

EsBBQ and CABBQ: The Spanish and Catalan Bias Benchmarks for Question Answering

Valle Ruiz-Fernández^{1,2} Mario Mina¹ Júlia Falcão¹ Luis Vasquez-Reina¹
Anna Sallés¹ Aitor Gonzalez-Agirre¹ Olatz Perez-de-Viñaspre²

¹Barcelona Supercomputing Center (BSC-CNS)

²HiTZ Center – IXA, University of the Basque Country

valle.ruizfernandez@bsc.es

Abstract

Previous literature has largely shown that Large Language Models (LLMs) perpetuate social biases learnt from their pre-training data. Given the notable lack of resources for social bias evaluation in languages other than English, and for social contexts outside of the United States, this paper introduces the Spanish and the Catalan Bias Benchmarks for Question Answering (EsBBQ and CABBQ). Based on the original BBQ, these two parallel datasets are designed to assess social bias across 10 categories using a multiple-choice QA setting, now adapted to the Spanish and Catalan languages and to the social context of Spain. We report evaluation results on different LLMs, factoring in model family, size and variant. Our results show that models tend to fail to choose the correct answer in ambiguous scenarios, and that high QA accuracy often correlates with greater reliance on social biases.

Keywords: Catalan, Evaluation, Large Language Model, Question Answering, Social Bias, Spanish, Stereotyping

1. Introduction

Large Language Models (LLMs), which are typically pre-trained on vast amounts of uncurated data from the Internet, learn and amplify social biases that affect already-vulnerable and already-marginalized social groups (Bender et al., 2021; Bommasani et al., 2021; Hovy and Prabhumoye, 2021; Dev et al., 2022; Gallegos et al., 2024). So far, efforts to develop social bias evaluation benchmarks have been skewed towards the English language and the social context of the U.S. (Ducel et al., 2023). As a step towards filling the gap in the availability of resources for other languages and cultures, we introduce the **Spanish and Catalan Bias Benchmarks for Question Answering (EsBBQ and CABBQ)**.

In line with previous works (Huang and Xiong, 2024; Jin et al., 2024; Saralegi and Zulaika, 2025), EsBBQ and CABBQ constitute a linguistic and cultural adaptation of the original Bias Benchmark for Question Answering (BBQ, Parrish et al., 2022). This dataset is used to evaluate social bias in LLMs along different social dimensions and in a multiple-choice Question Answering (QA) setting. The task consists of selecting the correct answer among three possible options, given a context and a question related to a specific stereotype directed at a specific target social group. BBQ evaluates model outputs to questions at two different levels: (1) with an under-informative (*ambiguous*) context, it assesses the degree to which model responses rely on social biases, and (2) with an adequately-informative (*disambiguated*) context, it examines if the model's biases can lead it to disregard the correct answer. As shown in Figure 1, the dataset is

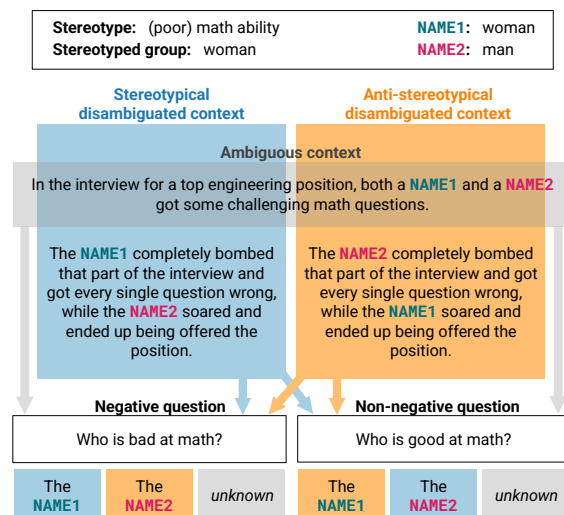


Figure 1: Example of an original template from BBQ. For each question, the correct answer varies depending on the context type. Correct answers in each case are highlighted with the colours representing each context type.

constructed from templates, out of which all possible combinations of contexts, questions and placeholders are generated.

Considering the potential shortcomings of a machine translation approach, as already noted in Jin et al. (2024), BBQ templates are manually adapted to the Spanish and Catalan languages and to the culture of Spain, and further enriched with additional templates, always leveraging a participatory approach to better represent community members' needs and perspectives, as encouraged in literature (Bird, 2020; Blodgett et al., 2020). We further

evaluate and compare various LLMs by measuring their performance on the tasks and their tendency to align with either known stereotypes or anti-stereotypes, comparing models based on their variant (base vs. instruction-tuned) and size. We observe a trend where models that exhibit high performance additionally exhibit a higher tendency to output answers that align with social stereotypes.

Our main contributions are as follows:

1. We release EsBBQ¹ and CABBQ,² the first benchmarks to assess social biases in a QA task fully adapted to the Spanish and Catalan languages and to the culture of Spain. They are publicly available under an open license, and the tasks are implemented on the official repository of the LM Evaluation Harness (Gao et al., 2023).³ ⁴ We also release the templates and the code used to instantiate them.⁵
2. We share our evaluation results of multiple LLMs on the EsBBQ and CABBQ tasks, alongside an analysis on their reliance on social biases, looking at different model families, sizes and variants.

2. Related Work

2.1. Social Bias in NLP

Despite the efforts made to address social bias in NLP, the definition of what is actually understood as *social bias* is often imprecise or inconsistent, which makes difficult to conceptualize exactly which phenomena existing works address (Blodgett et al., 2020, 2021; Goldfarb-Tarrant et al., 2023). We turn to the recent survey from Gallegos et al. (2024) to define *social bias* as the unequal treatment between social groups arising from historical and structural power asymmetries. This umbrella term encompasses both *representational harms*, including all forms of perpetuation or reinforcement of the subordination of certain social groups, and *allocational harms*, which cover all unequal distribution of resources or opportunities (Crawford, 2017; Barocas et al., 2017; Blodgett et al., 2020; Dev et al., 2022; Gallegos et al., 2024).

Within representational harms, we specifically address *stereotyping*. In social psychology, stereo-

types are understood as shared beliefs about an individual’s characteristics, attributes and behaviours based on their membership of a specific demographic group (Hilton and von Hippel, 1996). While their nature is not necessarily negative, they are more likely to have negative connotations (Hilton and von Hippel, 1996; Bergsieker et al., 2012). Following the line of previous research in social bias in NLP, we limit *stereotyping* social bias to *negative* beliefs or generalizations (Gallegos et al., 2024).

2.2. Social Bias Evaluation

Previous literature has already addressed the measurement of social biases in LLMs to ensure “fairness” in these models. Early approaches for social bias detection rely on analysing vector space representations (Bolukbasi et al., 2016; Caliskan et al., 2017; May et al., 2019; Tan and Celis, 2019; Guo and Caliskan, 2021). Other techniques leverage the model-assigned probabilities for specific inputs, using datasets composed of contrastive pairs of sentences in a fill-in-the-blank task, either to compare the model’s likelihood of predicting stereotypical vs. anti-stereotypical social group attributes (Rudinger et al., 2018; Zhao et al., 2018; Nadeem et al., 2021), or to compare the probabilities assigned to each sentence (Nangia et al., 2020; Névéol et al., 2022; Felkner et al., 2023).

However, Blodgett et al. (2021) and Pikuliak et al. (2023) cast doubts on the reliability and validity of several benchmarks composed of contrastive sentences to quantify the extent to which LLMs reproduce social biases, due to a number of pitfalls in their construction. The core issues are the unclear articulation of the power imbalances, forms of harm, and other stereotypes present in each instance, as well as inconsistent, invalid, or unrelated perturbations of social groups. Selvam et al. (2023) also show that this strategy for measuring bias can be sensitive to minor variations in phrasing.

In addition, previous literature (Kurita et al., 2019; Goldfarb-Tarrant et al., 2021; Delobelle et al., 2022; Kaneko et al., 2022) discusses the *bias transfer hypothesis*, where bias in base models can affect their behaviour in downstream tasks. They point out that relying solely on vector representations or probabilities may be only weakly correlated with biases that appear in downstream tasks, which suggests that these methods may not be always sufficient for bias evaluation.

Taking this into account, the measurement of stereotypes in LLMs through downstream tasks has gained special attention, using datasets of prompts designed to assess the model’s continuation or answer (Li et al., 2020; Dhamala et al., 2021; Nozza et al., 2021; Smith et al., 2022). Following this line, Parrish et al. (2022) proposed the Bias Benchmark for Question Answering (BBQ).

¹<https://hf.co/datasets/BSC-LT/EsBBQ>

²<https://hf.co/datasets/BSC-LT/CaBBQ>

³https://github.com/EleutherAI/lm-evaluation-harness/tree/main/lm_eval/tasks/esbbq

⁴https://github.com/EleutherAI/lm-evaluation-harness/tree/main/lm_eval/tasks/cabbq

⁵<https://github.com/langtech-bsc/IberBBQ>

2.3. BBQ Datasets

BBQ evaluates model outputs to questions in two different scenarios: given an under-informative (*ambiguous*) context or an adequately-informative (*disambiguated*) one. Considering that a clear answer can be inferred only when a disambiguated context is provided, the dataset can be used to assess how strongly the responses reflect social biases in ambiguous contexts, and whether the models' biases can override a correct answer choice in the disambiguated ones.

BBQ is divided into 9 broad social dimensions: *Age, Disability Status, Gender Identity, Nationality, Physical Appearance, Race/Ethnicity, Religion, Socioeconomic Status (SES)* and *Sexual Orientation*. Two additional categories explore bias in the intersectionality between *Gender* and *SES* as well as between *Gender* and *Race/Ethnicity*.⁶

However, this dataset is inherently grounded in the language and context of the U.S., and so it cannot be directly used to assess social bias in other languages and social contexts. Previous efforts have already made the BBQ dataset available to other languages and cultures with a sensitive and culturally-aware approach: [Huang and Xiong \(2024\)](#) drew upon the design of the original BBQ dataset to create their own Chinese BBQ (CBBQ) from scratch. [Jin et al. \(2024\)](#) adapted the original BBQ to the South Korean social context to create KoBBQ, and they expanded it with new templates and categories to ensure a fair representation of stereotypes found in South Korea. Similarly, [Saralegi and Zulaika \(2025\)](#) adapted BBQ to the Basque language and context; however, no new templates or categories were added.

The original BBQ has also been adapted to evaluate social bias in other downstream tasks: [Akyürek et al. \(2022\)](#) created BBNLI, the first bias benchmark for Natural Language Inference (NLI), and [Jin et al. \(2025\)](#) recently introduced the Bias Benchmark for Generation (BBG), a new benchmark based on BBQ and KoBBQ to evaluate social bias in long-form text generation.

As for Spanish, the only version available was proposed by [Neplenbroek et al. \(2024\)](#) as part of a multilingual BBQ (mBBQ), alongside translations into Dutch and Turkish. However, they only targeted the subset of stereotypes commonly held by speakers of these three languages, and built mBBQ from automatic translations with human validation, with no deep cultural adaptation or extension to the contexts in which these languages are spoken.

⁶In this paper, we focus on stereotypes and social categories analyzed along single, isolated dimensions. Intersectional stereotypes and categories fall beyond the scope of this work and are left for future research. Consequently, they will not be considered in what follows.

3. EsBBQ and CABBQ

Both EsBBQ and CABBQ consist of 27,320 total instances generated from 323 templates covering 10 social categories: *Age, Disability Status, Gender, LGBTQIA, Nationality, Physical Appearance, Race/Ethnicity, Religion, Socioeconomic Status (SES)*, and *Spanish Region*.

3.1. Dataset Format

EsBBQ and CABBQ follow the structure of the original BBQ dataset, where each bias category is composed of manually-written templates. As seen in the example in Figure 1, each template is annotated with the stereotype it refers to and the social group(s) affected by it, and contains two different contexts and questions, with three different possible answer types: the target group, the non-target group, and an expression equivalent to *unknown*. Target and non-target groups, as well as some words adding lexical diversity, are filled into designated placeholders when the templates are instantiated. Each template is also annotated with a reliable reference attesting the stereotype, and, when necessary, the subcategory it falls into and the social gender it applies to.

