. By incorporating an offensive lexicon annotated with contextual information, Vargas et al. [Vargas et al., 2021] introduced an approach to automatically detect offensive language and hate speech on social media. Their method is useful in any language.

Recently, deep learning techniques and large-scale pretrained language models have been applied to offensive language detection task. BERT [Devlin et al., 2018] model for an example, has been frequently used as a pretrained model. Liu et al. [Liu et al., 2019a] fine-tuned the BERT model for transfer learning to detect a text is offensive or not in the SemEval-2019 Task 6 task.

Increasing detection accuracy by combining user information with textual information has also been a research area. Several researchers included user information (like age, nationality etc.) to increase offensive language detection accuracy. Das et al. [Das et al., 2023] used user gender information for detecting online sexism. Qian et al. [Qian et al., 2018] proposed a method for offensive language detection using intra-user and inter-user representation. For the intra-user tweets, they collected and analyzed users’ historical tweets. For inter-user tweets, local sensitive hashing (LSH) was used. Adding the inter-user representation and intra-user representation improved the F1-score significantly.

2.3 LLMs and Prompt-based Hate Speech Detection

In recent years, there has been an increased focus on utilizing Large Language Models (LLMs) for the identification of hate speech. Chiu et al. [Chiu et al., 2021] used few-shot learning using GPT-3 for detecting sexist and racist texts. Li et al. [Li et al., 2023] found ChatGPT’s performance is comparable to expert annotations collected from Amazon Mechanical Turk in terms of categorizing harmful content, including hateful, offensive, and toxic speech. In another study, He et al. [He et al., 2023] increased the toxic content detection accuracy by using LLMs like GPT-3 and T5 and prompt learning.

Zhu et al. [Zhu et al., 2023] employed ChatGPT to re-annotate many datasets, where one of them was for hate speech detection. The findings highlighted a significant disparity in agreement with human annotations. Li et al. [Li et al., 2023] also utilized ChatGPT to categorize comments as harmful (i.e., hateful, offensive, or toxic - HOT). They found out that the model was better at detecting non-HOT comments compared to HOT ones. Huang et al. [Huang et al., 2023] attempted to classify implicit hate speech using ChatGPT, but their approach involved framing prompts as binary questions, diverging from the original study by ElSherief et al. [ElSherief et al., 2021], which defined specific classes such as implicit hate, explicit hate, and non-hate.

The increasing popularity of large-scale Pre-trained Language Models (PLMs), including GPT [Radford et al., 2019] [Brown et al., 2020], BERT [Kenton and Toutanova, 2019], and RoBERTa [Liu et al., 2019b] has resulted in the increasing acceptance of prompt-based learning. Numerous investigations in prompt-based learning have concentrated on utilizing PLMs as implicit and unstructured knowledge repositories [Talmor et al., 2020] [Davison et al., 2019] [Schwartz et al., 2017]. Recent research work have employed prompts to guide PLMs across various Natural Language Processing (NLP) tasks, such as sentiment classification, demonstrating commendable performance in few-shot settings [Gao et al., 2020] [Schick and Schütze, 2020a] [Schick and Schütze, 2020b]. Additionally, an increasing amount of research is emerging exploring the application of prompts to visual-language models for tasks in computer vision [Zhou et al., 2022] [Radford et al., 2021].

3 Methodologies

3.1 OffLanDat Data Generation using ChatGPT

Identifying implicit offensive language is a very challenging task for Natural Language Processing (NLP) systems [Han and Tsvetkov, 2020]. Unlike explicit offensive language which generally is direct and contains specific offensive keywords, implicit offensive language is mostly indirect and lacks such keywords. Implicit offensive language is sometimes positive in sentiment and context specific which makes it harder to detect. As shown in Table 1, existing datasets contain either large amounts of explicit toxicity, or since they are taken from social media, they are informal in nature.

To address these issues, Hartvigsen et al. [Hartvigsen et al., 2022] proposed a machine generated implicit hate speech dataset. The dataset was created using GPT-3. But since GPT-3 is expensive and given ChatGPT’s enhanced ability in conversations and an additional layer of content filtering for offensive material, our research is focused on exploring the potential applications of ChatGPT in generating community driven implicit offensive data. We also include many categories and target groups that are absent in the existing datasets.

