

Tackling Social Bias against the Poor: A Dataset and Taxonomy on Aporophobia

Georgina Curto¹, Svetlana Kiritchenko², Muhammad Hammad Fahim Siddiqui³,
Isar Nejadgholi², Kathleen C. Fraser²

¹United Nations University Institute in Macau, Macau SAR, China

²National Research Council Canada, Ottawa, Canada

³University of Ottawa, Ottawa, Canada

Correspondence: curto@unu.edu, svetlana.kiritchenko@nrc-cnrc.gc.ca

Abstract

Content Warning: This paper presents textual excerpts that may be offensive or upsetting.

Eradicating poverty is the first goal in the United Nations Sustainable Development Goals. However, *aporophobia*—the societal bias against people living in poverty—constitutes a major obstacle to designing, approving and implementing poverty-mitigation policies. This work presents an initial step towards operationalizing the concept of *aporophobia* to identify and track harmful beliefs and discriminative actions against poor people on social media. In close collaboration with non-profits and governmental organizations, we conduct data collection and exploration. Then we manually annotate a corpus of English tweets from five world regions for the presence of (1) direct expressions of *aporophobia*, and (2) statements referring to or criticizing *aporophobic* views or actions of others, to comprehensively characterize the social media discourse related to bias and discrimination against the poor. Based on the annotated data, we devise a taxonomy of categories of *aporophobic* attitudes and actions expressed through speech on social media. Finally, we train several classifiers and identify the main challenges for automatic detection of *aporophobia* in social networks. This work paves the way towards identifying, tracking, and mitigating *aporophobic* views on social media at scale.

1 Introduction

Poverty is a multidimensional phenomenon that affects 712 million people worldwide. Sub-Saharan Africa continues to be the region with the highest number of people (411.15 million) living under the poverty line of \$2.15 a day (The World Bank, 2024). Yet, poverty is not only a challenge for developing economies. In the United States, there are 37.9 million people living in poverty (Creamer et al., 2022), and in Europe, 4.6 million people

are at risk of poverty (Eurostat, 2023). While the traditional efforts to mitigate poverty have been losing effectiveness in the last decades (The World Bank, 2022), the capability approach to human development (Sen, 2001; Nussbaum, 2012; Comim et al., 2018) has become the main interdisciplinary alternative to the traditional economic frameworks on poverty mitigation.

The capability approach defines poverty as one’s lack of capabilities (where *capabilities* are the things that people can be and do if they choose so, like being in good health, getting married, being educated, or traveling) to conduct a meaningful life with dignity (Sen, 1979). Therefore, the focus of the approach is on increasing individuals’ well-being and agency. While classical utilitarian approaches center on the redistribution of wealth (Bentham, 2010; Mill, 2017), these often constitute a palliative strategy that fails to address the underlying dimensions of poverty (Sen, 1979). In turn, it has been argued that the rhetoric of equal opportunity (Rawls, 1971) could contribute to blaming the poor for their condition (Fishkin, 2016; Sandel, 2020). This is especially relevant in countries such as the United States and Canada, which embrace the narrative of the “land of opportunity” (Desmond, 2023; Curto et al., 2024), but in reality have low rates of social mobility (Chetty et al., 2014).

This paper aligns with the capability approach to human development and aims to provide language resources and new insights to tackle poverty by identifying, tracking and analyzing societal bias against people living in poverty, known as *aporophobia*. *Aporophobia* has been defined as the “rejection, aversion, fear and contempt for the poor” (Cortina, 2022). It increases the burden of poverty, impacting the well-being of this vulnerable group, and constitutes an obstacle to poverty mitigation. When the poor are blamed for their situation and considered undeserving of help, it is harder for

policymakers to approve and implement poverty-mitigation strategies (Arneson, 1997; Everatt, 2009; Nunn and Biressi, 2009).

The way society imagines and acts towards the poor is part of the poverty phenomenon, independently of how one decides to measure it (Sherman, 2001). Aporophobia is rooted in our beliefs and manifests through different degrees of attitudes and actions (Comim et al., 2020). In this study, we identify and characterize aporophobia by analyzing how it is expressed through language, leveraging natural language processing (NLP) techniques. We distinguish between *direct* expressions of personal aporophobic views and *reporting* on or discussing aporophobic views and actions of other people or institutions. While both types of expressions signify aporophobic bias in the society, this distinction can be informative for countering measures. ‘Direct’ expressions of aporophobia may directly be confronted through removal from public view (for severe expressions) or addressed with counterspeech. The ‘reporting’ instances, on the other hand, inform more complex mitigation strategies, often outside the scope of online platforms.

Since there are no available data resources annotated for aporophobia, we start by collecting and annotating English texts from various regions that refer to poor people. This provides us with a better understanding of the diversity of commonly expressed beliefs and behaviors regarding the poor. We focus on social media, and in particular X (formerly Twitter), as our data source for this paper.¹ While the discourse on this platform does not comprehensively represent the opinions of the world population, it constitutes a valuable resource for exploring vast amounts of data and conducting a preliminary analysis that can be complemented with other sources in further studies.

We analyze the data and create a taxonomy of aporophobic actions expressed through speech as the first resource to characterize aporophobia on social media. Then, we experiment with machine learning models to evaluate the viability of automatic aporophobia detection. The results of this evaluation are relevant for future work aimed at identifying and tracking aporophobia at scale, both regionally and globally. This would allow us to study the potential correlation between aporopho-

¹We note that while some aporophobic posts can constitute hate speech according to social media platforms’ policies on safe content, most aporophobic texts do not violate such safety guidelines.

bia and socio-economic indicators as well as the impact of various policies in different world regions.

Thus, the main contributions of this work are as follows:

- We present a novel data collection and annotation method, which relies on unsupervised topic modeling, to considerably increase the proportion of target (aporophobic) instances in a data sample to be labeled. We ensure fair representation of various geographical regions in the dataset by oversampling posts from the underrepresented regions, including Africa, South Asia, and Oceania.
- Grounded in cognitive science, philosophy of discrimination, and human development literature, as well as the inputs from specialized non-profits and governmental bodies, we manually annotate social media posts and devise a taxonomy of expressed aporophobic behaviors corresponding to various degrees of discriminative action. These categories are grouped under two types of speech: ‘Direct’ (expressing the speaker’s own aporophobic views) and ‘Reporting’ (stating or criticizing the views and behaviors of others).
- We release the first-ever dataset, DRAX (Direct and Reported Aporophobia on X), manually annotated for aporophobia and containing 1,816 English tweets from five world regions.²
- Using the annotated data, we investigate the feasibility of automatic aporophobia detection with various language models. Then, through an error analysis, we identify future directions for improving aporophobia detection models.

2 Related Work

In the last decades, an important body of work has been devoted to studying and mitigating discrimination against vulnerable groups, namely those defined on the grounds of sex, race, religion, political or other opinion, national or social origin, and other characteristics (United Nations, 1966; Council of Europe, 2010). Sadly, aporophobia has not received the attention it deserves in the literature as a distinct field of study (Curto et al., 2022). It was not

²The dataset is publicly available for research purposes: <https://svkir.com/projects/aporophobia-data.html>

until the 1990s that the term was coined by philosopher Adela Cortina (Cortina, 2022). The study of aporophobia unveils the prejudices on poverty, which has an impact on essential topics, including public policy (Comim et al., 2020), the judicial system (Terradillos Basoco, 2020) and the economic power dynamics (Desmond, 2023). More importantly, aporophobia affects the dignity of the disadvantaged as “the mere fact of being poor is itself cause for being isolated, left out, looked down upon, alienated, pushed aside, and ignored by those who are better off.” (Narayan and Petesch, 2002).

Interdisciplinary research in AI has attended to the detection and mitigation of biases with two ultimate goals: *ethical AI*, which aims to mitigate inherent biases in models (Curto et al., 2022; Blodgett et al., 2020; Lalor et al., 2022) and *AI for social impact*, which constitutes policy-relevant research where AI is used as a tool to detect and mitigate societal biases (Aguilera et al., 2024; Havaladar et al., 2024; Li et al., 2024). Our work distinctly focuses on the latter.

AI researchers, and specifically NLP scholars, have addressed social biases online by creating annotated datasets (Poletto et al., 2021), developing algorithms to detect biased language (Alkomah and Ma, 2022; Jahan and Oussalah, 2023) and to counter it (Fraser et al., 2023; Nejadgholi et al., 2024; Hassan and Alikhani, 2023), extending these models to accommodate multiple languages (Chhabra and Vishwakarma, 2023) and cultural contexts (Lee et al., 2023, 2024), and establishing ethical guidelines for deploying such technologies (Kiritchenko et al., 2021; Garg et al., 2023). However, most of these attempts focused on racism and sexism (Mansur et al., 2023) and, in some cases, bias against immigrants (Pitropakis et al., 2020), sexual minorities (Chakravarthi et al., 2022) or religious groups (Vidgen and Yasseri, 2020).

Only recently has aporophobia emerged as a topic in NLP research that revealed the pervasive yet under-studied issue of bias against poor people on social media (Kiritchenko et al., 2023). Research studies explored the association between crime and poverty on social media across eight geographically diverse countries (Curto et al., 2024), the correlation between aporophobia and increased wealth inequality (Aguilera et al., 2024), negative public attitudes towards people experiencing homelessness (Ranjit et al., 2024), and the impact of aporophobia on the stigmatization of specific groups (Brate et al., 2024). While these studies

solidified the necessity of confronting aporophobia using AI tools, our work is the first step in providing the necessary resources for training models to detect and track this form of bias at scale.

Creating resources to detect social biases online poses significant challenges (Yin and Zubiaga, 2021; Kovács et al., 2021). Firstly, the identification of biases needs to be rooted in cognitive science (Allport, 1954; Kahneman, 2011), the philosophy of discrimination (Young, 2022; Honneth, 1996; Taylor, 1931) and sociology (Fuchs, 2018; Kozłowski et al., 2019), among other disciplines, and involve domain experts. Secondly, while biases have been traditionally studied from a single-dimension perspective, they should be considered in a multi-axial approach (Hoffmann, 2019), for example, examining how aporophobia and xenophobia aggravate one another (Cortina, 2022).