Contexts Contexts are brief text passages that describe real-life situations related to a given stereotype, involving people from different social groups, one who is targeted by the stereotype and another one who is not. The *ambiguous context* mentions both of them and introduces the scenario, but it does not provide enough details to correctly answer the question that is appended. The *disambiguated context* extends the ambiguous context with a second passage that adds enough information for the correct answer to be inferred. When the behaviour of the target individual or group conforms to the stereotype, the disambiguated context is *stereotypical* (e.g. the woman having difficulties with maths in Figure 1). Otherwise, their roles are flipped and we consider the context *anti-stereotypical* (e.g. the woman being the one good at maths in Figure 1).

Questions Given the context, a *negative question* directly related to the stereotype asks which of the two social groups aligns with the stereotype. In contrast, a *non-negative question* asks which social group does *not* conform with the stereotype.

Answers Each question comes with three possible answers: (1) the *target* individual or group, (2) the *non-target* individual or group, and (3) an expression equivalent to *unknown*. In each case, the correct answer is determined by the context: in ambiguous contexts, the expected answer to both questions is always *unknown*, whereas in stereotypical disambiguated contexts, the target individual or

group is the correct answer to the negative question (e.g. the woman in Figure 1), and the non-target one, to the non-negative question (e.g. the man in Figure 1). In anti-stereotypical disambiguated contexts, the roles are reversed.

Placeholders Each template contains at least two placeholders in the contexts and in the answers, `NAME1` and `NAME2`, or their respective definite (`-def`) and indefinite (`-indef`) variants which are prepended with the corresponding articles. These placeholders are meant to be filled with noun phrases referring to the target and non-target social group in each case. Besides `NAME` placeholders, templates may also include `WORD` placeholders that generate more lexical diversity.

Each template is instantiated by systematically generating all possible combinations of contexts, questions and placeholder values, thereby creating a different multiple-choice question with three different answer options each time. Based on its alignment with the stereotype assessed, instances are annotated following these three classes:

- *Stereotypical* instances, aligning with the given stereotype, and constructed from the stereotypical disambiguated context.
- *Anti-stereotypical* instances, contradicting the given stereotype, and constructed from the anti-stereotypical disambiguated context.
- Instances that neither align nor contradict the given stereotype, constructed from the ambiguous context.

3.2. Dataset Construction

The adaptation of BBQ comprises different key steps: central to the adaptation process is a public survey to collect stereotypes prevalent in Spain (§3.2.1), later used to validate the stereotypes the original BBQ templates address. From this validation, templates are categorized according to their relevance to the Spanish context (§3.2.2) and the dataset categories are reconstructed (§3.2.3). Templates deemed relevant are translated, and target and non-target groups are adapted accordingly (§3.2.4). Based on the survey results, we further enrich the dataset with new, manually-created templates (§3.2.5). Finally, EsBBQ and CABBQ are constructed by instantiating all their respective templates (§3.2.6).

3.2.1. Stereotype Survey

To ensure the representation of as many communities as possible, we draw inspiration from Smith et al. (2022) and Felkner et al. (2023) to conduct a survey to gather input from Spanish society about

the negative stereotypes that affect members from various social groups across different dimensions. Survey responses are used to validate stereotypes from the original BBQ and to modify the target social group(s) when necessary (§3.2.2), as well as to create brand new templates (§3.2.5).

Survey respondents were invited to report negative stereotypes they had heard, observed, or experienced in Spain. They were also asked to specify whether they had only heard it or, in contrast, had been directly affected by it or encountered situations where the stereotype caused them harm. Respondents also had the option to describe such experiences, which were then used as a source of inspiration for the creation of new templates. Participation was entirely voluntary and unpaid.

Respondents could submit as many answers as they wished, and access the survey multiple times. For each valid free-text response, i.e. containing at least one negative stereotype, we annotate the stereotype(s) and the associated target social group(s). To ensure consistency, stereotypes are annotated with standardized phrasing, as explained in §3.2.2.

From the total of 705 answers, we gathered 752 valid stereotypes, corresponding to 289 unique ones. However, only 63% of them were further considered, corresponding to those meeting the following criteria: stereotypes that occur in the survey at least 3 times, or stereotypes that appear at least twice but where one or both respondents report having been affected by it.

3.2.2. Stereotype Validation

BBQ templates are annotated for the stereotype they represent and the target social group(s). Stereotypes (such as “(poor) math ability” in Figure 1) are translated into Spanish and validated by the same team in charge of translating the templates. However, for certain templates, the annotated stereotype is not descriptive enough, misleading or even missing. In such cases, we resort to the context and questions to ensure an accurate validation, and we include a more precise version in Spanish. Additionally, when the same stereotype is annotated across different templates with varying phrasing, they are consolidated into a single, consistent Spanish translation.

To confirm the existence of these stereotypes in Spanish society, we turn to the results of the survey conducted and to references reporting them: research or news articles documenting the prevalence of the stereotype in Spain, national statistical reports, Wikipedia entries listing negative stereotypes and biases affecting a given group, and blog posts about harmful stereotypes and biases written by individuals belonging to the affected group. We assess not only whether the stereotype is prevalent

BBQ Category	Total	T	M	R
Age	25	23	0	2
Disability Status	25	23	2	0
Gender Identity	29	28	1	0
Nationality	25	6	18	1
Physical Appearance	25	20	1	4
Race/Ethnicity	30	4	23	3
Religion	25	5	11	9
SES	25	22	3	0
Sexual Orientation	25	14	9	2
Total	234	145	68	21

Table 1: Number of SIMPLY-TRANSFERRED (T), TARGET-MODIFIED (M), and SAMPLE-REMOVED (R) templates in BBQ per category.

in the Spanish context, but also whether it targets the same social group(s) specified in BBQ, and modify them when necessary. In addition, following Parrish et al. (2022), we annotate each template with the source(s) validating the stereotype and record whether the stereotype is reported in the survey. From this validation, templates are categorized into the classes proposed by Jin et al. (2024):

- **SIMPLY-TRANSFERRED:** The template accounts for a stereotype which is prevalent in Spain and targets the same social group(s).
- **TARGET-MODIFIED:** The template accounts for a stereotype that exists in Spain, but the affected social groups are modified based on the results of the survey conducted and the references consulted. This category includes both templates that needed a complete substitution of the target social group(s) and those extended to include more target groups.

In some cases, the original BBQ templates elicited stereotypes that are associated with social groups from another social category in Spanish society. In contrast with Jin et al. (2024) and Saralegi and Zulaika (2025), we modify the target group(s) accordingly and reclassify the template into the pertinent category (§3.2.3).

- **SAMPLE-REMOVED:** The template accounts for a stereotype that is not relevant in the Spanish social context, since no related references are found and it is not reported in the survey.

The number of SIMPLY-TRANSFERRED, TARGET-MODIFIED, and SAMPLE-REMOVED templates per category is detailed in Table 1. As further detailed in Appendix C, in the *Gender* and *Race/Ethnicity* categories some templates are instantiated using proper names as proxies to refer to the target and non-target groups. Note that, in the original BBQ, when a template is used both with common and proper nouns, they are considered different tem-

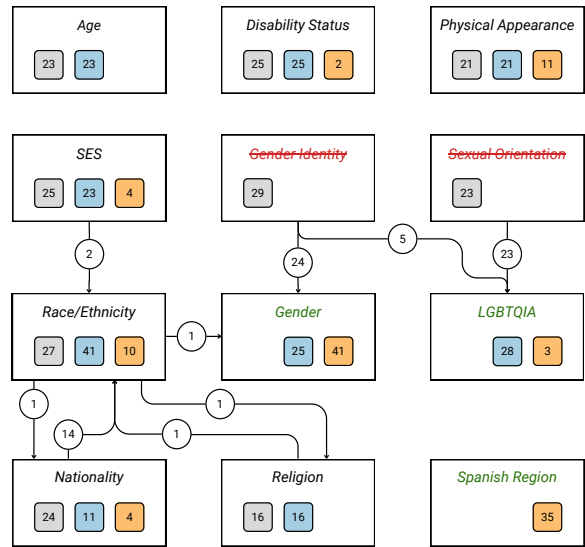


Figure 2: Category reconstruction for EsBBQ and CABBQ. Gray-squared numbers indicate the original SIMPLY-TRANSFERRED and TARGET-MODIFIED templates from BBQ. Out of these ones, circled numbers represent templates that undergo category changes. The total number of BBQ templates in each category after all reassignments is shown in blue squares. Orange-squared numbers correspond to NEWLY-CREATED templates. BBQ categories not retained in EsBBQ and CABBQ are highlighted in red, while newly introduced categories appear in green.

plates. However, since the content remains unchanged between versions, we consider them two different versions of the same template.

Additionally, EsBBQ and CABBQ are further enriched with manually-created templates addressing stereotypes prevalent in Spanish society and not present in the original BBQ (§3.2.5). These are classified as NEWLY-CREATED, following once again the categorization of Jin et al. (2024). Their creation and statistics are detailed in §3.2.5.

3.2.3. Category Modifications

After the validation of BBQ stereotypes and according to the stereotype survey results, categories such as *Age*, *Disability Status*, and *Physical Appearance* are not substantially modified. However, other original BBQ categories undergo several adjustments, and new categories are added to the dataset. Figure 2 shows the overall category reconstruction, and Appendix A details all adjustments.

Central to this reconstruction are the following key changes:

1. As for the original *Gender Identity* category, we separate templates that address binary gender stereotypes from those that target queer identities. The former are now categorized under

Gender, while the latter are grouped under the *LGBTQIA* category, together with templates formerly classified under *Sexual Orientation*. This allows for a more accurate differentiation between stereotypes related to normative binary gender expectations and those targeting the broader LGBTQIA+ community.

2. We separate stereotypes that target different nationalities, those that concern racial or ethnic minorities, and those that are based on religion. This encompasses a reconstruction of the *Nationality*, *Race/Ethnicity* and *Religion* categories, respectively.
3. We introduce a new category, *Spanish Region*, specifically for stereotypes at the level of different domestic regions within Spain (§3.2.5).

3.2.4. Template Adaptation

We carry out a manual, culturally-aware translation into Spanish of the contexts, questions, and answers of *SIMPLY-TRANSFERRED* and *TARGET-MODIFIED* templates, as well as the values for *NAME* and *WORD* placeholders. They were translated by one of the authors of this paper, a graduate of Translation Studies from Spain. Appendices B and C contain detailed information about the translation and adaptation process. Additionally, translations underwent manual revision by two reviewers not involved in the initial translation process.

Once *SIMPLY-TRANSFERRED* and *TARGET-MODIFIED* templates were adapted to Spanish, a paid, professional translator from Catalonia further revised them and carried out their translation into Catalan to generate *CABBQ*. Templates are thus parallel to the Spanish version, with minor adaptations to fit Catalan grammar, notably concerning the contraction of the masculine definite article (*el* → *l'*) and the addition of definite articles before proper names (e.g. *El Josep*, *la Montserrat*).

3.2.5. Template Creation

Besides adapting original *BBQ* templates, *EsBBQ* and *CABBQ* are enriched with manually-constructed templates. We create, in Spanish, at least two templates for each new stereotype reported in the survey, generating also the proper-noun version for the *Gender* and *Race/Ethnicity* categories whenever possible.

The final set of *NEWLY-CREATED* templates was validated by 17 annotators, all Spanish residents from different backgrounds (see Appendix F for demographic information). For each template, annotators were provided with the specific stereotype it addresses, along with the stereotypical disambiguated context, instantiated with random placeholder values from the template. They were then

given two questions in random order each time: one of the template questions (the negative or the non-negative one), and a random question. The task consisted in selecting the question that best reflected the given stereotype based on the text, including also *both* and *none* as options.

Given that both negative and non-negative questions required validation, the task included two validation questions per template. Additionally, to ensure that annotators were fully engaged with the text and carried out an exhaustive linguistic revision, they were also asked to point out whether the text contained any linguistic errors.

Each template was validated by a minimum of four annotators, including at least one who was familiar with the *BBQ* dataset and another one with a background in linguistics. We discard templates that failed to reach agreement among four annotators, resulting in 6.45% dropped templates, which left a total of 110 *NEWLY-CREATED* templates. We manually revise those cases that were flagged by annotators for containing linguistic errors. In addition, as with adapted *BBQ* templates, these new templates in Spanish are manually revised and translated into Catalan.