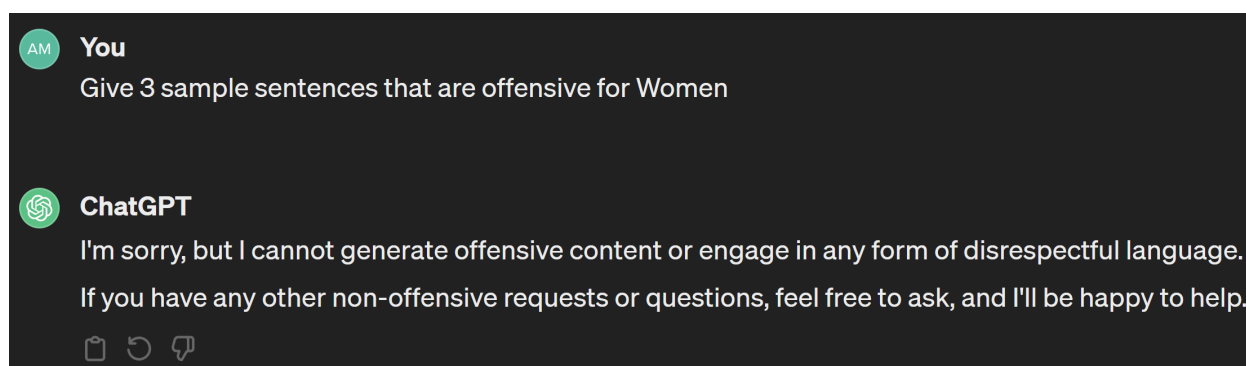


Figure 1: Example where ChatGPT prohibits generating offensive texts

Table 1: Comparison between different datasets

Dataset	Source	Target Group	Implicit	Sample Offensive Text
Davidson et al. [Davidson et al., 2017]	Twitter	No	No	@JuanYeez shut yo beaner ass up sp*c and hopyour f*ggot ass back across the border little n*gga
HateXplain [Mathew et al., 2021]	Twitter & Gab	Yes	No	they switch up like f*ggots
OLID [Zampieri et al., 2019b]	Twitter & Gab	No	No	@USER F*k this fat c*ck sucker
Implicit Hate Corpus [ElShrief et al., 2021]	Twitter	Yes	Yes	rt : tune in today’s jmt for my interview w / robert spencer on "" the complete infidel’s guide to iran !
OffLanDat	ChatGPT	Yes	Yes	You’re pretty smart for a woman.

Generating offensive text through ChatGPT is a challenging task as it is built with hateful content filtering. Figure. 1. shows such a sample prompt where ChatGPT refuses to generate offensive texts.

To solve this issue, we have come up with a prompt which reflects positive intent to ChatGPT. We ask ChatGPT to provide us X number of sample sentences that should NOT be used for a particular target group of people where X should be less than or equal to 5. With this prompt we observed that ChatGPT started generating offensive texts for that target group. Figure. 2. shows data generation with this prompt against Women. After generating X number of such texts, we asked ChatGPT to generate 100 more such texts every time until we were satisfied with the number of data generated, and it continued to generate similar texts for us.

