

# Assessing the Reliability of LLMs Annotations in the Context of Demographic Bias and Model Explanation

Hadi Mohammadi<sup>1\*†</sup> Tina Shahedi<sup>1\*</sup> Pablo Mosteiro<sup>1</sup>  
Massimo Poesio<sup>2,3</sup> Ayoub Bagheri<sup>1</sup> Anastasia Giachanou<sup>1</sup>

<sup>1</sup>Department of Methodology and Statistics, Utrecht University, The Netherlands

<sup>2</sup>Department of Information and Computing Sciences, Utrecht University, The Netherlands

<sup>3</sup>Queen Mary University of London, London, United Kingdom

⚠ The paper contains examples which are offensive in nature.

## Abstract

Understanding the sources of variability in annotations is crucial for developing fair NLP systems, especially for tasks like sexism detection where demographic bias is a concern. This study investigates the extent to which annotator demographic features influence labeling decisions compared to text content. Using a Generalized Linear Mixed Model, we quantify this influence, finding that while statistically present, demographic factors account for a minor fraction (8%) of the observed variance, with tweet content being the dominant factor. We then assess the reliability of Generative AI (GenAI) models as annotators, specifically evaluating if guiding them with demographic personas improves alignment with human judgments. Our results indicate that simplistic persona prompting often fails to enhance, and sometimes degrades, performance compared to baseline models. Furthermore, explainable AI (XAI) techniques reveal that model predictions rely heavily on content-specific tokens related to sexism, rather than correlates of demographic characteristics. We argue that focusing on content-driven explanations and robust annotation protocols offers a more reliable path towards fairness than potentially persona simulation.

## 1 Introduction

Reliable annotations are foundational to machine learning in NLP, guiding models toward accurate predictions. According to Uma et al. (2020), annotation involves humans labeling or transforming data inputs into "gold data", which guides machine learning practitioners in building their models. For instance, to create a gold dataset for a model that corrects grammatical errors, annotators might be asked to identify mistakes in a range of sample sentences. However, creating high-quality annotations

\*These authors contributed equally to this work.

†Corresponding author: h.mohammadi@uu.nl

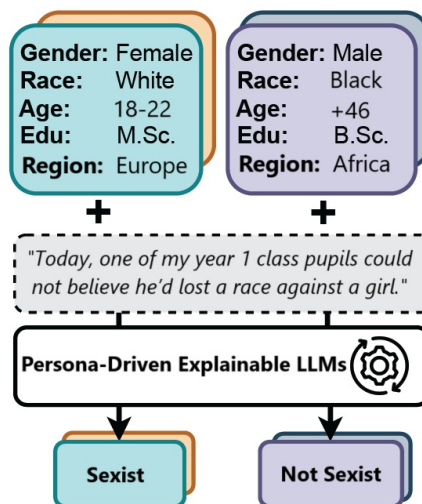


Figure 1: We instruct LLMs to replicate human annotations for subjective NLP tasks from different perspectives using persona prompting and XAI techniques. Our results show that simulated personas alone may not sufficiently capture human subjectivity. XAI analysis confirms that tweet content plays a more significant role in model decisions.

is not a straightforward task since it requires thoughtful consideration of the criteria that make annotations effective, consistent, and unbiased.

However, creating high-quality annotations is not a straightforward task since it requires thoughtful consideration of the criteria that make annotations effective, consistent, and unbiased. This raises the following question, what defines a robust annotation process? When it comes to evaluating annotation quality, several studies highlight Inter-Annotator Agreement (IAA), as defined by Krippendorff (2011), as a standard metric for labeled datasets (Pei and Jurgens, 2023; Plank et al., 2014). However, achieving high IAA is often challenging, particularly for subjective language tasks that rely on human judgments. For tasks like sexism detection, where subjectivity is inherent, addressing annotator agreement challenges is essential, as disagreements can significantly influence the performance of NLP models trained on this

data. In some cases, disagreement often arises from ambiguous sentences or vague label definitions, which can make it difficult for annotators to reach an agreement (Russell et al., 2008; Artstein and Poesio, 2008). Traditionally, aggregating judgments from multiple annotators to create a single "ground truth" for each data instance is widely used to address the inherent ambiguity and subjectivity in language interpretation. This approach is similar to initial methods for handling annotator disagreement, which focuses on estimating a "true" label. However, Pavlick and Kwiatkowski (2019) shows that even when annotators are provided with additional context, there is not always a single correct answer, and disagreements still persist.

Recent studies indicate a significant shift in how annotator disagreements are handled, particularly in subjective tasks involving human judgments (Pavlick and Kwiatkowski, 2019; Basile et al., 2021; Uma et al., 2021; Plank, 2022). Current research primarily focuses on developing models that can learn from these disagreements. While NLP researchers aim for consistency among annotators, some level of disagreement is both inherent and unavoidable in human annotation processes (Leonardelli et al., 2021). As Bless and Fiedler (2014) showed, annotators' demographic factors, personal perspectives, and differing value systems can lead to discrepancies in annotations.

Building on this foundation, researchers have systematically analyzed how the characteristics of annotators and the way tasks are framed can skew evaluation outcomes. For instance, Hosking et al. demonstrate that annotator assertiveness and the linguistic complexity of model outputs significantly bias judgments of factuality and consistency in crowdsourced error annotations. Their study finds that responses that sound more confident are judged as more accurate, even if they contain the same number of errors. Similarly, Kirk et al. (2024) reveal that factors such as cultural background, age, gender, and personal values lead to substantial variation in how responses are rated for helpfulness, creativity, and alignment with individual beliefs. These findings underscore the challenge of distinguishing true model performance from annotator-induced biases and motivate the need for more scalable and consistent annotation methods under controlled conditions.

This sparked researchers to explore the potential of GenAI models as substitutes for human annotators. Several studies have shown that large

language models (LLMs), when provided with demographic information, can imitate specific annotator groups by tailoring their outputs to reflect attributes such as gender, race, age, or education (Beck et al., 2024; Schäfer et al., 2024). However, LLMs often align more closely with certain demographics (e.g., younger, White, male) unless explicitly directed otherwise (Schäfer et al., 2024). To the best of our knowledge, the incorporation of XAI techniques to guide these models is still rare (Ralevski et al., 2024; He et al., 2024; Freedman et al., 2024). This creates critical gaps in evaluating how demographic biases impact annotation reliability and whether GenAI models, with XAI guidance, can effectively substitute human annotators, especially in subjective tasks such as sexism detection. For instance, Mohammadi et al. (2024) presents an explainability-enhanced sexism detection pipeline that bridges model predictions with token-level explanations, illustrating efforts to improve transparency in sexism detection.