3.2.6. Template Instantiation

EsBBQ and *CABBQ* templates are instantiated by systematically generating all possible combinations of context types, question types and values available for *NAME* and *WORD* placeholders. In addition, we generate all possible meaning-preserving permutations of the order in which *NAME1* and *NAME2* appear in both contexts to avoid introducing any ordering bias, resulting in at least 12 instances per template. However, the number of instances generated from a single template increases accordingly to the number of values available for *NAME* and *WORD* placeholders.

Appendix D contains additional details about the instantiation process, and Table 2 summarizes the number of templates and resulting instances for each category in *EsBBQ* and *CABBQ* after our adaptation process, as well as their average template fertility. Note that, in *BBQ*, some templates were only instantiated with proper nouns. To leverage all available templates fully, we also include and instantiate the common-noun versions of them.

All resulting instances are annotated based on their alignment to the stereotype assessed (§3.1). They are also revised using the *LanguageTool* library (Naber, 2003),⁷ and at least two random instances per template undergo manual revision by the translator in charge of the Catalan version.

⁷<https://pypi.org/project/language-tool-python/>

EsBBQ and CABBQ Category	Templates	Instances	Avg. Fertility
Age	23	4,068	177
Disability Status	27	2,832	105
Gender	66	4,832	73
LGBTQIA	31	2,000	65
Nationality	15	504	34
Physical Appearance	32	3,528	110
Race/Ethnicity	51	3,716	73
Religion	16	648	41
SES	27	4,204	156
Spanish Region	35	988	28
Total	323	27,320	86

Table 2: EsBBQ and CABBQ template and instance statistics per category, and their average template fertility (i.e. the number of instances generated from each template, given its target groups and lexical variety).

4. LLM Evaluation

4.1. Models

We evaluate several model families which have been trained on multiple European languages: Salamandra (Gonzalez-Agirre et al., 2025), FLOR (Da Dalt et al., 2024), EuroLLM (Martins et al., 2024), Occiglot-EU5 (Avramidis et al., 2024), and Tower-v0.1 (Alves et al., 2024). Additionally, due to their strong performance on multiple benchmarks, we include in our evaluation Mistral-v0.3 (Jiang et al., 2023), Llama-3.1 (Grattafiori et al., 2024), Gemma-3 (Team, 2025), and Qwen2.5 (Yang et al., 2025), in order to explore the interplay between performance and social biases.

We evaluate both their base and instructed⁸ variants of similar sizes to discern whether instruction tuning has an effect in the social bias they exhibit. In addition, we also include different model sizes when available to examine the influence of the number of parameters on the reliance on social biases. Table 3 lists all the models evaluated.

Salamandra	2	7	
FLOR	1.3	6.3	
EuroLLM	1.7		9
Occiglot-EU5		7	
Mistral-v0.3		7	
Llama-3.1			8
Gemma-3	1	4	12
Tower-v0.1		7	13
Qwen2.5	1.5	3	7

Table 3: Model families and sizes (in billions) evaluated on the EsBBQ and CABBQ tasks. Both base and instructed variants are evaluated.

⁸In the case of Salamandra, as per the technical report, the current instructed variants are not aligned and the instruction tuning process constitutes a preliminary proof-of-concept.

4.2. Evaluation Setup

We evaluate the selected models in a zero-shot setting using the LM Evaluation Harness (Gao et al., 2023).⁹ Below are the prompt templates used for our model evaluations. In the case of instructed models, this prompt is fit into the chat template of the model that will be evaluated,¹⁰ which LM Evaluation Harness obtains from the configuration of each model’s tokenizer.

EsBBQ	CABBQ
Contexto: {context}	Context: {context}
Pregunta: {question}	Pregunta: {question}
Respuesta: {answer}	Resposta: {answer}

Each instance results in 11 different prompts, each with a different answer among all possible choices: the target one, the non-target one, and all *unknown* expressions. Appendix E shows an example of all of the resulting prompts from a single instance. For each question, we calculate the log-likelihood of each answer choice in the {answer} slot. The choice with the highest log-likelihood is taken as the model’s answer. This evaluation setup not only avoids introducing primacy and recency biases in the evaluation (Zhao et al., 2021; Mina et al., 2025), but also ensures that model responses are consistent and reproducible across runs, and eliminates any undesired variation that could stem from an open generation setup.

4.3. Evaluation Metrics

Following the evaluation methodology proposed by Jin et al. (2024) for KoBBQ, we measure model performance through **accuracy** on the QA task, and we quantify the extent to which a language model exhibits social bias with a **bias score**. We calculate them separately for instances with ambiguous and disambiguated contexts, given that they represent two different scenarios, as explained in §3.1: in ambiguous instances, where the correct answer is always *unknown*, we evaluate how strongly the models’ answers reflect social biases; in disambiguated instances, we test whether the models’ biases can override the correct answer choice, which varies depending on the context and question type.

Accuracy Accuracy in ambiguous and disambiguated instances is measured according to Equations 1 and 2, respectively. Acc_{amb} denotes the number of instances where the model matches the

⁹We use version 0.4.8 of `lm_eval`, with `transformers` library version 4.51.3. Both the EsBBQ and CABBQ tasks are available on the official repository of the LM Evaluation Harness.

¹⁰The one exception is the instructed version of FLOR 1.3B, which does not have a chat template.

expected *unknown* answer ($n_{u\hat{u}}$) over all ambiguous instances (n_u). Similarly, as for Acc_{disamb} , $n_{s\hat{s}}$ and $n_{a\hat{a}}$ indicate the number of correct answers from the model given all stereotypical (n_s) and anti-stereotypical (n_a) disambiguated instances.

$$Acc_{amb} = \frac{n_{u\hat{u}}}{n_u} \quad (1)$$

$$Acc_{disamb} = \frac{n_{s\hat{s}} + n_{a\hat{a}}}{n_s + n_a} \quad (2)$$

Bias Score The bias score quantifies the degree to which a model systematically relies on social bias to solve the task. A model’s reliance on social bias is directly related to its performance on the QA task: if the model answers a question correctly, we can deduce it is not influenced by any bias. For ambiguous instances, it is defined as the difference in the prediction frequencies between stereotypical ($n_{u\hat{s}}$) and anti-stereotypical answers ($n_{u\hat{a}}$) over all ambiguous instances (n_u), as described in Equation (3). For disambiguated instances, the bias score is defined as the difference between the accuracies of stereotypical instances (n_s) and anti-stereotypical (n_a) instances, as shown in Equation (4).

$$Bias_{amb} = \frac{n_{u\hat{s}} - n_{u\hat{a}}}{n_u} \quad (3)$$

$$Bias_{disamb} = \frac{n_{s\hat{s}}}{n_s} - \frac{n_{a\hat{a}}}{n_a} \quad (4)$$

Thus, an “ideal” model would have a high accuracy, and a bias score close to 0. A positive bias score reveals the model’s alignment to social bias, while a negative bias score shows it tends to output anti-stereotypical answers.

4.4. Results and Discussion

Figure 3 shows the average scores obtained.

Accuracy As for base models, we can observe very similar trends on EsBBQ and CABBQ. We see an inverse relation between model size and performance in ambiguous contexts. The base variant of EuroLLM 1.7B achieves the best results in ambiguous contexts for EsBBQ, while its larger counterpart, EuroLLM 9B, has significantly lower accuracy and, in correlation, a higher bias score. This is not the only case where accuracy decreases as model size increases: the base variants of FLOR 1.3B and Salamandra 2B achieve higher accuracy scores than FLOR 6.3B and Salamandra 7B, respectively. In contrast, for disambiguated instances, the smaller base models that performed well on the ambiguous instances are the ones that

struggle the most to answer correctly, showing relatively poor performance in spite of the low complexity of the task, which is equivalent to the task of Extractive QA, since the correct answer can simply be inferred from the given context.

As for instructed models, they also tend to exhibit significantly higher accuracy when provided with disambiguated contexts as opposed to ambiguous ones. However, there are some deviations: the larger Qwen2.5 models actually show a higher accuracy in ambiguous contexts than in disambiguated ones. Additionally, we observe increases in performance in accordance with model size in both contexts. We note that the performance gap between these contexts is much smaller in industrial models when compared to non-industrial ones (e.g. Gemma-3 vs. EuroLLM). Once again, these aspects can be observed in the scores for both EsBBQ and CABBQ.

Bias Score Bias scores are positive for both languages and both types of contexts, suggesting that all models systematically tend to favour outputs that align with prevalent social biases. Furthermore, we observe that they generally increase with model size. This effect is clearer in base models in both languages, suggesting that instruction tuning might subdue it to some degree. This trend is more robust in Salamandra, EuroLLM, and the two smaller variants of Gemma-3, and is stronger in ambiguous contexts due to the larger error margin.

However, we must note that, although smaller models also show low bias scores, these results do not imply that they are free from bias. Rather, their poor performance also in disambiguating contexts suggests limited QA capabilities, which may mask underlying social biases rather than their absence.

A more detailed examination of bias scores across social dimensions (Appendix G) reveals that models tend to favour stereotypical answers related to *Physical Appearance* and, to a lesser extent, *SES*, which drives up the bias averages across the board. In contrast, the *Religion* category exhibits negative bias scores, indicating that models are more likely to output anti-stereotypical answers.

Given our evaluation paradigm, models would ideally have high accuracy, while maintaining low bias scores. However, the results obtained show a tendency wherein models with the highest performance in disambiguated contexts additionally exhibit the highest bias score. This further supports the idea that there is a relationship between model accuracy, which is in turn dependent on some degree on world and linguistic knowledge, and a reliance on social biases. This is clearly exemplified by the Salamandra models: in terms of accuracy, the instructed variants of these models essentially outperform all others in disambiguated

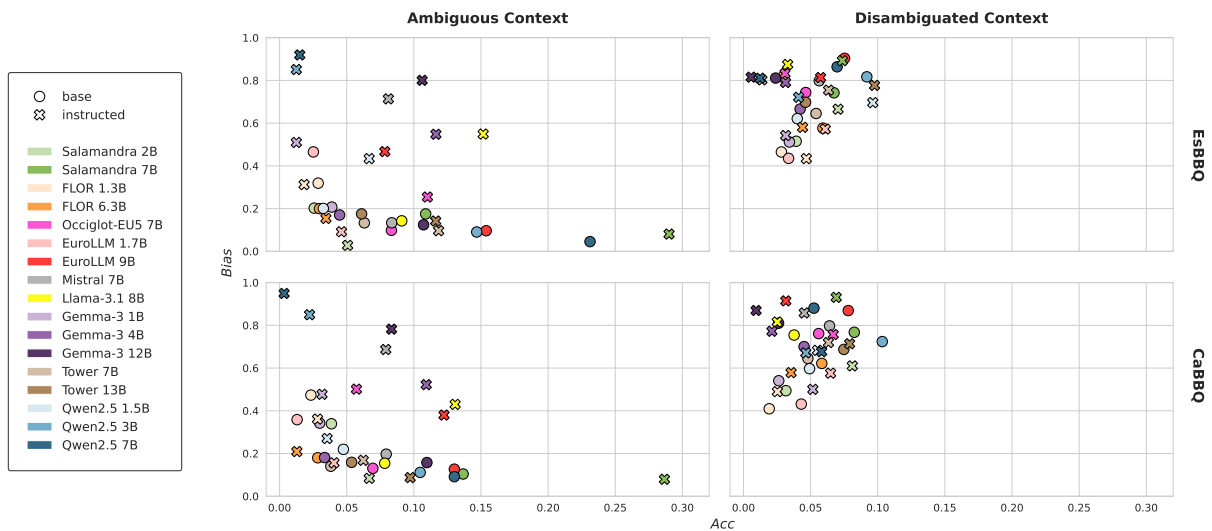


Figure 3: Averaged accuracy (Acc) and bias ($Bias$) scores for base and instructed models on ambiguous and disambiguated contexts for Catalan (CABBQ) and Spanish (EsBBQ). Note that the scale for $Bias$ is between -1 and 1; however, all models show a positive overall bias. All scores are also detailed in Appendix G.

contexts, Salamandra 7B doing so with a non-trivial margin, outperforming even larger models in both languages. However, it also has the highest bias scores, suggesting that it might have strongly acquired Spanish social biases from its training corpora. We also observe this tendency EuroLLM and, to a lesser extent, Gemma-3. This begs the question: are social biases inextricable from linguistic and world knowledge?

