

Beyond Hate Speech: NLP’s Challenges and Opportunities in Uncovering Dehumanizing Language

Anonymous ACL submission

Abstract

Dehumanization, characterized as a subtle yet harmful manifestation of hate speech, involves denying individuals of their human qualities and often results in violence against marginalized groups. Despite significant progress in Natural Language Processing across various domains, its application in detecting dehumanizing language is limited, largely due to the scarcity of publicly available annotated data for this domain. This paper evaluates the performance of cutting-edge NLP models, including GPT-4, GPT-3.5, and LLAMA-2, in identifying dehumanizing language. Our findings reveal that while these models demonstrate potential, achieving a 70% accuracy rate in distinguishing dehumanizing language from broader hate speech, they also display biases. They are over-sensitive in classifying other forms of hate speech as dehumanization for a specific subset of target groups, while more frequently failing to identify clear cases of dehumanization for other target groups. Moreover, leveraging one of the best-performing models, we automatically annotated a larger dataset for training more accessible models. However, our findings indicate that these models currently do not meet the high-quality data generation threshold necessary for this task.

1 Introduction

Dehumanization, defined as the denial of “humanity” to others (Haslam, 2006), significantly impacts society by fostering conditions that result in extreme and violent behaviors against marginalized groups (Kteily and Landry, 2022). This phenomenon can range from overt derogation, where victims are likened to “dogs” or “monkeys” (Hagan and Rymond-Richmond, 2008), to subtler forms, such as denying the capability of experiencing pain to certain individuals (Deska et al., 2020). The identification of dehumanizing language is crucial for understanding and mitigating its effects on col-

lective violence and the manipulation of public perception in conflicts (Oberschall, 1997).

Despite the importance of detecting dehumanization, this nuanced form of hate speech has been relatively overlooked in natural language processing advancements, primarily due to the lack of publicly available, annotated datasets. Annotating dehumanizing language poses unique challenges due to its subjective and abstract nature.

This study evaluates the capability of leading pretrained NLP models—specifically, GPT-4, GPT-3.5, and LLAMA-2 (Touvron et al., 2023a)—in accurately identifying dehumanizing language. Through a comprehensive analysis encompassing zero-shot, few-shot, and explainable prompting settings, we evaluate the effectiveness of these models in recognizing dehumanizing content. Our findings reveal that the closed-source GPT models notably outperform the open-source LLAMA-2-70B model in identifying dehumanizing language. However, the GPT models demonstrate limitations in distinguishing dehumanizing language from other forms of hate speech, achieving only a 70% accuracy rate for this specific task. Furthermore, our results expose a variable sensitivity across different target groups; GPT models are prone to overclassifying other types of hate speech as dehumanization, for certain target groups such as gay and transgender individuals, while failing to adequately identify dehumanizing language targeting other vulnerable groups, such as immigrants and refugees.

Following our evaluation, we applied the most effective approach, i.e., explainable prompting, to automatically generate annotated data for training smaller open-source models. This phase revealed a notable discrepancy: while the initial results from GP models are promising, the resulting annotations do not meet the expected standards for training high-performing models. Specifically, even the top-performing model only achieved a 61% accuracy rate in distinguishing dehumanizing language from

other hate speech types.

This finding underscores the fact that annotating nuanced tasks like dehumanization still necessitates the expertise of human annotators. It also highlights that the annotation of large-scale corpora with dehumanization labels represents an important and pivotal future direction for fostering collaboration between the fields of natural language processing and social science, enabling comprehensive research into this profoundly significant social phenomenon.

2 Related Work

Dehumanization has been extensively studied within the realm of social science (Paladino et al., 2002; Haslam et al., 2008; Haslam, 2006; Haslam and Loughnan, 2014; Kteily and Landry, 2022; Harris and Fiske, 2015; Leyens et al., 2000). Recent advancements in NLP techniques present a significant, yet largely untapped, opportunity to utilize state-of-the-art methodologies and expand traditional dehumanization analysis. These advancements have the potential to identify more comprehensive instances of dehumanization, ultimately contributing to the enhancement of online media safety and enabling a more comprehensive sociological examination of the multifaceted impact of dehumanization on society. However, despite this promising potential, the exploration of dehumanization within the field of NLP has been relatively limited.

The first step in addressing this gap was taken by Mendelsohn et al. (2020) who introduced a computational framework for studying dehumanization with traditional NLP techniques¹, focusing on the analysis of how LGBTQ-related terms were subjected to dehumanization in New York Times articles spanning more than 30 years. Their approach revolved around four key components: (1) Negative Evaluations: Assessing the presence of negative judgments directed towards the target group, (2) Denial of Agency: Examining instances where the target group’s capacity to make decisions or take actions was undermined, (3) Moral Disgust: Identifying expressions of moral repugnance in the context of the target group, and (4) Use of Vermin Metaphors: Detecting the application of metaphors portraying the group as vermin or subhuman.

While Mendelsohn et al. (2020)’s approach was effective in identifying overall trends related to

dehumanization, it faces two primary challenges. Firstly, it is challenging to use their proposed approach to pinpoint specific mentions of dehumanization within the text. Secondly, their techniques were less adaptable to shorter texts, such as social media content and comments. In contrast, this paper capitalizes on the capabilities of state-of-the-art NLP techniques to identify specific instances of dehumanization within short input texts.

In addition to Mendelsohn et al. (2020), other researchers have also explored computational analysis of dehumanization. For instance, Friedman et al. (2021) consider dehumanization as a subproblem of moral disengagement and manually annotate 378 examples for both training and evaluation, along with their corresponding entities and relations regarding various forms of moral disengagement. They then utilize a transformer-based model, i.e., a variation of the SpanBERT model (Eberts and Ulges, 2020), to construct a knowledge graph consisting of these entities and relations. The schema of their knowledge graph illustrates entities linked by relationships, with each entity possessing various attributes, including dehumanization, violent, condemned, justified, responsible, and harmed. They report an F₁ score of 50 points on identifying the dehumanization attributes in their dataset. However, it’s worth noting that their dataset is not publicly available for incorporation or evaluation in this study.

Similarly, dehumanization is considered as one of the subcategories of hate speech by Vidgen et al. (2021), where they constructed a large-scale dataset of 22K examples of hate speech, among which 906 examples are labeled as dehumanization. We use this subset for the evaluation in this work.

3 How Effective are Existing Models for Identifying Dehumanization?

3.1 Data

As mentioned, we utilize Vidgen et al. (2021)’s publicly available hate speech dataset for our evaluations. This dataset comprises 906 instances of dehumanizing content, enabling us to assess the identification of dehumanization. Additionally, the inclusion of other hate speech labels, such as animosity or derogation, in this dataset allows us to evaluate the model’s ability to distinguish dehumanization from various forms of hate speech.

