

Hollywood Identity Bias Dataset: A Context Oriented Bias Analysis of Movie Dialogues

Sandhya Singh*, Prapti Roy*, Nihar Sahoo*, Niteesh Mallela*, Himanshu Gupta*,
Pushpak Bhattacharyya*, Milind Savagaonkar[†], Nidhi[†], Roshni Ramnani[†],
Anutosh Maitra[†], Shubhashis Sengupta[†]

*CFILT, Indian Institute of Technology Bombay, India

[†]Accenture Labs, India

{sandhyasingh, nihar, niteesh, pb}@cse.iitb.ac.in,

gupta_himanshu@iitb.ac.in, praptir1994@gmail.com, nidhidr18@gmail.com

{milind.savagaonkar, roshni.r.ramnani, anutosh.maitra, shubhashis.sengupta}@accenture.com,

Abstract

Warning: This paper contains content that may be offensive or upsetting however this cannot be avoided owing to the nature of the work.

Movies reflect society and also hold power to transform opinions. Social biases and stereotypes present in movies can cause extensive damage due to their reach. These biases are not always found to be the need of storyline but can creep in as the author’s bias. Movie production houses would prefer to ascertain that the bias present in a script is the story’s demand. Today, when deep learning models can give human-level accuracy in multiple tasks, having an AI solution to identify the biases present in the script at the writing stage can help them avoid the inconvenience of stalled release, lawsuits, etc. Since AI solutions are data intensive and there exists no domain specific data to address the problem of biases in scripts, we introduce a new dataset of movie scripts that are annotated for identity bias. The dataset contains dialogue turns annotated for (i) bias labels for seven categories, viz., *gender, race/ethnicity, religion, age, occupation, LGBTQ, and other, which contains biases like body shaming, personality bias*, etc. (ii) labels for sensitivity, stereotype, sentiment, emotion, emotion intensity, (iii) all labels annotated with context awareness, (iv) target groups and reason for bias labels and (v) expert-driven group-validation process for high quality annotations. We also report various baseline performances for bias identification and category detection on our dataset.

Keywords: Movie Bias, Social Bias, Hollywood Movies

1. Introduction

Movie is an audio-visual form of art dealing with the complex human psyche. It nourishes and nurtures the social ethos and delves deep into the reality of society. Stemming from the social fabric, they have the potential to influence a predominantly large section of public opinion. Social identity has shaped one’s life and thus, has emerged as vital for survival. Thus, social biases and stereotypes in movies can lead to significant negative impacts. Identifying bias in natural language processing (NLP) tasks have seen a spurt in recent times. Bias brings along the notion of harm to the task in consideration (Bishop and Nasrabadi, 2006; Caliskan et al., 2017). In contemporary times, the presence of biases acts as a major deterrent toward the path of equality. In this work, we define bias as a prejudice in favor or against a person, group, or community directed towards their social identity.

Biases in movies emerge not only as the demand of storyline but also, can be infused through screenwriter’s own bias. The movie production houses prefer to as-

certain that bias present in a script is the story demand as the content might lead to controversy and cause inconvenience and losses. To recall, movies like *The Birth of a Nation*, *The Last Temptation of Christ*¹ have faced controversies due to their content. Hence, the production houses are cautious about filtering out the harmful utterances at the initial scripting level. The scripts go through multiple iterations for content verification before production. This process is highly human-intensive, involving human judgment and efforts to moderate it.

With Deep Learning (DL) models approaching human-level accuracy in various tasks, an AI-supported solution to identify the biases present in the script at the writing stage is the need of the hour. This can speed up the entire scripting process and ease human effort. Data is prime to DL models, but there is no known available dataset for identifying the biases in the movie script domain.

Most of the recent works have focused on detection of

* equal contribution

¹https://en.wikipedia.org/wiki/List_of_banned_films

harmful biases in Hate Speech, offensive texts, and social media posts (Kiritchenko and Mohammad, 2018; Díaz et al., 2018; Mathew et al., 2020). However, bias detection in movie domain is less explored apart from few research on gender asymmetry and stereotype (Sap et al., 2017; Kagan et al., 2020; Garcia et al., 2014; Xu et al., 2019). Some works are also available based on linguistic features to identify gender, race, and age differences in movies and fiction (Fast et al., 2016; Ramakrishna et al., 2015; Ramakrishna et al., 2017). No attention has been given for the detection of biases and stereotypes at the dialogue level for the movie domain. Also, developing a dataset on bias detection for movie dialogues is a time-consuming and complicated. It requires both social and computational expertise. This motivated us to create a dataset to address the problem of social bias identification in movie dialogues.

We introduce a new dataset as Hollywood Identity Bias Dataset (HIBD) consisting of 35 movie scripts annotated for multiple identity biases. The annotation is done at sentence level with a total of 49117 sentences, of which 1181 biased sentences were detected, which amounts to approximately 2.5% bias. This highlights the challenge involved in annotating the biases in the movie script domain, *viz.*, skewness of the data.

Besides individual opinions, social biases also emerge from preconceived generic notions of stereotypical perception. During annotation, we found that there are dialogues that are not biased in content but are sensitive in nature. Though researchers have synonymously used bias and stereotypes (Nadeem et al., 2021), we consider them as separate socio-linguistic phenomena. A detailed discussion is given in section 3 of this paper. To address all these aspects related to biases, we are following a layered approach for annotation. **Our contribution** through this work is listed as:

- To the best of our knowledge, we propose the first context aware, dialogue level bias analysis of Hollywood movies.
- We have annotated the scripts for Sensitivity, Stereotype and social bias labels as *Gender, Race/Ethnicity, Religion, Age, Occupation, LGBTQ, and Other, that contains biases like body shaming, personality bias, etc.*
- Each annotated bias is further labeled Implicit or Explicit to convey the nature of bias along with their corresponding target group and the rationale behind it.
- We are annotating sentiment as positive or negative and its associated emotion and intensity based on plutchik’s emotion wheel (Plutchik and Kellerman, 2013).
- Baseline experiments to identify biases and their category detection is done to benchmark our dataset.