Thirdly, collecting data to identify social biases is complex due to its scattered distribution on social media and the subtle, complex ways such beliefs are expressed (Vidgen and Derczynski, 2020). Random sampling results in a very small number of examples of the harmful class, which is insufficient for training classifiers (Vidgen et al., 2019b). Researchers have tackled this issue by employing targeted search techniques, using keywords and hashtags, and also boosted random sampling techniques (Madukwe et al., 2020; Naseem et al., 2021), which has been shown to introduce biases and over-reliance on keywords in trained classifiers (Wiegand et al., 2019). In this work, we first use generic keywords to collect data on the target population and then employ unsupervised topic modeling, which takes into account context beyond mere keyword matching, allowing us to concentrate on annotating data from the most relevant topics.

Finally, annotating such data is a demanding task, requiring precise definitions to capture nuanced expressions of harmful language and must address annotators’ disagreements, stemming from subjective interpretations across different cultures and personal experiences (Vidgen and Derczynski, 2020; Curry et al., 2024). While many studies have applied a binary classification, few works have attempted to create more detailed fine-grained categories (Salminen et al., 2018; Kirk et al., 2023). We annotate our dataset for three high-level categories and then devise a fine-grained taxonomy that is both grounded in the literature and on the findings from the dataset annotation.

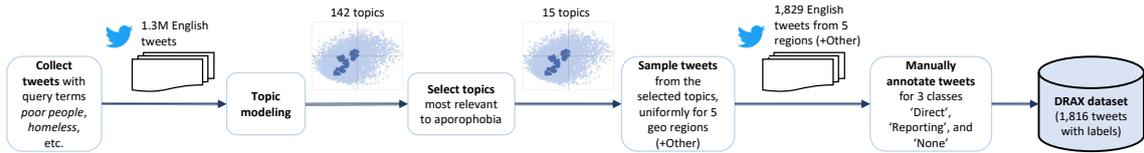


Figure 1: The process diagram for the novel methodology to collect and annotate data.

3 DRAX Dataset Collection and Annotation

3.1 Tweet Collection and Data Sampling

Through the Twitter API,³ we collected tweets in English between 25 August 2022 and 23 November 2022 using the following query terms: *the poor* (used as a noun as opposed to an adjective, as in ‘the poor performance’), *poor people*, *poor ppl*, *poor folks*, *poor families*, *homeless*, *on welfare*, *welfare recipients*, *low-income*, *underprivileged*, *disadvantaged*, *lower class*. (Further details on query term selection and tweet pre-processing are in Appendix A.1.) By using tweet location (‘place’ field), where available, or user location field, we grouped tweets into the following six regions: North America, Europe, Africa, South Asia, Oceania, and Other.⁴ Most of the collected tweets (over 62%) fall into the Other category; 26% come from North America, about 7% are from Europe, and only about 1% were written in Africa, South Asia, and Oceania.

Annotating a random sample of tweets for the presence of aporophobia would be inefficient since only a small percentage of tweets express aporophobic attitudes. On the other hand, sampling data instances that contain specific (harmful) keywords can significantly restrict and bias the collected data. Therefore, we propose a novel data collection method, based on topic modeling (depicted in Figure 1). Previous work showed that topic modeling is an effective tool in uncovering the subtle details within a raw dataset, taking into account the full context (Nejadgholi and Kiritchenko, 2020; Bourgeade et al., 2023; Piot et al., 2024).

First, we masked the query terms in the retrieved tweets, and applied unsupervised topic modeling, using BERTopic (Grootendorst, 2022), on a random sample of 600K tweets (see Appendix A.2 for

technical details). Overall, 142 topics were identified. Then, we manually analyzed the topic words and the most representative example tweets from the obtained topics and selected 15 topics highly relevant to the concept of aporophobia. The selected topics include subjects such as blame, crime, substance abuse, immigrants and refugees, and racism, among others. (The full list of topics with example tweets is available in Table A.1 in Appendix.) Finally, tweets from each of the 15 topics were randomly sampled to satisfy the following two conditions: (1) uniform distribution by region (equal amounts of tweets are sampled from each of the six geographical regions), and (2) uniform temporal distribution (equal amounts of tweets are sampled from each of the three months of the collected data). For topics with a broad relevant content (on crime, drug abuse, immigrants and refugees, and racism), around 250 tweets were sampled per topic; for the rest of the topics, around 100 tweets per topic were included. Since prevalence of topics differed across the regions, we could not always sample the exact number of tweets for each topic and region. Overall, this process resulted in 1,829 tweets for manual annotation.

3.2 Data Annotation

We manually annotated each of the selected textual instances (tweets) with one of the three categories: (1) ‘Direct Aporophobia’, defined as text expressing the speaker’s own aporophobic views, (2) ‘Reporting Aporophobia’, defined as text stating or criticizing the aporophobic views and behaviors of others, or (3) ‘None’ (none of the above). While ‘Reporting’ posts do not necessarily promote aporophobic views, both categories ‘Direct’ and ‘Reporting’ comprise instances referring to aporophobic attitudes or actions and present evidence on levels of aporophobic bias in the society.

In the first round of annotation, a team of two field experts, with multidisciplinary knowledge on human development, poverty studies, aporophobia, and NLP (authors of this paper), annotated 25 tweets from each selected topic and devised the

³Data was collected before the introduction of the paywall.

⁴‘Other’ group includes tweets with no geographical information in either field, tweets with location strings that could not be parsed into a valid location, and tweets originated from countries in other world regions, for which only very limited English data was available (e.g., Middle East).

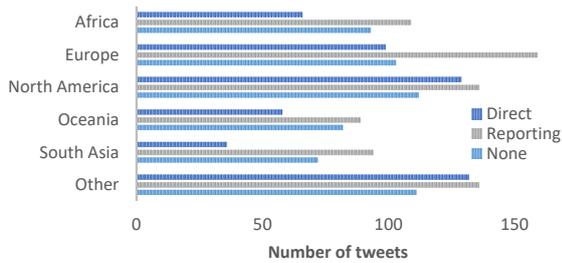


Figure 2: Class distribution in DRAX per geographical region.

classification schema and detailed annotation guidelines (available in Appendix C). Non-profits and governmental organizations specialized in poverty provided guidance during the annotation, and validated the approach and the resulting outcomes in a process that was performed in several iterations. Then, a third annotator (also an author of this paper) was trained using the guidelines and the annotated sample. Subsequently, all selected tweets were independently annotated by two annotators.

The overall inter-rater agreement was 71% (Cohen’s kappa of 0.57). Disagreements mostly emerged due to different interpretations of the annotated texts since tweets are short and sometimes lack the necessary context for precise interpretation of the intended meaning. Furthermore, aporophobia, as with any discriminatory phenomenon, is hard to define precisely for borderline cases. After individual annotations were completed, all the disagreements were discussed among the annotators and resolved in the final dataset. Thirteen tweets were judged as lacking sufficient context for annotation and were removed.

3.3 Characteristics of the DRAX Dataset

The resulting dataset contains 1,816 annotated tweets. We refer to it as DRAX (Direct and Reported Aporophobia on X). There are 520 (29%) instances labeled as ‘Direct Aporophobia’, 723 (40%) instances labeled as ‘Reporting Aporophobia’, and 573 (32%) instances labeled as ‘None’.

Figure 2 shows the class distribution per geographical region.⁵ For each region, we annotated 202–377 tweets. One can observe that tweets from North America and those grouped as ‘Other’ in terms of location, contain the highest proportion

⁵In Table A.2 in Appendix, we break this distribution down per country. Not all countries in a region are represented equally. For most regions, the overwhelming majority of tweets originate from the largest English-speaking country, such as the U.S. for North America and the U.K. for Europe.

of ‘Direct’ aporophobic actions expressed through speech (~35%).⁶ In contrast, subsets from South Asia, Africa, and Oceania include the lowest proportion of tweets in the ‘Direct’ category (18%, 25%, 25%, respectively). The data from Europe presents the highest amount of tweets ‘Reporting’ aporophobic actions of others.

The topics with the highest proportion of ‘Direct’ aporophobic statements are those referring to drug addiction and mental health issues, and immigrants and refugees (see Figure A.1 and Table A.1 in Appendix). Other topics with a high proportion of tweets in the ‘Direct’ category refer to crime, homeless encampments, smell, alcohol addiction, and fear. Such messages often stereotype poor people, and especially the homeless group, as substance addicts and criminals, or express the general attitudes of fear and contempt toward the group. Another topic contains texts communicating the views of rejecting immigrants and refugees since they do not bring any resources and depend on the state’s support.

Topics with a high rate of ‘Reporting Aporophobia’ refer to racism, crime, hatred, the military, laws and courts, laws and regulations, and blaming the poor. Messages in these topics often criticize the governments and people in power for taking advantage and discriminating against poor people through unfair enforcement of laws and regulations, and blaming all social and economic issues on the lower socio-economic class. Black communities are seen as the most targeted since race-based discrimination results in both social and economic disadvantages.

4 Taxonomy of Aporophobic Actions Expressed Through Speech

Based on the qualitative analysis of the DRAX dataset, the conceptual framework for the nature of prejudices, bias and discrimination (Allport, 1954; Kahneman, 2011; Taylor, 1931; Honneth, 1996; Blodgett et al., 2020; Fuchs, 2018) and the concept of aporophobia (Cortina, 2022; Comim et al., 2020), we devise a taxonomy of aporophobic actions expressed through speech (depicted in Figure 3).

⁶We note a very similar class distribution for the tweets from North America and those grouped as ‘Other’. Looking at the content of tweets in the ‘Other’ group, we conjecture that a large portion of tweets with no identifiable geo-location originate from the U.S.