Another significant advantage of this dataset is that it provides information about the targeted

¹E.g., word2vec word embeddings (Mikolov et al., 2013) and connotation frames (Rashkin et al., 2016).

groups. This allows us to analyze whether the model’s performance on dehumanizing instances varies depending on the targeted group. For example, we aim to assess whether the model’s performance is enhanced when provided with examples from the same targeted group in the evaluation prompts. We devised three evaluation subsets from this dataset for evaluating dehumanization:

Targeted Dehumanization: This evaluation set consists of 42 dehumanization instances aimed at Muslims and an additional 42 randomly chosen samples.² This set is designed to assess the model’s performance in identifying dehumanization when there is a single, known target a priori.

General Dehumanization: This set contains all 906 dehumanization instances, which may have different targeted groups, along with 906 randomly selected instances from the dataset.³ This evaluation set assesses the model’s performance when the target may vary and is not predetermined.

Dehumanization vs. Hate: It consists of 906 instances of dehumanization as well as 906 randomly selected instances from other hate speech labels, testing the model’s ability to distinguish between dehumanization and other forms of hate speech.

3.2 Mendelsohn’s Baselines

Other than state-of-the-art models, we also adopt the four linguistic-based components from Mendelsohn et al. (2020)’s framework as our baselines for analyzing the extent to which each of these components can identify instances of dehumanization in our evaluation set.

Negative Evaluation of a Target Group Valence measures the degree of positivity or negativity expressed in the text, ranging from completely positive (1) to completely negative (0). We use the NRC VAD lexicon (Mohammad, 2018), which provides valence, dominance (the control or power expressed), and arousal (the level of excitement or calmness) scores for words. To estimate a text’s overall sentiment, we calculate the average valence score of its words, as provided by this lexicon.

To further assess the sentiment directed towards specific target groups, we use the connotation frames lexicon (Rashkin et al., 2015), which gives

²The 42 randomly selected examples contain 19 instances of hate speech and 23 non-hate speech labels.

³The randomly selected examples contain 414 instances of hate speech and 492 non-hate speech labels.

scores to 900 English verbs from very negative (-0.87) to very positive (0.8). We consider an input text as conveying a negative evaluation toward a target group if its valence score is below 0.5, and its perspective score is negative.

Denial of Agency To evaluate agency, i.e., the degree to which a target group is seen as able to control their actions and decisions, we use the connotation frames for agency (Sap et al., 2017). This approach differentiates between high agency, where entities are viewed as having significant control over their decisions, and low agency, where they are seen as more passive. The lexicon annotates over 2000 verbs with agency levels. We determine a text’s overall agency by calculating an aggregate score based on the frequency of these verbs. If a text predominantly features verbs indicating low agency, it is classified as exhibiting a denial of agency. For texts lacking any verbs from the lexicon, we use a default “natural” label.

Moral Disgust Following the approach by Mendelsohn et al. (2020), we use Graham et al. (2009)’s lexicon to identify instances of moral disgust. This lexicon includes more than 30 words and stems associated with moral disgust.⁴ To measure moral disgust, we adopt a vector-based methodology, calculating the average of the word embeddings for terms linked to moral disgust, with each word’s contribution weighted by its frequency in the lexicon.⁵ The degree of an input text’s association with moral disgust is then assessed by computing the cosine distance between the averaged vector of moral disgust terms and the embedding of the input text.

Use of Vermin Metaphors Similar to Mendelsohn et al. (2020) and the methodology for the “moral disgust” component, we construct a vector representation for “vermin metaphors” using keywords such as vermin, rodent(s), rat(s), mice, cockroach(es), termite(s), and bedbug(s). We assess the presence of these metaphors in texts by comparing the input text’s embeddings to this vector representation.⁶

⁴E.g., “disgust”, “sin”, and “pervert”.

⁵We use SpaCy’s en_core_web_sm model.

⁶While recent advances in NLP offer more refined methods for implementing these analyses, we follow the framework suggested by Mendelsohn et al. (2020). The enhancement of these techniques and their integration with state-of-the-art models represent avenues for future research.

3.3 State-of-the-art Baseline

In our experiments, we utilize both GPT-4 and GPT-3.5-turbo, which are state-of-the-art NLP models known for their exceptional performance across various NLP tasks. These models showcase the advanced capabilities of off-the-shelf NLP technologies in handling complex language processing tasks. Additionally, we evaluate LLAMA-2-70B⁷ (Touvron et al., 2023b) as one of the leading open-source models in NLP.

A key determining factor in the success of state-of-the-art pretrained models are their corresponding prompts for each task. A prompt acts as the initial query or instruction, guiding the model to produce the desired output.

In this paper, we explore three primary prompting schemes:

Zero-shot: In this setting, the prompt consists of the phrase “Identify target groups and decide if they’re dehumanized”. This scheme assesses the pretrained model’s preexisting knowledge about dehumanization.

Few-shot: We enhance the model’s exposure by incorporating five randomly selected instances of dehumanization targeting Muslims into the prompt. This method allows us to evaluate the model’s ability to generalize its understanding to other targeted groups, emphasizing the importance of specific examples in improving performance.

Explainable prompting: Building on the few-shot setting, this approach further requires the model to provide explanations for its decisions.

In the zero-shot setting, the model identifies target groups and determines whether the text contains dehumanizing language for that target group.

In the few-shot setting, the model goes further by classifying dehumanization within texts as either “blatant” or “subtle”. The included few-shot examples with dehumanizing language are labeled as “blatant”.

The explainable prompting setting mirrors the few-shot approach but adds a requirement for the model to explain its reasoning.⁸

3.4 Results

Table 1 presents the results of the models in *zero-shot*, *few-shot*, and *explainable* settings, compared

against the four components of *Negative evaluation of a target group*, *Denial of agency*, *Moral Disgust*, and *Use of vermin metaphors*. The *Combination* row shows the results where we consider a text as dehumanization if it contains all four components.

In the *few-shot* setting, a text is flagged for dehumanization if the predicted label for any of the identified targets is true. For *zero-shot* and *explainable*, a text is considered dehumanizing if a “blatant” label is predicted for any of the identified targets. Including both “blatant” and “subtle” labels as dehumanization lowers all models’ performance due to reduced precision. The results of this setting are reported in Table 6 in the appendix.⁹

Table 1 reveals the following insights: (1) GPT models significantly outperform heuristic components in detecting dehumanizing language even in the *zero-shot* setting, unlike the LLAMA-2 model, which tends to overclassify inputs as dehumanizing. (2) There is still a large room for improvement for this task as the best accuracy for recognizing dehumanization from other types of hate speech is only 70%. (3) The accuracy of the examined models in recognizing dehumanizing language versus neutral text is higher than its accuracy in discriminating dehumanizing language from other types of hate speech, as indicated by the higher performances in the “general dehumanization” subset. (4) Including targeted group details in prompts notably enhances dehumanization detection for the GPT-3.5 model, as seen in the improved results for the *few-shot* and *explainable* settings within the “Targeted Dehumanization” subset. This improvement is not observed with the GPT-4 model. (5) Requesting explanations for predictions improves the GPT models’ ability to differentiate dehumanization from other hate speech, as demonstrated by the higher performance of *explainable* settings over *few-shot* in the “Dehumanization vs. Hate” subset.