In this study, we use data from the EXIST 2024 challenge (Plaza et al., 2024), a shared task on sexism detection in social networks—a highly subjective task.<sup>1</sup> Our primary goal is to assess annotation reliability and examine how demographic biases influence annotator decisions. Using a Generalized Linear Mixed Model (GLMM), we analyze both fixed and random effects, revealing that demographic variables account for nearly 8% of the variance in labeling behavior, suggesting the presence of demographic biases in human judgments. We also evaluate LLM performance by simulating annotation and classification under various prompting scenarios, model configurations, and temperature settings. Our methodology compares state-of-the-art models across open-source frameworks and proprietary APIs, exploring how prompt modifications affect outcomes. To improve explainability, we employ SHAP values to reveal the influence of specific tokens on predictions across demographic groups. By integrating SHAP analyses into persona prompting, we examine how demographic attributes shape predictions. Results show that combining SHAP with persona prompting enhances both interpretability and reliability of LLM-generated annotations.

**Bias Statement** This paper examines how demographic factors, such as gender, ethnicity, education, and region, may influence both human and

<sup>1</sup><https://nlp.uned.es/exist2024/>

LLM annotations in detecting sexist content on social media. We focus on potential representational harms, wherein certain demographic groups' viewpoints or sensitivities to biased language might be underrepresented or misjudged. By highlighting differences in labeling behaviors across diverse annotator backgrounds, we aim to reduce the risk that an NLP system trained on these annotations will inadvertently perpetuate stereotypes or unfairly discount certain cultural experiences. We take the normative stance that all groups deserve unbiased and respectful treatment in both data collection and model predictions. Our ultimate goal is to ensure that technology, especially in sensitive tasks like sexism detection, does not exacerbate inequalities or reinforce harmful narratives.

## 2 Related Work

Recent studies have explored how annotators' personal backgrounds, experiences, and identities influence labeling outcomes, particularly in subjective tasks (Pei and Jurgens, 2023). However, findings in this area are mixed. Some studies report significant correlations between demographic features and annotation results (Excell and Al Moubayed, 2021), while others observe minimal statistically significant differences, especially regarding gender (Biester et al., 2022). These conflicting results highlight the complexity of the relationship between annotator characteristics and labeling decisions. Contrasts are particularly evident in tasks such as identifying sexist content, offensive language, and political ideologies, where an individual's personal experiences and group affiliations can significantly influence their perception and categorization of content (Kamruzzaman et al., 2024). The diversity of findings underscores the need for ongoing research to better understand the intricate interplay between annotator attributes and labeling outcomes. This understanding is crucial for developing more robust and inclusive NLP models that can effectively incorporate diverse perspectives in the annotation process. While some studies attempt to enhance data quality by analyzing disagreements among annotators, systematic investigations into how annotators' demographic biases affect annotation results remain limited (Gupta et al., 2024).

### 2.1 Persona Prompting for LLMs Annotations

One promising approach to NLP annotation tasks involves using GenAI Models, such as GPT-4,

which have been explored for automating annotation tasks due to their advanced language understanding capabilities (Manikandan et al., 2023). Furthermore, LLMs have shown potential in simulating diverse human perspectives by integrating demographic features into prompts (Hu and Collier, 2024). This technique, known as "persona prompting", has been effectively utilized to model human behavior and facilitate role-playing scenarios (Beck et al., 2024). For instance Hu and Collier (2024) examined how demographic, social, and behavioral persona variables influence LLM predictions and highlighted the importance of considering personal attributes in subjective NLP tasks. The success of LLMs in this domain has sparked discussions about their potential to replace human subjects in research contexts, particularly in annotation tasks (Dillion et al., 2023; Grossmann et al., 2023).

However, this raises concerns about identity misrepresentation and the flattening of group nuances (Wang et al., 2024). Moreover, persona prompting is not without its challenges. LLMs may carry inherent biases from their training data, potentially affecting annotation quality (Bender et al., 2021; Pavlovic and Poesio, 2024). Recent studies highlight these limitations, noting that LLMs often replicate societal biases or fail to adequately capture the nuances of minority perspectives (Hu et al., 2025; Pavlovic and Poesio, 2024). These issues emphasize the need for nuanced techniques to evaluate and mitigate the extent to which LLMs can accurately simulate human-like predictions.

### 2.2 LLMs Annotations' Interpretability

XAI techniques, particularly SHAP (SHapley Additive exPlanations), have become powerful tools for improving model interpretability by attributing importance to input features (Zhao et al., 2024). In NLP, SHAP effectively identifies influential tokens driving classification decisions and uncovers potential model biases (Ribeiro et al., 2016). Recent advances have expanded XAI's role in tasks such as sentiment classification, bias detection, toxic language identification, and inference (He et al., 2024). He et al. (2024) introduced a two-step framework using GPT-3.5, where the model first generates explanations and then annotates data through prompting. This approach has achieved performance comparable to or exceeding human annotators in tasks like Question Answering (QA) and Word-in-Context (WiC), demonstrating the potential of LLMs for annotation. Similarly, Ralevski

et al. (2024) applied GPT-3.5 and GPT-4 for annotating housing instability using chain-of-thought prompting. While LLMs are not yet suitable for full automation due to challenges such as bias, they show strong potential for computer-assisted annotation, reducing the time and cost of manual efforts.

### 3 Experimental Setup

#### 3.1 Dataset

We used data from the EXIST 2024 challenge (Plaza et al., 2024), which comprises datasets sourced from Twitter (now X). The labeled dataset contains tweets in both English and Spanish, with the training set comprising 6920 tweets in both languages (3260 in English, and 3660 in Spanish). For simplicity, we focus exclusively on Task 1 which involves binary classification of tweets to determine whether they express content related to sexism. Each tweet in the dataset was annotated by six individuals, who also provided demographic information across five categories: gender, age, ethnicity, education, and country. Specifically, gender was recorded as male or female; age was grouped into three categories (18–22, 23–45, and 46+); ethnicity included Asian, Black, White, Latino, Middle Eastern, Multiracial, and Other; education levels ranged from less than high school to doctorate; and annotators came from 45 countries. To simplify the analysis, these countries were categorized into five regions: Europe, America, Africa, Asia, and the Middle East. This grouping reduced the total number of unique demographic combinations from 266 to 117. We then eliminated combinations with rare representations, which we explain in detail in the next section.

#### 3.2 Generalized Linear Mixed Model

We ran a GLMM to examine how annotators’ demographic features affect labeling decisions. The model accounts for clustering of labels within tweets by incorporating random effects, ensuring that demographic influences are estimated independently of tweet-specific characteristics and individual differences. In our dataset, tweets and annotators serve as grouping variables, forming a crossed random effects structure: each tweet is labeled by multiple annotators, and each annotator labels multiple tweets. Also, tweets are hierarchically nested within languages. To account for both crossed and nested random effects, the following mixed-effects

logistic regression model is specified.<sup>2</sup>

$$E(\text{label}_{ij} | \mathbf{b}) = \text{logit}^{-1}(\mathbf{X}_{ij}\boldsymbol{\beta} + \mathbf{Z}_{ij}\mathbf{b})$$