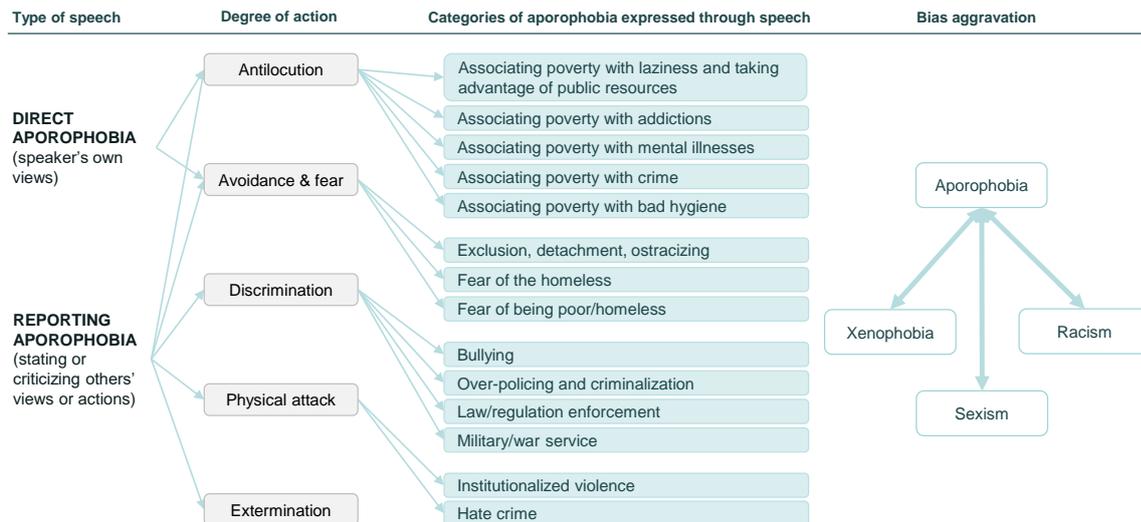


Figure 3: Taxonomy of categories for three levels of classification of aporophobia: Type of speech, Degree of action, Categories of aporophobia expressed through speech.

The first level of classification “Type of Speech” (Figure 3, left column) defines whether the text constitutes ‘Direct’ aporophobia or whether it is ‘Reporting’ the aporophobic views or actions of others, as defined in Section 3.2. Here, note that statements refuting stereotypes, such as “*most poor people aren’t violent*” are also included in the ‘Reporting’ category. Despite the fact that these sentences negate the stereotypes, they at the same time acknowledge and may indirectly reinforce them (Beukeboom and Burgers, 2019). Both ‘Direct’ and ‘Reporting’ categories reveal the levels of aporophobic bias in the society. However, these two types of expressions may require different mitigating strategies.

The second level of classification “Degree of Action” (Figure 3, second left column) corresponds to Allport’s categorization for the different degrees of negative action resulting from prejudices: antilocution (i.e., negative verbal remarks), avoidance and fear, discrimination, physical attack, and extermination (Allport, 1954). It allows classifying aporophobic instances based on their severity: from ‘Antilocution’ in sentences such as “*poor people tend to abuse alcohol*” to the most severe category of ‘Extermination’, evidenced in sentences such as “*poor people should be euthanized*”. We consider that the texts are a form of speech, and as such can *directly* express a person’s aporophobic attitudes through antilocution or fear (e.g., statements expressing prejudice) or *report* on aporophobic attitudes and behaviors of others. However, behavior such as discrimination, physical attack, and ex-

termination are actions that occur out in the real world, and therefore can only be *reported* on social media. For that reason, the former two categories (antilocution and fear) appear under both direct and reported aporophobia, but the latter three categories (discrimination, physical attack, and extermination) only appear under the reporting category.

The third level of classification “Categories of Aporophobia Expressed Through Speech” (Figure 3, middle column) corresponds to the different themes of aporophobia expressed through language, based on the analysis of the DRAX dataset. We see themes on negative attitudes towards poor people (antilocution) due to the stereotypical associations of this group with substance abuse, mental illnesses, crime, bad hygiene, and the general belief that poverty is the result of laziness and lack of agency to hold one’s life in control (e.g., “*he had so many chances but he told me that he would rather do drugs and be homeless*”). Another common theme is fear and ostracism of poor people and especially people in a situation of homelessness (e.g., “*homeless encampments are taking over streets and parks, residents are in desperation for the unsafe conditions*”). Many posts report on the tendency of people in power to blame any negative circumstances on poor people (bullying) as in the statement “*It’s not easy to blame politicians and their friends. It’s easier to blame the poor and their families.*” Furthermore, criminalization of poverty and over-policing as well as disproportionate law and regulation enforcement on people lacking resources have been actively criticized in posts from

across the world (e.g., “*Typical, the law is meant only for the poor.*”, “*American justice system: rich people get warnings and probation, poor people go to prison.*”). We also see messages reporting incidents of physically attacking homeless persons (e.g., “*Two arrested for beating and looting a homeless old man*”). Appendix C provides additional examples for each of the categories.

It is important to emphasize that most common themes associated with poverty (addictions, crime, laziness) put the blame for living in poverty on the poor themselves. However, in reality, the overwhelming majority of the poor population in North America (the region with the highest representation in the dataset and the highest level of ‘Direct’ aporophobic instances in DRAX) are either born into poverty or became poor due to circumstances beyond their control, such as disabilities, divorce, illness, old age, or low wages (United Nations, 2018; Desmond, 2023).

Bias Aggravation: The taxonomy (Figure 3, on the right) also highlights the often intersectional nature of bias, when aporophobia can be aggravated by as well as affect other types of discrimination. We particularly found numerous examples where xenophobia and racism are linked with aporophobia. For example, a recurrent argument that appears on social media is the “need to take care of our homeless first” as an argument to reject migrants and ethnic minority groups (e.g., “*We don’t want them! Can’t help our homeless, but sure let’s put a roof over migrants heads*”). Another example is tweets from the United States that reveal how the poor are often assumed to be people of color and immigrants (e.g., “*only Black people are on welfare*”). These preliminary findings unveil existing prejudices, since in fact 44% of the population in poverty in the U.S. are white, followed by 28.4% Hispanic, of any race (Shrider, 2023). Correlations between sexism and aporophobia also constitute an interesting area of study in NLP. Social science literature shows that non-market care work is often excluded from mainstream economics analysis, which can marginalize women and undervalue their (unpaid) contributions to the community (Folbre, 2021).

5 Automatic Aporophobia Detection

In this section, we investigate the viability of automatic aporophobia detection in text data. For this purpose, we experiment with a range of state-of-

the-art NLP techniques, using the DRAX dataset, and establish benchmark results. Further, we analyze the model performance and the different types of errors to point out future research directions for the effective automatic detection of aporophobia.

We approach the task of automatic aporophobia detection as a three-class classification problem (‘Direct’, ‘Reporting’, or ‘None’), and employ pre-trained language models from the BERT family as well as generative large language models (LLMs). Specifically, we leverage BERT (Devlin et al., 2019), BERTweet (Nguyen et al., 2020), RoBERTa (Liu et al., 2019), and DistilBERT (Sanh et al., 2019)—four of the most prominent BERT-family models known for their high performance on various text classification tasks. We fine-tune these models on DRAX to optimize their ability to identify and categorize aporophobic actions expressed through language. Additionally, we perform zero-shot and few-shot experiments using two open-source LLMs, Llama 3.1 405B Instruct⁷ (Llama Team, 2024) and Mixtral 8x22B Instruct,⁸ and three OpenAI’s generative LLMs, GPT-3.5 Turbo, GPT-4o, and GPT-4o mini.⁹ These experiments aim to assess the models’ ability to detect aporophobia with minimal or no task-specific training. In Appendix B.3, we also present the results of the experiments using two existing models trained to detect general toxicity and hate speech. These experiments reveal that the existing models, though effective in general toxicity detection, fall short in accurately identifying aporophobic content, thus highlighting the need for specialized models in this domain.

We split the DRAX dataset into a training and a test subset chronologically, using data from the first two months for training and the last month for testing.¹⁰ This approach allows us to simulate a real-world scenario where models are trained on past data and deployed on newer, unseen data. Table 1 shows the number of instances in each subset.

We fine-tune models from the BERT family on the training portion and evaluate the performance on the test subset. For generative LLMs, we crafted 20 distinct prompts, each providing a brief defini-

⁷<https://huggingface.co/meta-llama/Llama-3.1-405B-Instruct>

⁸<https://huggingface.co/mistralai/Mixtral-8x22B-Instruct-v0.1>

⁹<https://platform.openai.com/docs/models>

¹⁰To ensure the reliability of our results, we also performed a 3-fold cross-validation and obtained similar results.

Data subset	Direct	Reporting	None	Total
Training set	347	494	389	1,230
Test set	173	229	184	586
Overall	520	723	573	1,816

Table 1: The training and test dataset statistics.

tion of the three classes (‘Direct,’ ‘Reporting,’ and ‘None’). The models are instructed to output only one of these labels. For few-shot prompting, we added nine paraphrased examples from the training set, three for each class label. We tested these prompts on a validation set (stratified by region and label), and selected the best-performing prompt for each model for evaluation on the test set. The overall best-performing prompt (few-shot GPT-4o) is listed in Appendix (Table A.5). For performance evaluation, we employ standard evaluation metrics: accuracy and support-weighted average precision, recall, and F1-score. For fine-tuned models, we repeat experiments 3 times with different fixed seeds, and report the average results.¹¹ The results are presented in Table 2.

Among the tested models, RoBERTa achieved the highest overall performance with an F1-score of 64%. While not very high, such performance may already be adequate for the task of tracking major changes in societal attitudes towards people living in poverty in response to local or global events, such as wars, natural disasters, or new policies. For more fine-grained analyses, further improvements on classification models are needed.

Below, we summarize the key insights gained from manual inspection of the best performing classifier’s errors.

Highly nuanced contexts: Classifiers struggle to effectively incorporate the nuances of the context. Surprisingly, we observe that our best-performing classifier makes the most mistakes in detecting the ‘None’ class, labeling more than 50% of it as ‘Direct’ or ‘Reporting’ aporophobia. Deeper investigations of errors uncover additional challenges. First, the classifiers struggle to interpret the nuances in semantic associations between entities mentioned in the text and negative traits or behaviors (who is accused of what). For example, “*the local homeless shelter director and the CVS pharmacist were busted for dealing opioids at the shelter*” is misclassified as ‘Direct’. Second, statements that challenge stereotypical assumptions are not always rec-

¹¹The results across the runs are stable, each within 1-2% of the average.