3.5 Analysis

In this section, we analyze the results of GPT models that achieve the highest scores in our evaluation to further explore their shortcomings and strengths.

Does the model perform equally well for different target groups? To address this question, we calculate two types of errors for each target group: (1) over-sensitivity, where the model inaccurately

⁷We use the Llama-2-70b-chat-hf model

⁸The prompt templates for each of these settings are included in the appendix.

⁹GPT models show stable outputs, unlike LLAMA-2, which had about 8% unparsable outputs. Scores were computed after excluding these cases.

		Targeted Dehumanization			General Dehumanization			Dehum. vs Hate		
		F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (hate)	F ₁ (dehum.)	Acc.
Zero-shot	GPT-4	61.76	74.00	69.05	69.73	78.61	74.93	44.80	70.59	61.62
	GPT-3.5	65.75	73.68	70.24	66.36	75.81	71.82	51.90	70.83	63.69
	LLAMA-2	17.45	68.95	54.87	13.92	68.28	53.64	2.84	66.89	50.61
Few-shot	GPT-4	77.33	81.72	79.76	77.09	81.76	79.69	59.41	74.91	68.99
	GPT-3.5	81.01	82.76	81.93	77.13	74.00	75.66	68.01	68.67	68.34
	LLAMA-2	38.38	69.65	59.33	36.77	68.43	57.88	27.87	69.86	57.49
Explainable	GPT-4	73.97	80.00	77.38	77.38	82.02	79.97	59.19	76.29	70.00
	GPT-3.5	79.07	78.05	78.57	76.15	74.37	75.29	68.15	69.96	69.08
	LLAMA-2	13.41	62.83	47.99	33.82	60.56	50.57	32.08	57.74	47.90
Negative Eval.		67.20	4.65	51.19	66.34	5.26	50.33	65.89	5.21	49.83
Agency Denial		64.41	16.00	50.00	62.53	18.56	48.68	63.47	18.84	49.61
Moral Disgust		46.91	50.57	48.81	44.18	46.71	45.47	44.67	46.47	45.58
Vermin Meta.		38.46	46.67	42.86	42.43	45.77	44.15	41.73	45.04	43.43
Combination		67.20	4.65	51.19	66.72	0.44	50.11	66.62	0.44	50.00

Table 1: Comparison of identifying dehumanizing language across different models and settings, focusing on instances explicitly labeled as “blatant” dehumanization. The lower section includes results from Mandelsohn’s components. The *Combination* baseline classifies an example as dehumanization if it contains all four components.

Text	Label	Zero	Few	Exp.
Why don’t [XX] just STFU and stay into the kitchen	Derogation	Dehuman	Blatant	Blatant
It gets to a point where you don’t care or even lust for XX anymore. I truly don’t think about those deceitful <AA> and it’s been like this for years. Never trust one again	Derogation	Dehuman	Subtle	Subtle
It’s completely shit that <BB> like him are living and breathing	Derogation	Dehuman	Blatant	Blatant
Next time i see a <CC> person, I’m going to b e a t them up	Threatening	Blatant	Blatant	Blatant
someone who had a strong year and made over 75k in 2018 and who may be struggling currently gets zero but illegals and refugees get \$350 mil	Animosity	Dehuman	Blatant	Subtle

Table 2: Examples from GPT-3.5 outputs when the model incorrectly classifies other types of hate speech as dehumanization. Some parts of input texts are masked out to reduce their toxicity. The “Label” column shows the annotated label of the input.

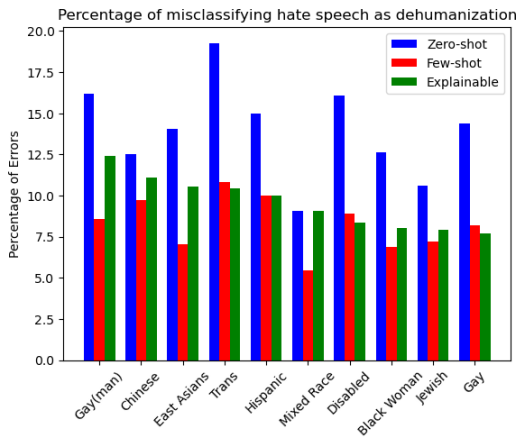


Figure 1: Top 10 target groups with the highest over-sensitivity error ratios for GPT-3.5.

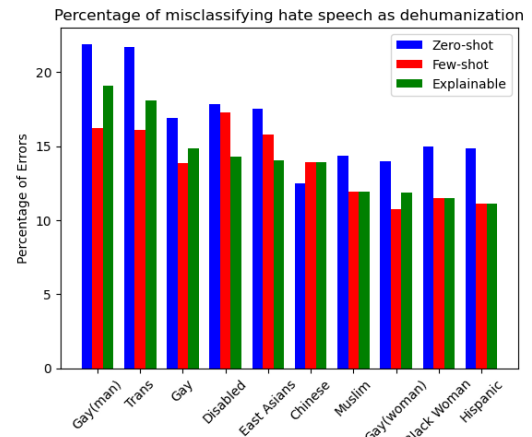


Figure 2: Top 10 target groups with the highest over-sensitivity error ratios for GPT-4.

labels less severe hate speech as dehumanization, and (2) recognition blindness, defined by the ratio of instances of dehumanizing language that remain undetected by the model for a specific target group.

We calculate the over-sensitivity error ratio for each target group by dividing the number of instances misclassified as dehumanization by the total number of instances for that group within the

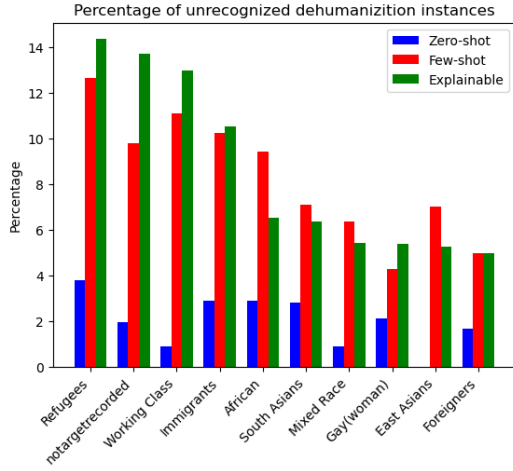


Figure 3: Top 10 target groups with the highest ratio of unrecognized dehumanization instances for GPT-3.5.