In the model,  $\text{label}_{ij}$  is the binary response variable indicating whether the label for the  $i$ -th tweet by the  $j$ -th annotator is YES or NO. The design matrix  $\mathbf{X}_{ij}$  includes fixed effects for annotator demographic features, with  $\boldsymbol{\beta}$  representing their corresponding coefficients. Random effects are modeled as  $\mathbf{Z}_{ij}\mathbf{b}$ , capturing variation among tweets nested within languages and annotators. The random effects vector  $\mathbf{b}$  follows a multivariate normal distribution  $\sim N(0, \mathbf{G})$ . A logistic inverse link function,  $\text{logit}^{-1}(\cdot)$ , is used to model the binary outcome. This model evaluates demographic biases while accounting for tweet-level variability and annotator differences. Following prior studies (Pei and Jurgens, 2023), we excluded rare demographic features (i.e., representing less than 2% of annotators), such as the “Middle Eastern” ethnicity with only three annotators. Consequently, 69 out of 725 annotators were removed. We also excluded unique demographic combinations represented by only one annotator unless present in both languages. This resulted in 56 unique demographic combinations, detailed in Appendix A, Table 3. To address demographic and label-class imbalances, we assigned weights to each observation based on the inverse frequency of its demographic attributes and label class. The raw weight ( $W_{\text{raw}}$ ) for each observation was calculated as:

$$W_{\text{raw}} = \prod_{\text{features}} \frac{1}{f_{\text{group}}} \times \frac{1}{f_{\text{label}}}$$

Here,  $f_{\text{group}}$  denotes the relative frequency of a demographic category, and  $f_{\text{label}}$  the label class frequency. This approach, commonly used in survey weighting to address sample imbalances (Groves et al., 2011). For computational stability, raw weights were normalized to  $[0, 1]$  using  $W_{\text{norm}} = \frac{W_{\text{raw}}}{\max(W_{\text{raw}})}$  and then scaled for use in the mixed-effects model. As shown in Appendix C, Figure 5, the top ten demographic combinations with the highest weight contributions are identified across both YES and NO labels. For instance, female annotators aged 23–45, identifying as Black, holding a bachelor’s degree, and residing in Africa, provide the most balanced weighted input.

#### 3.3 BERT Model and SHAP Values

To classify texts as sexist or non-sexist, we use the Bidirectional Encoder Representations from Trans-

<sup>2</sup>In R notation,  $\text{label}_{ij} \sim \text{Annotators}' \text{ demographic factors} + (1 | \text{lang/id\_EXIST}) + (1 | \text{annotator\_id})$

formers (BERT) multilingual model. BERT captures word context by considering both left and right surroundings in a sentence (Devlin et al., 2019). The multilingual version is particularly suited to our dataset, which contains texts in two languages. During training, we fine-tune the BERT multilingual model using standard procedures. We use the Adam optimizer with a learning rate of  $3 \times 10^{-5}$  and a batch size of 128. The maximum sequence length is set to 512 tokens to handle longer texts. Binary cross-entropy is used as the loss function for this binary classification task. The model is trained for up to 10 epochs, with early stopping based on validation loss to prevent overfitting and ensure good generalization (Brownlee, 2018). To incorporate explainability into our methodology, we use SHAP values, following the approach by (Mohammadi et al., 2024). SHAP values quantify each token’s contribution to the model’s prediction, highlighting the most influential parts of the text. The SHAP value for each token  $t$ ,  $S_t$ , is computed by measuring the change in the model’s output when the token is included versus omitted across all possible subsets of input tokens. The SHAP value  $S_t$  for token  $t$  is computed as:

$$S_t = \sum_{T' \subseteq T \setminus \{t\}} \frac{|T'|!(|T|-|T'|-1)!}{|T|!} [f(T' \cup \{t\}) - f(T')]$$

Where  $T$  is the set of all tokens in the input text,  $T'$  is a subset of  $T$  excluding token  $t$ , and  $f(\cdot)$  represents the model’s prediction function. To find the most influential tokens, we calculate the SHAP importance  $SI_t$  for each token  $t$  by averaging the absolute SHAP values across all instances  $N_t$  where the token appears, considering only the cases where the model’s prediction matches the true label:

$$SI_t = \frac{1}{N_t} \sum_{i=1}^{N_t} |S_t(i)| \cdot \mathbb{I}(y_i = \hat{y}_i)$$

Here,  $S_t(i)$  is the SHAP value of token  $t$  in instance  $i$ ,  $y_i$  is the true label,  $\hat{y}_i$  is the predicted label, and  $\mathbb{I}(\cdot)$  is the indicator function. After that, we normalize the SHAP importance scores to compute the importance ratio for each token:  $IR_t = \frac{SI_t}{\sum_{k \in T} SI_k}$ . Tokens are ranked by importance ratios, and cumulative importance is calculated as  $CI_k = \sum_{i=1}^k IR_i$  to select the most influential tokens such that  $CI_k \leq T_c$ . We set the threshold  $T_c = 0.95$  to retain tokens contributing to 95% of the total importance. These top tokens are identified per class and incorporated into the GenAI prompts by bolding them, guiding the generative model to focus on critical parts of the text. Integrating SHAP enhances classifier transparency, revealing key factors driving decisions. Crucially,

analyzing high-importance tokens helps determine whether the model relies on meaningful indicators of sexism or spurious correlations. These tokens are then used in GenXAI and GenPXAI scenarios, which will be described in more detail in the section 3.5, to guide LLMs, allowing us to assess whether highlighting content-relevant features improves annotation reliability.

### 3.4 Large Language Models

We experiment with a range of LLMs, including local open-source models and cloud-based proprietary APIs, including OpenAI-based models (GPT-4o and GPT-4o mini)<sup>3</sup> and LLaMA-based models (LLaMA 3.2 3B and LLaMA 3.3 70B).<sup>4</sup> We evaluate how factors such as model size and architecture (e.g., OpenAI vs LLaMA variants) influence sexism detection across different scenarios and prompts.

### 3.5 GenAI Scenarios

We evaluate four main scenarios, each designed to probe the effect of additional context or highlighting on the model’s responses. These scenarios, inspired by a previous study (Kamruzzaman et al., 2024). The first scenario, **GenAI**, involves using a generative model to classify texts without additional context or guidance. The prompt instructs the model to make a straightforward classification based solely on the input text. In the second scenario, **Persona-Driven GenAI (GenP)**, a persona is added to the prompt to assess its impact on model performance. The persona provides specific demographic characteristics the model should adopt, aiming to influence its perspective and potentially enhance sensitivity to sexist content. The third scenario, **Explainable GenAI (GenXAI)**, incorporates SHAP values to highlight influential parts of the text using bold formatting, guiding the model’s attention to key sections. This tests whether focusing on important tokens improves classification accuracy. In the fourth scenario, **Persona-Driven Explainable GenAI (GenPXAI)**, we combine persona guidance with SHAP-based highlighting to assess the joint effect of perspective adoption and attention emphasis on model performance.

### 3.6 Personas

Personas are constructed based on demographic attributes such as gender, age, ethnicity, education,

<sup>3</sup><https://openai.com/index/hello-gpt-4o/> & <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>

<sup>4</sup><https://huggingface.co/meta-llama/Llama-3.2-3B> & <https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct>

and region, corresponding to the 56 unique demographic combinations listed in Appendix A, Table 3. They shape the perspective from which the model evaluates text, potentially reflecting cultural, linguistic nuance or demographic sensitivities.