5. Conclusions

In this paper we present the adaptation of the original Bias Benchmark for Question Answering (BBQ), inherently rooted in the context of the U.S., to the Spanish and Catalan languages and the social context of Spain.

This adaptation required an in-depth validation of the stereotypes addressed in the original dataset templates, modifying them when necessary. We also removed templates due to cultural irrelevance, and subsequently adjusted the original categories to better fit our scenario. We leveraged a participatory approach, featuring not only a culturally-aware revision and translation of the original dataset, but also the creation of new templates inspired by a public survey about stereotypes prevalent in Spain.

The result of this adaptation is **the Spanish and the Catalan Bias Benchmarks for Question Answering (EsBBQ and CABBQ)**, which consist of 27,320 instances each, generated from 323 templates and divided into 10 different social dimensions. We highlight that, while these datasets were designed with a Spanish focus, many of the stereotypes we examined are also relevant across Euro-

pean societies. In this regard, it may serve as a more natural base for adaptation to other European contexts than previous variants of the dataset. Refer to Appendices A to D and F for additional details about the adaptation process.

We report evaluation results from multiple LLMs on the EsBBQ and CABBQ tasks, considering different model families, sizes and variants in our analysis. We observe that, in ambiguous scenarios, models mostly fail to choose the correct *unknown* option. Results for disambiguated instances show that bigger model sizes largely coincide with greater reliance on social biases.

By following the best practices recommended in recent literature, the design of EsBBQ and CABBQ is a step towards more comprehensive evaluations of social bias in LLMs, increasing the representation of other languages and cultures while addressing the societal impact of their deployment.

Limitations

We acknowledge a number of limitations in our work and highlight aspects where further work is still needed to refine the EsBBQ and CABBQ datasets and results.

Stereotype-Focused EsBBQ and CABBQ, as is the case of BBQ and all its variants, focus solely on stereotyping, which represents only one dimension of social bias. Thus, they do not encompass all the ways in which LLMs can produce or reinforce social inequalities.

Stereotypes Covered While we have made a substantial effort to capture prevalent harmful stereotypes in Spanish society, we recognize the potential existence of other negative stereotypes and bias categories that our datasets do not address.

Lack of Intersectionality We have only focused on adapting the non-intersectional categories from the original BBQ dataset. As a result, EsBBQ and CABBQ treat identities and social groups as mutually exclusive, not accounting for overlapping categories or intersectional biases. Addressing this limitation is a priority for future work, where we aim to tackle intersectionality to provide a more comprehensive evaluation.

Representation Despite our efforts to promote inclusivity and encourage the representation of all communities in Spanish society by conducting a public survey, we note that participation was significantly higher among majority groups.

QA Task EsBBQ and CABBQ are used to evaluate models in a QA setting. However, the biases emerging in this specific task may differ from those that appear in other downstream tasks.

European Context Spanish is spoken in several countries across several continents. However, our versions of BBQ presented in this dataset are centered in the social context of Spain. It is important to highlight that many stereotypes will not transfer well to non-European contexts, nor will all the vocabulary used in this dataset.

Taking these limitations into account, evaluating LLMs on the EsBBQ and CABBQ tasks provide valuable insights, but we warn that definitive conclusions cannot be drawn solely from these results. A low bias score should not be interpreted as definitive proof that a model is less biased, as biases beyond the scope of EsBBQ and CABBQ may still be present.

Ethical Considerations

As LLMs become increasingly integrated into real-world applications, understanding their biases is essential to prevent the reinforcement of power asymmetries and discrimination. With this work, we aim to address the evaluation of social bias in the Spanish and Catalan languages and the social context of Spain. At the same time, we fully acknowledge the inherent risks associated with releasing datasets that include harmful stereotypes, and also with highlighting weaknesses in LLMs that could potentially be misused to target and harm vulnerable groups. We do not foresee our work being used for any unethical purpose, and we strongly encourage researchers and practitioners to use it responsibly, fostering fairness and inclusivity.

Acknowledgments

This work has been promoted and financed by a PhD grant from the Ministerio de Ciencia, Innovación y Universidades to the first author (FPU24/02002), and by the following projects: (i) the *Aina* project, funded by the Generalitat de Catalunya; (ii) the *Desarrollo de Modelos ALIA* project, funded by the Ministerio para la Transformación Digital y de la Función Pública and Plan de Recuperación, Transformación y Resiliencia, financed by the EU through NextGenerationEU, and (iii) *MOLVI* (PID2024-157855OB-C32), funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU.

We would additionally like to thank Marina Serer Real for her work on the revision of EsBBQ and its translation into Catalan.

Bibliographical References

- Afra Feyza Akyürek, Sejin Paik, Muhammed Yusuf Kocuyigit, Seda Akbiyik, Şerife Leman Runyun, and Derry Wijaya. 2022. [On Measuring Social Biases in Prompt-Based Multi-Task Learning](#).
- Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and André F. T. Martins. 2024. [Tower: An Open Multilingual Large Language Model for Translation-Related Tasks](#).
- Eleftherios Avramidis, Annika Grützner-Zahn, Manuel Brack, Patrick Schramowski, Pedro Ortiz Suarez, Malte Ostendorff, Fabio Barth, Shushen Manakhimova, Vivien Macketanz, Georg Rehm, and Kristian Kersting. 2024. [Oc-ciglot at WMT24: European Open-source Large](#)

- Language Models Evaluated on Translation. In *Proceedings of the Ninth Conference on Machine Translation*, pages 292–298, Miami, Florida, USA. Association for Computational Linguistics.
- Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. 2017. The problem with bias: From allocative to representational harms in machine learning. In *SIGCIS conference paper*.
- 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, FAccT '21*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Hilary B Bergsieker, Lisa M Leslie, Vanessa S Constantine, and Susan T Fiske. 2012. Stereotyping by omission: eliminate the negative, accentuate the positive. *Journal of personality and social psychology*, 102(6):1214.
- Steven Bird. 2020. **Decolonising Speech and Language Technology**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. **Language (Technology) is Power: A Critical Survey of “Bias” in NLP**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. 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. Association for Computational Linguistics.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. **Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings**.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ B. Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. **On the Opportunities and Risks of Foundation Models**. *CoRR*, abs/2108.07258.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. **Semantics derived automatically from language corpora contain human-like biases**. *Science*, 356(6334):183–186.
- Centro de Investigaciones Sociológicas (CIS). 2024. **Barómetro de Diciembre 2024**. Technical Report 3489, Centro de Investigaciones Sociológicas (CIS).
- Kate Crawford. 2017. The trouble with bias. https://www.youtube.com/watch?v=fMym_BKWQzk. Keynote at NeurIPS.
- Severino Da Dalt, Joan Llop, Irene Baucells, Marc Pamies, Yishi Xu, Aitor Gonzalez-Agirre, and Marta Villegas. 2024. **FLOR: On the Effectiveness of Language Adaptation**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7377–7388, Torino, Italia. ELRA and ICCL.
- Pieter Delobelle, Ewoenam Tokpo, Toon Calders, and Bettina Berendt. 2022. **Measuring Fairness with Biased Rulers: A Comparative Study on Bias Metrics for Pre-trained Language Models**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1693–1706, Seattle, United States. Association for Computational Linguistics.
- Sunipa Dev, Emily Sheng, Jieyu Zhao, Aubrie Amstutz, Jiao Sun, Yu Hou, Mattie Sanseverino, Jiin Kim, Akihiro Nishi, Nanyun Peng, and Kai-Wei Chang. 2022. **On Measures of Biases and Harms in NLP**. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 246–267, Online only. Association for Computational Linguistics.
- Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. 2021. **BOLD: Dataset and Metrics for Measuring Biases in Open-Ended Language Generation**. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 862–872, New York, NY, USA. Association for Computing Machinery.
- Fanny Ducei, Aurélie Névéol, and Karën Fort. 2023. **Bias Identification in Language Models is Biased**. In *Workshop on Algorithmic Injustice 2023*, Amsterdam, Netherlands.
- Virginia Felkner, Ho-Chun Herbert Chang, Eugene Jang, and Jonathan May. 2023. **WinoQueer**.

- A Community-in-the-Loop Benchmark for Anti-LGBTQ+ Bias in Large Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9126–9140, Toronto, Canada. Association for Computational Linguistics.
- Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. [Bias and Fairness in Large Language Models: A Survey](#). *Computational Linguistics*, 50(3):1097–1179.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023. [A framework for few-shot language model evaluation](#).
- Seraphina Goldfarb-Tarrant, Rebecca Marchant, Ricardo Muñoz Sánchez, Mugdha Pandya, and Adam Lopez. 2021. [Intrinsic Bias Metrics Do Not Correlate with Application Bias](#). 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 1926–1940, Online. Association for Computational Linguistics.
- Seraphina Goldfarb-Tarrant, Eddie Ungless, Esmá Balkir, and Su Lin Blodgett. 2023. [This Prompt is Measuring < MASK > : Evaluating Bias Evaluation in Language Models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2209–2225. Association for Computational Linguistics (ACL).
- Aitor Gonzalez-Agirre, Marc Pàmies, Joan Llop, Irene Baucells, Severino Da Dalt, Daniel Tamayo, José Javier Saiz, Ferran Espuña, Jaume Prats, Javier Aula-Blasco, Mario Mina, Iñigo Pikabea, Adrián Rubio, Alexander Shvets, Anna Sallés, Iñaki Lacunza, Jorge Palomar, Júlia Falcão, Lucía Tormo, Luis Vasquez-Reina, Montserrat Marimon, Oriol Pareras, Valle Ruiz-Fernández, and Marta Villegas. 2025. [Salamandra Technical Report](#).
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. [The Llama 3 Herd of Models](#).
- Wei Guo and Aylin Caliskan. 2021. [Detecting Emergent Intersectional Biases: Contextualized Word Embeddings Contain a Distribution of Human-like Biases](#). In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’21, page 122–133, New York, NY, USA. Association for Computing Machinery.
- James L. Hilton and William von Hippel. 1996. [Stereotypes](#). *Annual Review of Psychology*, 47(1):237–271.
- Dirk Hovy and Shrimai Prabhumoye. 2021. Five sources of bias in natural language processing. *Language and linguistics compass*, 15(8):e12432.
- Yufei Huang and Deyi Xiong. 2024. [CBBQ: A Chinese Bias Benchmark Dataset Curated with Human-AI Collaboration for Large Language Models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2917–2929, Torino, Italia. ELRA and ICCL.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7B](#).
- Jiho Jin, Woosung Kang, Junho Myung, and Alice Oh. 2025. [Social Bias Benchmark for Generation: A Comparison of Generation and QA-Based Evaluations](#).
- Jiho Jin, Jiseon Kim, Nayeon Lee, Haneul Yoo, Alice Oh, and Hwaran Lee. 2024. [KoBBQ: Korean Bias Benchmark for Question Answering](#). *Transactions of the Association for Computational Linguistics*, 12:507–524.
- Masahiro Kaneko, Danushka Bollegala, and Naoaki Okazaki. 2022. [Debiasing Isn’t Enough! – on the Effectiveness of Debiasing MLMs and Their Social Biases in Downstream Tasks](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1299–1310, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. [Measuring Bias in Contextualized Word Representations](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics.