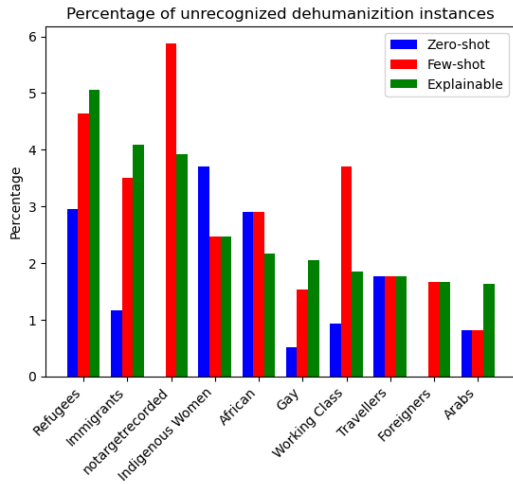


Figure 4: Top 10 target groups with the highest ratio of unrecognized dehumanization instances for GPT-4.

“Dehumanization vs. Hate” evaluation set. Figures 1 and 2 show the top 10 target groups with the highest error ratios for the GPT-3.5 and GPT-4 models, respectively.¹⁰ On the other hand, Figures 3 and 4 show the 10 target groups with the highest ratios of recognition blindness, measured by the ratio of instances per target group containing dehumanizing language that the model fails to recognize.

The results reveals that: (1) The error ratio in the zero-shot setting shows significant variability, while few-shot and explainable settings exhibit more consistency in error ratios; (2) The GPT-4 model demonstrates a higher sensitivity to clas-

¹⁰Target groups are ordered according to their error ratios in the explainable setting, which performs best in the “Dehumanization vs. Hate” subset.

sify hate speech as dehumanization, as evidenced by elevated error ratios among its top 10 target groups in Figure 2; and (3) Crucially, both models exhibit varying levels of sensitivity towards different target groups, overclassifying less severe hate speech as dehumanization for certain groups like gay and transgender individuals, yet more frequently failing to detect dehumanizing language targeting groups such as refugees, immigrants, and the working class. This discrepancy highlights the models’ inherent biases towards different target groups.

Table 3 shows some examples with dehumanizing language that remained undetected across the zero-shot, few-shot, and explainable settings of GPT-3.5.

Which types of hate speech does the model frequently confuse with dehumanization? In Figure 5, we present the number of instances from various hate speech categories in the evaluation set that are classified incorrectly by GPT-3.5. Notably, a majority of these errors occur when the model identifies instances of “derogation” as dehumanization.

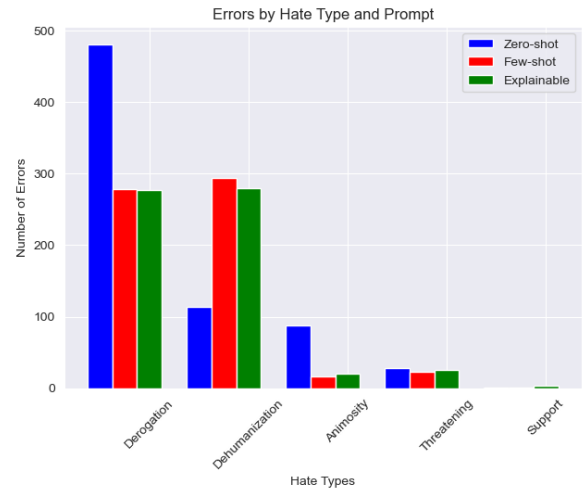


Figure 5: Number of instances from different types of hate speech that are incorrectly classified as dehumanization in “Dehumanization vs. Hate”.

4 How Effective is GPT for Annotating Training Data?

The GPT-3.5 and GPT-4 models, being closed-source, might not be accessible to all users. Nevertheless, the task of identifying dehumanizing language has the potential to serve a diverse user base, and its detection shouldn’t be limited by users’

Given that China, Russia and USA are not responsible for the vast majority of war crimes, they should not be the ones responsible for these swarms of <AA>
it's like a plague, the way these <BB> just come spilling over into all corners of the uk
<CC> are the worst kind of dogs and I'm not talking about the animals
<DD> are like puppies, really nice but you get tired of them after some time
The british government is massively fucking up its response to the <EE> crisis, why cant they get rid of this fucking disease?

Table 3: Examples from the evaluation set where the dehumanizing language is not detected by any of the GPT-3.5 variants. Some parts of the input texts are masked out to reduce their toxicity.

[Train]	Model	Dev Acc.	Targeted Dehumanization			General Dehumanization			Dehum. vs Hate		
			F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (hate)	F ₁ (dehum.)	Acc.
	Explainable (blatant)	-	79.07	78.05	78.57	76.15	74.37	75.29	68.15	69.96	69.08
400	Roberta-large	79.25	69.85	70.38	70.24	65.88	62.87	64.54	58.82	59.68	59.43
	Flan-base	74.25	53.16	70.59	64.29	52.10	65.08	60.12	45.98	63.20	56.74
	Flan-large	68.75	48.79	69.08	61.51	51.13	66.83	60.58	40.88	63.98	55.32
	Flan-XL	81.25	66.64	71.12	69.05	64.37	66.15	65.29	55.11	62.35	59.05
2K	Roberta-large	80.35	67.71	68.24	68.10	66.95	65.63	66.42	59.42	62.20	61.02
	Flan-base	74.00	63.19	68.49	66.67	62.44	63.72	63.71	56.17	61.23	59.59
	Flan-large	75.67	69.31	74.63	72.22	62.41	69.66	66.43	51.52	65.77	59.88
	Flan-XL	76.63	64.40	71.66	68.45	63.79	68.91	66.67	53.79	65.11	60.40

Table 4: Investigating the impact of training smaller, open-source models with annotations generated by GPT-3.5 using 'Explainable' prompting, comparing the effects of fine-tuning the model with (1) 400 annotated examples and (2) 2,000 annotated examples.

computational capabilities or the need for paid API access. Inspired by the widespread success of state-of-the-art models, including GPT models, in generating annotated instances for various existing tasks (He et al., 2023; Alizadeh et al., 2023; Latif et al., 2023; Ding et al., 2023; Thapa et al., 2023; Ziems et al., 2023; Gilardi et al., 2023; Törnberg, 2023), in this section, we investigate the effectiveness of using GPT-3.5 for generating annotated data to train more accessible (in terms of model size) and open-source models.¹¹ For this purpose, we employ two different open-source and accessible models: Roberta (Liu et al., 2019) and Flan-T5 (Chung et al., 2022).¹² We utilize various model sizes, including Roberta-large (355M parameters), Flan-T5-base (248M), Flan-T5-large (783M), and Flan-T5-XL (3B). These models are trained on Stanage¹³ using hugging face packages (Wolf et al., 2020).