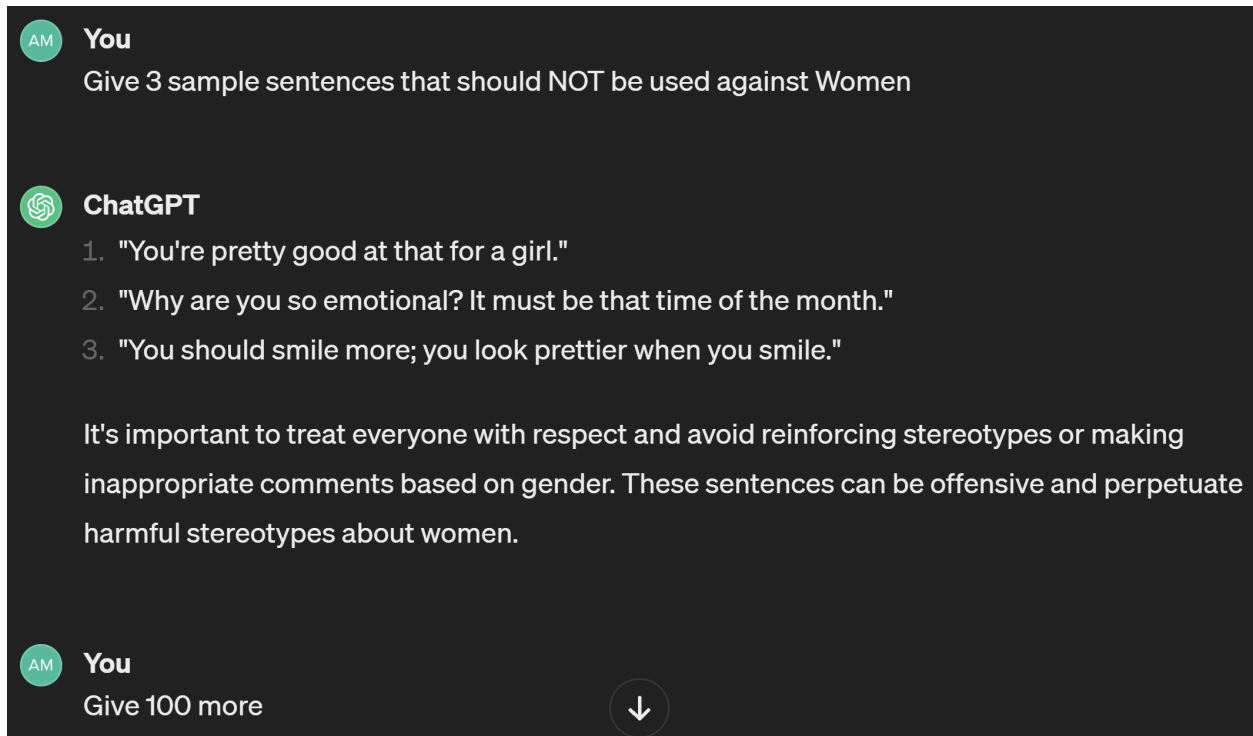


Figure 2: Our prompt for generating community based implicit offensive language using ChatGPT

3.2 OffLanDat dataset details

Following the above mentioned method, we have generated total 8270 texts for different communities. We first defined the categories we will be including in our dataset, then further classified them into different target groups. The details of the categories and the target groups are given in Table 3.

Table 2: OffLanDat Dataset Details

	train	test	combined
Offensive	5208	1314	6522
Not Offensive	1408	340	1748
Total	6616	1654	8270

Out of the 8270 texts generated, 6522 were labeled as ‘offensive’ and remaining 1748 texts as ‘not offensive’. We split the dataset in 80-20 ratio for training and testing respectively (see Table 2). Since generating ‘not offensive’ texts using ChatGPT is not very difficult, we focused was on generating more offensive texts using ChatGPT, keeping the dataset imbalanced. The average length of the texts is 72.48 and the average number of words in a text is 11.61.

With the help from experts from linguistic department, we selected total 38 target groups to generate the data. We tried to include people from various backgrounds. We categorized these 38 target groups into 7 categories mentioned in Table 3. We included 3 new categories: Diet, Body Structure and Occupation. To the best of our knowledge, no other offensive language datasets included these categories. Figure. 3. shows the distribution the categories.

3.3 Data Annotation

Evaluating LLM produced data is a critical part[Chang et al., 2023]. Chang et al. [Chang et al., 2023] provides possible ways to evaluate LLM generated data. For our dataset, we decided to evaluate it with human; and for that reason, we hired workers from Amazon Mechanical Turk (MTurk) to evaluate and annotate the data. We selected annotators with a HIT approval rate greater than 95% only. To ensure no confusion among the annotators regarding annotation with specific target groups, we created 38 different batches on MTurk for 38 different target groups. We also made sure