Model	Acc.	P	R	F1
<i>Fine-tuned models</i>				
BERT	0.62	0.62	0.62	0.61
BERTweet	0.64	0.63	0.64	0.63
DistilBERT	0.63	0.63	0.63	0.62
RoBERTa	0.64	0.64	0.64	0.64
<i>Generative LLMs</i>				
Llama 3.1 405B - zero shot	0.57	0.61	0.59	0.56
Mixtral 8x22B - zero shot	0.63	0.65	0.61	0.60
GPT-3.5 Turbo - zero shot	0.49	0.49	0.50	0.49
GPT-4o mini - zero shot	0.53	0.53	0.53	0.53
GPT-4o - zero shot	0.60	0.58	0.60	0.59
Llama 3.1 405B - few shot	0.60	0.61	0.61	0.60
Mixtral 8x22B - few shot	0.61	0.65	0.59	0.59
GPT-3.5 Turbo - few shot	0.54	0.56	0.54	0.54
GPT-4o mini - few shot	0.55	0.56	0.55	0.55
GPT-4o - few shot	0.63	0.69	0.63	0.63

Table 2: The performance of fine-tuned and generative large language models on the DRAX dataset. The evaluation metrics are overall accuracy (Acc.), and support-weighted average precision (P), recall (R), and F1-score (F1). The highest numbers in each group are in bold.

ognized as such and misclassified as ‘Direct’, e.g., “*Perception. Homeless are often wrongly assumed to be drug addicts, drinkers, ex-offenders etc.*” Another type of error has to do with the direction of the association. For example, the belief that all homeless people are drug addicts is aporophobic, but the viewpoint that drug addiction can lead to homelessness is not. Yet, we observe some instances where the model makes such mistakes, for example misclassifying the text “*Take heroin and you’ll find yourself ruined, homeless, fallen in the gutters*” as ‘Direct’ (instead of ‘None’).

Impact of tone: The model tends to interpret texts written in a more formal tone as more objective (and classify them as ‘Reporting’) and more informal texts as expressing subjective views (and classify them as ‘Direct’). For example, the post “*Very hard to take care of the poor with no policing. No one wants to invest in the poorer areas as they will lose their businesses due to crime.*” is misclassified as ‘Reporting’ (instead of ‘Direct’), whereas the post “*They’re not mutually exclusive Paddy. The govt should be able to house migrants and the homeless.*” is misclassified as ‘Direct’ (instead of ‘Reporting’).

Impact of topic bias: A number of topics are skewed towards one of the classes, leading to higher performance on the majority class for the topic. For example, the model struggles to distinguish between ‘Direct’ aporophobic actions towards refugees and innocuous statements about this

group, due to ‘Direct’ aporophobia being the dominant class in the topic that discusses this group (57%). For example, the message “*So sad and contradictory, with so many refugees and homeless people around the world*” is misclassified as ‘Direct’ aporophobia (instead of ‘None’).

Generalization across geographical regions: The best performance is achieved on tweets from Oceania (F-score of 68%) and North America and South Asia (both with F-score of 67%), and the worst performance is obtained on data from Europe and ‘Other’ groups (F-score of 61%). Further experiments with the training data sampled only from the dominant regions (North America and Other) show a substantial decrease in results for Africa and Oceania, demonstrating the importance of balanced data sampling across various geographical regions (see Appendix B.4 for the detailed results).

Overall, the above observations suggest that specialized language models capable of handling more nuanced contexts may be of help here, yet will substantially increase the computational costs due to further training and utilizing bigger language models. Additionally, larger and more linguistically and semantically diverse training data can mitigate some of the issues. Synthetic data generation techniques can potentially offer the required scale and diversity of the data while reducing the costs of manual data selection and annotation. Finally, machine learning techniques for handling class imbalance (e.g., downsampling the dominant class) may reduce the impact of topic bias. By creating and releasing the annotated dataset, we provide the resources and encourage further work on this challenging NLP task.

6 Discussion and Future Work

This article examines aporophobia expressed through online speech and provides new insights to better understand the multidimensionality of poverty. In particular, we create the first annotated dataset and the first taxonomy for the phenomenon of aporophobia in social networks. These resources present preliminary evidence of how the cultural contexts in different regions affect the beliefs and degrees of negative action towards people living in poverty. While North America has the highest percentage of tweets (~35%) expressing ‘Direct’ aporophobia, posts from European countries express a higher percentage of ‘Reporting’ aporophobic actions. Future research will further explore

this geographical comparative, to provide insights on whether the higher levels of ‘Direct’ aporophobia in North America could be aligned with the tradition of meritocracy (Sandel, 2020) and the narrative of the “the American dream”. This line of work has important implications for policy making (United Nations, 2018), since caricatured narratives towards the poor could reinforce stereotypes and constitute an obstacle for addressing the shortcomings of social protections.

Future work will extend the analysis to cover other languages and data sources, including multimodal data combining text with images and video. Moreover, based on the existing dataset and taxonomy, further studies can examine aporophobia in other types of written texts, such as legal or corporate texts, bringing new insights on aporophobia in diverse real-world scenarios. In addition, while the labeling in the current dataset includes only the categories of ‘Direct’, ‘Reporting’ and ‘None’, future work can incorporate the annotation of all the categories of aporophobia identified in the taxonomy. Further, analyzing and characterizing aporophobia as an aggravator of other types of bias (including racism, xenophobia, and sexism) is an important avenue for future studies. We will also work on extending and improving NLP models for automatic aporophobia detection, with the ultimate goal of tracking and informing on the trends of aporophobia online over time.

This line of research aims to complement other lines of work on aporophobia, such as agent-based modeling simulations, that inform on the impact of poverty-mitigation policies under discussion at local parliamentary level (Aguilera et al., 2024). The capacity of NLP techniques to analyze and track complex social phenomena, such as aporophobia, over vast amounts of real-world data, combined with the capabilities of agent-based modeling to represent social impact contexts constitute a promising new set of tools to tackle poverty by acting on social biases.

Limitations

As with any analysis of an intricate social challenge through computational techniques, this study has limitations. There are 2.6 billion people worldwide who do not have internet access (International Telecommunications Union (ITU), 2023), and not all regions, genders, or identity groups are equally represented online (Chan et al., 2021). Among on-

line users, we capture data only from users of the social media platform X, which is predominantly used in the United States (Barbieri et al., 2020) and represents specific socioeconomic demographics in terms of age, gender, ethnicity, etc. (Mislove et al., 2011). Thus, many demographics are underrepresented, and even if present, they might shy away from expressing their real beliefs. Further, the data is collected only in English and, mainly, the largest English-speaking countries are represented. In addition, the user posts are collected using a pre-specified set of terms in standard English that may exclude related terms in regional dialects. In general, data collection from social media may introduce different types of biases, such as selection bias, platform-induced behavior and public image curation bias. Therefore, the findings from this study may not generalize to population at large and should be considered preliminary.

The data annotation has been performed by three annotators with diverse cultural backgrounds and familiarity with local events and societal expectations in different regions of the world. Still, misinterpretation of content and annotation bias may be present.

Finally, it must be emphasized that, while allowing processing large amounts of data, NLP techniques can only analyze what people say, as opposed to what they think. Complementary data collection methods, such as the generation of an aporophobia-specific Harvard Implicit Association Test (Xu et al., 2014), workshops and interviews with the specialized non-profits and government officials, as well as surveys among the affected population will be considered to obtain additional quantitative and qualitative data.

Ethics Statement

Detecting, tracking, and mitigating the negative impact of aporophobia in the online social discourse poses a number of risks and ethical issues, discussed at length in previous works in the context of abusive and toxic language detection (Hovy and Spruit, 2016; Vidgen et al., 2019a; Cortiz and Zubiega, 2020; Kiritchenko et al., 2021). These issues include tension between freedom of speech and respect for equality and dignity, biased data sampling and data annotation, dual use of technology, and many others. The research, design and deployment of such technology should comply with trustworthy AI principles of transparency, justice and fairness,

non-maleficence, responsibility, and privacy (Jobin et al., 2019). This should include clear communication about the models' limitations, potential biases, and intended use cases. We would also like to emphasize the importance of involving key stakeholders—including non-profit organizations and underserved communities themselves—in the development process to ensure that the technology is aligned with broader social values and equipped with safeguards against potential harm.

Acknowledgments

We would like to thank the teams in Arrels Fundacio, Caritas and Motels4Now, among many other contributors, for their inputs and guidance as domain experts. Their expertise has been key to this research and has allowed us to contextualize what aporophobia means in practice for the affected population.

References

- Alba Aguilera, Nieves Montes, Georgina Curto, Carles Sierra, and Nardine Osman. 2024. Can poverty be reduced by acting on discrimination? an agent-based model for policy making. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pages 22–30.
- Fatimah Alkomah and Xiaogang Ma. 2022. A literature review of textual hate speech detection methods and datasets. *Information*, 13(6):273.
- Gordon W Allport. 1954. *The nature of prejudice*. Basic Books.
- Richard J. Arneson. 1997. Egalitarianism and the Undeserving Poor. *The Journal of Political Philosophy*, 5(4):327–350.
- Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. [TweetEval: Unified benchmark and comparative evaluation for tweet classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.
- Jeremy Bentham. 2010. *An introduction to the principles of morals and legislation*. White Dog Publishing, Pennsylvania.
- Camiel J Beukeboom and Christian Burgers. 2019. How stereotypes are shared through language: A review and introduction of the social categories and stereotypes communication (SCSC) framework. *Review of Communication Research*, 7:1–37.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is](#)