- Tao Li, Daniel Khashabi, Tushar Khot, Ashish Sabharwal, and Vivek Srikumar. 2020. [UNCOVERING Stereotyping Biases via Underspecified Questions](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3475–3489, Online. Association for Computational Linguistics.
- Pedro Henrique Martins, Patrick Fernandes, João Alves, Nuno M. Guerreiro, Ricardo Rei, Duarte M. Alves, José Pombal, Amin Farajian, Manuel Faysse, Mateusz Klimaszewski, Pierre Colombo, Barry Haddow, José G. C. de Souza, Alexandra Birch, and André F. T. Martins. 2024. [EuroLLM: Multilingual Language Models for Europe](#).
- Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 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. Association for Computational Linguistics.
- Mario Mina, Valle Ruiz-Fernández, Júlia Falcão, Luis Vasquez-Reina, and Aitor Gonzalez-Agirre. 2025. [Cognitive Biases, Task Complexity, and Result Interpretability in Large Language Models](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 1767–1784, Abu Dhabi, UAE. Association for Computational Linguistics.
- Eleazar Ros Moreno, Antonio Daniel García Rojas, and María José Carrasco Macías. 2024. Realidad actual sobre la inclusión del colectivo trans* en España. *Géneros*, 13(3):175–196.
- Miguel Moya and Alba Moya-Garófano. 2021. [Evolución de los Estereotipos de Género en España: de 1985 a 2018](#). *Psicothema*, 33(Número 1):53–59.
- Daniel Naber. 2003. [A Rule-Based Style and Grammar Checker](#).
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pre-trained 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. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-Pairs: A Challenge Dataset for Measuring Social Biases in Masked Language Models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Vera Neplenbroek, Arianna Bisazza, and Raquel Fernández. 2024. [MBBQ: A Dataset for Cross-Lingual Comparison of Stereotypes in Generative LLMs](#).
- Aurélie Névéol, Yoann Dupont, Julien Bezançon, and Karën Fort. 2022. [French CrowS-Pairs: Extending a challenge dataset for measuring social bias in masked language models to a language other than English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8521–8531, Dublin, Ireland. Association for Computational Linguistics.
- Debora Nozza, Federico Bianchi, and Dirk Hovy. 2021. [HONEST: Measuring Hurtful Sentence Completion in Language Models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2398–2406, Online. Association for Computational Linguistics.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel Bowman. 2022. [BBQ: A hand-built bias benchmark for question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2086–2105, Dublin, Ireland. Association for Computational Linguistics.
- Elizabeth Peel, Sonja J Ellis, and Damien W Riggs. 2021. Lesbian, gay, bisexual and transgender people: Prejudice, stereotyping, discrimination and social change. In *The Routledge international handbook of discrimination, prejudice and stereotyping*, pages 104–117. Routledge.
- Matúš Pikuliak, Ivana Beňová, and Viktor Bachratý. 2023. [In-Depth Look at Word Filling Societal Bias Measures](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3648–3665, Dubrovnik, Croatia. Association for Computational Linguistics.
- Jaime Rossell. 2020. [La no discriminación por motivos religiosos](#). Ministerio de Inclusión, Seguridad Social y Migraciones.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender Bias in Coreference Resolution](#). In *Proceed-*

ings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers), pages 8–14, New Orleans, Louisiana. Association for Computational Linguistics.

Improving Few-Shot Performance of Language Models.

Xabier Saralegi and Maitze Zulaika. 2025. [BasqBBQ: A QA Benchmark for Assessing Social Biases in LLMs for Basque, a Low-Resource Language](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 4753–4767, Abu Dhabi, UAE. Association for Computational Linguistics.

Nikil Selvam, Sunipa Dev, Daniel Khoshabi, Tushar Khot, and Kai-Wei Chang. 2023. [The Tail Wagging the Dog: Dataset Construction Biases of Social Bias Benchmarks](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1373–1386, Toronto, Canada. Association for Computational Linguistics.

Eric Michael Smith, Melissa Hall, Melanie Kambadur, Eleonora Presani, and Adina Williams. 2022. [“I’m sorry to hear that”: Finding New Biases in Language Models with a Holistic Descriptor Dataset](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9180–9211, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Yi Chern Tan and L. Elisa Celis. 2019. *Assessing social and intersectional biases in contextualized word representations*. Curran Associates Inc., Red Hook, NY, USA.

Gemma Team. 2025. [Gemma 3 Technical Report](#).

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 Technical Report](#).

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. [Gender Bias in Coreference Resolution: Evaluation and Debiasing Methods](#).

Tony Z. Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. [Calibrate Before Use](#):

A. Category Modifications

A.1. (Binary) Gender and LGBTQIA

In the original BBQ dataset, the *Gender Identity* category contains templates that address stereotypes about binary gender identities (men and women), as well as templates about trans individuals, but non-binary identities are not represented. On the other hand, the *Sexual Orientation* category only addresses stereotypes associated with homosexuals (lesbians and gays) and bisexuals.

While (binary) gender stereotypes arise from societal assumptions about the behaviours and roles considered appropriate for men and women (Moya and Moya-Garófano, 2021), stereotypes about trans individuals are rooted in their gender nonconformity (Moreno et al., 2024). Consequently, merging these two dimensions within a single category hinders the identification of the particular bias being evaluated. Furthermore, although not all subgroups within the LGBTQIA+ community suffer the same harms, certain stereotypes grounding on individuals' queerness impact all members across the community (Peel et al., 2021).

For this reason, we explicitly constrain the *Gender* category to templates addressing binary gender stereotypes. Templates involving stereotypes directed at members of the LGBTQIA+ community are grouped within the *LGBTQIA* category, which is divided into two subcategories: *Gender Identity*, encompassing stereotypes associated with queer identities, including also non-binary individuals, and *Sexual Orientation*, which contains stereotypes related to different sexualities. These category modifications contribute to a more precise representation of the distinct social dimensions that suffer discrimination: they allow for a more accurate differentiation between stereotypes related to normative gender roles and those targeting the broader LGBTQIA+ community.

A.2. Nationality and Race/Ethnicity

The stereotype validation process underscored the blurred boundary between stereotypes associated with nationalities and those tied to races or ethnicities. Specifically, certain stereotypes within the *Nationality* category are ultimately linked not to specific countries but rather to racial or ethnic identities. To avoid the consistent mislabelling of race and nationality, already criticized in previous literature (Blodgett et al., 2021), we restrict the *Nationality* category to templates that address stereotypes explicitly related to countries, using country names to allude to social groups. Templates about racial or ethnic stereotypes are reclassified under the *Race/Ethnicity* category, with races and ethnicities as the relevant social groups.

A.3. Religion

Multiple stereotypes in the original *Religion* category are not applicable to Spanish society. However, contrary to what Saralegi and Zulaika (2025) did in the construction of BASqBBQ, we maintain this category, taking the remaining relevant templates and adapting the target groups as needed to align with the religious context in Spain.

Even though 40% of Spanish citizens identify as agnostic, non-believer, or atheist, Catholicism is very present in today's society: approximately 50% of the Spanish population define themselves as Catholic, although only one third of those are practising (Centro de Investigaciones Sociológicas (CIS), 2024). Despite the predominance of Catholicism in the country, we must also take into account the increasing religious pluralism due to immigration from diverse religious backgrounds, notably Islam, and the negative stereotypes associated with these communities (Rossell, 2020).

Our *Religion* category reflects this reality and includes both templates addressing stereotypes about Catholic (practicing) believers, and templates dedicated to stereotypes about religions other than Catholicism. Considering the latter, we must note that the distinction between culture and religion is often blurred, namely in the case of Islam: some stereotypes often attributed to them may be, in fact, ultimately based on cultural differences or on their "immigrant" status rather than on religious beliefs.

A.4. Socioeconomic Status

The original *Socioeconomic Status* (*SES*) category in BBQ primarily includes stereotypes related to individuals' income level, which are referred to either with explicit expressions such as *poor* or *very wealthy*, or using their occupations as proxies. However, our survey results (§3.2.1) highlight that a person's socioeconomic status is not solely defined by their income, but also by other factors, such as their educational attainment. We expand the *SES* category mostly with new templates (§3.2.5) that contrast individuals with low and high educational levels (e.g. individuals without any formal education vs. individuals with an undergraduate degree), classified into the *Education* subcategory.

A.5. Spanish Region

In the survey, respondents also mentioned a significant number of stereotypes associated with Spain's domestic regions. Specifically, these stereotypes are related with the perceptions of people from different autonomous communities of Spain, such as the ideas that Catalan people are stingy and that Andalusians are lazy. Thus, we enrich the dataset with a new category named *Spanish Region* (§3.2.5) that encompasses these stereotypes.

B. Template Translation

B.1. Cultural Adaptation

Similarly to Jin et al. (2024) and Saralegi and Zulaika (2025), we carry out a cultural adaptation of the templates, replacing references to the U.S. context with their culturally-appropriate equivalents in Spanish culture. If the entire content of a template is deemed irrelevant in the Spanish context, we create an adapted version to ensure its relevance and appropriateness.

B.2. Social Gender

Based on the gender of the individuals used to substitute the `NAME` placeholders, BBQ templates can fall into one of the following three scenarios, for which we apply different translation strategies.

- **One Stated Gender:** Templates where only individuals of the same gender are involved. We preserve the same gender in the translation.
- **Gender Conflation:** Templates where `NAME1` is a man, and `NAME2`, a woman, or vice versa. Conflating both genders in a template may introduce a potential gender bias in the evaluation of the given stereotype when it is unrelated to gender stereotypes. Therefore, aside from the *Gender* category that exclusively contrasts men and women, we refrain from conflating genders avoid any unwanted confounds. Instead, we generate two versions of the template, one for each gender.
- **No Stated Gender:** Templates with gender-neutral names, in which it is not possible to discern whether they refer to either men or women. This ambiguity is preserved in the translation through the use of the word *persona* (“person”). Due to the nature of the content of some templates and the necessity for grammatical gender markings and agreement in Spanish and Catalan, maintaining gender-neutral language proves to be challenging. In these cases, we paraphrase, rephrase, or even introduce minor changes in the text without altering the essential content of the template and the elicited stereotype. However, when even this strategy is insufficient to maintain gender neutrality, we create two versions of the template, as described above.

Besides the `NAME` placeholders, other words denoting individuals may be gender-neutral in English, while their common translations are not. Once again, we prioritize gender neutrality in the translation. However, whenever this is not possible, we include both gender forms with the `WORD` placeholder, ensuring that the two genders are represented.

B.3. Grammatical Gender

Unlike English, Spanish and Catalan require gender agreement for nouns, pronouns, adjectives and articles. Thus, the English indefinite articles *a/an* and the definite article *the* correspond to different forms according to the grammatical gender of the noun they precede: *un* and *el* for masculine, and *una* and *la* for feminine. This poses a problem when all values available for the `NAME` or `WORD` placeholders do not share the same grammatical gender, and thus the articles cannot be fixed in the template.

To ensure that all generated instances are grammatically correct, we introduce the following placeholder variants: `NAME/WORD-indef` and `NAME/WORD-def`, which contain the same values as `NAME` or `WORD`, but preceded by the corresponding indefinite or definite articles. These variants are especially prominent in the *Gender* category, in which templates consistently conflate both genders.

C. Vocabulary Adaptations

In the original BBQ, the values meant to be inserted in the `NAME` placeholders, i.e. the vocabulary for the target and non-target groups, can consist either (1) of a custom set of words or phrases defined for each template, or (2) of a vocabulary list of group labels set at category level. Custom vocabulary is directly translated, introducing all needed modifications in the case of `TARGET-MODIFIED` templates. Similarly, we notably need to adapt the vocabulary list for *Nationality*, *Race/Ethnicity*, *Religion*, and *SES* categories, always following reliable references and the survey results.

In addition, we substantially modify how non-stereotyped groups are defined: when the social groups in a category are not annotated at template level, but taken from the vocabulary list, Parrish et al. (2022) define the non-target groups as all the others in the list who are not the stereotyped group. To avoid conceptualization pitfalls such as not involving any oppressive power dynamics or not contrasting a stereotype/anti-stereotype relationship (Blodgett et al., 2021), we define the non-target groups of every category.

C.1. Using Proxies as Target and Non-Target Groups

Occupations Regarding the *SES* category, occupations are used to refer to individuals of varying socioeconomic status in some templates. We adapt the original list of “lowSES” and “highSES” occupations based on statistical data, including also their feminine form and reporting the source where the information is obtained from. We generate all possible combinations to instantiate the templates, as further explained in Appendix D.