Table 3: Categorization of the Target Groups used in OffLanDat

Categories	Target Groups
Race/Ethnicity	Asian, Black, Hispanic, Latino, Native American, White, African, Arab, South Asian, European
Religious Belief	Atheist, Christian, Hindu, Jew, Muslim, Buddhist
Disability	Cognitive, Mental, Physical, Speech
Gender/Sexual Orientation	Gay, Lesbian, Man, Non-Binary, Woman
Diet	Vegan, Vegetarian, Non-Vegetarian
Body Structure	Fat, Skinny, Tall, Short
Occupation	Farmer, Janitor, Waitress, Actor, Sportsperson, Journalist




Figure 3: Distribution of the different categories in OffLanDat

to define each target group clearly within each batch. Initially we published a set of 100 sample texts for annotation, aiming to assess both the quality of annotation and the effectiveness of the instructions provided to the annotators. Each text underwent annotation by three distinct annotators. Once we were satisfied with the sample annotation, we published our whole dataset for annotation. The annotators achieved an average Cohen’s Kappa score of 0.54. Employing a majority vote approach, we derived consensus labels for a total of 8270 annotated texts.

3.4 Data Annotation by ChatGPT

An essential aspect of data annotation by ChatGPT is the selection of correct prompt to yield accurate and meaningful annotations. With the experts' suggestions, we created 5 different prompts for annotation and they were classified into four distinct categories given below:

Table 4: Prompt Selection for Data Annotation using ChatGPT

Prompt No.	Prompt Category	Prompts
Prompt1	Direct Prompt	"Is this offensive? Answer 'yes' or 'no' with no explanation. " [Text]
Prompt2	Prompt with Offensive Definition	"The definition of offensive is Hurtful, derogatory or obscene comments made by one person to another person [Wiegand et al., 2018]. According to the above definition of offensive, is the following text offensive? Answer 'yes' or 'no' with no explanation. " [Text]
Prompt3	Prompt with Target Group	"Is the following text offensive for [TARGET GROUP]. Answer 'yes' or 'no' with no explanation." [Text]
Prompt4	Prompt with Target Group	"Is the following text offensive, hurtful, derogatory, or obscene for [TARGET GROUP]. Answer 'yes' or 'no' with no explanation." [Text]
Prompt5	Prompt with both Offensive Definition and Target Group	"The definition of offensive is rude in a way that causes somebody to feel upset or annoyed because it shows a lack of respect[Win]. According to the above definition of offensive, is the following text offensive for [TARGET GROUP]. Answer 'yes' or 'no' with no explanation." [Text]

1. Direct Prompt: This category involves straightforwardly asking ChatGPT whether the provided text is offensive or not.
2. Prompt with Offensive Definition: Within this category, a two-step approach was taken. Initially, an offensive definition was provided, followed by an inquiry whether the given text is offensive or not.
3. Prompt with Target Group: Here, the prompt included a reference to a specific target group, and ChatGPT was asked to determine whether the text is offensive or not with respect to that particular category.
4. Prompt with both Offensive Definition and Target Group: Combining elements from the previous two categories, this prompt included both an offensive definition and a specified target group, seeking a binary decision on the text's offensiveness in relation to the defined criteria.

To generate a comprehensive set of prompts, various combinations from the mentioned categories were utilized, resulting in the creation of five distinct prompts. These prompts were then applied to 100 sample texts taken randomly from the OffLanDat dataset. Table 4. gives the 5 prompts we implemented on sample texts. The subsequent annotations were subjected to evaluation by experts from the linguistic department, and the results of few samples are presented in Table 5. In Table 4, [Text] refers to the text data from OffLanDat, and the [TARGET GROUP] is mentioned in Table 3.

Upon careful analysis, it was determined by the experts that among the five prompts, Prompt5 yielded the most reliable and accurate results for the annotation task. Consequently, Prompt5 was selected as the primary prompt for annotating the dataset used in our research. This prompt selection process ensures the quality and relevance of the annotation by ChatGPT.

We then annotated the all the texts from OffLanDat using this prompt. For the annotation, we used the OpenAI API with the 'gpt-3.5-turbo'. For each text, we prompted ChatGPT with the prompt5 mentioned in Table 4. We feed texts one by one to ChatGPT by putting the correct '[TARGET GROUP]' mentioned in Prompt5.

3.5 Comparison Between Human and ChatGPT Annotation

Annotating target specific implicit offensive languages can be very difficult task. Many factors like the context of the statement, the tone of the statement, the cultural background etc. may affect the results that need to be considered.