- **English:** "You are a {gender} individual, aged {age}, who identifies as {ethnicity}, has a {study\_level}, and currently resides in {region}. You have the cultural and personal background of someone with these demographics."
- **Spanish:** "Eres una persona {gender}, de {age} años, que se identifica como {ethnicity}, posee un nivel de estudios {study\_level}, y actualmente reside en {region}. Tienes el trasfondo cultural y personal de alguien con estas características demográficas."

### 3.7 Important Tokens

For scenarios involving GenXAI and GenPXAI, we rely on previously computed important tokens from SHAP values. We highlight the top tokens by wrapping them in bold formatting (**\*\*token\*\***) to draw the model’s attention. This approach aims to help the model focus on terms that are most indicative of sexism.

### 3.8 Majority Voting

Majority voting is used to assign hard labels, while probabilities are used for soft labels. This provides a robust benchmark for evaluating automated methods. To simulate multiple annotators, the model generates six responses per text under each scenario and temperature setting. These six outputs represent "virtual annotators," and majority voting is applied to produce a single prediction per text. This simulates inter-annotator variability and offers a more robust estimate of the model’s stance, similar to human annotation aggregation.

## 4 Results and Discussion

**Do demographic biases mainly drive labeling differences, or does tweet content play a larger role?** To investigate this question, we first fit a flat logistic regression model with annotator demographic features as fixed effects. This provides a baseline assessment of demographic influence without accounting for tweet-specific or annotator-level variability. We then extend the analysis using a mixed-effects logistic regression model, incorporating crossed random intercepts for annotators and nested random effects for tweets within languages. This approach captures both annotator variability and tweet-specific differences while retaining demographic features as fixed effects.

Our findings show that incorporating tweet-level and annotator-level variability in the mixed-effects model substantially improves performance over the flat model. The mixed model achieves higher accuracy (73.73% vs 48.76%) and F1 score (75.77% vs 45.09%), along with better fit indicated by lower AIC and BIC values and a higher AUC. A kappa value of (47.06%) and an intraclass correlation coefficient (ICC) of 92.3% highlight the importance of accounting for tweet-specific differences, which the flat model ignores. Notably, the random effect for tweets shows high variance (33.72), indicating that tweet content is the main source of labeling variability. The annotator random effect (5.54) also contributes meaningfully, while the language effect (0.30) has minimal influence. These findings confirm the mixed-effects model as a more accurate and nuanced approach for understanding the labeling process.

Table 1: Comparison of Flat Model and Mixed-Effects Model Coefficients. Significant codes: '\*\*\*'very strong( $p < 0.001$ ), '\*\*'strong( $0.001 \leq p < 0.01$ ), '\*'moderate( $0.01 \leq p < 0.05$ ), '.'weak( $0.05 \leq p < 0.1$ ), '-'very weak( $0.1 \leq p < 1$ ).

Variable	Coef_Flat	P_Flat >  z	Coef_Mixed	P_Mixed >  z
(Intercept) <sup>1</sup>	0.274	***	-0.328	-
Female	0.020	***	0.055	-
23-45	0.206	***	0.027	-
46+	-0.089	***	0.111	-
Black	0.214	***	1.704	.
Latino	-0.237	***	-0.770	*
High school	-0.255	***	-0.465	*
Master	-0.506	***	0.048	-
Africa	-0.732	***	-2.865	**
America	0.178	***	0.370	-

<sup>1</sup> The reference group is male annotators aged 18–22 from Europe who hold a bachelor’s degree and identify as white.

Table 1 compares the coefficients of the flat logistic regression and mixed-effects models for each demographic feature. The flat model assumes independence among observations, ignoring the dataset’s hierarchical structure. As a result, it attributes all variability to fixed effects and residual error, potentially leading to biased coefficient estimates. For example, the flat model suggests females are slightly more likely to label YES than males, but it fails to account for content-specific variability, leading to a misleading interpretation. In contrast, the mixed-effects model incorporates random effects for tweet-level and language-level variability, showing that gender does not significantly influence labeling. This aligns with [Biester et al. \(2022\)](#), who found no significant gender-based differences in annotation behavior across various NLP tasks. Based on these findings, we use the mixed-effects model for further analysis, as it

offers a more robust and accurate framework for interpreting demographic impacts.

#### 4.1 Random Effects Interpretation

The odds ratio (OR)<sup>5</sup> for English tweets (OR = 0.84) indicates they are less likely to be labeled as sexist compared to Spanish tweets (OR = 1.95). Among the 347 annotators labeling Spanish tweets, 223 (64.27%) are from Spanish-speaking countries, while only 73 out of 302 (24.17%) annotators labeling English tweets are from English-speaking countries. Although we assume annotators are fluent in the language they label, regional residency may influence familiarity with cultural nuances and idiomatic expressions, affecting labeling decisions. Additionally, the grammatical structure of Spanish—being a gendered language—may make gender biases more explicit than in English. This aligns with [Lomotey \(2015\)](#), who emphasize the impact of grammatical gendering in Spanish. Thus, the observed differences in labeling may reflect both linguistic and cultural factors. Also, prior studies have found that classifiers achieve higher sexism-detection performance in English than in Spanish, likely due to the greater abundance of English-language training resources ([Fivez et al., 2024](#)).

#### 4.2 Fixed Effects Interpretation

While the OR for females is slightly above 1, suggesting women may be more attuned to gender bias, gender does not significantly influence labeling decisions. Male and female annotators exhibit similar behavior, supported by a 74% agreement in majority labeling, indicating consistency across genders. Similarly, although older annotators show slightly higher ORs, suggesting greater sensitivity to sexist content, no significant differences are observed across age groups, indicating age is not a decisive factor in labeling behavior. In contrast, ethnicity significantly affects labeling. Black annotators are more likely to label tweets as sexist (OR = 5.50), while Latino annotators are less likely compared to White annotators (OR = 0.46). These findings align with [Tahaei and Bergler \(2024\)](#) and [Kwarteng et al. \(2023\)](#), which highlight the heightened sensitivity of Black annotators, particularly Black women, due to lived experiences with intersectional discrimination. The lower likelihood among Latino annota-

<sup>5</sup>The odds ratio ( $OR = e^{\beta}$ ) refers to how the odds of the outcome ( $label = yes$ ) change when a predictor variable changes, while all other variables are held constant.

tors may reflect cultural norms. Regarding education, no significant differences are found between annotators with bachelor’s and master’s degrees. However, those with only a high school degree are significantly less likely to label tweets as sexist (OR = 0.63). Geographical location also plays a key role. Annotators from Africa are much less likely to label tweets as sexist (OR = 0.06), supporting findings from [Tahaei and Bergler \(2024\)](#) that emphasize the influence of country of origin and linguistic background on annotation behavior.