- power: A critical survey of “bias” in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Tom Bourgeade, Patricia Chiril, Farah Benamara, and Véronique Moriceau. 2023. [What did you learn to hate? a topic-oriented analysis of generalization in hate speech detection](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3495–3508, Dubrovnik, Croatia. Association for Computational Linguistics.
- Ryan Brate, Marieke Van Erp, and Antal van den Bosch. 2024. A bayesian quantification of aporophobia and the aggravating effect of low–wealth contexts on stigmatization. In *Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024)*, pages 234–243.
- Ricardo JGB Campello, Davoud Moulavi, and Jörg Sander. 2013. Density-based clustering based on hierarchical density estimates. In *Proceedings of the Pacific–Asia Conference on Knowledge Discovery and Data Mining*, pages 160–172. Springer.
- Bharathi Raja Chakravarthi, Ruba Priyadarshini, Thenmozhi Durairaj, John Philip McCrae, Paul Buiteelaar, Prasanna Kumaresan, and Rahul Ponnusamy. 2022. Overview of the shared task on homophobia and transphobia detection in social media comments. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 369–377.
- Alan Chan, Chinasa T. Okolo, Zachary Turner, and Angelina Wang. 2021. The Limits of Global Inclusion in AI Development. In *Proceedings of the AAAI Workshop on Reframing Diversity in AI*.
- Raj Chetty, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, David Autor, Gary Becker, David Card, David Dorn, John Friedman, James Heckman, Nathaniel Hilger, Richard Hornbeck, Lawrence Katz, Sara Lalumia, Adam Looney, Pablo Mitnik, Jonathan Parker, Laszlo Sandor, Gary Solon, Danny Yagan, Sarah Abraham, Alex Bell, Shelby Lin, Alex Olssen, Evan Storms, Michael Stepner, and Wentao Xiong. 2014. [Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States](#). *The Quarterly Journal of Economics*, 129(4):1553–1623.
- Anusha Chhabra and Dinesh Kumar Vishwakarma. 2023. A literature survey on multimodal and multilingual automatic hate speech identification. *Multi-media Systems*, 29(3):1203–1230.
- F Comim, S Fenel, and P B Anand. 2018. *New Frontiers of the Capability Approach*. Cambridge University Press.
- Flavio Comim, Mihály Tamás Borsi, and Octasiano Valerio Mendoza. 2020. [The Multi-dimensions of Aporophobia](#). MPRA Paper 103124, University Library of Munich, Germany.
- Adela Cortina. 2022. *Aporophobia. Why we reject the poor instead of helping them*. Princeton University Press.
- Diogo Cortiz and Arkaitz Zubiaga. 2020. Ethical and technical challenges of AI in tackling hate speech. *The International Review of Information Ethics*, 29.
- Council of Europe. 2010. [European Convention on Human Rights](#).
- John Creamer, Emily Shrider, Kalee Burns, and Frances Chen. 2022. *Poverty in the US: 2022*. Technical report, United States Census Bureau.
- Amanda Cercas Curry, Gavin Abercrombie, and Zeerak Talat. 2024. Subjective isms? on the danger of conflating hate and offence in abusive language detection. In *Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024)*, pages 275–282.
- G. Curto, M. Jojoa Acosta, F. Comim, and B. Garcia-Zapirain. 2022. Are the poor being discriminated against on the Internet? A machine learning analysis using Word2vec and GloVe embeddings to identify aporophobia. *AI & Society*.
- Georgina Curto, Svetlana Kiritchenko, Kathleen Fraser, and Isar Nejadgholi. 2024. [The crime of being poor: Associations between crime and poverty on social media in eight countries](#). In *Proceedings of the Sixth Workshop on Natural Language Processing and Computational Social Science (NLP+CSS 2024)*, pages 32–45, Mexico City, Mexico. Association for Computational Linguistics.
- Matthew Desmond. 2023. *Poverty, by America*. Crown.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Eurostat. 2023. [Database - Income and living conditions](#).
- David Everatt. 2009. The Undeserving Poor: Poverty and the Politics of Service Delivery in the Poorest Nodes of South Africa. *Politikon*, 35(3):293–319.
- Joseph Fishkin. 2016. [Bottlenecks: A New Theory of Equal Opportunity](#). Oxford University Press.
- Nancy Folbre. 2021. *The Rise and Decline of Patriarchal Systems. An Intersectional Political Economy*. Verso, London - New York.

- Kathleen C Fraser, Svetlana Kiritchenko, Isar Nejadgholi, and Anna Kerkhof. 2023. What makes a good counter-stereotype? evaluating strategies for automated responses to stereotypical text. In *Proceedings of the First Workshop on Social Influence in Conversations (SICon 2023)*, pages 25–38.
- Christian Fuchs. 2018. [Capitalism, Patriarchy, Slavery, and Racism in the Age of Digital Capitalism and Digital Labour](#). *Critical Sociology*, 44(4-5):677–702.
- Tanmay Garg, Sarah Masud, Tharun Suresh, and Tanmoy Chakraborty. 2023. Handling bias in toxic speech detection: A survey. *ACM Computing Surveys*, 55(13s):1–32.
- Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- Sabit Hassan and Malihe Alikhani. 2023. Discgen: A framework for discourse-informed counterspeech generation. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 420–429.
- Shreya Havaldar, Salvatore Giorgi, Sunny Rai, Thomas Talhelm, Sharath Chandra Guntuku, and Lyle Ungar. 2024. NoBuilding Knowledge-Guided Lexica to Model Cultural Variation Title. *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Volume 1:.
- Anna Lauren Hoffmann. 2019. [Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse](#). *Information, Communication & Society*, 22(7):900–915.
- Axel Honneth. 1996. *The Struggle for Recognition*. Polity Press.
- Dirk Hovy and Shannon L. Spruit. 2016. [The social impact of natural language processing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.
- International Telecommunications Union (ITU). 2023. Measuring Digital Development Facts and Figures.
- Md Saroar Jahan and Mourad Oussalah. 2023. A systematic review of hate speech automatic detection using natural language processing. *Neurocomputing*, 546:126232.
- Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. [The global landscape of AI ethics guidelines](#). *Nature Machine Intelligence*, 1:389–399.
- Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Svetlana Kiritchenko, Georgina Curto Rex, Isar Nejadgholi, and Kathleen C. Fraser. 2023. [Aporophobia: An overlooked type of toxic language targeting the poor](#). In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 113–125, Toronto, Canada. Association for Computational Linguistics.
- Svetlana Kiritchenko, Isar Nejadgholi, and Kathleen C Fraser. 2021. Confronting abusive language online: A survey from the ethical and human rights perspective. *Journal of Artificial Intelligence Research*, 71:431–478.
- Hannah Rose Kirk, Wenjie Yin, Bertie Vidgen, and Paul Röttger. 2023. [SemEval-2023 Task 10: Explainable Detection of Online Sexism](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*. Association for Computational Linguistics.
- György Kovács, Pedro Alonso, and Rajkumar Saini. 2021. Challenges of hate speech detection in social media: Data scarcity, and leveraging external resources. *SN Computer Science*, 2(2):95.
- Austin C. Kozlowski, Matt Taddy, and James A. Evans. 2019. [The geometry of culture: Analyzing the meanings of class through word embeddings](#). *American Sociological Review*, 84(5):905–949.
- John Lalor, Yi Yang, Kendall Smith, Nicole Forsgren, and Ahmed Abbasi. 2022. [Benchmarking intersectional biases in NLP](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3598–3609, Seattle, United States. Association for Computational Linguistics.
- Nayeon Lee, Chani Jung, Junho Myung, Jiho Jin, Jose Camacho-Collados, Juho Kim, and Alice Oh. 2024. Exploring cross-cultural differences in english hate speech annotations: From dataset construction to analysis. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4205–4224.
- Nayeon Lee, Chani Jung, and Alice Oh. 2023. Hate speech classifiers are culturally insensitive. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 35–46.
- Bryan Li, Samar Haider, and Chris Callison-Burch. 2024. This Land is Your, My Land: Evaluating Geopolitical Bias in Language Models through Territorial Disputese. *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Volume 1:.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv*, abs/1907.11692.

- AI Meta Llama Team. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Varvara Logacheva, Daryna Dementieva, Sergey Ustyantsev, Daniil Moskovskiy, David Dale, Irina Krotova, Nikita Semenov, and Alexander Panchenko. 2022. [ParaDetox: Detoxification with parallel data](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6804–6818, Dublin, Ireland. Association for Computational Linguistics.
- Kosisochukwu Madukwe, Xiaoying Gao, and Bing Xue. 2020. In data we trust: A critical analysis of hate speech detection datasets. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 150–161.
- Zainab Mansur, Nazlia Omar, and Sabrina Tiun. 2023. Twitter hate speech detection: A systematic review of methods, taxonomy analysis, challenges, and opportunities. *IEEE Access*, 11:16226–16249.
- John Stuart Mill. 2017. *On Liberty*. Amazon Classics, Seattle.
- Alan Mislove, Sune Lehmann, Yong-Yeol Ahn, Jukka-Pekka Onnela, and J.Niels Rosenquist. 2011. Understanding the Demographics of Twitter Users. *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*.
- Deepa Narayan and Patti Petesch, editors. 2002. *Voices of the Poor. From Many Lands*. Oxford University Press & The World Bank.
- Usman Naseem, Imran Razzak, and Peter W Eklund. 2021. A survey of pre-processing techniques to improve short-text quality: a case study on hate speech detection on twitter. *Multimedia Tools and Applications*, 80:35239–35266.
- Isar Nejadgholi, Kathleen C Fraser, Anna Kerkhof, and Svetlana Kiritchenko. 2024. Challenging negative gender stereotypes: A study on the effectiveness of automated counter-stereotypes. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3005–3015.
- Isar Nejadgholi and Svetlana Kiritchenko. 2020. On cross-dataset generalization in automatic detection of online abuse. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 173–183.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Heather Nunn and Anita Biressi. 2009. The undeserving poor. *Soundings*, 41.
- Martha C Nussbaum. 2012. *Creating Capabilities*. Harvard University Press, Cambridge, Massachusetts and London, England.
- Paloma Piot, Patricia Martín-Rodilla, and Javier Parapar. 2024. Metahate: A dataset for unifying efforts on hate speech detection. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 2025–2039.
- Nikolaos Pitropakis, Kamil Kokot, Dimitra Gkatzia, Robert Ludwiniak, Alexios Mylonas, and Miltiadis Kandias. 2020. Monitoring users’ behavior: anti-immigration speech detection on twitter. *Machine Learning and Knowledge Extraction*, 2(3):11.
- Fabio Poletto, Valerio Basile, Manuela Sanguinetti, Cristina Bosco, and Viviana Patti. 2021. Resources and benchmark corpora for hate speech detection: a systematic review. *Language Resources and Evaluation*, 55:477–523.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Jaspreet Ranjit, Brihi Joshi, Rebecca Dorn, Laura Petry, Olga Koumoundouros, Jayne Bottarini, Peichen Liu, Eric Rice, and Swabha Swayamdipta. 2024. Oath-frames: Characterizing online attitudes towards homelessness with llm assistants. *arXiv preprint arXiv:2406.14883*.
- J Rawls. 1971. *A Theory of Justice*. Oxford University Press, Oxford.
- Joni Salminen, Hind Almerkhi, Milica Milenković, Soon-gyo Jung, Jisun An, Haewoon Kwak, and Bernard Jansen. 2018. Anatomy of online hate: developing a taxonomy and machine learning models for identifying and classifying hate in online news media. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Michael J. Sandel. 2020. *The tyranny of merit*. Penguin Random House.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. In *Proceedings of the 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing*.
- A Sen. 2001. *Development as freedom*. Oxford University Press.
- Amartya Sen. 1979. Equality of What? *Stanford University*.
- S Sherman. 2001. *Imagining Poverty: quantification and the decline of paternalism*. Ohio State University Press.
- Em Shrider. 2023. Poverty Rate for the Black Population Fell Below Pre-Pandemic Levels. Technical report, United States Census Bureau.