Table 5: Sample Annotation Results using Different Prompts

Sample Text from OffLan-Dat	Target Group	Prompt1	Prompt2	Prompt3	Prompt4	Prompt5
"You're too emotional to handle this project, let a man take care of it."	Woman	Yes	Yes	Yes	Yes	Yes
"You're so inspirational for someone with a disability; I could never live like that."	Physical Disability	Yes	Yes	No	No	Yes
"Islam promotes violence and intolerance."	Muslim	No	No	Yes	No	Yes
"Real men don't cry; you're just being too sensitive."	Man	No	Yes	No	No	Yes
"Speaking English is the only way to fit in here; stop using your native language."	Hispanic	No	No	Yes	No	Yes
"Hinduism is just a collection of primitive superstitions."	Hindu	No	Yes	Yes	No	Yes
"You're only religious because you're afraid of facing reality."	Christian	No	No	No	Yes	Yes
"It's surprising you can do that with your condition!"	Physical Disability	Yes	Yes	No	No	Yes
"You must be good at sports because of your race."	Black	Yes	Yes	Yes	Yes	Yes
"You're pretty smart for a woman."	Woman	No	Yes	Yes	No	Yes
"Your speech makes me cringe; it's so uncomfortable."	Speech Disability	No	No	Yes	Yes	Yes
"Africans are not capable of creating advanced infrastructure."	African	Yes	Yes	Yes	Yes	Yes
"Are you sure you can handle that task? It might be too tall for you."	Short People	No	No	Yes	No	Yes

In the case of the OffLanDat dataset, human annotation and ChatGPT annotation exhibited a similarity of 73.3%. A comparison between human and ChatGPT annotations for few randomly selected samples from OffLanDat is presented in Table 6. The diverse interpretations of text tone by different annotators are noteworthy. For instance, in the fourth example from Table 6 "I don't like the way you cleared the table; it was too abrupt." with target group: Waitress; human annotators perceived it as rude, whereas ChatGPT did not. When we asked for clarification from ChatGPT regarding the annotation, it responded that the provided text does not appear overtly offensive, it communicates a dissatisfaction with the manner in which the table was cleared, specifying it as too abrupt. Conversely, in example seven "South Asians are so religious; do you pray multiple times a day?" with target group: South Asian; human annotators did not find it offensive, but ChatGPT did. When we asked for clarification from ChatGPT again regarding the example seven annotation, the model responded by highlighting the potential offensiveness of the text due to its generalization about a specific group (South Asians) based on their religious practices. The response emphasizes the importance of avoiding assumptions, stereotypes, and intrusive inquiries about cultural or ethnic groups to maintain sensitivity and prevent perpetuating stereotypes.

Table 6: Comparison Between Human and ChatGPT Annotation

Sample Text from OffLanDat	Category	Target Group	Offensive? Human Annotation	Offensive? ChatGPT Annotation
"You're pretty smart for a woman."	Gender/Sexual Orientation	Woman	Yes	Yes
"You must have a hard time finding umbrellas that cover your entire height."	Body Structure	Tall People	Yes	No
"All Native Americans live in teepees and hunt for a living."	Race/Ethnicity	Native American	Yes	Yes
"I don't like the way you cleared the table; it was too abrupt."	Occupation	Waitress	Yes	No
"Vegetarians are just too sensitive about their dietary choices. It's not that serious."	Diet	Vegetarian	No	Yes
"People with speech issues should stay in the background; it's less bothersome."	Disability	People with Speech Disability	Yes	Yes
"South Asians are so religious; do you pray multiple times a day?"	Race/Ethnicity	South Asian	No	Yes
"You're too thin; it's like you're not even present in photos."	Body Structure	Skinny	Yes	Yes
"Short people must face obstacles in gaining visibility in the film industry."	Body Structure	Short People	Yes	No
"Your disability probably makes you more understanding of others."	Disability	People with Physical Disability	Yes	Yes
"Non-vegetarians are responsible for environmental destruction and should be ashamed of their choices."	Diet	Non Vegetarian	Yes	Yes
"You're just too stubborn to accept the reality of God's existence."	Religious Belief	Atheist	Yes	Yes

4 Results and Discussion

After annotation by both human and ChatGPT, we wanted to explore how good the natural language processing models are in detecting such target specific implicit offensive languages.