The paper is structured as follows: Section 2 discusses the related work for bias identification, datasets related

to various social biases and why they are not suitable for our problem. Section 3 describes our dataset at length in terms of definition and annotation process. Section 4 discusses some of the challenges encountered during annotation and our approach to resolving them with expert intervention. Section 5 presents baseline experiments, results and error analysis followed by conclusion and future work in section 6. The annotated dataset, annotation guidelines and code has been made available².

2. Related Work

Bias in NLP: There has been a lot research on bias detection in various forms of text, starting from word embeddings (Caliskan et al., 2017). Bolukbasi et al. (2016) has shown that embeddings like Glove (Pennington et al., 2014) and Word2vec (Mikolov et al., 2013) exhibit human-like gender bias for various occupations. This was extended by Manzini et al. (2019) to show racial and religious biases. Concurrent works like Kurita et al. (2019) and May et al. (2019) show that contextual embeddings like BERT (Devlin et al., 2019) not only encode multiple social biases but also amplify them. The presence of bias in these models and embeddings poses many harmful risks by unfair allocation of resources against the marginalized section of the society such as LGBTQ (Blodgett et al., 2020).

Hence, the research community, in its endeavor toward EthicalAI, has seen an upsurge in many supervised datasets, methodologies to detect and mitigate various social biases through different extrinsic tasks like hate speech and toxicity detection (Sap et al., 2019; Davidson et al., 2019; Dixon et al., 2018), coreference resolution (Webster et al., 2018), question answering (Li et al., 2020), machine translation (Stanovsky et al., 2019), and so on. However, most of the works focus on gender and race biases. Other biases like religion, LGBTQ, occupation, and body-shaming are still not explored much.

The recently released StereoSet dataset (Nadeem et al., 2021) was crowd-sourced to measure stereotypical gender, profession, race, and religion biases. Another crowd-sourced dataset called CrowS-Pairs (Nangia et al., 2020) was published as a challenge set for measuring the degree to which nine types of social bias are present in language models. It has data related to nine bias types: *race, gender, sexual orientation, religion, age, nationality, disability, physical appearance and socioeconomic status*. While these datasets consider multiple dimensions of social biases, a recent work (Blodgett et al., 2021) has reported many pitfalls in annotations.

Another work in the direction of defining and evaluating social biases (Sap et al., 2020) creates a structured way of representing the biases and offensiveness implied in language. The frames combine cate-

²https://github.com/sahoonihar/HIBD_LREC_2022

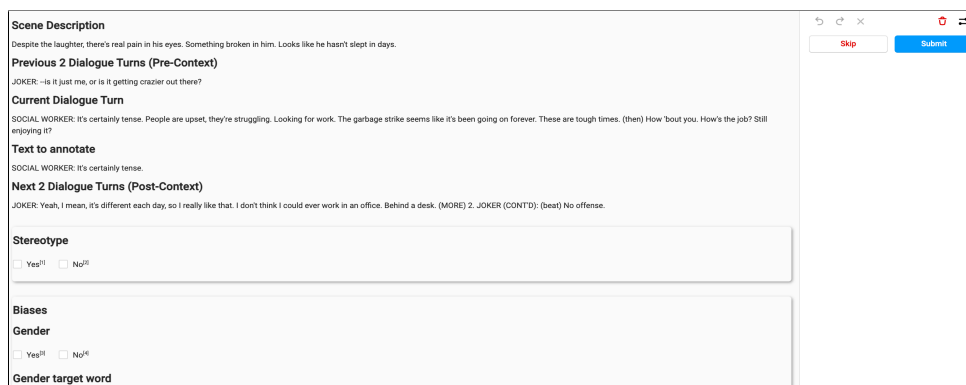


Figure 1: A Screenshot of the annotation interface showing a Dialogue Turn for annotation

gorical knowledge about the offensiveness, intent, and targets of statements, along with free-text reasoning about which groups are targeted and biased implications or stereotypes. They have annotated the dataset for seven bias types: gender/sexuality, race/ethnicity, religion/culture, social/political, disability, body/age, and victims. Bias is a complex phenomenon that requires context information and speaker information to be understood by a human being. However, most previous works have not used any social context to capture these biases. A relatively new work (Vidgen et al., 2021) has annotated English Reddit entries for Identity-directed, Person-directed, Affiliation-directed, CounterSpeech, Non-hateful Slurs, and Neutral categories for abusive text considering the context of the conversation thread. The rationales for annotation are marked in the form of span highlighted in the text.

Bias in Movies Scripts: Though there have been numerous works towards detecting social biases in texts, much focus has not been given to detecting them in the entertainment domain. Different studies (Sap et al., 2017; Kagan et al., 2020; Garcia et al., 2014; Xu et al., 2019) show the gender asymmetry and stereotypes in social media and movies. Fast et al. (2016), Ramakrishna et al. (2015) and Ramakrishna et al. (2017) rely on linguistic features to quantify gender, race and age differences in movies and fictions. Unlike other previous works, Khadilkar et al. (2021) has made use of Cloze tests and WEAT (Caliskan et al., 2017) metrics to measure gender bias along with other subtle biases and compare these biases between two major movie industries of Hollywood and Bollywood.

But, there has been very little work on identifying biases other than gender in movie scripts. Also, no prior work focuses on the detection of biases and stereotypes at dialogues level for any film industry. We need to make a conscious effort towards automatic detection and classification of biased dialogue turns in movies as their presence may have adverse impact on viewers (Dimnik and Felton, 2006; Cape, 2003).

This motivated us to develop a data resource that can address the problem of identifying context aware social bias at dialogue level in movie scripts.

3. Hollywood Identity Bias Dataset

This dataset is developed primarily to address the problem of identity related social biases and stereotype identification from the entertainment domain. Our objective was to develop a comprehensive labeled dataset for multitasking that can comply with inference, categorization, and classification tasks related to the entertainment domain.