- Charles Taylor. 1931. *Multiculturalism and "the politics of recognition"*. Princeton University Press.
- Juan Terradillos Basoco. 2020. *Aporofobia y plutofilia : la deriva jánica de la política criminal contemporánea*. J. M. Bosch.
- The World Bank. 2022. [Poverty and Shared Prosperity 2022: Correcting Course](#). Technical report, World Bank.
- The World Bank. 2024. Poverty and Inequality Platform (PIP).
- United Nations. 1966. International Covenant on Civil and Political Rights.
- United Nations. 2018. Report of the Special Rapporteur on extreme poverty and human rights on his mission to the United States of America.
- Bertie Vidgen and Leon Derczynski. 2020. Directions in abusive language training data, a systematic review: Garbage in, garbage out. *Plos one*, 15(12):e0243300.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019a. [Challenges and frontiers in abusive content detection](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy. Association for Computational Linguistics.
- Bertie Vidgen, Helen Margetts, and Alex Harris. 2019b. How much online abuse is there. *Alan Turing Institute*, 11.
- Bertie Vidgen and Taha Yasseri. 2020. Detecting weak and strong islamophobic hate speech on social media. *Journal of Information Technology & Politics*, 17(1):66–78.
- Michael Wiegand, Josef Ruppenhofer, and Thomas Kleinbauer. 2019. [Detection of Abusive Language: the Problem of Biased Datasets](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 602–608, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kaiyuan Xu, Brian Nosek, and Anthony G. Greenwald. 2014. [Data from the Race Implicit Association Test on the Project Implicit Demo Website](#). *Journal of Open Psychology Data*, 2(1):e3.
- Wenjie Yin and Arkaitz Zubiaga. 2021. Towards generalisable hate speech detection: a review on obstacles and solutions. *PeerJ Computer Science*, 7:e598.
- Iris Marion Young. 2022. *Gender, justice, and the politics of difference*. Routledge innovators in political theory. Routledge, Abingdon, Oxon ; New York, NY.
- Xuhui Zhou, Maarten Sap, Swabha Swayamdipta, Yejin Choi, and Noah Smith. 2021. [Challenges in automated debiasing for toxic language detection](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3143–3155, Online. Association for Computational Linguistics.

A Details on the Dataset Creation

A.1 Tweet Collection and Pre-processing

The initial query term list was constructed from social psychology literature and various web sources, and extended with synonyms from thesauruses. Then, terms resulted in few retrieved tweets or large amount of unrelated tweets in a preliminary search were removed. The single term *poor* was excluded since it often appears in unrelated contexts (e.g., ‘poor taste’, ‘poor results’). We also excluded derogatory terms, like ‘tramp’, ‘beggar’, or ‘trailer trash’, since in our preliminary search we found these terms being disproportionately used as personal insults and not directly related to people living in poverty (for example, ‘*You are a trailer trash queen who has no business in congress*’).

Re-tweets, duplicate tweets, tweets with external URLs, and tweets with more than five hashtags were discarded. Further, we removed tweets originated from accounts with the word ‘bot’ in their user name or screen name.

For geographical location we used Twitter fields ‘place’ (tweet location) and ‘user location’. Only about 2% of the collected tweets had exact tweet location (‘place’) specified. On the other hand, about 60% of the tweets had user location field filled. The user location field is a free-form text, so we applied simple string matching to extract the most frequently mentioned countries. When country was not mentioned, we also tried to match U.S. states (or their two-letter abbreviations), Canadian provinces (or their two-letter abbreviations), and major cities in the U.S., U.K., and Canada. Since we collected tweets in English, most of the tweets originated from countries where English is commonly spoken. We grouped the tweets with known geographical location into the following regions: North America (U.S. and Canada), Europe (mostly U.K. and Ireland), Africa (Nigeria, South Africa, Kenya, Uganda, Ghana), South Asia (India, Pakistan, and Philippines), and Oceania (Australia and New Zealand). Tweets from other parts of the world or with unknown location were grouped as ‘Other’.

Topic	Topic words	Examples
5	drug, addicts, mental, drugs, mentally, ill, addiction, health, addicted, addict	Leave the drugs and alcohol out and you won't be homeless, will you? [DIRECT]
6	black, white, color, racist, minorities, blacks, whites, race, racism, brown	Most republicans think only Blacks are on welfare [REPORTING]
10	immigrants, migrants, illegals, illegal, asylum, refugees, seekers, hotels, border, immigration	We are overrun with refugees and not protecting our own homeless [DIRECT]
14	crime, police, cops, criminals, jail, crimes, prison, arrest, commit, criminal	Poor people just steal other people's property instead of getting a job and earning money [DIRECT]
38	hate, hating, hates, hatred, say, despise, just, hated, really, people	Tories just hate the poor [REPORTING]
49	stealing, steal, stole, stolen, thieves, theft, money, steals, thief, millions	It's disgusting to steal from the poor [NONE]
56	war, military, wars, army, soldiers, fight, join, recruitment, peace, die	Men in the military are from most disadvantaged families that have no other option [REPORTING]
67	crime, gas, prices, high, inflation, highest, taxes, higher, record, border	Too much crime, too much pollution, too much homeless [DIRECT]
88	encampments, encampment, city, near, park, street, clear, looks, large, cleared	Clean the homeless encampments now! [DIRECT]
91	smell, piss, bathroom, smells, toilet, bathrooms, toilets, smelling, restrooms, pee	Homeless swarm public transportation, forcing everyone to breathe in their body odor [DIRECT]
96	justice, court, lawyers, lawyer, legal, courts, judges, judge, attorney, law	The objective for the creation of justice centres is to widen access of the poor to justice through legal aid [NONE]
100	drunk, beer, drink, drinking, alcohol, cigarette, drunks, drinks, liquor, smoking	I'd say, alcohol is the drink of the poor and frustrated, not the rich [DIRECT]
106	law, laws, apply, rule, applies, rules, order, break, enforced, legal	Typical: one law for the rich, one law for the poor [REPORTING]
118	fear, scared, scary, anxiety, afraid, terrified, scare, terrifying, fears, mongering	The scariest thing in the world is to be homeless, right? [DIRECT]
139	blame, blaming, fault, blamed, blames, problematic, problems, instead, mistakes, victim	Let's admit that we don't like the poor, blame them for their conditions, but have done nothing to help [REPORTING]

Table A.1: Fifteen topics manually selected as the most relevant to aporophobia. The topics are ordered by the amount of tweets assigned to the topic. The topic words (from BERTopic) are the words that tend to appear frequently in the topic of interest, and less frequently in the other topics. Example tweets are manually selected as the most representative of the topic. Example tweets are paraphrased to preserve the anonymity of the users. The class label for each example is as annotated in the dataset.

A.2 Topic Modeling with BERTopic

BERTopic (Grootendorst, 2022) is a flexible toolkit for supervised, semi-supervised, and unsupervised topic modeling. It utilizes a density-based clustering technique HDBSCAN (Campello et al., 2013), which produces clusters of arbitrary shapes and leaves some documents that do not fit any of the identified topics as outliers. The obtained topics/clusters are then represented with topic words, which are identified using class-based TF-IDF (c-TF-IDF). The ‘topic words’ are defined as the words that tend to appear frequently in the topic of interest, and less frequently in the other topics.

We ran BERTopic in the unsupervised mode with the following parameters. For converting text to numerical representations, we used the sentence transformers method based on the *all-MiniLM-L6-v2* pre-trained embedding model.¹² For the vectorizer model, we used the CountVectorizer method,¹³ and removed English stopwords and terms that appeared in less than 5% of the sentences ($min_df = 0.05$). For the HDBSCAN clustering algorithm, we specified the minimum size of the clusters as $min_cluster_size = 500$. For all the other parameters, the default settings of the BERTopic package were used.

BERTopic identified 142 topics and left about 42% of tweets unclustered. We manually analyzed the topic words and the most representative example tweets from the obtained topics and selected 15 topics highly relevant to the concept of aporophobia. The selected topics, along with the topic words and example tweets, are listed in Table A.1.

A.3 Human Annotators

The three annotators of the DRAX dataset are authors of this paper. Two of them identify as females, and one as male. The ages vary from 20s to 40s. All three have received higher education at Western institutions, but have different cultural backgrounds. They have extensive knowledge on social biases, aporophobia, and NLP.

A.4 Data Distribution

Table A.2 shows the data distribution per country. Note that not all countries in a region are

¹²https://www.sbert.net/docs/pretrained_models.html

¹³https://scikit-learn.org/stable/modules/generated/sklearn.feature_extraction.text.CountVectorizer.html

Country	Direct	Reporting	None	Total
Ghana	4	6	2	12
Kenya	7	22	11	40
Nigeria	17	26	27	70
South Africa	34	50	47	131
Uganda	4	5	6	15
Africa (total)	66	109	93	268
France	1	1	3	5
Germany	3	3	0	6
Ireland	9	17	6	32
United Kingdom	86	138	94	318
Europe (total)	99	159	103	361
Canada	5	16	8	29
United States	124	120	104	348
North America (total)	129	136	112	377
Australia	49	68	63	180
New Zealand	9	21	19	49
Oceania (total)	58	89	82	229
India	28	56	50	134
Pakistan	5	28	17	50
Philippines	3	10	5	18
South Asia (total)	36	94	72	202
Other (total)	132	136	111	379

Table A.2: The number of annotated tweets in the DRAX dataset per region and country.

represented equally. For most regions, the overwhelming majority of tweets originate from the largest English-speaking country, such as the U.S. for North America and the U.K. for Europe.