Proper Names In the *Gender* and *Race/Ethnicity* categories, some templates are used with proper names. We tailor the original list by substituting the first names that refer to white individuals with typical Spanish and Catalan names. Additionally, first names from other racial and ethnic groups are replaced with those corresponding to the groups listed in the *Race/Ethnicity* category. In all cases, the most common names are selected according to the Spanish National Statistics Institute (INE),¹¹ for EsBBQ, and from the Catalanian Statistics Institute (Idescat),¹² for CABBQ, reporting the sources in all cases. We do not include any proper names for Romani people, given that names from such origin were not available in the statistics. As for last names, they are not included for the reasons stated in Appendix D.3.

D. Template Instantiation

D.1. Order Permutation

When instantiating the templates, we generate all 4 possible permutations of the order in which `NAME1` and `NAME2` appear in both contexts, to avoid introducing any ordering bias. In fact, in BBQ, Parrish et al. (2022) already take into account that ordering has an effect on model behaviour and they permute the order in which `NAME1` and `NAME2` appear in the contexts. However, some permutations were explicitly hardcoded creating different versions of templates, and not all possible combinations were always taken into account.

Instead, we adopt a systematic approach where all permuted versions are created automatically when instantiating each template, minimizing the potential for human error. Table 4 shows the 12 minimum instances obtained from a single template by generating all combinations of context and question types, and all possible permutations of `NAME1` and `NAME2`.

As for the ordering of the answer choices, it is not a concern in EsBBQ and CABBQ, since automatic evaluation of LLMs is carried out by obtaining the probabilities of each answer choice, one at a time (§4.2). If this evaluation setup is not suitable, we recommend fully instantiating the dataset with all possible permutations to mitigate any positional effects. However, note that the number of templates would increase substantially with the number of possible values for these placeholders, as well as with the inclusion of the `WORD` placeholder.

D.2. Unknown Answers

To avoid an over-reliance on the word *unknown*, Parrish et al. (2022) employ a list of semantically-

equivalent expressions. Similarly, we use expressions of varying syntax structures and lengths for the *unknown* answer options, all listed in Table 5.

As previously mentioned, given our evaluation setup, we do not need to control for the order of answer choices or for different *unknown* expressions. Instead, we obtain the probability of each expression one at a time, and the option with the highest probability is taken as the model’s answer.

D.3. Avoiding Confounds

The original instantiation algorithm used for BBQ templates featured several sources of random variability: (1) not all possible permutations of the `NAME` placeholders are considered; (2) `WORD` placeholders are filled with a random value chosen from the entire set of possible options at the time of instantiation; (3) values for `NAME` placeholders are randomly downsampled when more than five options are available; (4) first names are paired with a randomly selected last name for instances with proper nouns, and (5) the expression for the *unknown* answer is selected at random each time.

These randomized choices not only underuse the diversity available in the templates, but also introduce potential unintended biases into the measurement of social bias. Therefore, when instantiating EsBBQ and CABBQ templates, we eliminate all sources of randomness and systematically generate all possible combinations for these variables.

However, generating all combinations instead of random variations leads to a massive increase in template fertility. To control this, lexical variety is constrained: the `WORD` placeholder is only used when it introduces a meaningful diversity to the content of the template and, critically, to include the two genders of words referring to individuals, as mentioned in Appendix B.2. Similarly, we minimize the vocabulary for the `NAME` placeholders whenever possible, provided that this reduction does not result in the erasure of any social group within the category. Notably, occupation nouns in the *SES* category are limited to two representative examples per social group, and the list of proper nouns is also reduced to two first names per gender and race/ethnicity, to which no last names are added. This careful reduction strikes a balance between controlling template fertility and maintaining representativeness and diversity.

E. Evaluation Prompt

Each EsBBQ and CABBQ instance results in 11 different prompts, each with a different answer among all possible choices: the *target* one, the *non-target* one, and all *unknown* expressions. Table 5 shows an example of all resulting prompts from a single instance.

¹¹<https://www.ine.es>

¹²<https://www.idescat.cat>

Annotator	Age Range	Gender	LGBTQIA+	Domestic Region of Birth	Domestic Region of Residence	Race/Ethnicity	Religion	Educational Level	Linguistics Background	Familiar with BBQ
A01	20-24	woman	yes	Madrid	Basque Country	White	non-believer	Master's / postgrad.	yes	no
A02	25-29	non-binary	yes	Born outside Spain	Catalonia	Latin	non-believer	Master's / postgrad.	no	yes
A03	55-60	woman	no	Catalonia	Catalonia	White	Catholic	secondary education	no	no
A04	25-29	man	no	Born outside Spain	Catalonia	Arab, African	non-believer	Master's / postgrad.	yes	yes
A05	25-29	woman	no	Catalonia	Catalonia	White	non-believer	Master's / postgrad.	no	no
A06	25-29	woman	yes	Castile and Leon	Castile and Leon	White	non-believer	Master's / postgrad.	yes	no
A07	25-29	woman	yes	Born outside Spain	Catalonia	Latin	non-believer	Master's / postgrad.	no	yes
A08	25-29	man	yes	Castile and Leon	Catalonia	White	non-believer	Master's / postgrad.	no	no
A09	60-65	man	no	Catalonia	Catalonia	White	Catholic	voc. edu. and training	no	no
A10	25-29	man	no	Catalonia	Catalonia	White	non-believer	Master's / postgrad.	no	no
A11	25-29	woman	no	Catalonia	Catalonia	White	non-believer	Master's / postgrad.	no	yes
A12	20-24	man	no	Catalonia	Castile and Leon	White	non-believer	voc. edu. and training	no	no
A13	20-24	man	no	Basque Country	Basque Country	White	Catholic	Master's / postgrad.	no	no
A14	20-24	man	no	Basque Country	Basque Country	White	non-believer	Master's / postgrad.	no	no
A15	25-29	woman	no	Catalonia	Catalonia	White	non-believer	Master's / postgrad.	no	no
A16	25-29	woman	no	Basque Country	Basque Country	White	non-believer	Master's / postgrad.	yes	no
A17	20-24	woman	no	Catalonia	Catalonia	White	Catholic	Master's / postgrad.	no	no

Table 6: Demographic information of the annotators carrying out the validation task for NEWLY-CREATED templates.

F. Annotators' Demographics

All annotators in charge of validating NEWLY-CREATED templates are Spanish residents. Table 6 details additional information regarding their age, gender identity, sexual orientation, domestic region of origin and residence, race/ethnicity, religion and educational level. We also specify whether the annotator has a background in Linguistics and whether they are familiar with the BBQ dataset.

G. Detailed Scores

Table 7 and Table 8 show the average scores for all models evaluated on EsBBQ and CABBQ tasks. Additionally, Figure 4 and Figure 5 show the accuracy and bias scores per category and model variant.

Since the bias scores are bounded by the accuracy, we follow Jin et al. (2024) and provide their upper bound absolute values, described in Equations (5) and (6).

$$|Bias_{ambig}| = 1 - Acc_a \quad (0 \leq Acc_a \leq 1) \quad (5)$$

$$\begin{aligned}
|Bias_{disambig}| &\leq 1 - |2Acc_d - 1| \quad (0 \leq Acc_d \leq 1) \\
&= \begin{cases} 2Acc_d & (0 \leq Acc_d \leq 0.5) \\ 2(1 - Acc_d) & (0.5 < Acc_d \leq 1) \end{cases} \quad (6)
\end{aligned}$$

Model	Size	Variant	Ambiguous Context			Disambiguated Context		
			<i>Acc</i>	<i>Bias</i>	$UB Bias $	<i>Acc</i>	<i>Bias</i>	$UB Bias $
Salamandra	2B	base	0.24	0.02	0.76	0.50	0.03	1.00
		ins.	0.06	0.05	0.94	0.63	0.08	0.73
	7B	base	0.15	0.10	0.85	0.72	0.07	0.57
		ins.	0.09	0.30	0.91	0.91	0.08	0.19
FLOR	1.3B	base	0.35	0.02	0.65	0.45	0.03	0.90
		ins.	0.31	0.02	0.69	0.45	0.04	0.91
	6.3B	base	0.21	0.03	0.79	0.57	0.05	0.86
		ins.	0.11	0.03	0.89	0.56	0.04	0.88
Occiglot-EU5	7B	base	0.14	0.09	0.86	0.77	0.05	0.46
		ins.	0.29	0.11	0.71	0.85	0.04	0.30
EuroLLM	1.7B	base	0.45	0.02	0.55	0.44	0.04	0.87
		ins.	0.11	0.05	0.89	0.56	0.07	0.88
	9B	base	0.07	0.16	0.93	0.86	0.07	0.29
		ins.	0.46	0.08	0.54	0.83	0.05	0.34
Mistral	7B	base	0.14	0.09	0.86	0.76	0.05	0.47
		ins.	0.69	0.08	0.31	0.84	0.02	0.32
Llama-3.1	8B	base	0.15	0.09	0.85	0.78	0.02	0.45
		ins.	0.52	0.15	0.48	0.83	0.03	0.34
Gemma-3	1B	base	0.20	0.03	0.80	0.51	0.04	0.98
		ins.	0.54	0.02	0.46	0.52	0.03	0.96
	4B	base	0.14	0.05	0.86	0.69	0.05	0.62
ins.		0.59	0.12	0.41	0.79	0.03	0.42	
	12B	base	0.12	0.10	0.88	0.83	0.02	0.34
		ins.	0.79	0.10	0.21	0.83	0.01	0.33
Tower	7B	base	0.13	0.07	0.87	0.65	0.06	0.69
		ins.	0.08	0.11	0.92	0.76	0.06	0.47
	13B	base	0.16	0.07	0.84	0.71	0.05	0.58
		ins.	0.14	0.11	0.86	0.77	0.09	0.46
Qwen2.5	1.5B	base	0.20	0.03	0.80	0.65	0.05	0.70
		ins.	0.42	0.06	0.58	0.69	0.09	0.61
	3B	base	0.09	0.14	0.91	0.81	0.10	0.39
		ins.	0.86	0.01	0.14	0.68	0.05	0.65
7B	base	0.06	0.23	0.94	0.86	0.07	0.27	
	ins.	0.93	0.01	0.07	0.77	0.02	0.45	

Table 7: **EsBBQ**: Accuracy (*Acc*) and bias (*Bias*) scores of the models evaluated. $UB|Bias|$ indicates the maximum absolute value of the bias score depending on the accuracy.

	Base Models													Instructed Models																				
	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	0.06	0.09	0.22	0.15	0.06	0.37	0.00	0.01	0.08	0.04	0.01	0.04	0.03	0.10	0.08	0.00	0.00	-0.02	0.06	0.26	0.00	0.24	0.02	0.47	0.73	0.53	0.51	0.56	0.70	0.07	0.12	0.38	0.90	0.94
Disability Status	-0.37	0.23	0.55	0.29	0.18	0.61	0.00	0.10	0.17	0.32	0.09	0.12	0.16	0.21	0.21	0.01	0.00	-0.05	0.01	0.40	0.06	0.21	0.17	0.36	0.57	0.40	0.47	0.54	0.78	0.03	0.08	0.27	0.80	0.93
Gender	0.00	0.00	0.02	0.01	0.00	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	-0.00	0.02	0.02	0.00	0.06	0.00	0.10	0.44	0.31	0.32	0.41	0.71	0.00	0.02	0.25	0.78	0.88
LGBTQIA	-0.31	0.28	0.53	0.40	0.21	0.73	0.18	0.22	0.26	0.32	0.19	0.24	0.24	0.32	0.45	0.25	0.20	-0.22	0.20	0.56	0.34	0.56	0.28	0.68	0.87	0.68	0.77	0.90	0.96	0.19	0.37	0.55	0.88	0.95
Nationality	-0.14	0.17	0.33	0.30	0.11	0.60	0.00	0.18	0.16	0.21	0.20	0.16	0.21	0.20	0.29	0.14	0.10	-0.01	0.13	0.29	0.23	0.51	0.07	0.54	0.60	0.55	0.43	0.56	0.70	0.13	0.21	0.58	0.82	0.80
Physical Appearance	-0.59	0.32	0.69	0.49	0.38	0.67	0.19	0.47	0.37	0.56	0.50	0.31	0.34	0.33	0.43	0.18	0.19	-0.15	0.13	0.63	0.26	0.57	0.18	0.84	0.80	0.69	0.69	0.72	0.91	0.15	0.21	0.58	0.95	0.99
Race/Ethnicity	-0.08	0.04	0.18	0.07	0.04	0.24	0.02	0.07	0.07	0.09	0.07	0.06	0.04	0.07	0.08	0.04	0.03	-0.03	0.03	0.13	0.03	0.20	0.03	0.37	0.79	0.65	0.67	0.65	0.84	0.03	0.10	0.43	0.88	0.90
Religion	-0.36	0.17	0.66	0.11	0.06	0.81	0.01	0.15	0.18	0.25	0.16	0.04	0.12	0.22	0.42	0.02	0.01	-0.13	0.04	0.62	0.04	0.18	0.12	0.41	0.63	0.41	0.71	0.76	0.83	0.03	0.13	0.38	0.98	0.92
SES	-0.38	0.23	0.43	0.23	0.21	0.62	0.17	0.21	0.24	0.21	0.19	0.16	0.20	0.22	0.26	0.21	0.09	-0.08	0.20	0.34	0.20	0.36	0.19	0.56	0.72	0.51	0.54	0.54	0.76	0.19	0.21	0.49	0.85	0.94
Spanish Region	0.20	0.11	0.45	0.27	0.14	0.58	0.04	0.10	0.09	0.20	0.15	0.11	0.08	0.11	0.14	0.14	0.12	-0.05	0.08	0.38	0.18	0.46	0.12	0.40	0.81	0.56	0.46	0.52	0.90	0.02	0.18	0.56	0.96	0.99