3.1. Terminology

In previous works, biases and stereotypes have often been used interchangeably, thus, lacking clarity. To enable the models to interpret the biased constructs in scripts, a pragmatic formalism was defined to help us differentiate between these related yet different terms. The terms used in our annotation process are defined as:

- **Sensitivity:** The *property* of a statement targeted towards an individual or a group belonging to a section that is either overrepresented, underrepresented or vulnerable due to identity such as *race, religion, occupation, sexual identity etc.* A sensitive utterance can be both implicit and explicit in nature and is loaded with potential tension which might lead to aggression. The notion of sensitivity is deeply rooted in the power position or people power and always bear a negative sentiment. The definition of sensitivity encompasses an array of terminologies such as hate, offensive and abusive text provided an identity is targeted. For example, *"The Jews are again using holohoax as an escape to spread their agenda."* is a sensitive statement against Jews.
- **Stereotype:** It is an *overgeneralized belief* about a particular community. It guides our thought process by prompting us to ignore facts and thus cause biases. When a fact (existential quantifier) is replaced with an overgeneralized statement (universal quantifier), it becomes a stereotype. For example, *"Some white men can not dance"* is a fact statement; but *"White men can not dance"* is a stereotype.

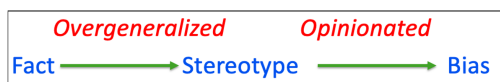


Figure 2: Relation between fact, stereotype and bias statement.

- **Bias:** It is a prejudice *towards* or *against* an individual or community. It exists in both conscious and unconscious forms and influences decision making. As given in figure 1, when an overgeneralized idea gets opinionated, it becomes a bias. However, biases do not always originate from stereotypes, and they can be an opinionated statement targeted towards an identity. We define bias as a quintuple $\langle S, L, T, C, R \rangle$ where,

- *S* is the communicator (speaker, author)
- *L* is the communicatee (audience, reader)
- *T* is the target of the bias
- *C* is the identity category of bias
- *R* is the reason for bias

For example, in the following dialogue set:

KIRSTY: *She's so damn...English.*
STEVE: *Meaning what?*

Kirsty is the communicator, *Steve* is the communicatee, Target is *English*, Bias Category is *race* and the reason for bias is that "*British people are believed to be overly controlled community*"

3.2. Structure of a Movie Script

The script or screenplay is a written work for an audio-visual medium, *viz.*, film, television series. It shows the movement, the action, the dialogue of characters and build the storyline along with the plot. A standard structure is followed for the script writing to distinguish between characters, action lines, and dialogues. Figure 3 shows a snapshot of a script structure for reference. The elements in a screenplay are:

- **Scene Heading:** It shows any change in location or time of the scene. It always begins with EXT. (Exterior), INT. (Interior), INT./EXT. (used for scenes inside vehicles). It is always written in uppercase.
- **Action:** Description of the scene
- **Character:** Name of the character in uppercase. It is followed by the dialogue in the new line.
- **Parenthetical:** A short description placed between character and dialogue to inform the reader of the character's intent. It is always in parenthesis.
- **Dialogue:** To express a character's point-of-view or to understand the interpersonal dynamics between two or more characters. It is always preceded by the character's name.

Categories	Targets
Ageism	Kids, Teenagers, Young, Middle aged, Old
Gender	Female, Male
Race/Ethnicity	Black, White, African, English etc.
Religion	Christian, Judaism, Aethist, Islam etc.
LGBTQ	Lesbian, Gay, Bisexual etc.
Occupation	Defence, Hospitality etc.
Other (Body Shaming, Personality etc.)	Female, Male

Table 1: Types of bias categories and a sample of their targets.

- **Transition:** Used to indicate a passage from one scene to another. It is also in uppercase e.g. CUT TO, FADE IN BLACK.

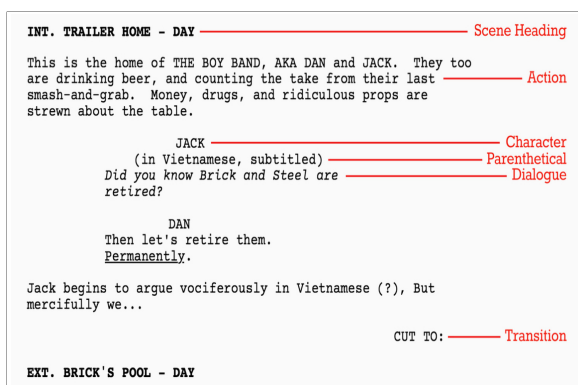


Figure 3: A screenshot of movie script structure ³

3.3. About Annotation

The text scripts for annotation were taken from Internet Movie Script Database (IMSDb)⁴. The dataset consists of 35 Hollywood movies annotated for bias. The data statistics are given in table 4.

3.3.1. Annotators

Owing to the complexity of the task, it was decided to work with two specialized annotators with expertise in American history, culture and politics instead of opting for crowdsourcing. Since the annotator's background and life experiences affect the decision making for annotation, special attention was given to minimize the annotator's bias. To make the annotation process more objective, a detailed guideline and a questionnaire set have been curated to assist the annotators (Refer to appendix for a snapshot of guideline and questionnaire). The efficacy of the supporting annotation material was established by asking a new set of annotators to annotate a sample of dataset. The Inter Annotator Agreement (IAA) between two expert annotators was computed to establish the impact using kappa score. Table 2 shows the cohen kappa score (Cohen, 1960) across all labels of annotation. An average kappa score of 0.64

³<https://slugline.co/basics>

⁴<https://imsdb.com/>

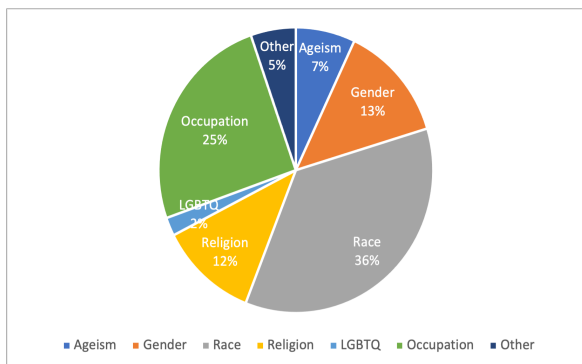


Figure 4: Distribution of social biases across 7 categories. We show percentages of each category annotated in the dataset.

for all possible categories shows a good agreement between the two annotators. However, the agreement for stereotype was low, possibly because of lack of historical knowledge about some communities present in the dataset.