Figure A.1 presents the class distribution per topic. Topics with the highest proportion of ‘Direct’ aporophobic statements are 5 (drug addiction and mental health issues) and 10 (immigrants and refugees). Other topics with a high proportion of the ‘Direct’ category include 14 (crime), 67 (crime and other issues), 88 (homeless encampments), 91 (smell), 100 (alcohol addiction), and 118 (fear). Topics with a high rate of ‘Reporting’ aporophobia are 6 (racism), 14 (crime), 38 (hatred), 56 (military), 96 (laws and courts), 106 (laws and regulations), and 139 (blame).

B Additional Details and Experiments with NLP Models

B.1 Fine-tuned Language Models

In this paper, we fine-tune several models from the BERT family to tackle the challenging task of aporophobia classification on social media. The first model, the original BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2019), serves as the foundation for many

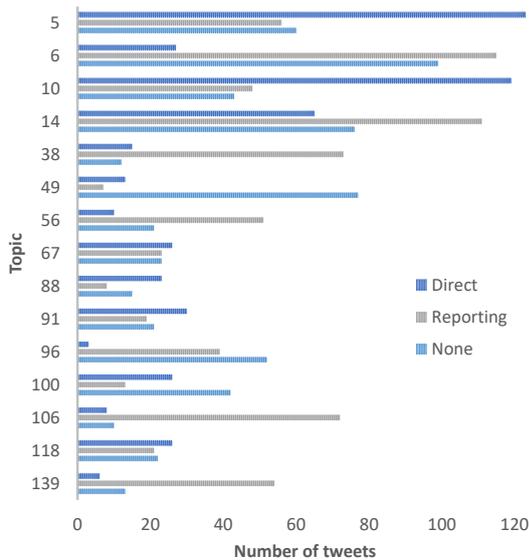


Figure A.1: Class distribution per topic in the DRAX dataset.

modern NLP models. BERT, with its 110 million parameters, is well-suited for tasks like ours. By leveraging pre-training on large text corpora, BERT can be fine-tuned effectively for domain-specific applications. Next, we used RoBERTa (Robustly Optimized BERT Pretraining Approach) (Liu et al., 2019), an advanced variant of BERT, which has the same architecture as BERT but uses a byte-level BPE as a tokenizer (same as GPT-2 (Radford et al., 2019)) and uses a different pretraining scheme. It has 125 million parameters and focuses solely on masked language modeling, which lead to improved performance in our classification experiments. We also fine-tune DistilBERT (Sanh et al., 2019), a distilled version of BERT designed to be more efficient without significantly compromising on accuracy. DistilBERT, with its reduced architecture of 66 million parameters, offers nearly the same language processing capabilities as BERT while being faster and more resource efficient, an important factor when deploying models at scale. Finally, we utilize BERTweet (Nguyen et al., 2020), a variant of BERT specifically pre-trained on English tweets. With the same architecture and 110 million parameters as BERT, BERTweet is particularly adept at handling the informal language found on social media platforms like Twitter. This specialization makes BERTweet an excellent candidate model for our task.

The tweet pre-processing involved removing user mentions and URLs. The experiments were conducted using Google Colab Pro, leveraging

their L4 GPU for faster training and inference. We utilized the models provided by the transformers library. The following hyperparameters were used for each model, representing the best-performing configurations after experimenting with different combinations:

- **BERT, RoBERTa, DistilBERT:** Batch size = 4, Optimizer = Adam, Epochs = 4
- **BERTweet:** Batch size = 8, Optimizer = Adam, Epochs = 4

All reported results are the average of three independent runs with seed values of 42, 62, and 82 to ensure robustness and mitigate the impact of random initialization.

B.2 Generative LLMs

We used both open-source and closed-source models for our zero-shot and few-shot experiments with generative LLMs. For open-source models, we used Llama 3.1 405B Instruct, and Mixtral 8x22B Instruct. The Llama 3.1 405B is Meta’s flagship language model with 405 billion parameters. It features a 128,000-token context window, enabling it to handle extensive textual inputs effectively. The Mixtral 8x22B is a Sparse Mixture-of-Experts (SMoE) model developed by Mistral AI, comprising 141 billion parameters, with only 39 billion active during inference. This architecture offers significant cost efficiency and performance advantages. The model supports a 64,000-token context window, facilitating precise information recall.

In addition, we experimented with three commercial models via the OpenAI chat completions API¹⁴: GPT-3.5 Turbo, GPT-4o mini, and GPT-4o. The GPT-3.5 Turbo 0125 model, with a 16,385-token context window, has been optimized for higher accuracy over GPT-3. GPT-4o mini, a smaller yet powerful model with a 128,000-token context window, has demonstrated more capabilities than GPT-3.5 Turbo and has shown efficiency on lightweight tasks. Finally, GPT-4o, the flagship OpenAI model with the same 128,000-token context window, was developed for handling more complex tasks.

The following parameters were used for all models: temperature $t = 0.7$, the maximum number of output tokens $max_tokens = 10$, and the number of chat completion choices $n = 1$.

¹⁴<https://platform.openai.com/docs/api-reference/chat/create>

Model	Acc.	P	R	F1
ToxDect	0.63	0.59	0.63	0.60
RoBERTa Toxicity	0.67	0.63	0.67	0.64
RoBERTa fine-tuned on DRAX	0.79	0.80	0.79	0.80

Table A.3: Performance of two off-the-shelf RoBERTa-based toxicity models and the binary RoBERTa model fine-tuned on DRAX. The evaluation metrics are overall accuracy (Acc.), and support-weighted average precision (P), recall (R), and F1-score (F1).

Each of these models are tested to explore their performance in classifying aporophobic content with varying levels of prior information provided through prompts. The highest performance on DRAX was obtained with GPT-4o using a few-shot prompt listed in Table A.5.

B.3 Experiments with Off-the-Shelf Toxicity Models

We experiment with two off-the-shelf toxicity detection models: ToxDect¹⁵ (Zhou et al., 2021) and RoBERTa Toxicity Classifier¹⁶ (Logacheva et al., 2022). Both models are based on the RoBERTa model and were fine-tuned specifically for toxicity classification tasks. We selected these models because they are publicly available on Huggingface, were trained on social media data, and are relatively recent. Additionally, their popularity on Huggingface, as indicated by a high number of downloads, further motivated their use, ensuring they are well-known within the community and suitable for our task.

Given that both toxicity classifiers are binary classifiers returning ‘Toxic’ or ‘Non-toxic’ category labels, we adapt our dataset for a binary classification setting. Specifically, we merge the classes ‘Reporting’ and ‘None’ as ‘Non-toxic’, and treat the ‘Direct’ class as ‘Toxic’. For comparison, we also fine-tune a RoBERTa model on the DRAX training subset (modified for the binary classification setting). The results are shown in Table A.3. We observe that the RoBERTa model fine-tuned on the DRAX dataset substantially outperforms the pre-trained toxicity classifiers. This result indicates that while pre-trained toxicity models are generally good at identifying harmful speech, they are not optimized for the specific nuances of aporophobic content, highlighting the need for tailored models

¹⁵<https://huggingface.co/Xuhui/ToxDect-roberta-large>

¹⁶https://huggingface.co/s-nlp/roberta_toxicity_classifier This model is licensed under the OpenRAIL++ License.

in this context. This conclusion is in line with the observations made by Kiritchenko et al. (2023) that existing toxicity and hate speech detection models and data resources are not effective for aporophobia detection.

B.4 The Impact of Uniform Sampling across Regions

To evaluate the impact of our oversampling strategy (i.e., uniform sampling from all six regional groups), we repeated the experiments with the training data sampled only from the dominant regions (North America and Other). We fine-tuned RoBERTa model on the training subset originated from North America and Other, and evaluate the performance of the model on the full test set. The results are shown in Table A.4. Since the training set is now smaller than before (two regions vs. six), the overall performance decreased. The largest losses in the weighted averaged F-score are observed for Oceania (11%) and Africa (7%), whereas the performance for North America decreased only by 3%. Moreover, the results for the ‘Direct’ category (the main category of interest) are the lowest for Africa and Oceania: F-score of 52% and 53% respectively, compared to 65% for North America and 66% for Other, with these differences driven predominantly by the differences in precision. These results demonstrate the importance of balanced data sampling across various geographical regions.

Region	Training on North America & Other				Full Training Set
	Direct	Reporting	None	Overall	Overall
Africa	0.52	0.65	0.47	0.56	0.63
Europe	0.56	0.67	0.45	0.57	0.61
North America	0.65	0.68	0.58	0.64	0.67
Oceania	0.53	0.74	0.38	0.57	0.68
South Asia	0.62	0.76	0.45	0.63	0.67
Other	0.66	0.61	0.32	0.55	0.61
All	0.61	0.68	0.45	0.58	0.64

Table A.4: Performance of the RoBERTa-based model fine-tuned on the training data pertained to the dominant regions (North America and Other) and the RoBERTa-based model fine-tuned on the full training data. The evaluation metric is support-weighted average F1-score.