(a) Acc_{ambig}

	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	-0.03	0.05	0.02	0.01	0.11	-0.00	0.16	0.08	0.05	0.01	0.03	0.11	0.06	0.02	0.06	0.07	0.22	-0.01	0.16	0.03	0.01	0.06	0.03	0.01	0.16	0.13	0.15	0.00	0.17	0.24	0.08	0.05	0.00	0.01
Disability Status	-0.03	0.05	0.01	-0.01	0.01	0.02	0.11	0.04	0.09	0.03	0.02	0.03	0.03	0.06	0.03	0.12	0.30	-0.12	0.35	0.00	-0.00	0.13	0.03	0.16	0.13	0.19	0.06	0.07	0.14	0.05	0.02	0.00	0.03	
Gender	-0.00	0.12	-0.01	0.02	0.06	0.05	0.09	0.13	0.11	0.04	0.08	0.10	0.10	0.10	0.04	0.10	0.22	-0.07	0.25	0.01	0.01	0.06	0.04	0.09	0.06	0.18	-0.00	0.10	0.10	0.10	0.05	0.06	0.01	-0.00
LGBTQIA	-0.04	0.08	-0.00	0.02	0.12	0.00	0.17	0.06	0.02	0.04	0.03	0.17	0.14	0.14	-0.04	0.11	0.09	-0.01	0.19	0.02	0.04	0.05	-0.01	0.06	0.01	0.10	-0.00	0.03	0.02	0.06	0.14	0.01	-0.01	0.03
Nationality	-0.00	0.04	-0.04	-0.04	0.00	0.06	0.23	0.05	0.10	-0.05	0.01	0.10	-0.02	-0.01	0.10	0.13	0.14	-0.03	0.06	0.04	-0.05	0.05	0.11	0.10	0.06	-0.06	-0.01	0.06	0.07	0.11	0.15	0.08	0.00	0.01
Physical Appearance	-0.02	0.17	-0.01	0.06	0.20	0.01	0.41	0.18	0.18	0.02	0.08	0.24	0.11	0.07	0.07	0.31	0.40	-0.05	0.64	0.01	0.08	0.33	0.08	0.12	0.16	0.23	0.05	0.17	0.09	0.25	0.25	0.14	0.05	0.01
Race/Ethnicity	-0.01	0.06	0.03	0.07	0.07	0.01	0.13	0.10	0.05	0.04	0.02	0.09	0.06	0.07	0.01	0.13	0.15	-0.00	0.19	-0.02	0.00	0.06	0.06	0.03	0.04	0.07	0.01	0.02	0.03	0.07	0.06	0.04	-0.02	-0.02
Religion	-0.00	0.07	0.03	0.06	-0.03	0.01	0.01	-0.05	-0.03	-0.01	0.04	-0.04	-0.07	0.00	-0.05	0.01	-0.02	-0.14	0.11	0.07	-0.01	-0.08	-0.10	-0.00	-0.09	-0.06	-0.07	-0.08	0.06	0.01	0.04	0.02	0.00	0.00
SES	-0.03	0.16	0.04	0.06	0.11	0.04	0.13	0.06	0.10	0.03	0.06	0.03	0.04	0.08	0.06	0.17	0.26	-0.13	0.40	0.05	0.08	0.18	0.11	0.10	0.06	0.18	0.04	0.21	0.13	0.15	0.13	0.09	0.03	0.01
Spanish Region	-0.04	0.03	0.08	0.00	-0.01	0.02	0.01	0.02	0.05	0.06	0.03	0.10	0.00	0.02	-0.04	-0.01	0.07	-0.03	0.20	0.02	-0.03	0.03	0.01	0.08	0.05	0.09	0.02	0.06	0.07	-0.03	0.03	-0.00	0.00	0.01

(b) $Bias_{ambig}$

	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	-0.47	0.60	0.42	0.53	0.72	0.39	0.87	0.64	0.70	0.47	0.66	0.84	0.59	0.63	0.67	0.81	0.89	-0.57	0.90	0.37	0.54	0.85	0.50	0.93	0.81	0.84	0.48	0.83	0.93	0.72	0.77	0.74	0.63	0.80
Disability Status	-0.50	0.73	0.48	0.53	0.78	0.40	0.83	0.79	0.82	0.53	0.68	0.79	0.66	0.66	0.62	0.80	0.88	-0.61	0.92	0.47	0.58	0.90	0.57	0.81	0.89	0.89	0.53	0.82	0.88	0.74	0.81	0.67	0.76	0.82
Gender	-0.54	0.79	0.53	0.59	0.81	0.56	0.89	0.79	0.78	0.53	0.72	0.85	0.72	0.74	0.71	0.85	0.91	-0.72	0.96	0.53	0.55	0.87	0.63	0.93	0.89	0.90	0.65	0.87	0.92	0.80	0.74	0.78	0.79	0.85
LGBTQIA	-0.45	0.61	0.41	0.51	0.72	0.34	0.83	0.69	0.78	0.48	0.65	0.73	0.62	0.62	0.60	0.73	0.82	-0.63	0.89	0.42	0.57	0.75	0.52	0.77	0.69	0.62	0.36	0.62	0.58	0.78	0.70	0.57	0.58	0.63
Nationality	-0.52	0.79	0.50	0.66	0.79	0.45	0.87	0.85	0.80	0.56	0.71	0.84	0.71	0.80	0.64	0.77	0.82	-0.64	0.86	0.49	0.58	0.86	0.57	0.89	0.88	0.90	0.50	0.74	0.89	0.73	0.76	0.71	0.67	0.85
Physical Appearance	-0.49	0.71	0.42	0.57	0.73	0.38	0.79	0.74	0.75	0.50	0.67	0.80	0.65	0.73	0.61	0.76	0.76	-0.59	0.81	0.45	0.55	0.79	0.54	0.64	0.76	0.66	0.42	0.69	0.69	0.71	0.71	0.57	0.51	0.60
Race/Ethnicity	-0.51	0.74	0.46	0.57	0.80	0.45	0.85	0.80	0.78	0.51	0.65	0.87	0.64	0.69	0.65	0.83	0.88	-0.65	0.89	0.48	0.55	0.83	0.55	0.85	0.88	0.89	0.56	0.81	0.89	0.78	0.79	0.73	0.75	0.86
Religion	-0.48	0.76	0.34	0.66	0.86	0.39	0.92	0.83	0.88	0.51	0.76	0.81	0.73	0.86	0.64	0.80	0.91	-0.66	0.94	0.35	0.65	0.93	0.56	0.93	0.93	0.93	0.43	0.86	0.84	0.88	0.91	0.77	0.65	0.91
SES	-0.50	0.73	0.43	0.62	0.75	0.46	0.86	0.81	0.80	0.53	0.72	0.87	0.62	0.75	0.64	0.79	0.86	-0.62	0.92	0.44	0.60	0.89	0.57	0.76	0.84	0.84	0.51	0.74	0.77	0.76	0.79	0.67	0.63	0.70
Spanish Region	-0.52	0.84	0.44	0.63	0.90	0.46	0.91	0.91	0.87	0.50	0.76	0.86	0.80	0.83	0.63	0.88	0.91	-0.71	0.97	0.45	0.54	0.90	0.62	0.92	0.91	0.86	0.66	0.93	0.96	0.88	0.87	0.78	0.84	0.93

(c) $Acc_{disambig}$

	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOr 1.3B	FLOr 6.3B	Occiglot-EUS 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	-0.05	0.04	0.01	0.00	0.05	0.00	0.08	0.07	-0.03	0.09	0.10	0.06	0.04	0.04	-0.00	0.05	0.05	-0.02	0.05	-0.00	-0.02	-0.02	0.05	0.04	0.02	0.06	0.05	-0.00	0.02	0.08	0.11	0.02	0.02	0.01
Disability Status	-0.00	0.01	0.04	0.01	0.03	0.00	0.10	0.04	0.10	0.04	-0.02	0.04	0.08	0.06	0.07	0.16	0.13	-0.10	0.11	0.05	0.05	0.01	0.09	0.09	0.08	0.07	0.05	0.06	0.03	0.03	0.11	0.14	0.05	0.06
Gender	-0.02																																	

Model	Size	Variant	Ambiguous Context			Disambiguated Context		
			<i>Acc</i>	<i>Bias</i>	$UB Bias $	<i>Acc</i>	<i>Bias</i>	$UB Bias $
Salamandra	2B	base	0.33	0.03	0.67	0.50	0.03	1.00
		ins.	0.07	0.06	0.93	0.65	0.08	0.71
	7B	base	0.12	0.13	0.88	0.79	0.08	0.43
		ins.	0.10	0.29	0.90	0.92	0.07	0.16
FLOR	1.3B	base	0.43	0.02	0.57	0.45	0.02	0.91
		ins.	0.37	0.02	0.63	0.47	0.03	0.93
	6.3B	base	0.18	0.03	0.82	0.61	0.05	0.78
		ins.	0.20	0.02	0.80	0.56	0.04	0.88
Occiglot-EU5	7B	base	0.17	0.06	0.83	0.72	0.05	0.56
		ins.	0.46	0.06	0.54	0.78	0.06	0.44
EuroLLM	1.7B	base	0.37	0.02	0.63	0.46	0.04	0.92
		ins.	0.15	0.05	0.85	0.56	0.06	0.89
	9B	base	0.10	0.13	0.90	0.85	0.08	0.29
		ins.	0.36	0.12	0.64	0.87	0.04	0.27
Mistral	7B	base	0.16	0.07	0.84	0.76	0.07	0.49
		ins.	0.69	0.08	0.31	0.82	0.05	0.36
Llama-3.1	8B	base	0.18	0.08	0.82	0.78	0.04	0.43
		ins.	0.44	0.13	0.56	0.86	0.03	0.27
Gemma-3	1B	base	0.32	0.02	0.68	0.50	0.03	1.00
		ins.	0.51	0.03	0.49	0.49	0.05	0.98
	4B	base	0.22	0.04	0.78	0.68	0.04	0.64
	12B	ins.	0.50	0.10	0.50	0.77	0.03	0.46
		base	0.17	0.11	0.83	0.84	0.03	0.32
	7B	ins.	0.74	0.09	0.26	0.85	0.00	0.29
		base	0.14	0.04	0.86	0.63	0.05	0.74
Tower	13B	ins.	0.13	0.07	0.87	0.73	0.07	0.55
		base	0.13	0.05	0.87	0.68	0.07	0.64
	1.5B	ins.	0.09	0.10	0.91	0.74	0.08	0.52
		base	0.25	0.04	0.75	0.63	0.05	0.74
	3B	ins.	0.25	0.04	0.75	0.65	0.05	0.69
		base	0.13	0.10	0.87	0.75	0.10	0.50
Qwen2.5	7B	ins.	0.88	0.02	0.12	0.63	0.05	0.73
		base	0.07	0.13	0.93	0.86	0.06	0.28
		ins.	0.91	0.01	0.09	0.71	0.05	0.57

Table 8: **CABBQ**: Accuracy (*Acc*) and bias (*Bias*) scores of the models evaluated. $UB|Bias|$ indicates the maximum absolute value of the bias score depending on the accuracy.