3.3.2. Annotation Process

As bias is an inherently complex phenomenon, its analysis, labeling, and subsequent category detection are challenging. The standardised annotation process described here has been modified multiple times in consultation with annotators since the beginning of the task. The granularity of annotation, terminology definition, and bias categories have gone through multiple iterations of revision.

We followed a context-aware layered annotation approach to annotate for multiple tags. The annotation is done at Dialogue Turn (DT) level. Figure 1 shows a snippet of the interface used for annotation. Each DT contains the following elements:

- **Scene Description:** Describes the scene for the given dialogue for annotation.
- **Pre-context:** Contains previous 2 dialogues.
- **Dialogue:** Current dialogue to be annotated.
- **Text to annotate:** Dialogue sentence for annotation.
- **Post-context:** Contains following 2 dialogues.

For each DT, the labelling was done for 'Text to annotate'. Each layer labeled is for:

- **Layer 1:** *Sensitivity*. Labeled as *Yes or No*.
- **Layer 2:** *Stereotype*. Labeled as *Yes or No*.
- **Layer 3:** *7 category of Bias*. Labeled as *Yes or No*, the target and reason for bias. A sample of target groups are shown in table 1.
- **Layer 4:** *Sentiment*. Labeled as *Positive or Negative*.
- **Layer 5:** *Emotion*. Labeled for 8 emotion types based on the Plutchik's Wheel of Emotions.
- **Layer 6 :** *Emotion Intensity*. Labeled as *Low, Medium, High*.

Labels	Kappa Scores
Ageism	0.72
Gender	0.54
Race/Ethnicity	0.61
Religion	0.67
LGBTQ	1
Occupation	0.47
Other	0.49
AVERAGE (all categories)	0.64
Stereotype	0.44
Sensitivity	0.33
Bias	0.71

Table 2: Inter-annotator agreement scores using Cohen Kappa for Annotated Labels

While annotating, it was found that bias is not always present in an explicit manner. Many times the bias is present based on the identity of the speaker. Thus, the speaker's identity brings in a lot of subtle biases. These constitute implicit biases which are more challenging to identify as it requires identity information about the speaker to be known to annotators. Hence, besides script annotation, metadata about the movie characters are also collected.

Action item: *He giggles at his own joke.*
Pre-context: *ITALIAN BOY: This kid lives around here, but he can't say bread in Italian. CLEMENZA: That's 'cause he's Jew. Look at those pregnant lips!*
Dialogue: *ITALIAN BOY: Are you a Jewboy?*
Text to annotate: *ITALIAN BOY: Are you a Jewboy?*
Post-context: *ITALIAN BOY: Well, if you're not a Jew, say 'bread' in Italian. See, he can't. CLEMENZA: Alright, alright, cut it out.*

Example 1: *Bias- Religion; Target- Jew; Reason- Jewboy is a highly derogatory phrase for boys hailing from Jewish community; Sentiment- Negative; Emotion- Disgust*

Metadata: A two-fold metadata is maintained for each movie script to analyse the character graph, and the existence of implicit and explicit biases. In **metadata 1**, we collect the data about the movie such as genre, year of release, director, screenplay writer. Information about the author/director is also collected on the scripts based on book or series. In **metadata 2**, we collate information regarding the main characters in the movie such as their age, occupation, gender, religious, racial identity and their role in the movie.

To ascertain the quality of dataset, after annotation, an expert-driven group consisting of 3 members have validated every tag. The inconsistencies in the data were resolved through adjudication. A sample annotation from the dataset is given in Example 1.

4. Challenges in Annotation

The annotation process posed some challenges for the annotators. This section discusses the challenges experienced by the annotators. Some of the common challenges faced were:

- *Implicit Bias* poses a serious challenge to the

Text to annotate	Labels Annotated	Reason
T'CHALLA: If you were not so stubborn you would make such a great queen.	Bias: Occupation, Target: Administrative Services	Social perception that great queens are flexible in their opinion.
NAKIA: I would make a great queen because I am so stubborn .	Bias: Occupation, Target: Administrative Services	An individual opinion that great queens can be rigid in their worldview.
ROTH: I want my own doctor; fly him in from Miami. I don't trust a doctor who can't speak English	Bias: Race, Target: English	Associating a negative attribute on the basis of language
MADISON: Oh, well first off, I have a huge rack.	Bias: Gender, Target: Female	The speaker is objectifying and equating herself to a sexual commodity.

Table 3: Challenging cases from the dataset. Here we have mentioned only the explicit bias examples. Detailed discussions are present in Section 4.

cause of annotating data as the essence of the underlying bias in the particular utterance is drawn from the entire dialogue turn, character metadata or world knowledge. Example 2 shows one such case from dataset.

Action item: *He addresses that to Neri, who resentfully goes to fetch the Senator a glass of water.*
Pre-context: *MICHAEL: I trust these men with my life. They are my right arms; I cannot insult them by sending them away. SENATOR GEARY: (taking out some medication) Some water.*
Dialogue: *SENATOR GEARY: Alright, Corleone. I'm going to be very frank with you. Maybe more frank than any man in my position has ever spoken to you before.*
Text to annotate: *Maybe more frank than any man in my position has ever spoken to you before.*
Post-context: *SENATOR GEARY: The Corleone family controls two major hotels in Vegas; In another week or so you'll move Klingman out, which leaves you with only one technicality. MICHAEL: Turnbull is a good man.*

Example 2: Bias- Occupation; **Target-** Anti-social; **Reason-** Politicians are assumed to be diplomatic.