Our analysis shows that tweet-specific characteristics have a substantial impact on annotation outcomes, outweighing the influence of annotator demographics. While demographic features such as ethnicity, region, and education exhibit some significant associations with labeling tendencies, our mixed-effects model indicates that these effects are secondary to the inherent properties of the tweets. With an intraclass correlation coefficient (ICC) of 92%, the majority of the variance in labeling outcomes is attributed to tweet-level variability, with language contributing only a minor additional source of variation. The remaining 8% of the variance is explained by demographic variables and residual error. These findings suggest that, although demographic biases are not the dominant source of variability, they still play a meaningful role and should not be overlooked.

#### 4.3 BERT Model Interpretation

We employed a multilingual BERT model for binary sexism classification, fine-tuning it on 90% of the dataset using class weights and early stopping. Evaluated on the remaining 10% (ensuring representation of all demographic combinations), the model achieved test accuracies of 77% in English and 79% in Spanish, demonstrating consistent cross-lingual performance. To interpret the model’s decisions, particularly for classifying tweets as sexist (YES), we utilized SHAP values. Calculating normalized mean SHAP importance for tokens in correctly classified YES instances revealed insights into feature attribution.

As shown in [Figure 2](#), while a relatively small number of tokens capture roughly 50% of the cumulative importance, explaining near-total importance (e.g., 95%) necessitates considering a significantly larger lexicon, a trend particularly pronounced in Spanish. This suggests reliance on both core indicators and a broader range of terms for comprehensive detection. Examining the most influential

tokens provides further clarity.




Figure 2: Threshold vs. Number of Selected Tokens in both English and Spanish.

Figure 3 displays the top 20 tokens by SHAP importance. In English, terms like *slut*, *women*, *girls*, and *wife* dominate, highlighting the model’s focus on overtly gendered and potentially insulting language. Similarly, in Spanish, tokens such as *masculino*, *mujeres*, *feminist*, *mujer*, *mach*, and *sexual* are highly ranked, indicating a strong reliance on explicit gendered terms and references to sexual characteristics or ideologies.




Figure 3: Top 20 tokens by SHAP importance in (a) English and (b) Spanish.

This analysis confirms that both language models heavily weigh content features directly related to sexism. While the specific influential tokens differ due to linguistic variations, the underlying mechanism points towards content-based classifi-

cation. The distribution of influence also varies slightly, with the top 50 tokens accounting for 40% of importance in English versus 45% in Spanish (Appendix D, E), suggesting a slightly more concentrated reliance on key terms in Spanish.<sup>6</sup>

#### 4.4 GenAI Scenarios Results

We evaluate our approach on a 10% random sample of the dataset, comprising 326 English texts and 366 Spanish texts, covering all demographic groups. We measure performance using accuracy and F1-score. Table 2 presents the performance metrics for all GenAI scenarios across four models, LLaMA 3.2 3B, LLaMA 3.3 70B, OpenAI GPT-4o, and OpenAI GPT 4o-mini, in both English and Spanish.

Table 2: Performance metrics for all scenarios (see section 3.5). Numbers represent the scenarios: 1.GenAI, 2.GenP, 3.GenXAI, and 4.GenPXAI.

Accuracy	English				Spanish			
	1	2	3	4	1	2	3	4
LM 3B	0.50	0.47	<b>0.59</b>	0.53	0.43	0.43	0.48	<b>0.50</b>
LM 70B	<b>0.66</b>	0.64	0.65	0.64	<b>0.64</b>	0.58	0.57	0.58
GPT-4o	0.76	0.75	0.73	<b>0.78</b>	0.75	0.77	0.72	<b>0.77</b>
4o-mini	<b>0.79</b>	0.78	0.77	<b>0.79</b>	0.81	0.80	<b>0.82</b>	0.79

F1-score	English				Spanish			
	1	2	3	4	1	2	3	4
LM 3B	0.51	0.47	0.53	<b>0.53</b>	0.43	0.43	0.45	<b>0.47</b>
LM 70B	<b>0.66</b>	0.60	0.62	0.58	<b>0.62</b>	0.51	0.49	0.47
GPT-4o	0.74	0.74	0.71	<b>0.77</b>	0.74	0.76	0.70	<b>0.76</b>
4o-mini	0.78	0.78	0.77	<b>0.79</b>	0.81	0.80	<b>0.82</b>	0.79

Overall, OpenAI GPT 4o-mini and GPT-4o perform better, while LLaMA 3.2 3B tends to perform worse, and LLaMA 3.3 70B is in between. The English subset often shows a baseline advantage for the more capable models, while the Spanish subset sometimes benefits more from certain prompting strategies. Differences across scenarios help reveal the impact of introducing personas and focusing attention on important tokens (XAI). Critically, assessing the utility of demographic personas (Scenario 2, GenP), we observe that it often provides no significant improvement over the baseline GenAI (Scenario 1) and occasionally leads to worse performance (e.g., LLaMA 3B and 70B models show decreased accuracy or F1-score in English, and LLaMA 70B sees a notable drop in F1-score in Spanish when personas are added). Even for the higher-performing GPT models, the gains from persona prompting alone are minimal or absent (e.g., GPT-4o mini accuracy slightly decreases in both languages). This suggests that simply layering demographic characteristics onto the prompt does not

<sup>6</sup>An exploratory analysis of unique token diversity across annotator demographic groups, detailed in Appendix B.



reliably enhance the LLM’s ability to replicate nuanced human judgments for this task, questioning the value of such personas for improving annotation reliability.

Focusing on XAI (Scenario 3, GenXAI), highlighting important tokens identified by SHAP often helps smaller models (e.g., LLaMA 3.2 3B shows a marked improvement in accuracy in English going from 0.50 to 0.59, and in Spanish from 0.43 to 0.48) and provides a solid baseline, suggesting benefit from focusing the model on content features deemed important by an explainability analysis. For larger models, the effect of XAI alone is mixed, sometimes resulting in slight performance dips compared to the baseline (e.g., GPT-4o). For larger models, the combined approach (Scenario 4, GenPXAI) sometimes yields the highest scores (e.g., GPT-4o achieves its peak accuracy and F1 in both languages, and 4o-mini peaks in English). However, the improvement of GenPXAI over GenXAI is often marginal or inconsistent. For instance, with GPT-4o mini in Spanish, the GenXAI scenario (0.82 Acc, 0.82 F1) actually slightly outperforms the combined GenPXAI scenario (0.79 Acc, 0.79 F1). This pattern raises questions about whether the persona component in GenPXAI adds substantial value beyond the guidance provided by the content-focused XAI highlighting. The data suggests that directing the model’s attention to relevant textual features (XAI) might be the more robust and impactful strategy, rather than attempting to simulate demographic perspectives through personas, whose contribution appears less certain. In summary, these results indicate that while baseline GenAI models already achieve strong performance on this task, the addition of demographic persona information offers questionable and inconsistent benefits for improving annotation reliability in this context. Guiding the model’s attention using XAI based on content features appears more consistently helpful, particularly when paired with capable models, suggesting that focusing on the text itself through explainability methods is a more promising path forward than relying on potentially superficial persona simulation.