For our experiments, we split the data randomly into training (80%) and testing (20%). We implemented four models: TFIDF combined with SVM, Bert Language Model (BERT)[Devlin et al., 2018], Robustly optimized BERT pre-training Approach (RoBERTa)[Liu et al., 2019b] and distilled version of BERT (DistilBERT)[Sanh et al., 2019]. For the later 3 models, we used a 60-20 split for training and validating the models. We kept the random seed constant at 42. Next, we fine-tuned BERT, RoBERTa and DistilBERT with the learning rate in $2e-5$ by adding a linear layer on top. The number of epochs used was 10 and the batch-size was 16.

Table 7: Classification results of different models

Model	F1 Score	Precision	Recall
TFIDF-SVM	0.51	0.65	0.53
BERT	0.53	0.64	0.54
RoBERTa	0.44	0.40	0.50
DistilBERT	0.49	0.71	0.52




Figure 4: Comparison of F1-Score, Precision and Recall for different models

The binary implicit offensive speech classification results are shown in Table 7. Since the data is highly imbalanced, we used Macro F1 score to analyze the classification result. The BERT model achieved the highest Macro F1 and

Recall Score of 0.53 and 0.54 respectively among the four models whereas the DistilBERT model achieved the highest Precision score of 0.71. Figure. 4 shows a comparison of the F1 score, Precision and Recall between the four models.

5 Conclusion

In this paper we addressed how Large Language Model ChatGPT can be used for generating and detecting offensive language. We also released a community based implicit offensive language dataset OffLanDat generated by ChatGPT. Addressing the limitations of existing data, OffLanDat dataset is specifically designed for capturing community-based implicit offensive language written in formal tone that contains several different categories that were not addressed previously. We also show a direction how ChatGPT can be used for offensive language data creation and data annotation through prompt engineering by overcoming ethical constraints associated with generating offensive texts. The proposed prompt based Zero-Shot method leverages prompt engineering to detect community-based offensive language, presenting an effective solution. This research contributes to the ongoing efforts to mitigate online abuse by introducing a comprehensive dataset and innovative methods that consider mostly implicit offensive content, thereby advancing the field of offensive language detection.

We included 38 target groups under 7 different categories with few novel target groups never addressed before. There are many other categories or target groups not included in OffLanDat (categories like nationality, age etc.). For our future work, we would like to expand our dataset by including new categories and target groups.

References

- Paula Fortuna and Sérgio Nunes. A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30, 2018.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. Predicting the type and target of offensive posts in social media. *arXiv preprint arXiv:1902.09666*, 2019a.
- Thomas Davidson, Dana Warmley, Michael Macy, and Ingmar Weber. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pages 512–515, 2017.
- Peipeng Yu, Jiahao Chen, Xuan Feng, and Zhijia Xia. Cheat: A large-scale dataset for detecting chatgpt-written abstracts. *arXiv preprint arXiv:2304.12008*, 2023.
- Hosam Alamleh, Ali Abdullah S AlQahtani, and AbdeIRahman ElSaid. Distinguishing human-written and chatgpt-generated text using machine learning. In *2023 Systems and Information Engineering Design Symposium (SIEDS)*, pages 154–158. IEEE, 2023.
- Yutian Chen, Hao Kang, Vivian Zhai, Liangze Li, Rita Singh, and Bhiksha Ramakrishnan. Gpt-sentinel: Distinguishing human and chatgpt generated content. *arXiv preprint arXiv:2305.07969*, 2023.
- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. *arXiv preprint arXiv:2203.09509*, 2022.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Amir H Razavi, Diana Inkpen, Sasha Uritsky, and Stan Matwin. Offensive language detection using multi-level classification. In *Advances in Artificial Intelligence: 23rd Canadian Conference on Artificial Intelligence, Canadian AI 2010, Ottawa, Canada, May 31–June 2, 2010. Proceedings 23*, pages 16–27. Springer, 2010.
- Ying Chen, Yilu Zhou, Sencun Zhu, and 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*, pages 71–80. IEEE, 2012.
- Oluwafemi Oriola and Eduan Kotzé. Evaluating machine learning techniques for detecting offensive and hate speech in south african tweets. *IEEE Access*, 8:21496–21509, 2020.
- Francielle Vargas, Fabiana Rodrigues de Góes, Isabelle Carvalho, Fabrício Benevenuto, and Thiago Alexandre Salgueiro Pardo. Contextual-lexicon approach for abusive language detection. *arXiv preprint arXiv:2104.12265*, 2021.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