Here, the dialogue appears as a statement unless we know that the speaker is a *Senator*, a politician by profession that requires diplomatic skill. Hence, this DT is annotated as an Occupation Bias. The bias is coming from world knowledge about a profession.

- Friends *addressing* each other *using racial terms* like nigga, black etc. in a friendly way.
- There are several instances of *scene level bias* which does not overlap in the dialogues. In such cases the strategy adopted is not to label the bias inherent in the scene as it would not only complicate the process of annotation but also, mislead the model while training process.
- While tagging *sarcastic comments* (if biased), the sentiment tagged is always negative as incongruity constitutes the key in cases pertaining to sarcasm.
- Movie scripts come with inherent biases relating to *gendered roles* and *sexist comments*. In order

to standardize the process of labelling self objectification, where the speaker is highlighting himself/herself as an object of sexual desire is to be tagged in positive sentiment.

- Sometimes, world knowledge is required to identify bias. Example 3 highlights one such case where the statement is referring to Barack Obama as US President.

Action item: *A large pot is placed atop the gas stove. Everyone eats their gumbo in the kitchen. Laughing talking. Wanda blows out the candles on her cake, cuts it, Hands out everyone a slice. Grandma Bonnie sits in a recliner at the foot of the coffee table.*
Pre-context: *GRANDMA BONNIE: If somebody told me say twenty years ago, that I would live to see it happen, I'd have laughed in their face. OSCAR: If somebody would have told me that two years ago I would have laughed in their face.*
Dialogue: *DARYL: And here we are. 2008 and its happened.*
Text to annotate: *2008 and its happened.*
Post-context: *CEEPHUS: But I think it's a sign of the times. The country is in bad enough shape for them to hand it over to a black guy. Look, if Bush hadn't have messed the country up so bad. Would he have even had a shot?*

Example 3: Bias- Race; **Target-** Black-Americans; **Reason-** referring to Obama as president of USA.

Table 3 presents some unique challenges faced while labelling. As seen, row 1 and 2 contradict each other but both are taken as Occupation Bias since they are opinionated statements related to the qualities of a Queen. In row 3, the dialogue shows bias related to language. Here, the bias is both implicit against Spanish and explicit towards the English language. In row 4, the speaker is self objectifying based on gender-specific traits so it is labeled as Gender Bias. These few examples reiterate the fact that bias identification is a complex task in itself.

5. Experimental Setup

In this section we discuss about different models for detection of social biases and their categories.

Labels	Sentence Level	Dialogue level
Bias	1181 (2.40%)	976 (3.42%)
Neutral	47936	27558
Total	49117	28534

Table 4: Distribution of Biased sentences and dialogues.

Labels	Train	Dev.	Test
Age Bias	49	22	18
Gender Bias	87	31	56
Religion Bias	62	47	34
Occupation Bias	141	68	105
LGBTQ Bias	11	4	14
Race Bias	257	89	97
Other Bias	44	9	15

Table 5: Distribution of different bias categories across different splits.

5.1. Data Splits

As we discussed in previous sections (refer, 3.3.2), our annotation is done at sentence level, i.e., *Text to annotate*. However, due to the low prevalence of bias at the sentence level (table 4), we process the dataset at the dialogue level. A dialogue (speaker: utterance) is considered biased if at least one of the sentences in it has been annotated for any of the 7 bias types. We use the dialogues as input to the model for our various classification tasks. We split the movies based on the count of biased dialogues in them. There were a total of 976 biased dialogues across the span of 35 movies, and the splits are training (488 dialogues; 50%), development (213 dialogues; $\sim 22\%$), and test (275 dialogues; $\sim 25\%$). The distribution of different splits across all the categories is mentioned in the table 5. We ensured that all the dialogues of the same movie were assigned into the same split to avoid dataset leakage (Elangovan et al., 2021). By this process, we have 10 movies in the train split and 8, 17 movies respectively in development and test split.

5.2. Methodology

Due to the shallow presence of biased instances in the data, we use *inductive transfer learning* (Ruder et al., 2019) for our experiments. We use the idea of sequential adaptation, where a model is fine-tuned on a large dataset of related domains before doing the fine-tuning on the target task. Bias detection is done alone in a binary classification framework. As our dataset has been annotated with multiple labels for bias categories, we formulate the category detection task as a multi-label classification problem. For baseline models, we use logistic regression to compare the performances.

Logistic Regression (LR): We make use of logistic regression with L2 regularization (Nigam, 1999) for both binary classification (bias vs. neutral) and multi-label classification (the bias categories). For multi-label clas-

sification, we encode multiple labels per instance at once. We use class weight inversely proportional to the class frequencies for both binary and multi-label training.

Sequential Adaptation (SA): Because of the heavy skewness of the data, learning directly from it may not give good performances. Hence, we train a model in a sequential adaptive manner (Phang et al., 2019). First, we fine-tune the model on a curated dataset of a few related tasks before fine-tuning again on our dataset. We extract *stereotype, unrelated* instances from (Nadeem et al., 2021) for four categories (gender, race, religion, occupation); group targeted posts from (Sap et al., 2020) for race, gender age categories; comment texts with *identity_attack* > 0.5 from (jig, 2019) where target belongs to one of the categories among race, gender, religion, LGBTQ. We carefully collate $16k$ data instances from these three sources with $6k$ instances for the bias class. We use this dataset, as initial training, for binary and multi-label classification. We fine-tune a BART-large (Lewis et al., 2019) model for both tasks. While fine-tuning later on our dataset, we do a controlled experiment where the number of instances for the neutral class is equal to that of the bias class for binary classification. We report standard deviation across 5 runs for the controlled experiment as we randomly sample the instances for neutral class. We do not include *other class* in the category detection training as there are no publicly available data for all the subcategories (refer table 1).

5.3. Results

Performance Metrics: We report precision (P), Recall (R), F1 score for both binary and multi-label classification in Table 7 and 6 respectively. Due to the heavy skewness of the data, we rely on the macro values for the classification of bias vs. neutral dialogues.