## 5 Conclusion

This study evaluated the reliability of LLM annotations for sexism detection, focusing on the roles of annotator demographics and model explainability. Mixed-effects modeling showed that demographic

factors, while sometimes statistically significant, accounted for only 8% of the variance in human labels, tweet content and individual differences were the main drivers. We tested the use of demographic personas to guide LLMs but found this strategy had limited, inconsistent, and sometimes negative effects on performance. SHAP analysis confirmed that content drove model decisions. These findings suggest that bias mitigation should focus less on broad demographic corrections and more on content and individual-level understanding. Simulated personas may oversimplify complexity and risk reinforcing stereotypes. This limitation is underscored by evidence that LLM often exhibits uniform stylistic patterns (Mohammadi et al., 2025), showing that current models cannot fully emulate the diverse differences of human annotators. Instead, explainability tools that highlight content-relevant features offer a more promising path toward fairness and reliability in NLP. Future research should explore richer ways to capture diverse perspectives and improve content-based guidance in LLM annotations.

## 6 Limitations and Future Work

Although our analysis suggests that demographic factors account for only a fraction of the variability in the labeling, our findings may not generalise to other languages or cultural contexts. Future work should examine a wider range of datasets and linguistic settings to better assess the robustness and cross-cultural applicability of our approach. Our persona-driven prompts and explainability techniques rely on relatively broad demographic categories, which cannot capture the full richness of individual identities or personal experiences. Additionally, LLMs can exhibit hidden biases derived from their training data, and our prompts may not always surface or mitigate these biases.

### Reproducibility

All codes and experiment notebooks are available on GitHub.<sup>7</sup>

### Acknowledgments

We thank the organizers and annotators for their contributions in EXIST 2024 and acknowledge support from the focus area Applied Data Science (ADS) funding from Utrecht University.

<sup>7</sup>[https://github.com/mohammadi-hadi/Explainable\\_Annotations\\_Reliability](https://github.com/mohammadi-hadi/Explainable_Annotations_Reliability)

## References

- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational linguistics*, 34(4):555–596.
- Valerio Basile, Michael Fell, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, Massimo Poesio, Alexandra Uma, and 1 others. 2021. We need to consider disagreement in evaluation. In *Proceedings of the 1st workshop on benchmarking: past, present and future*, pages 15–21. Association for Computational Linguistics.
- Tilman Beck, Hendrik Schuff, Anne Lauscher, and Iryna Gurevych. 2024. Sensitivity, performance, robustness: Deconstructing the effect of sociodemographic prompting. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2589–2615.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Laura Biester, Vanita Sharma, Ashkan Kazemi, Naihao Deng, Steven Wilson, and Rada Mihalcea. 2022. Analyzing the effects of annotator gender across nlp tasks. In *Proceedings of the 1st Workshop on Perspectivist Approaches to NLP@ LREC2022*, pages 10–19.
- Herbert Bless and Klaus Fiedler. 2014. *Social cognition: How individuals construct social reality*. Psychology Press.
- Jason Brownlee. 2018. A gentle introduction to early stopping to avoid overtraining neural networks. *Machine Learning Mastery*, 7.
- 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.
- Danica Dillion, Niket Tandon, Yuling Gu, and Kurt Gray. 2023. Can ai language models replace human participants? *Trends in Cognitive Sciences*, 27(7):597–600.
- Elizabeth Excell and Noura Al Moubayed. 2021. Towards equal gender representation in the annotations of toxic language detection. *GeBNLP 2021*, page 55.
- Pieter Fivez, Walter Daelemans, Tim Van de Cruys, Yury Kashnitsky, Savvas Chamezopoulos, Hadi Mohammadi, Anastasia Giachanou, Ayoub Bagheri, Wessel Poelman, Juraj Vladika, and 1 others. 2024. The clin33 shared task on the detection of text generated by large language models. *Computational Linguistics in the Netherlands Journal*, 13:233–259.
- Gabriel Freedman, Adam Dejl, Deniz Gorur, Xiang Yin, Antonio Rago, and Francesca Toni. 2024. Argumentative large language models for explainable and contestable decision-making. *arXiv preprint arXiv:2405.02079*.
- Igor Grossmann, Matthew Feinberg, Dawn C Parker, Nicholas A Christakis, Philip E Tetlock, and William A Cunningham. 2023. Ai and the transformation of social science research. *Science*, 380(6650):1108–1109.
- Robert M Groves, Floyd J Fowler Jr, Mick P Couper, James M Lepkowski, Eleanor Singer, and Roger Tourangeau. 2011. *Survey methodology*. John Wiley & Sons.
- Vipul Gupta, Pranav Narayanan Venkit, Shomir Wilson, and Rebecca J Passonneau. 2024. Sociodemographic bias in language models: A survey and forward path. In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 295–322.
- Xingwei He, Zhenghao Lin, Yeyun Gong, A-Long Jin, Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan Duan, and Weizhu Chen. 2024. Annollm: Making large language models to be better crowdsourced annotators. In *NAACL (Industry Track)*.
- Tom Hosking, Phil Blunsom, and Max Bartolo. Human feedback is not gold standard. In *The Twelfth International Conference on Learning Representations*.
- Tiancheng Hu and Nigel Collier. 2024. Quantifying the persona effect in llm simulations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10289–10307.
- Tiancheng Hu, Yara Kyrychenko, Steve Rathje, Nigel Collier, Sander van der Linden, and Jon Roozenbeek. 2025. Generative language models exhibit social identity biases. *Nature Computational Science*, 5(1):65–75.
- Mahammed Kamruzzaman, Hieu Nguyen, Nazmul Hasan, and Gene Louis Kim. 2024. "a woman is more culturally knowledgeable than a man?": The effect of personas on cultural norm interpretation in llms. *arXiv preprint arXiv:2409.11636*.
- Hannah Rose Kirk, Alexander Whitefield, Paul Rottger, Andrew M Bean, Katerina Margatina, Rafael Mosquera-Gomez, Juan Ciro, Max Bartolo, Adina Williams, He He, and 1 others. 2024. The prism alignment dataset: What participatory, representative and individualised human feedback reveals about the subjective and multicultural alignment of large language models. *Advances in Neural Information Processing Systems*, 37:105236–105344.
- Klaus Krippendorff. 2011. Computing krippendorff’s alpha-reliability.