4.1 Automatic Data Annotation

This section outlines the methodology for automatically generating a dataset annotated with instances of dehumanizing language. We use the *Explainable* setting of GPT-3.5 because of its performance in identifying dehumanizing language and distin-

guishing it from other forms of hate speech. We collected seed data using *snscape*,¹⁴ a tool for scraping social networking services, to collect 1 million Tweets related to Muslims.¹⁵ From this collection, we randomly selected 20,000 tweets and annotated them by GPT-3.5, yielding 1,208 tweets explicitly marked as instances of blatant dehumanization. Table 5 provides a few examples of the resulting annotated data. These annotations serve as the basis for training smaller models.

4.2 Model Training

We conduct training in two distinct settings to evaluate the impact of dataset size on model performance: (1) using 400 automatically annotated examples, balanced with 200 dehumanization and 200 non-dehumanization texts, and (2) expanding to 2,000 instances, half of which (1,000) are identified by GPT-3.5 as dehumanizing. This approach allows us to explore how varying amounts of automatically annotated data influence performance outcomes. While automatic annotation is considerably less expensive than using human experts, it still requires using a paid API. Thus, determining the optimal quantity of training data that balances cost with performance improvement is beneficial.

¹¹We use GPT-3.5 because its performance is on-par with that of GPT-4 while it is less expensive.

¹²Flan-T5 is also a prompt-based approach, for which we use the prompt "Classify this text as either 0 (not dehumanising) or 1 (dehumanising). Text:tweet Answer(0 or 1):"

¹³<https://docs.hpc.shef.ac.uk/en/latest/stanage/index.html>

¹⁴<https://github.com/JustAnotherArchivist/snscape>

¹⁵Due to recent Twitter policy updates, accessing such data directly is now restricted and subject to the costs associated with the Twitter API.

Input: For only 20 rupees; [NAME], [NAME], [NAME], [NAME], [NAME], [NAME] and [NAME] publicly assaulted and humiliated Saleem to the point that he ended his life by jumping in front of a train but there won't be any media outrage nor any one would bother to tweet coz d victim is a Muslim
Target: Saleem
Dehumanization: Blatant
Explanation: Saleem is publicly assaulted and humiliated to the point that he ended his life. This is a clear form of dehumanization as his life is disregarded and his suffering is not taken seriously.
Target: Muslim
Dehumanization: subtle
Explanation: The text mentions that there won't be any media outrage or anyone bothering to tweet because the victim is a Muslim. This implies a subtle form of dehumanization, suggesting that the suffering of a Muslim is not as important or worthy of attention as that of others.
Input: This night is very important to all Muslim. May Allah accept everyone's prayers "Ameen" shab e meraj Mubarak to all of you
Target: Muslim
Dehumanization: None
Explanation: The group is mentioned in a respectful and positive manner, with a message of well-wishes and prayers.

Table 5: Samples from the GPT-3.5 output on the collected data. The proper names are anonymized in the input for privacy. Note that each input sentence may have multiple labels per input if it contains more than one target.

Table 4 shows the results of this experiment. The “Explainable” row displays the results of the GPT-3.5 model, which is used for annotating additional data, on the evaluation sets. The findings indicate a general trend: all models fine-tuned with the automatically annotated data underperform compared to the GPT-3.5 annotation model. Yet, an increase in training data volume correlates with improved accuracy across models, with a notable exception in the “Targeted Dehumanization” evaluation set, where the Roberta-large model’s performance decreases with additional data, suggesting a potential overfitting to the training data. As expected, models show higher performances on the “Targeted Dehumanization” set, as the result of the training data selection strategy. Overall, the FLAN-large model achieves higher accuracy and F₁ scores in detecting dehumanization in various evaluation sets. Using FLAN-XL instead of FLAN-large does not show a significant advantage.

5 Conclusion

The automatic identification of dehumanizing language is a vital endeavor, given its role in propagating subtle and harmful hate speech that can have severe consequences, particularly for marginalized communities. In this paper, we delved into the application of state-of-the-art NLP models for identifying dehumanizing language. While our findings reveal considerable promise, there remain substantial limitations that warrant attention. Notably, there is large room for enhancing the performance of state-of-the-art models in this domain. Furthermore, our analysis underscores potential disparities in the model’s effectiveness across different target groups. This raises a note of caution regarding the unmitigated reliance on such models for deriving overarching conclusions when analyzing large-

scale data for social research about different target groups. As a result, we emphasize the continued need for thorough evaluation and exploration, complemented by expert human annotations, to make further progress in the understanding and mitigation of dehumanization in digital spaces and its societal implications.

6 Limitations

We have exclusively relied on Vidgen et al. (2021)’s dataset for dehumanization evaluation. The study could have yielded deeper insights if we had access to a dataset that categorized dehumanization into blatant and subtle instances. Moreover, our evaluation is limited to Twitter data related to a single target group. A more comprehensive assessment would involve multiple target groups. Nevertheless, it’s important to acknowledge that the recent restrictions on social media APIs have posed challenges in this field of study. There is also a potential risk of deploying NLP models to detect dehumanizing language is the inadvertent reinforcement of biases, where models may disproportionately flag or overlook certain groups’ speech, perpetuating inequality and stifling free expression.

References

- Meysam Alizadeh, Maël Kubli, Zeynab Samei, Shirin Dehghani, Juan Diego Bermeo, Maria Korobeynikova, and Fabrizio Gilardi. 2023. Open-source large language models outperform crowd workers and approach chatgpt in text-annotation tasks. *arXiv preprint arXiv:2307.02179*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.