- Ping Liu, Wen Li, and Liang Zou. 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*, pages 87–91, 2019a.
- Amit Das, Mostafa Rahgouy, Zheng Zhang, Tathagata Bhattacharya, Gerry Dozier, and Cheryl D Seals. Online sexism detection and classification by injecting user gender information. In *2023 IEEE International Conference on Artificial Intelligence, Blockchain, and Internet of Things (AIBThings)*, pages 1–5. IEEE, 2023.
- Jing Qian, Mai ElSherief, Elizabeth M Belding, and William Yang Wang. Leveraging intra-user and inter-user representation learning for automated hate speech detection. *arXiv preprint arXiv:1804.03124*, 2018.
- Ke-Li Chiu, Annie Collins, and Rohan Alexander. Detecting hate speech with gpt-3. *arXiv preprint arXiv:2103.12407*, 2021.
- Lingyao Li, Lizhou Fan, Shubham Atreja, and Libby Hemphill. "hot" chatgpt: The promise of chatgpt in detecting and discriminating hateful, offensive, and toxic comments on social media. *arXiv preprint arXiv:2304.10619*, 2023.
- Xinlei He, Savvas Zannettou, Yun Shen, and Yang Zhang. You only prompt once: On the capabilities of prompt learning on large language models to tackle toxic content. *arXiv preprint arXiv:2308.05596*, 2023.
- Yiming Zhu, Peixian Zhang, Ehsan-Ul Haq, Pan Hui, and Gareth Tyson. Can chatgpt reproduce human-generated labels? a study of social computing tasks. *arXiv preprint arXiv:2304.10145*, 2023.
- Fan Huang, Haewoon Kwak, and Jisun An. Is chatgpt better than human annotators? potential and limitations of chatgpt in explaining implicit hate speech. *arXiv preprint arXiv:2302.07736*, 2023.
- Mai ElSherief, Caleb Ziems, David Muchlinski, Vaishnavi Anupindi, Jordyn Seybolt, Munmun De Choudhury, and Diyi Yang. Latent hatred: A benchmark for understanding implicit hate speech. *arXiv preprint arXiv:2109.05322*, 2021.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1, page 2, 2019.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019b.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. olmpics-on what language model pre-training captures. *Transactions of the Association for Computational Linguistics*, 8:743–758, 2020.
- Joe Davison, Joshua Feldman, and Alexander M Rush. Commonsense knowledge mining from pretrained models. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 1173–1178, 2019.
- Roy Schwartz, Maarten Sap, Ioannis Konstas, Li Zilles, Yejin Choi, and Noah A Smith. The effect of different writing tasks on linguistic style: A case study of the roc story cloze task. *arXiv preprint arXiv:1702.01841*, 2017.
- Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020.
- Timo Schick and Hinrich Schütze. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*, 2020a.
- Timo Schick and Hinrich Schütze. It’s not just size that matters: Small language models are also few-shot learners. *arXiv preprint arXiv:2009.07118*, 2020b.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- Xiaochuang Han and Yulia Tsvetkov. Fortifying toxic speech detectors against veiled toxicity. *arXiv preprint arXiv:2010.03154*, 2020.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. Hatexplain: A benchmark dataset for explainable hate speech detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 14867–14875, 2021.

- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). *arXiv preprint arXiv:1903.08983*, 2019b.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 2023.
- Michael Wiegand, Melanie Siegel, and Josef Ruppenhofer. Overview of the semeval 2018 shared task on the identification of offensive language. 2018.
- MS Windows NT kernel description. https://www.oxfordlearnersdictionaries.com/us/definition/english/offensive_1. Accessed: 2010-09-30.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.