- Joseph Kwarteng, Gregoire Burel, Aisling Third, Tracie Farrell, and Miriam Fernandez. 2023. Understanding misogynoir: A study of annotators’ perspectives. In *Proceedings of the 15th ACM Web Science Conference 2023*, pages 271–282.
- Elisa Leonardelli, Stefano Menini, Alessio Palmero Aprosio, Marco Guerini, and Sara Tonelli. 2021. Agreeing to disagree: Annotating offensive language datasets with annotators’ disagreement. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10528–10539.
- Benedicta Adokarley Lomotey. 2015. On sexism in language and language change—the case of peninsular spanish. *Linguistik online*, 70(1):167–183.
- Hariharan Manikandan, Yiding Jiang, and J Zico Kolter. 2023. Language models are weak learners. *Advances in Neural Information Processing Systems*, 36:50907–50931.
- Hadi Mohammadi, Anastasia Giachanou, and Ayoub Bagheri. 2024. A transparent pipeline for identifying sexism in social media: Combining explainability with model prediction. *Applied Sciences*, 14(19):8620.
- Hadi Mohammadi, Anastasia Giachanou, Daniel L Oberski, and Ayoub Bagheri. 2025. Explainability-based token replacement on llm-generated text. *arXiv preprint arXiv:2506.04050*.
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Maja Pavlovic and Massimo Poesio. 2024. The effectiveness of llms as annotators: A comparative overview and empirical analysis of direct representation. *LREC-COLING 2024*, page 100.
- Jiaxin Pei and David Jurgens. 2023. When do annotator demographics matter? measuring the influence of annotator demographics with the popquorn dataset. In *The 17th Linguistic Annotation Workshop (LAW-XVII)@ ACL 2023*, page 252.
- Barbara Plank. 2022. The “problem” of human label variation: On ground truth in data, modeling and evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682.
- Barbara Plank, Dirk Hovy, Anders Sogaard, and 1 others. 2014. Learning part-of-speech taggers with inter-annotator agreement loss. In *Proceedings of EACL*. Association for Computational Linguistics.
- Laura Plaza, Jorge Carrillo-de Albornoz, Enrique Amigó, Julio Gonzalo, Roser Morante, Paolo Rosso, Damiano Spina, Berta Chulvi, Alba Maeso, and Víctor Ruiz. 2024. Exist 2024: sexism identification in social networks and memes. In *European Conference on Information Retrieval*, pages 498–504. Springer.
- Alexandra Ralevski, Nadaa Taiyab, Michael Nossal, Lindsay Mico, Samantha N Piekos, and Jennifer Hadlock. 2024. Using large language models to annotate complex cases of social determinants of health in longitudinal clinical records. *medRxiv*.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144.
- Bryan C Russell, Antonio Torralba, Kevin P Murphy, and William T Freeman. 2008. Labelme: a database and web-based tool for image annotation. *International journal of computer vision*, 77:157–173.
- Johannes Schäfer, Aidan Combs, Christopher Bagdon, Jiahui Li, Nadine Probol, Lynn Greschner, Sean Pappay, Yarik Menchaca Resendiz, Aswathy Velutharambath, Amelie Wüthrl, and 1 others. 2024. Which demographics do llms default to during annotation? *arXiv preprint arXiv:2410.08820*.
- Narjes Tahaei and Sabine Bergler. 2024. Analysis of annotator demographics in sexism detection. In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 376–383.
- Alexandra Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2020. A case for soft loss functions. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 173–177.
- Alexandra N Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470.
- Angelina Wang, Jamie Morgenstern, and John P Dickerson. 2024. Large language models cannot replace human participants because they cannot portray identity groups. *CoRR*.
- Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology*, 15(2):1–38.

## A Annotators' demographic combination

The total number of unique demographic combinations after removing those with rare representations.

Table 3: Unique Demographic Combinations

#	Possible Combination	# of es Ann	# of en Ann
1	F, 18-22, Black, Bachelor, Africa	0	4
2	F, 18-22, Black, High school, Africa	0	3
3	F, 18-22, Latino, Bachelor, America	19	1
4	F, 18-22, Latino, High school, America	15	4
5	F, 18-22, Latino, High school, Europe	1	1
6	F, 18-22, White, Bachelor, America	2	0
7	F, 18-22, White, Bachelor, Europe	15	18
8	F, 18-22, White, High school, Europe	7	25
9	F, 23-45, Black, Bachelor, Africa	0	9
10	F, 23-45, Black, High school, Africa	0	2
11	F, 23-45, Latino, Bachelor, America	34	11
12	F, 23-45, Latino, High school, America	6	0
13	F, 23-45, Latino, Master, America	2	0
14	F, 23-45, White, Bachelor, America	1	1
15	F, 23-45, White, Bachelor, Europe	7	20
16	F, 23-45, White, High school, Europe	1	3
17	F, 23-45, White, Master, Europe	9	14
18	F, 46+, Black, Bachelor, Africa	0	4
19	F, 46+, Latino, Bachelor, America	12	0
20	F, 46+, Latino, Bachelor, Europe	3	0
21	F, 46+, Latino, High school, America	2	1
22	F, 46+, Latino, Master, America	6	1
23	F, 46+, White, Bachelor, America	3	2
24	F, 46+, White, Bachelor, Europe	11	9
25	F, 46+, White, High school, Africa	0	3
26	F, 46+, White, High school, America	2	2
27	F, 46+, White, High school, Europe	4	16
28	F, 46+, White, Master, America	2	0
29	F, 46+, White, Master, Europe	7	6
30	M, 18-22, Black, Bachelor, Africa	0	2
31	M, 18-22, Black, High school, Africa	0	2
32	M, 18-22, Latino, Bachelor, America	10	2
33	M, 18-22, Latino, Bachelor, Europe	1	2
34	M, 18-22, Latino, High school, America	17	7
35	M, 18-22, Latino, High school, Europe	3	2
36	M, 18-22, Latino, Master, Europe	2	0
37	M, 18-22, White, Bachelor, Europe	17	11
38	M, 18-22, White, High school, Europe	11	25
39	M, 18-22, White, Master, Europe	0	3
40	M, 23-45, Black, Bachelor, Africa	0	7
41	M, 23-45, Black, Master, Africa	0	2
42	M, 23-45, Latino, Bachelor, America	8	5
43	M, 23-45, Latino, Bachelor, Europe	1	2
44	M, 23-45, Latino, Master, America	2	0
45	M, 23-45, Latino, Master, Europe	2	0
46	M, 23-45, White, Bachelor, Europe	24	10
47	M, 23-45, White, High school, Europe	4	10
48	M, 23-45, White, Master, Europe	18	15
49	M, 46+, Latino, Bachelor, America	8	3
50	M, 46+, Latino, Master, America	2	0
51	M, 46+, White, Bachelor, Africa	0	2
52	M, 46+, White, Bachelor, America	5	5
53	M, 46+, White, Bachelor, Europe	21	14
54	M, 46+, White, High school, America	0	2
55	M, 46+, White, High school, Europe	12	15
56	M, 46+, White, Master, Europe	8	5

## B Unique Token Analysis by Demographic Group

To further explore potential secondary demographic influences, we analyzed the distribution of unique token counts within tweets annotated by different demographic groups. This exploratory analysis aimed to identify potential variations in linguistic engagement or lexical diversity associ-

ated with annotator characteristics. As shown in Figure 4, we observed some variation across categories in both English and Spanish. For instance, certain groups exhibited broader ranges of unique tokens, potentially hinting at subtle cultural or linguistic factors influencing how they engage with the text. However, consistent with our primary findings, these observed differences appear secondary to the overwhelming influence of the tweet content itself on the annotation process and model interpretation.