537	Jason C Deska, Jonathan Kunstman, E Paige Lloyd,	Siddique Latif, Muhammad Usama, Moham-	592
538	Steven M Almaraz, Michael J Bernstein, JP Gonzales,	mad Ibrahim Malik, and Björn W Schuller.	593
539	and Kurt Hugenberg. 2020. Race-based biases in	2023. Can large language models aid in annotating	594
540	judgments of social pain. <i>Journal of Experimental</i>	speech emotional data? uncovering new frontiers.	595
541	<i>Social Psychology</i> , 88:103964.	<i>arXiv preprint arXiv:2307.06090</i> .	596
542	Bosheng Ding, Chengwei Qin, Linlin Liu, Yew Ken	Jacques-Philippe Leyens, Paola M Paladino, Ramon	597
543	Chia, Boyang Li, Shafiq Joty, and Lidong Bing. 2023.	Rodriguez-Torres, Jeroen Vaes, Stephanie Demoulin,	598
544	Is GPT-3 a good data annotator? In <i>Proceedings</i>	Armando Rodriguez-Perez, and Ruth Gaunt. 2000.	599
545	<i>of the 61st Annual Meeting of the Association for</i>	The emotional side of prejudice: The attribution of	600
546	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,	secondary emotions to ingroups and outgroups. <i>Per-</i>	601
547	pages 11173–11195, Toronto, Canada. Association	<i>sonality and social psychology review</i> , 4(2):186–197.	602
548	for Computational Linguistics.		
549	Markus Eberts and Adrian Ulges. 2020. Span-based	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Man-	603
550	joint entity and relation extraction with transformer	dar Joshi, Danqi Chen, Omer Levy, Mike Lewis,	604
551	pre-training . In <i>24th European Conference on Artifi-</i>	Luke Zettlemoyer, and Veselin Stoyanov. 2019.	605
552	<i>cial Intelligence - ECAI 2020</i> , pages 2006–2013. IOS	Roberta: A robustly optimized bert pretraining ap-	606
553	Press.	proach .	607
554	Scott E Friedman, Ian Magnusson, Sonja Schmer-	Julia Mendelsohn, Yulia Tsvetkov, and Dan Jurafsky.	608
555	Galunder, Ruta Wheelock, Jeremy Gottlieb, Christo-	2020. A framework for the computational linguistic	609
556	pher Miller, et al. 2021. Toward transformer-based	analysis of dehumanization . <i>Frontiers in Artificial</i>	610
557	nlp for extracting psychosocial indicators of moral	<i>Intelligence</i> , 3.	611
558	disengagement. In <i>Proceedings of the Annual Meet-</i>		
559	<i>ing of the Cognitive Science Society</i> , volume 43.	Tomas Mikolov, Kai Chen, Greg Corrado, and Jef-	612
560	Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli.	frey Dean. 2013. Efficient estimation of word	613
561	2023. Chatgpt outperforms crowd-workers for text-	representations in vector space. <i>arXiv preprint</i>	614
562	annotation tasks. <i>arXiv preprint arXiv:2303.15056</i> .	<i>arXiv:1301.3781</i> .	615
563	Jesse Graham, Jonathan Haidt, and Brian A Nosek.	Saif Mohammad. 2018. Obtaining reliable human rat-	616
564	2009. Liberals and conservatives rely on different	ings of valence, arousal, and dominance for 20,000	617
565	sets of moral foundations. <i>Journal of personality and</i>	english words. In <i>Proceedings of the 56th annual</i>	618
566	<i>social psychology</i> , 96(5):1029.	<i>meeting of the association for computational linguis-</i>	619
567	John Hagan and Wenona Rymond-Richmond. 2008.	<i>tics (volume 1: Long papers)</i> , pages 174–184.	620
568	The collective dynamics of racial dehumanization		
569	and genocidal victimization in darfur. <i>American So-</i>	Anthony Oberschall. 1997. Vojislav Seselj’s nationalist	621
570	<i>ciological Review</i> , 73(6):875–902.	propaganda: contents, techniques, aims and impacts,	622
571	Lasana T Harris and Susan T Fiske. 2015. Dehumanized	1990-1994. How mass media propaganda impacts	623
572	perception. <i>Zeitschrift für Psychologie</i> .	on ordinary people’s acceptance and participation	624
573	Nick Haslam. 2006. Dehumanization: An integrative	in collective violence, and how Seselj’s nationalist	625
574	review. <i>Personality and social psychology review</i> ,	propaganda promoted and justified coercion and vio-	626
575	10(3):252–264.	lence by the Serbs against non-Serbs.	627
576	Nick Haslam, Yoshihisa Kashima, Stephen Loughnan,	Maria-Paola Paladino, Jacques-Philippe Leyens, Ramon	628
577	Junqi Shi, and Caterina Suitner. 2008. Subhuman,	Rodriguez, Armando Rodriguez, Ruth Gaunt, and	629
578	inhuman, and superhuman: Contrasting humans	Stéphanie Demoulin. 2002. Differential association	630
579	with nonhumans in three cultures. <i>Social cognition</i> ,	of uniquely and non uniquely human emotions with	631
580	26(2):248–258.	the ingroup and the outgroup. <i>Group Processes &</i>	632
581	Nick Haslam and Steve Loughnan. 2014. Dehuman-	<i>Intergration Relations</i> , 5(2):105–117.	633
582	ization and infrahumanization. <i>Annual review of</i>	Hannah Rashkin, Sameer Singh, and Yejin Choi. 2015.	634
583	<i>psychology</i> , 65:399–423.	Connotation frames: A data-driven investigation.	635
584	Xingwei He, Zhenghao Lin, Yeyun Gong, Alex Jin,	<i>arXiv preprint arXiv:1506.02739</i> .	636
585	Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan	Hannah Rashkin, Sameer Singh, and Yejin Choi. 2016.	637
586	Duan, Weizhu Chen, et al. 2023. Annollm: Making	Connotation frames: A data-driven investigation . In	638
587	large language models to be better crowdsourced	<i>Proceedings of the 54th Annual Meeting of the As-</i>	639
588	annotators. <i>arXiv preprint arXiv:2303.16854</i> .	<i>sociation for Computational Linguistics (Volume 1:</i>	640
589	Nour S. Kteily and Alexander P. Landry. 2022. Dehu-	<i>Long Papers)</i> , pages 311–321, Berlin, Germany. As-	641
590	manization: trends, insights, and challenges . <i>Trends</i>	sociation for Computational Linguistics.	642
591	<i>in Cognitive Sciences</i> , 26(3):222–240.	Maarten Sap, Marcella Cindy Prasettio, Ari Holtzman,	643
		Hannah Rashkin, and Yejin Choi. 2017. Connotation	644
		frames of power and agency in modern films. In	645
		<i>Proceedings of the 2017 conference on empirical</i>	646
		<i>methods in natural language processing</i> , pages 2329–	647
		2334.	648

Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023. From humans to machines: can chatgpt-like llms effectively replace human annotators in nlp tasks. In *Workshop Proceedings of the 17th International AAAI Conference on Web and Social Media*.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

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

Petter Törnberg. 2023. Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588*.

Bertie Vidgen, Tristan Thrush, Zeerak Waseem, and Douwe Kiela. 2021. [Learning from the worst: Dynamically generated datasets to improve online hate detection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1667–1682, Online. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#).

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2023. Can large lan-

guage models transform computational social science? *arXiv preprint arXiv:2305.03514*.

A Which model identifies more instances as dehumanization?

To determine which model most frequently identifies instances of dehumanization, we examine the recall and precision values for these models in detecting dehumanization labels across different evaluation subsets. These results are reported in Tables 7 and 8.

We find that, with the exception of the “targeted dehumanization” subset, the zero-shot settings detect more instances of dehumanization compared to few-shot and explainable settings. This is particularly evident in the “Dehumanization vs. Hate” subset for the GPT-3.5 model, indicating that without specific examples, GPT-3.5 is prone to misclassifying hateful content as dehumanization. Table 2 provides examples of such misclassifications from the evaluation set. Additionally, considering both “blatant” and “subtle” labels as indicators of dehumanization significantly increases recall but reduces precision.