System prompt
You are an AI trained to classify tweets based on their content into one of three categories: “Direct”, “Reporting”, or “None” with respect to aporophobia.
User prompt
Classify the tweet into one of these categories: “Direct”, “Reporting” or “None” Aporophobia involves negative sentiments and actions directed at the poor, including rejection, fear, and contempt. Direct: The tweet shows the author’s own aporophobic attitudes, such as derogatory remarks, negative stereotypes, or support for exclusionary practices. Reporting: The tweet mentions or critiques aporophobic behaviors or attitudes of others without expressing the author’s own negative views. None: The tweet does not exhibit any aporophobic content or references. Examples: Tweet, Classification “Let’s be honest though, illegals are at least normal people trying to get a better life; homeless are usually mentally unwell junkies that would strip the copper out of the furniture for more crack” , Direct “No, we gonna get the streets cleaned up, stop crime and do something about the homeless invasion. NM is a shithole atm.”, Direct “Put the homeless in jail and start work camps and chain gangs.” , Direct “Being poor does not take away your moral agency. Most poor people manage to not commit violent crime.” , Reporting “So a theft automatically equates to a homeless person?” , Reporting “Those kinda people never do time. Prison is for poor people.” , Reporting “Local charities are working hard to provide shelters and job training programs for the homeless to help improve their lives.” , None “Our city has launched a new initiative aiming to offer better healthcare access to the poor, ensuring everyone receives the care they need.”, None “A recent study highlights the importance of community support in uplifting homeless individuals, demonstrating significant improvements in their well-being through local outreach programs.” , None Tweet: “Tweet” Limit the response to only one category from “Direct”, “Reporting” or “None” Select the category:

Table A.5: Best-performing few-shot prompt for GPT-4o.

C Annotation Guidelines

Aporophobia is defined as “rejection, aversion, fear and contempt for the poor” (Cortina, 2022). In or-

der to annotate aporophobia in textual instances of social media, we make distinction between (1) direct instances of aporophobic beliefs and attitudes expressed by the speaker (toxic language) and (2)

instances of aporophobic attitudes the speaker reports about others. Therefore, we have three main categories to be annotated: ‘**Direct Aporophobia**’, ‘**Reporting Aporophobia**’, and ‘**None**’.

We note that aporophobia can manifest through the different degrees of action resulting from prejudice (Allport, 1954):

1. **Antilocution or verbal rejection:** when an in-group freely purports negative images of an out-group (negative stereotypes, jokes, negative statements).
2. **Avoidance and fear:** when members of the in-group actively avoid people in the out-group (expressing feelings of fear and the desire to avoid any contact).
3. **Discrimination:** when the prejudiced person makes active detrimental distinctions, by denying the out-group opportunities and services. Note that segregation is considered an institutionalized form of discrimination, enforced either legally or by tradition.
4. **Physical attack:** Prejudices can also lead to acts of violence or semi-violence, such as forcibly evicting families from their homes or neighborhoods, or physically attacking persons in a situation of homelessness.
5. **Extermination:** Lynchings, massacres are the ultimate degree of violent expression of prejudice. Even though we do not expect to see expressions of extermination directly linked to aporophobia on social media, we acknowledge that aporophobia can be an element in the prepared ground of previous hostility towards a particular group that is attacked and where other types of discrimination are intertwined.

Note that the categories ‘antilocution’ and ‘avoidance and fear’ can be expressed directly by the speaker on social media as well as reporting the attitudes of others, therefore, they can be labeled as either ‘Direct Aporophobia’ or ‘Reporting Aporophobia’. In contrast, the categories ‘discrimination’, ‘physical attack’ and ‘extermination’, when they appear in textual instances, can only be considered ‘Reporting Aporophobia’.

In reporting statements, we often find *benevolent* instances where the speaker opposes the stereotype, such as in the following examples: “*Most poor*

people manage to not commit violent crime.”,¹⁷ “*I don’t think you have to do drugs to be homeless*”. We consider these as instances of ‘Reporting Aporophobia’ because, despite having a good intention, such messages indirectly acknowledge and reinforce the stereotype (Beukeboom and Burgers, 2019).

These guidelines are based on our position that associating poor/homeless people with negative behavior, like addiction or crime, even if based on factual information, perpetuates the negative stereotypes and hinders social and economic measures to reduce poverty. We draw parallels with racism and sexism where stereotypical associations based on current standings (e.g., more doctors are male) are considered not acceptable for a just society.

DIRECT APOROPHOBIA

Antilocution: includes, but is not limited to the following topics:

- Associating poverty with laziness and taking advantage of public resources: “*Homeless population don’t pay taxes and get lots of free-bees*”; “*Poor people always blame others for their misfortunes.*”; “*Rich people see opportunities and focus on rewards. Poor people see obstacles and focus on the risks*”; “*Most poor people don’t want to try new careers or new business opportunities*”.
- Associating the poor with addiction: “*Do you realize that most of these folks are homeless BECAUSE they’re addicts and not the other way around?*”; “*If you don’t subsidize drug users and the homeless issue will get cut in half overnight*”
- Associating the poor with mental illness: “*Most homeless are mentally ill, so put them into a home for the mentally ill. If there is a shortage of institutional homes, build more. Problem solved.*”; “*Let’s be honest, illegals are at least normal people, homeless are usually mentally unwell junkies that would do anything for more crack.*”
- Associating the poor with crime, which can be divided into:
 - Overestimating the correlation between poverty and crime: “*Most of the students*

¹⁷All example tweets in these guidelines are paraphrased to preserve the anonymity of the users.

of such institutions are from poor families who are even willing to commit state organized crimes”; “But crime is a route that many poor people take just to survive.”

- The criminalization of poverty: *“put the homeless in jail and start work camps”; “ENFORCE THE LAW ON THESE CRIMINAL “HOMELESS” who refuse services.”; “Prison would be a lot more comfortable than cold pavements and begging for cash.”*
- Associating poverty with bad hygiene: *“you can stink like the poor”; “smells like poor people in here”; “I can’t get the smell of Poor People out of my hands”*

Avoidance and fear: includes, but is not limited to the following topics:

- Exclusion, detachment and ostracizing: *“We gonna get the streets cleaned up and end the homeless invasion”; “That is the place to move!! No homeless and no migrants!!”*
- Fear of poor / homeless people: *“I check every room in my house for a homeless man who could be lurking in my carpet”; “homeless-phobia is a big reason things have changed.”*
- Fear of being poor / homeless: *“Being homeless is the scariest thing ever.”* Note: expressing concern about being in a situation of poverty is understandable since one might not have the necessary resources to cover one’s needs and conduct a meaningful life with dignity. We consider specific expressions in this category when we believe they also imply social stigma or ostracism.

REPORTING APOROPHOBIA

Antilocution: includes, but is not limited to the following topics:

- Associating poverty with laziness and taking advantage of public resources: *“Me or anyone on social welfare is not a burden. Our situations are not unique, this could happen to anyone.”; ““Should we throw more taxpayers money at them?” You seem to be blaming the entire ‘disadvantaged’ community”; “it’s the ‘being homeless is their own fault and I should*

be allowed to shoot them with a gun’ part that turns people off.”

- Associating the poor with addiction: *“not everyone is homeless due to addiction”; “the guy uses that hateful term ‘open air drug markets’ to refer to homeless encampments”; “It’s wrong to assume people become homeless due to drugs, I’ve never touched drugs in my life, but was homeless for a year after my divorce.”*
- Associating the poor with mental illness: *“This man is talking about how homeless people should just all be thrown into insane asylums.”*
- Associating the poor with crime, which can be divided into:
 - Overestimating the correlation between poverty and crime: *“So a theft automatically equates to a homeless person?”*
 - The criminalization of poverty: *“if it were up to you, you’d fill up every jail with the homeless”; “For them, the law is for the perpetual enslavement of the poor”; “The reality is that American law and order is brutal on the poor.”*
- Associating poverty with bad hygiene: *“If a homeless person shits on the street because they can’t access a bathroom, it’s an excuse to evict and attempt to exterminate”; “Unless you want people to think the poor are dirty and smelly...which they are not!”*

Avoidance and fear: includes, but is not limited to the following topics:

- Exclusion, detachment and ostracizing: *“The obsession with trying to “improve” the lives of the disadvantaged, without those being discussed having a say.”; “Humane society takes dog from homeless person arguing it deserves a better life and leaves person on street”; “No one left behind” Except the poor, the disabled and the asylum seekers”*
- Fear of poor / homeless people: *“Nah. It’s just political opportunism. ABC has staked their entire campaign on fear of poor people”; “Hating homeless addicts is literally based on a childish fear of them”; “There is no crime or homeless problem. There is only fear-mongering.”; “If you find homeless*

people 'scary and uncomfortable', give them homes."

Discrimination: includes, but is not limited to the following topics:

- Bullying: *"Blaming the poor is a great tactic to get people on your side."; "They want to grind homeless addicts under the boot or ship them somewhere else"*
- Over-policing and criminalization: *"All he does is harass homeless people and the local chapter/other leftist orgs"; "it is a horribly broken system of police training environment that prioritizes policing the poor, leading to brutality and overreactions"*
- Law / regulation enforcement: *"Those kinds of people never do time. Prison is for poor people."; "It is more profitable to try, convict, and incarcerate poor people"; "Too many poor people in jail for minor offenses and it has nothing to do with community safety."*
- Military / war service: *"A lot of poor people joined the military in the day as Black and White ghetto kids often had poor nutrition at home"; "They are actually admitting that the military depends on recruiting poor people."*

Physical attack: includes examples such as *"Police charged for pouring oil on the homeless"; "Two arrested for beating and looting a homeless old man around midnight 2 days back."; "That is a clever little plan to pay young people to beat up homeless people"; "Maybe concentrate on stopping the homeless being abused on the streets."*

Bias Aggravation

It is important to highlight that aporophobia can act as an aggravator of other types of discrimination (namely, racism, xenophobia, and sexism) and in many instances these different types of discrimination appear intertwined. For example, a recurrent argument that appears on social media is the "need to take care of our homeless first" as an argument to reject migrants and ethnic minority groups. We consider the following examples as 'Direct Aporophobia' because, in this case, poverty is part of the argument to reject ethnic minority groups and migrants (i.e., if these population groups were rich, they would not be rejected with the argument that they are competing for resources with the local

population). In other words, when the minority group is poor, this contributes to racism or xenophobia. Some examples are as follows: *"We don't want them! Can't help our homeless, but sure let's put a roof over migrants heads. No thanks, send them back where they came from or dump them elsewhere."; "They spend millions to house, feed and educate immigrants all across America in ways which they do not do for the Homeless Americans in their own cities"; "BLACK people are more likely to be homeless and need financial support."*