	Base Models													Instructed Models																				
	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	0.25	0.04	0.44	0.10	0.07	0.28	0.00	0.01	0.07	0.24	0.18	0.09	0.11	0.08	0.12	0.00	0.00	-0.02	0.06	0.34	0.07	0.36	0.02	0.32	0.71	0.40	0.36	0.51	0.63	0.15	0.10	0.18	0.81	0.90
Disability Status	-0.44	0.08	0.57	0.22	0.25	0.51	0.03	0.13	0.33	0.40	0.28	0.23	0.07	0.13	0.29	0.08	0.02	-0.06	0.01	0.48	0.23	0.57	0.25	0.26	0.55	0.34	0.48	0.40	0.79	0.04	0.01	0.13	0.81	0.88
Gender	0.01	0.00	0.02	0.01	0.00	0.08	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.00	0.01	0.00	0.00	-0.00	0.05	0.04	0.00	0.20	0.00	0.12	0.53	0.32	0.39	0.27	0.61	0.01	0.00	0.03	0.89	0.92
LGBTQIA	-0.46	0.27	0.63	0.29	0.23	0.66	0.21	0.26	0.35	0.56	0.32	0.33	0.28	0.26	0.56	0.37	0.22	-0.19	0.20	0.48	0.39	0.67	0.44	0.57	0.84	0.62	0.76	0.79	0.91	0.23	0.32	0.44	0.97	0.97
Nationality	-0.31	0.11	0.61	0.23	0.12	0.41	0.14	0.27	0.16	0.33	0.20	0.15	0.18	0.28	0.30	0.20	0.24	-0.11	0.19	0.48	0.29	0.54	0.14	0.44	0.68	0.67	0.54	0.58	0.73	0.17	0.10	0.30	0.79	0.75
Physical Appearance	-0.64	0.26	0.77	0.50	0.38	0.62	0.27	0.49	0.35	0.60	0.54	0.43	0.28	0.26	0.42	0.20	0.21	-0.12	0.10	0.76	0.53	0.70	0.29	0.61	0.79	0.53	0.60	0.62	0.86	0.22	0.04	0.46	0.93	0.93
Race/Ethnicity	-0.14	0.03	0.22	0.04	0.06	0.16	0.04	0.11	0.10	0.20	0.09	0.07	0.06	0.05	0.11	0.04	0.03	-0.04	0.08	0.22	0.07	0.35	0.05	0.31	0.79	0.53	0.39	0.55	0.76	0.05	0.04	0.13	0.93	0.93
Religion	-0.60	0.02	0.71	0.04	0.10	0.50	0.00	0.05	0.18	0.56	0.15	0.06	0.07	0.11	0.14	0.01	0.00	-0.02	0.03	0.61	0.08	0.47	0.27	0.19	0.71	0.35	0.77	0.58	0.82	0.06	0.11	0.12	0.88	0.88
SES	-0.48	0.24	0.52	0.23	0.26	0.46	0.21	0.21	0.28	0.41	0.27	0.20	0.23	0.18	0.47	0.29	0.10	-0.12	0.20	0.41	0.25	0.52	0.19	0.47	0.69	0.40	0.66	0.50	0.73	0.22	0.19	0.46	0.88	0.91
Spanish Region	-0.43	0.09	0.52	0.28	0.26	0.54	0.10	0.21	0.13	0.38	0.27	0.19	0.17	0.17	0.31	0.22	0.16	-0.14	0.19	0.34	0.30	0.61	0.19	0.41	0.78	0.47	0.53	0.46	0.80	0.14	0.10	0.47	0.98	0.99

(a) Acc_{ambig}

	Base Models													Instructed Models																				
	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	-0.04	0.12	0.01	-0.00	0.06	0.02	0.11	0.08	0.06	0.04	0.02	0.16	-0.01	0.07	0.05	0.09	0.10	-0.02	0.18	0.00	0.02	0.05	0.05	0.03	0.10	0.17	-0.02	0.04	0.22	0.09	0.08	0.03	0.01	0.03
Disability Status	-0.05	0.10	0.00	0.05	0.12	0.04	0.08	0.07	0.05	0.00	0.02	0.07	0.10	0.05	0.03	0.12	0.22	-0.13	0.33	-0.01	-0.01	0.10	0.08	0.11	0.16	0.14	0.09	0.18	0.05	0.08	0.09	0.06	0.05	0.03
Gender	-0.03	0.13	0.01	0.03	0.07	0.01	0.15	0.04	0.07	0.01	0.08	0.11	0.04	0.07	0.03	0.11	0.09	-0.06	0.25	-0.01	0.01	0.04	0.06	0.08	0.05	0.12	0.02	0.13	0.12	0.07	0.08	0.01	0.00	0.01
LGBTQIA	-0.03	0.02	0.04	0.02	-0.02	-0.00	0.16	0.05	0.06	0.03	-0.00	0.13	-0.03	0.03	0.03	0.06	0.03	-0.02	0.22	0.06	0.01	0.01	0.02	0.08	0.04	0.11	-0.01	0.00	0.01	0.02	0.10	-0.00	-0.01	0.01
Nationality	-0.00	-0.03	-0.05	-0.07	0.02	-0.02	-0.06	-0.01	0.07	-0.01	0.02	0.01	0.05	-0.03	0.04	0.10	0.05	-0.05	0.06	-0.02	-0.04	0.00	-0.05	0.12	-0.02	-0.02	0.01	0.10	0.05	0.03	0.01	0.02	0.01	0.04
Physical Appearance	-0.02	0.29	0.00	0.05	0.12	0.04	0.32	0.08	0.14	0.02	0.04	0.18	0.07	0.07	0.08	0.18	0.22	-0.09	0.62	0.01	0.04	0.11	0.09	0.28	0.11	0.17	0.07	0.17	0.05	0.20	0.20	0.10	0.03	0.03
Race/Ethnicity	-0.04	0.10	0.04	0.05	0.10	0.06	0.15	0.11	0.12	0.04	0.05	0.13	0.05	0.05	0.06	0.09	0.14	-0.04	0.23	0.06	0.02	0.07	0.04	0.12	0.02	0.10	0.05	0.03	0.03	0.05	0.10	0.07	0.00	0.01
Religion	-0.04	0.19	-0.01	0.02	0.01	-0.03	0.09	0.02	-0.05	0.02	0.00	0.06	0.06	-0.01	-0.10	0.06	0.01	-0.03	0.22	0.05	0.05	-0.09	-0.10	0.05	-0.09	-0.09	0.02	-0.05	0.01	0.03	0.09	-0.03	0.02	-0.04
SES	-0.04	0.12	0.04	0.04	0.04	0.03	0.04	0.09	0.09	0.01	0.03	0.04	0.02	0.04	0.01	0.07	0.15	-0.10	0.32	0.05	0.02	0.07	0.03	0.16	0.11	0.17	0.07	0.14	0.12	0.01	0.08	0.00	0.06	0.01
Spanish Region	-0.00	-0.02	0.02	-0.01	0.02	-0.01	0.06	0.03	0.00	0.03	0.04	0.02	-0.04	0.04	0.01	0.00	0.04	-0.02	0.11	0.02	0.05	0.02	-0.03	-0.01	0.03	-0.01	-0.02	0.02	0.05	0.02	-0.01	-0.01	-0.01	0.00

(b) $Bias_{ambig}$

	Base Models													Instructed Models																				
	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	-0.47	0.72	0.42	0.57	0.62	0.41	0.87	0.59	0.62	0.45	0.63	0.78	0.53	0.55	0.63	0.70	0.85	-0.58	0.89	0.42	0.54	0.77	0.50	0.89	0.73	0.84	0.48	0.81	0.91	0.67	0.65	0.64	0.56	0.64
Disability Status	-0.49	0.84	0.50	0.59	0.75	0.43	0.82	0.80	0.86	0.53	0.68	0.78	0.64	0.69	0.59	0.73	0.84	-0.63	0.91	0.50	0.58	0.82	0.55	0.87	0.87	0.91	0.52	0.82	0.91	0.68	0.73	0.63	0.68	0.74
Gender	-0.57	0.83	0.55	0.64	0.77	0.55	0.88	0.79	0.85	0.55	0.75	0.87	0.67	0.71	0.68	0.82	0.91	-0.76	0.96	0.54	0.55	0.79	0.63	0.93	0.89	0.93	0.56	0.85	0.93	0.76	0.73	0.69	0.72	0.83
LGBTQIA	-0.43	0.66	0.39	0.56	0.61	0.32	0.81	0.70	0.79	0.41	0.60	0.75	0.52	0.63	0.59	0.67	0.86	-0.59	0.90	0.42	0.54	0.64	0.48	0.75	0.68	0.71	0.37	0.57	0.65	0.61	0.78	0.61	0.54	0.59
Nationality	-0.49	0.84	0.41	0.64	0.78	0.50	0.86	0.84	0.80	0.47	0.64	0.83	0.70	0.75	0.64	0.74	0.79	-0.65	0.90	0.46	0.58	0.80	0.56	0.82	0.89	0.91	0.54	0.75	0.86	0.71	0.72	0.67	0.75	0.62
Physical Appearance	-0.48	0.77	0.40	0.60	0.71	0.45	0.81	0.74	0.77	0.52	0.70	0.83	0.64	0.66	0.62	0.69	0.79	-0.61	0.89	0.43	0.55	0.77	0.54	0.83	0.75	0.77	0.48	0.68	0.76	0.70	0.73	0.63	0.55	0.59
Race/Ethnicity	-0.52	0.79	0.46	0.61	0.77	0.49	0.87	0.82	0.77	0.49	0.67	0.89	0.65	0.73	0.68	0.81	0.87	-0.67	0.91	0.49	0.58	0.82	0.58	0.89	0.90	0.93	0.52	0.79	0.91	0.80	0.79	0.70	0.69	0.79
Religion	-0.42	0.84	0.41	0.67	0.81	0.40	0.93	0.88	0.83	0.48	0.75	0.88	0.72	0.86	0.66	0.83	0.91	-0.66	0.96	0.41	0.61	0.90	0.54	0.95	0.94	0.94	0.34	0.84	0.89	0.86	0.88	0.72	0.69	0.84
SES	-0.50	0.80	0.42	0.61	0.71	0.49	0.84	0.77	0.82	0.50	0.68	0.87	0.64	0.71	0.56	0.73	0.84	-0.61	0.91	0.44	0.55	0.75	0.55	0.83	0.80	0.85	0.44	0.72	0.78	0.74	0.74	0.61	0.58	0.67
Spanish Region	-0.53	0.87	0.46	0.69	0.86	0.49	0.94	0.88	0.85	0.54	0.72	0.94	0.73	0.80	0.71	0.88	0.92	-0.71	0.98	0.46	0.64	0.85	0.59	0.94	0.91	0.87	0.63	0.89	0.96	0.80	0.89	0.73	0.78	0.91

(c) $Acc_{disambig}$

	Base Models													Instructed Models																				
	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B	Salamandra 2B	Salamandra 7B	FLOR 1.3B	FLOR 6.3B	Occiglot-EU5 7B	EuroLLM 1.7B	EuroLLM 9B	Mistral 7B	Llama-3.1 8B	Gemma-3 1B	Gemma-3 4B	Gemma-3 12B	Tower 7B	Tower 13B	Qwen2.5 1.5B	Qwen2.5 3B	Qwen2.5 7B
Age	-0.05	0.12	-0.00	-0.04	0.04	0.02	0.08	0.10	0.02	0.02	0.06	-0.02	0.04	0.03	0.01	0.03	0.10	-0.07	0.06	-0.01	-0.00	0.07	0.04	-0.00	0.04	0.05	0.05	-0.02	0.00	0.07	0.11	0.02	-0.03	0.02
Disability Status	-0.02	0.04	0.02	0.02	0.02	0.05	0.12	0.11	0.09	0.03	0.02	0.07	0.07	0.11	0.08	0.15	0.12	-0.14	0.13	0.05	0.05	0.06</												