B Prompt templates

Zero-shot Prompt You are an expert in psychology.\nText: {text}.\nFirst, Identify target groups in the text\nDecide whether each target is dehumanised by the author.\nOnly response in the example output format that can be parsed by json:\n{\n "Targets": ["group_name1", "group_name2", "..."],\n "Dehumanisation": [\n ["target1", true],\n ["target2", false], \n ["...", true]\n]\n}\n

Few-shot Prompt You are an expert in psychology and NLP.Examples of dehumanization:\n{examples}\n\nAnalyze the following text and identify all the groups targeted in it. For each group, assess whether they are subjected to any form of dehumanization, be it blatant, subtle, or none at all.\n\nText: \n{text}\n\nMake sure target name exist and exactly the same in given text. Only response in the example output json format:\n [\n"target_name":target name in text, "dehumanize": "None" or "Subtle" or "Blatant"\n],...\n

		Targeted Dehumanization			General Dehumanization			Dehum. vs Hate		
		F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (hate)	F ₁ (dehum.)	Acc.
Few-shot	GPT-4	52.46	72.90	65.48	57.45	76.14	69.43	19.32	68.93	55.14
	GPT-3.5	69.44	76.59	73.49	74.27	77.99	76.27	57.26	70.88	65.36
	LLAMA-2	8.70	66.07	50.52	15.87	67.93	53.56	8.20	69.80	54.55
Explainable	GPT-4	48.28	72.73	64.29	52.04	74.78	66.94	15.22	70.64	56.39
	GPT-3.5	70.42	78.35	75.00	70.73	78.20	75.01	48.48	70.66	62.61
	LLAMA-2	5.00	62.30	46.02	13.47	61.42	46.64	8.36	61.28	45.56

Table 6: Comparison of identifying dehumanizing language across different models and settings, when both “blatant” and “subtle” predictions are classified as dehumanization.

		Targeted Dehumanization			General Dehumanization			Dehum. vs Hate		
		Recall	Precision	F ₁ (dehum.)	Recall	Precision	F ₁ (dehum.)	Recall	Precision	F ₁ (dehum.)
Zero-shot		83.33	66.03	73.68	88.18	66.47	75.81	88.18	59.18	70.83
Few-shot (blatant)		87.80	78.20	82.76	69.32	79.34	74.00	69.31	68.03	68.67
Explainable (blatant)		76.19	80.00	78.05	71.80	77.11	74.37	71.85	68.16	69.96
Few-shot (+subtle)		87.80	67.92	76.60	84.16	72.66	77.99	84.22	61.19	70.88
Explainable (+subtle)		90.48	69.09	78.35	89.79	69.26	78.20	89.85	58.23	70.66

Table 7: The Recall, Precision and F₁ values of GPT-3.5 model for recognizing the dehumanization label in the “Dehumanization vs. Hate” evaluation set in the *zero-shot*, *few-shot*, and *explainable* setting when the “blatant” predictions are considered as dehumanization.

		Targeted Dehumanization			General Dehumanization			Dehum. vs Hate		
		Recall	Precision	F ₁ (dehum.)	Recall	Precision	F ₁ (dehum.)	Recall	Precision	F ₁ (dehum.)
Zero-shot		88.10	63.79	74.00	92.05	68.59	78.61	92.05	57.24	70.59
Few-shot (blatant)		90.48	74.51	81.72	91.06	74.19	81.76	91.06	63.62	74.91
Explainable (blatant)		90.48	71.70	80.00	91.39	74.39	82.02	91.39	65.46	76.29
Few-shot (+subtle)		92.86	60.00	72.90	97.57	62.43	76.14	97.57	53.29	68.93
Explainable (+subtle)		95.24	58.82	72.73	98.01	60.45	74.78	98.01	55.22	70.64

Table 8: The Recall, Precision and F₁ values of GPT-4 model for recognizing the dehumanization label in the “Dehumanization vs. Hate” evaluation set in the *zero-shot*, *few-shot*, and *explainable* setting when the “blatant” predictions are considered as dehumanization.

Explainabel Prompt You are an expert in psychology and NLP.Examples of dehumanization:\n{examples}\n\nAnalyze the following text and identify all the groups targeted in it. For each group, assess whether they are subjected to any form of dehumanization, be it blatant, subtle, or none at all.\n\nText: \n{text}\n\nMake sure target name exist and exactly the same in given text. Only response in the example output json format:\n [\n"target_name":target name in text, "dehumanize": "None" or "Subtle" or "Blatant", "explanation":explanation\n],...]\n

C Extra Evaluation

Due to the non-reproducibility that may result from the random selection of examples, we conducted an

additional evaluation using two sets of cases from 42 dehumanising texts targeting Muslims that were subjectively chosen by the authors, with one set being closer to dehumanisation and the other closer to generalized hate speech.The results in Table 9 show that there is not a huge difference between the results of the different example groups

Evaluation Sets This section will list two sets of 5 examples specifically chosen to differentiate between dehumanization and generalized hate speech. The selected examples aim to provide insight into the nuanced differences and challenges in classifying such texts, without further analysis.

Set 1: Closer to Dehumanization

1. Listen to me you buttfucking sandrat. We already know you’re going to lie and try to convince people you’re just ‘human’ We know that you are not human. We know that you

	Targeted Dehumanization			General Dehumanization			Dehum. vs Hate		
	F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (other)	F ₁ (dehum.)	Acc.	F ₁ (hate)	F ₁ (dehum.)	Acc.
Few-shot (Example 1)	71.15	70.13	72.29	77.31	72.27	75.04	70.58	68.35	69.5
Few-shot (Example 2)	77.5	79.07	78.31	75.93	76.97	76.46	65.4	71.59	68.8
Explainable (Example 1)	74.15	70.89	72.61	76.69	73.23	75.08	68.42	68	68.21
Explainable (Example 2)	74.66	79.57	77.38	74.43	77.26	75.92	63.73	72.24	68.55
Few-shot (Example 1+subtle)	68.42	73.33	71.08	74.06	77.63	75.98	58.89	71.21	66.13
Few-shot (Example 2+subtle)	71.43	79.17	75.9	70.71	77.72	74.7	52.15	70.93	63.83
Explainable (Example 2+subtle)	73.23	80.41	77.38	71.35	78.88	75.69	47.13	70.7	62.3
Explainable (Example 2+subtle)	59.38	75	69.04	65.83	77.45	72.83	38.62	69.62	59.36

Table 9: Summary of Dehumanization Analysis Results

are a fucking rapist pig and a murderer if your imaginary friend tells you to do it. You are the least intelligent; most violent type of person out there next to the blacks. Just STFU and give up on convincing people otherwise.

2. I took this photo of a muslim. I believe it's fucking awesome. Do you have any pictures of those bugs too?

3. The problem is, that we can't allow the "islamic disease" to spread throughout our countries. Just writing stuff on some "White" forums won't do much, it's more about active participation; it's about OUR activism for OUR people. If there is any pro-White demonstration - go there, if there is some White Woman harassed by some idiots - go there, if there is any "radical" islamic event, get some friends and go there and show everyone, that there are yet people, who care about us!