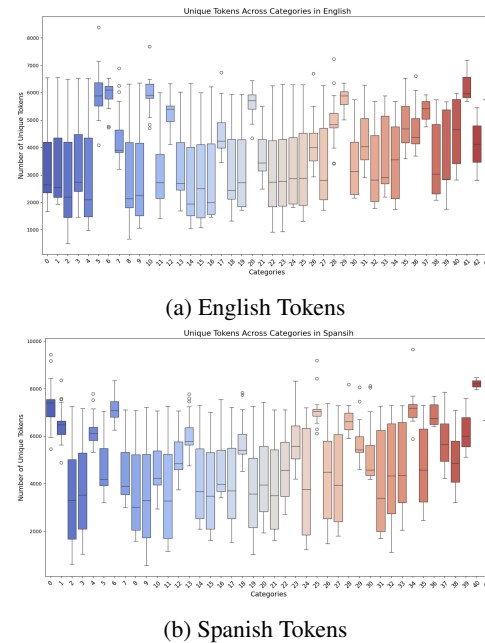


Figure 4: Distribution of unique tokens per tweet across various annotator demographic categories in (a) English and (b) Spanish. This exploratory analysis hints at subtle variations but confirms the secondary nature of these effects compared to content.

## C Top ten demographic combination




Figure 5: Different demographic combinations that have the highest weight contributions across both label classes

## D Complete Lists of Important Tokens

Here are the tokens identified by the model that contribute to classifying tweets as sexist, along with their importance scores.

Table 4: 50 Top important English Tokens

Token	SHAP	Ratio	Cum.	Token	SHAP	Ratio	Cum.
slut	0.4041	0.0246	0.0246	feminist	0.1017	0.0062	0.2818
women	0.3928	0.0239	0.0485	periods	0.0991	0.0060	0.2878
girls	0.3561	0.0217	0.0702	pro	0.0974	0.0059	0.2938
fem	0.3324	0.0202	0.0905	her	0.0972	0.0059	0.2997
Wife	0.2896	0.0176	0.1082	ok	0.0935	0.0057	0.3054
scholar	0.2858	0.0174	0.1256	She	0.0924	0.0056	0.3110
woman	0.2807	0.0171	0.1427	boys	0.0896	0.0054	0.3165
onde	0.2559	0.0156	0.1583	ti	0.0871	0.0053	0.3218
ches	0.2278	0.0138	0.1722	Like	0.0853	0.0052	0.3270
teaching	0.2264	0.0138	0.1860	mbo	0.0837	0.0051	0.3321
stitute	0.1735	0.0105	0.1966	ips	0.0836	0.0051	0.3372
pregnant	0.1682	0.0102	0.2068	ts	0.0820	0.0050	0.3422
gang	0.1624	0.0099	0.2167	coverage	0.0808	0.0049	0.3472
men	0.1430	0.0087	0.2255	really	0.0806	0.0049	0.3521
biggest	0.1382	0.0084	0.2339	wife	0.0776	0.0047	0.3568
bl	0.1249	0.0076	0.2415	dies	0.0773	0.0047	0.3615
girl	0.1182	0.0072	0.2487	finger	0.0768	0.0046	0.3662
Women	0.1156	0.0070	0.2558	trophy	0.0759	0.0046	0.3708
bit	0.1155	0.0070	0.2628	dressed	0.0747	0.0045	0.3754
pen	0.1073	0.0065	0.2694	ina	0.0742	0.0045	0.3799
financial	0.1021	0.0062	0.2756	Why	0.0739	0.0045	0.3844
female	0.0734	0.0044	0.3889	comment	0.0733	0.0044	0.3934
dress	0.0702	0.0042	0.3977	sex	0.0672	0.0041	0.4017
male	0.0669	0.0040	0.4058	husband	0.0668	0.0040	0.4099
ehan	0.0654	0.0039	0.4139	ouse	0.0649	0.0039	0.4179

Table 5: 50 Top important Spanish Tokens

Token	SHAP	Ratio	Cum.	Token	SHAP	Ratio	Cum.
apa	0.1573	0.0063	0.3787	feminist	0.3258	0.0132	0.1557
ones	0.1489	0.0060	0.3848	mujer	0.3184	0.0129	0.1686
ios	0.1478	0.0059	0.3907	lab	0.3151	0.0127	0.1814
var	0.1476	0.0059	0.3967	vas	0.3123	0.0126	0.1941
novia	0.1416	0.0057	0.4025	hombre	0.3026	0.0122	0.2063
bian	0.1415	0.0057	0.4082	mach	0.2965	0.0120	0.2184
golf	0.1414	0.0057	0.4140	dama	0.2881	0.0116	0.2301
male	0.1393	0.0056	0.4196	tú	0.2822	0.0114	0.2415
marido	0.1384	0.0056	0.4252	bia	0.2508	0.0101	0.2517
tant	0.1289	0.0052	0.4305	Od	0.2485	0.0100	0.2618
laga	0.1269	0.0051	0.4356	sexual	0.2453	0.0099	0.2717
ñas	0.1242	0.0050	0.4406	fem	0.2309	0.0093	0.2811
ellas	0.1235	0.0050	0.4457	femenino	0.2263	0.0091	0.2903
amo	0.1227	0.0049	0.4506	doctor	0.2237	0.0090	0.2993
aca	0.1179	0.0047	0.4554	princesa	0.2231	0.0090	0.3084
loc	0.1080	0.0043	0.4598	nen	0.2200	0.0089	0.3173
ball	0.1023	0.0041	0.4640	masculin	0.2189	0.0088	0.3262
nar	0.0781	0.0234	0.0234	Mujeres	0.2137	0.0086	0.3349
masculino	0.4012	0.0162	0.0397	niña	0.2028	0.0082	0.3431
prend	0.3953	0.0160	0.0557	bella	0.1890	0.0076	0.3508
mach	0.3804	0.0154	0.0712	ton	0.1839	0.0074	0.3582
zo	0.3665	0.0148	0.0860	niños	0.1807	0.0073	0.3656
mujeres	0.3642	0.0147	0.1008	ment	0.1670	0.0067	0.3723
mans	0.3615	0.0146	0.1155	novi	0.3394	0.0137	0.1292
señor	0.3266	0.0132	0.1425	sÃ	0.1003	0.0040	0.4680

## E Cumulative importance of the top 50 tokens

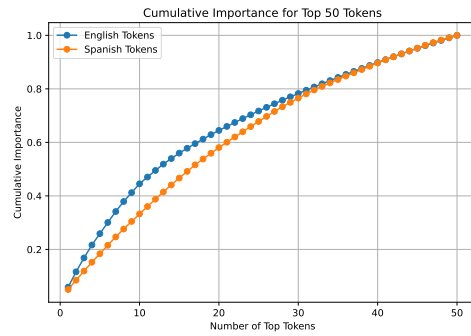




Figure 6: The cumulative importance of the top 50 tokens in both English and Spanish.

## F LLMs Performance Comparison



(a) English



(b) Spanish

Figure 7: Comparing TPR and FNR across models, scenarios, and languages.