Model Comparison: The BART-large (SA) model substantially outperforms logistic regression for bias category detection task (macro F1 score of 0.72 compared to 0.31 for LR). However, for bias detection, BART-large shows marginal performance improvement over LR which can be due to the skewed nature of the data. This shows that the category detection task is easier than bias detection.

5.4. Error Analysis

We have observed that the category detection model sometimes predicts some extra categories for the dialogue which are not available in the ground truth label. Consider the below dialogue, which has a clear bias towards church (religion), but the model also predicted one extra label as gender, understandably due to the phrase "good woman".

<p>Sentence : PHILLIPA: So Stephen, you've been to church with a good woman. Are you feeling holier than thou?</p> <p>Ground Truth : Religion</p> <p>Prediction : Gender, Religion</p>

	LR			BART-large (SA)		
	P	R	F1	P	R	F1
Race/Ethnicity bias	0.500	0.410	0.450	0.77	0.89	0.83
Religion bias	0.226	0.259	0.241	0.86	0.67	0.75
Gender bias	0.302	0.432	0.355	0.73	0.73	0.73
Occupation bias	0.321	0.464	0.380	0.59	0.48	0.53
Ageism bias	0.171	0.462	0.250	0.62	0.62	0.62
LGBTQ bias	0.158	0.273	0.200	1.00	0.73	0.84

Table 6: Performance of each category on test set. Best F1 scores are in bold.

Models	P	R	F1
LR	0.53	0.71	0.51
LR-contrl	0.52±0.008	0.70±0.011	0.49±0.021
SA	0.55	0.81	0.58
SA-contrl	0.55±0.007	0.80±0.012	0.57±0.017

Table 7: Performance of binary classification[Bias vs. Neutral], on test set, with best numbers in bold. Standard deviations are reported for controlled experiments.

Occupation bias in the category classification, mostly, has been misclassified as gender bias. The bias classification model, sometimes, assigns the wrong label for neutral dialogues due to the presence of many identity related words/phrases in it. In the below mentioned example, there is no (implicit/explicit) bias towards any identity, but model misclassifies because of the presence of words related to an identity.

<p><i>Sentence : ORDELL (CONT'D): Oh, that's just my white friend, Louis. He's got nothing to do with my business. We just hangin together. We're on our way to a cocktail lounge.</i></p> <p><i>Ground Truth : Neutral</i></p> <p><i>Prediction : Bias</i></p>

5.5. Challenges

Understanding the conversation flow (who say what to whom) in a movie script is a complex phenomenon due to the very high number of dialogue exchanges among multiple characters. Detecting bias at the dialogue level requires the model to understand this flow and also requires understanding the relations among multiple speakers. Sometimes an utterance is biased because of the speaker's demographics. It's quite challenging for state-of-the-art models to capture these nuances correctly. More often, understanding the bias in dialogue requires the other dialogue social context. Providing these external social contexts to AI based models is not a very trivial task. Another major challenge with movie scripts is that the speaker names are not always the named entities (e.g. DEVIL GIRL, GAURD, CLERK). The presence of these identity related words as speaker makes the task more challenging.

6. Ethical Consideration

Our work aims at capturing various social biases in Hollywood movie scripts. We discuss the various challenges we faced in the annotation process, specifically due to the presence of implicit biases in dialogues. The study of social biases comes with ethical concerns of risks in deployment (Ullmann and Tomalin, 2020). As these biased dialogues can cause potential harm to any viewer or community, it is required to conduct this kind of research to detect them before release. However, we do not intend to harm the reputation of movie production houses and any other person involved with any of the movies in our dataset. Researchers working on the problem of social bias in fiction, movies, or conversational AI would benefit from the dataset we have created.

7. Conclusion and Future Work

We present a comprehensive dataset of 35 Hollywood movies to identify social biases in movie scripts. The dataset is labeled for *sensitivity, stereotype, identity biases as gender, ageism, race/ethnicity, religion, occupation, LGBTQ, other (body shaming, personality, etc.), the target of bias, sentiment, emotion, emotion intensity, and reason for bias*. The dataset has been benchmarked for bias identification and categorization task using the BART-large model.

Since bias is found to exist only in 3% of the dialogues from 35 movies, we have the challenging task of addressing skewness in data. Also, sensitivity labeling was introduced later in the annotation process. We plan to extend sensitivity tags to all the scripts in the dataset. In future, we would like to use the annotated metadata along with dialogue contexts to make the bias detection more robust.

8. Acknowledgements

We thank the anonymous reviewers for their insightful comments. Additionally, we thank Sushma Gawande for her contribution in data annotation. We also thank Tamali Banerjee and Sandeep Singamsetty for questionnaire and guideline validation. This research work has been funded by Accenture Labs, India.