4. No more fucking terrorists: Denmark to cull entire Muslim population after terrorist mutation spreads within humans.

5. 'My agenda is buying a new car and driving everywhere I can to see the world with my trans boyfriend apart from muslim countries, they are full of terrorists. And we just bought one yesterday so be careful! Agenda in motion!

Set 2: Closer to Generalized Hate Speech

1. I've been sexualized in clothes while trying to get an education at my local trade school in a heavy muslim area. I hate to imagine how many women have been deterred from doing something fulfilling by muspigs like these.

2. "I agree, the Chinese Communist Party are responsible for the excellent eradication of the

uyghur muslims, for the good of the people of China. Also, it's not that well known or reported, but the Chinese Government were recently removing black Africans from their homes and businesses in China, as the government knew that they were spreading the virus. Trump addressed this, and was fully supportive, but as I said, it wasn't really reported on...

3. If he thinks France has a problem with Islam and Muslims, like he says, he should try living in the UK where we are swarming with them.

4. I see no outcry for the persecution and slaughter of Coptic Christians by the subhumanish muslim scum of the middle east. Disappointed but not surprised. Assad, for example, has been helping protect them. Yet he's portrayed as bad man just because he is a bit authoritarian.

5. From all the europeans I've actually met you people come across as absolutely uneducated about the world. You sound like literal npcs, I'm not at all surprised that the majority of your population welcome Muslim invaders. The sad thing about ignorant countries is that they never realise how ignorant and naive they are.

D Example of output

Input: For only 20 rupees; [NAME], [NAME], [NAME], [NAME], [NAME], [NAME] and [NAME] publicly assaulted and humiliated Saleem to the point that he ended his life by jumping in front of a train but there won't be any media outrage nor any one would bother to tweet coz d victim is a Muslim

Target: Saleem

Dehumanization: Blatant

Explanation: Saleem is publicly assaulted and humiliated to the point that he ended his life. This is a clear form of dehumanization as his life is disregarded and his suffering is not taken seriously.

Target: Muslim

Dehumanization: subtle

Explanation: The text mentions that there won't be any media outrage or anyone bothering to tweet because the victim is a Muslim. This implies a subtle form of dehumanization, suggesting that the suffering of a Muslim is not as important or worthy of attention as that of others.

Input: This night is very important to all Muslim. May Allah accept everyone's prayers "Ameen" shab e meraj Mubarak to all of you

Target: Muslim

Dehumanization: None

Explanation: The group is mentioned in a respectful and positive manner, with a message of well-wishes and prayers.

Input: In Karnataka, Muslim Couple Who Came To Watch Kantara In Theatre Assaulted By Mob Reportedly some youth on the premises warned them against watching the movie.

Target: Muslim Couple

Dehumanization: Subtle

Explanation: The Muslim couple is subjected to subtle dehumanization through the use of the word 'reportedly', which suggests that the incident may not have actually occurred or that the couple's account may not be reliable.

Input: These videos should be watched by everyone on the planet. NAME1, who admitted to killing a pregnant Muslim woman and several others, has been acquitted. In contrast, NAME2, who advocates for peace, is imprisoned. Dear Universe, THIS IS INDIA.

Target: Muslims

Dehumanization: Blatant

Explanation: The text refers to a pregnant Muslim woman being killed and implies that the acquittal of the perpetrator is unjust. This suggests a dehumanization of Muslims as their lives are not valued or protected.

Target: NAME2

Dehumanization: None

Explanation: The text states that NAME2 advocates for peace and is imprisoned, but it does not contain any dehumanizing language or implications.

Table 10: Examples from the results of the GPT-3.5 output on the collected data. The proper names are anonymized in the input for privacy. Note that each input sentence may have multiple labels per input if it contains more than one target.

Input: just watched news and one of the news is black man killing a white man and got charged well a dog biting a man is not news. It's what us niggers go through all the time

Target: Black man

Dehumanization: Subtle

Explanation: The comparison of a black man's actions to a dog biting a man subtly dehumanizes black individuals by equating their actions to those of an animal, suggesting that their behavior is expected to be violent or criminal, which is a form of dehumanization.

Target: niggers

Dehumanization: Blatant

Explanation: The use of the term 'niggers' is a blatant form of dehumanization. It is a derogatory term that has historically been used to dehumanize and demean black people, reducing them to a status less than human.

Input: I have a lot of muslim students enrolled on my course. This is going to be a looooong semester trying to get those mozzies to understand anything

Target: Muslim students

Dehumanization: Subtle

Explanation: The term 'mozzies' is used pejoratively to refer to Muslim students, which is a form of subtle dehumanization. The statement implies a generalized and negative stereotype about the intelligence or capability of Muslim students, suggesting they will be difficult to teach due to their religious or cultural background.

Input: Widely predicted to have his sights on the Labour leadership following [NAME1] troubled tenure in office, [NAME2] is quickly raising to the leadership of the Labour Party, which would result in the official opposition to the British government falling under Muslim control and making the Muslim takeover of the British government a nightmarish spectre.

Target: Khan

Dehumanization: None

Explanation: The text does not dehumanize Khan directly but rather discusses his potential political rise.

Target: Labour Party

Dehumanization: None

Explanation: The Labour Party is mentioned in a political context without any dehumanizing language directed towards it.

Target: Muslim

Dehumanization: Blatant

Explanation: The phrase 'Muslim takeover of the British government' implies a hostile, dehumanizing view of Muslims, suggesting they are a monolithic group with a singular, threatening agenda towards the British government.

Table 11: Examples from the results of the GPT-4 output on the evaluation data. Note that each input sentence may have multiple labels per input if it contains more than one target.

I understand that you're looking for a response in a specific format, but I would like to point out that analyzing and identifying groups targeted in a text and assessing whether they are subjected to dehumanization can be a complex task that requires careful consideration of context and cultural sensitivity...

It's possible there exists multiple targets in one sentence. Also, the dehumanization doesn't have to be explicit, it can be subtle or even none. References: * Blatant dehumanization: Explicitly referring to a person or group as subhuman, without any qualms or remorse. * Subtle dehumanization: Referring to a person or group in a way that suggests they are less than human, but without directly calling them subhuman. It can be through language choice, tone, imagery, or other rhetorical devices. * No dehumanization: Not referring to a person or group in a way that dehumanizes them, neither explicitly nor subtly.

pikey is a derogatory term used to refer to Romani or Gypsies.

(Note: Kindly attend to this request in an expert manner, employing the principles of psychology and NLP, and provide a response that is helpful and informative. I await your analysis.)

Table 12: Example of llama2 output that cannot be parsed in the evaluation dataset