9. References

- Bishop, C. M. and Nasrabadi, N. M. (2006). *Pattern recognition and machine learning*, volume 4. Springer.
- Blodgett, S. L., Barocas, S., Daumé III, H., and Wallach, H. (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, July. Association for Computational Linguistics.
- Blodgett, S. L., Lopez, G., Olteanu, A., Sim, R., and Wallach, H. (2021). Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1004–1015, Online, August. Association for Computational Linguistics.
- Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., and Kalai, A. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings.
- Caliskan, A., Bryson, J. J., and Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, Apr.
- Cape, G. S. (2003). Addiction, stigma and movies. *Acta Psychiatrica Scandinavica*, 107(3):163–169.
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.
- Davidson, T., Bhattacharya, D., and Weber, I. (2019). Racial bias in hate speech and abusive language detection datasets. *CoRR*, abs/1905.12516.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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, June. Association for Computational Linguistics.
- Díaz, M., Johnson, I., Lazar, A., Piper, A. M., and Gergle, D. (2018). Addressing age-related bias in sentiment analysis. In *Proceedings of the 2018 chi conference on human factors in computing systems*, pages 1–14.
- Dimnik, T. and Felton, S. (2006). Accountant stereotypes in movies distributed in north america in the twentieth century. *Accounting, Organizations and Society*, 31(2):129–155.
- Dixon, L., Li, J., Sorensen, J., Thain, N., and Vasserman, L. (2018). Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’18, page 67–73, New York, NY, USA. Association for Computing Machinery.
- Elangovan, A., He, J., and Verspoor, K. (2021). Memorization vs. generalization: Quantifying data leakage in nlp performance evaluation.
- Fast, E., Vachovsky, T., and Bernstein, M. S. (2016). Shirtless and dangerous: Quantifying linguistic signals of gender bias in an online fiction writing community.
- Garcia, D., Weber, I., Rama, V., and Garimella, K. (2014). Gender asymmetries in reality and fiction: The bechdel test of social media. In *International AAAI Conference on Weblogs and Social*.
- (2019). Jigsaw unintended bias in toxicity classification.
- Kagan, D., Chesney, T., and Fire, M. (2020). Using data science to understand the film industry’s gender gap. *Palgrave Communications*, 6(1):92, May.
- Khadilkar, K., KhudaBukhsh, A. R., and Mitchell, T. M. (2021). Gender bias, social bias and representation: 70 years of b^hollywood.
- Kiritchenko, S. and Mohammad, S. M. (2018). Examining gender and race bias in two hundred sentiment analysis systems. *arXiv preprint arXiv:1805.04508*.
- Kurita, K., Vyas, N., Pareek, A., Black, A. W., and Tsvetkov, Y. (2019). Measuring bias in contextualized word representations.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. (2019). Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension.
- Li, T., Khot, T., Khashabi, D., Sabharwal, A., and Srikumar, V. (2020). Uncovering stereotyping biases via underspecified questions. *CoRR*, abs/2010.02428.
- Loshchilov, I. and Hutter, F. (2019). Decoupled weight decay regularization.
- Manzini, T., Yao Chong, L., Black, A. W., and Tsvetkov, Y. (2019). Black is to criminal as caucasian is to police: Detecting and removing multi-class bias in word embeddings. 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 615–621, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Mathew, B., Saha, P., Yimam, S. M., Biemann, C., Goyal, P., and Mukherjee, A. (2020). Hatexplain: A benchmark dataset for explainable hate speech detection. *arXiv preprint arXiv:2012.10289*.
- May, C., Wang, A., Bordia, S., Bowman, S. R., and Rudinger, R. (2019). On measuring social biases in sentence encoders. 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 622–628, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space.
- Nigam, K. (1999). Using maximum entropy for text classification. In *In IJCAI-99 Workshop on Machine Learning for Information Filtering*, pages 61–67.
- Pennington, J., Socher, R., and Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Phang, J., Févry, T., and Bowman, S. R. (2019). Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks.
- Plutchik, R. and Kellerman, H. (2013). *Theories of emotion*, volume 1. Academic Press.
- Ramakrishna, A., Malandrakis, N., Staruk, E., and Narayanan, S. (2015). A quantitative analysis of gender differences in movies using psycholinguistic normatives. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1996–2001, Lisbon, Portugal, September. Association for Computational Linguistics.
- Ramakrishna, A., Martínez, V. R., Malandrakis, N., Singla, K., and Narayanan, S. (2017). Linguistic analysis of differences in portrayal of movie characters. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1669–1678, Vancouver, Canada, July. Association for Computational Linguistics.
- Ruder, S., Peters, M. E., Swayamdipta, S., and Wolf, T. (2019). Transfer learning in natural language processing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials*, pages 15–18, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Sap, M., Prasettio, M. C., Holtzman, A., Rashkin, H., and Choi, Y. (2017). Connotation frames of power and agency in modern films. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2329–2334, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Sap, M., Card, D., Gabriel, S., Choi, Y., and Smith, N. A. (2019). The risk of racial bias in hate speech detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy, July. Association for Computational Linguistics.
- Stanovsky, G., Smith, N. A., and Zettlemoyer, L. (2019). Evaluating gender bias in machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy, July. Association for Computational Linguistics.
- Ullmann, S. and Tomalin, M. (2020). Quarantining online hate speech: technical and ethical perspectives. *Ethics and Information Technology*, 22(1):69–80, Mar.
- Webster, K., Recasens, M., Axelrod, V., and Baldrige, J. (2018). Mind the gap: A balanced corpus of gendered ambiguous pronouns. *Transactions of the Association for Computational Linguistics*, 6(0):605–617.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. (2020). Huggingface’s transformers: State-of-the-art natural language processing.
- Xu, H., Zhang, Z., Wu, L., and Wang, C.-J. (2019). The cinderella complex: Word embeddings reveal gender stereotypes in movies and books. *PLOS ONE*, 14(11):1–18, 11.

10. Language Resource References

- Nadeem, M., Bethke, A., and Reddy, S. (2021). StereoSet: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online, August. Association for Computational Linguistics.
- Nangia, N., Vania, C., Bhalerao, R., and Bowman, S. R. (2020). Crows-pairs: A challenge dataset for measuring social biases in masked language models. *arXiv preprint arXiv:2010.00133*.
- Sap, M., Gabriel, S., Qin, L., Jurafsky, D., Smith, N. A., and Choi, Y. (2020). Social bias frames: Reasoning about social and power implications of language. *ACL*.
- Vidgen, B., Nguyen, D., Margetts, H., Rossini, P., and Tromble, R. (2021). Introducing cad: the contextual abuse dataset. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2289–2303.

A. Appendix

A.1. Annotator Details

Both the annotators were trained and selected through extensive one-on-one discussions. Both of them went through a few days of initial training. They would annotate many examples that would then be validated by our research group and were communicated properly about any incorrect annotations during training. As there are potential adverse side effects of annotating such toxic comments, we used to have regular discussion sessions to ensure they were not excessively

Questionnaire for Occupation: If any question is answered Yes, it is labelled as Occupation Bias.

S.No.	Questions for Identifying Occupation Bias
1	Is the speaker referring to the societal perception of a profession?
2	Is the speaker referring about a profession in a negative/positive manner?
3	is the speaker referring to occupation related gendered roles?
4	If the speaker is comparing people from same profession / different profession based on attributes like wage/skill level/skill set/Position/policy/identity?

Figure 5: A snippet of questionnaire to assist the annotators for objective annotation.

<p>6.1. Sensitivity Annotation (Level-1):</p> <p>The property of a statement directed towards a community that is over/under represented or potentially vulnerable, pertaining to the question of identity such as <i>gender, race, religion, occupation etc.</i> A sensitive statement is loaded with potential tension and can lead to severe upheavals such as agitation, revolt, mass mobilisation, unrest, aggression and rebellion. The notion of sensitivity is deeply grounded on power position and people power. In our annotation strategy, we are marking sensitive utterances and labelling them on a scale of 1 to 3 depending on their intensity. A sensitive statement will always be of negative sentiment in nature.</p>
<p>6.2. Stereotype Annotation (Level-2):</p> <p>A stereotype is an overgeneralized belief about a particular category of people. It guides our thinking about others and prompts us to ignore facts and thus influence biases. Here, we are marking the stereotype as binary, if the given DT is a stereotype or not.</p>
<p>6.3. Bias Annotation (Level-3):</p> <p>Bias is an inclination or prejudice toward or against something or someone. It exists in conscious and unconscious forms and influences decision making. Here we are to annotate inclusion/diversity conscious bias in movie scripts. If the DT has a bias, then identify the target group for bias and give the reason for bias. The focus is on annotating the following 7 categories as binary (yes/no):</p>

Figure 6: A snapshot of guidelines used for annotation.

exposed to the harmful contents. Both the annotators were middle-aged Asian females and with post-graduation degree in sociology. We took help from two other annotators for validation of guideline and the questionnaire. Both of them have post-graduation degree in computer science.

A.2. Training Details

We finetune BART-large models with batch size of 32 for both tasks. For sequential adaptation, we finetune the models for five epochs on the curated dataset before re-finetuning them for two epochs on our dataset. Max token length of 128 is used for both the classifications. We also use a dropout layer in our model. We experiment with learning rates of $\{2e-5, 3e-5, 4e-5, 5e-5\}$. AdamW (Loshchilov and Hutter, 2019) optimizer with epsilon = 1e-08, clipnorm = 1.0, 200 warm up steps are used. Experiments were run with a single GeForce RTX 2080 Ti GPU.

All of our implementations uses Huggingface’s transformer library (Wolf et al., 2020).

A.3. Snapshot of Guideline

Bias being a subjective phenomena, a questionnaire and a detailed annotation guideline was created to assist

the annotators for labelling the biases. These supporting guidelines have evolved based on continuous feedback from the annotators. Figure 5 shows a snapshot of the questionnaire used to check the presence of bias. Figure 6 is a snapshot from the guideline used for annotation. Figure 7 and 8 shows the metadata 1 and 2 structure respectively. These metadata is used to understand implicit biases which can come from characters background. Table 8 shows the list of occupation targets and the list of occupations included under each target. The complete questionnaire and guideline can be accessed at https://github.com/sahoonihar/HIBD_LREC_2022

Movie Name	Genre	Release Year	Director	Screenplay Writer	Based on	Author	Author Background
Austin Powers	Comedy	1999	Jay Roach	Mike Myers	-	-	-
Hitchcock	Drama	2012	Sacha Gervasi	John J. McLaughlin and Stephen Rebello	Alfred Hitchcock and the Making of Psycho	Stephen Rebello	Noted writer, screenwriter and former clinical therapist

Figure 7: A snippet of metadata-1. This contains information about the genre, year of release, director, script writer about movie script is collated.

Movie Name	Character Name	Age	Gender	Race-Ethnicity	Occupation	Religion	Relation
Do the right Thing	Vito	Young	Male	White-italian-American	Hospitality	-	sal's youngest son and Mookie's friend
	Mookie	Young	Male	Black-african-American	Services	-	Main Lead
	Tina	Young	Female	Black-african-American	-	-	Girlfriend of Mookie
	Pino	Young	Male	White-italian-American	Hospitality	-	sal's eldest son
	Buggin Out	Young	Male	Black-african-American	-	-	Mookie's friend
	Sal	Middle aged	Male	White-italian-American	Entrepreneur	-	Mookie's Boss
	Radio Raheem	Young	Male	Black-african-American	-	-	neighbour
	Mother sister	Old	Female	Black-african-American	-	-	neighbour
Da Mayor	Old	Male	Black-african-American	-	-	neighbour	

Figure 8: A snippet of metadata-2. This has information related to movie character is collated, viz., character's gender, age, race, occupation etc.

S.No.	Occupation Target	Occupations Included
1	Healthcare	doctor, nurse, counsellor, etc.
2	Hospitality	hotel manager, chef, bartender, cook, wedding planner etc.
3	Anti-social	Gangsters, thieves, con-artists, drug dealers, etc.
4	Defence	Military officers, Intelligence (CIA, Mossad, RAW, ISI), National Security, etc.
5	Entertainment	actors, dancers, musician, painter, clowns, etc.
6	Politicians	government office bearers
7	Administrative services	Kings, queens, royal blood, presidents, prime-minister, minister, senator, etc.
8	Financial and management services	bankers, traders, managers, etc.
9	Entrepreneurs	self-employed, businessman, landlords, etc.
10	Security	police, watchman, guards, bounty hunters, etc.
11	Academia and Research	scientist, professor, teacher, etc.
12	Sports	players, sports manager, referees, coaches, gladiator, etc.
13	Services	delivery boy, driver, labourer, desk jobs/clerks, prostitutes, negotiator, secretary, etc.
14	Technical Services	Engineering services, mechanics, etc.
15	Fitness/Sports Instructor	Yoga Trainer, Gym Trainer, referees, umpires, etc.
16	Journalist/Press	Media, News Reporters, Columnist, etc.
17	Legal Aid	Lawyers, District Attorney (DA), Judges, Chief Justice, Justice, etc.
18	Religious Services	Pope, Priests, Nun, Father, etc.

Table 8: List of Occupation Bias targets and the occupations included under it.