

MAFIA: Multi-Adapter Fused Inclusive LanguAge Models

Warning: This paper has content which may be offensive, or upsetting, however this cannot be avoided owing to the nature of the work

Anonymous ACL submission

Abstract

Pretrained Language Models (PLMs) are widely used in NLP for various tasks. Recent studies have identified various biases that such models exhibit and have proposed methods to correct these biases. However, most of the works address a limited set of bias dimensions independently such as gender, race, or religion. Moreover, the methods typically involve finetuning the full model in order to maintain the performance on the downstream task. In this work, we aim to *modularly* debias a pretrained language model across *multiple* dimensions. Previous works extensively explored debiasing PLMs by using limited US-centric counterfactual data augmentation (CDA). We use structured knowledge and a large generative model to build a diverse CDA across multiple bias dimensions in a semi-automated way. We highlight how existing debiasing methods do not consider interactions between multiple societal biases and propose a debiasing model that exploits the synergy amongst various societal biases and enables multi-bias debiasing simultaneously. An extensive evaluation on multiple tasks and languages demonstrates the efficacy of the approach.

1 Introduction

PLMs are growing in power and prominence across numerous NLP tasks. Their reach has expanded beyond academia, reaching general users through services like code assistance and chatbots (Li et al., 2023; Köpf et al., 2023). Despite the extraordinary performance of these models on their respective tasks, several works have identified the harmful social biases picked up by these models as an artifact of their pretraining on web-scale corpus consisting of unmoderated user-generated content (Manzini et al., 2019; Webster et al., 2020; Nadeem et al., 2021, *inter alia*).

While most previous works focus on (binary) gender biases, other societal biases, such as race

and religion, are less studied in the context of PLMs. Moreover, these biases are often intertwined with each other, creating complex and nuanced forms of discrimination. We define intersectional biases as the biases that arise from the combination of different attributes, such as gender, race, and religion. In this work, we focus on building debiasing techniques that can model and mitigate gender (including non-binary), race, religion, profession and intersectional biases, which are often ignored in previous works.

The community has developed a gamut of methods to measure and mitigate biases in LLMs (Bordia and Bowman, 2019; Liang et al., 2020; Ravfogel et al., 2020; Webster et al., 2020; Schick et al., 2021; Lauscher et al., 2021; Kumar et al., 2023). The majority of these methods finetune *all* the parameters of a language model to debias it towards a particular bias dimension such as gender or race, and the escalating size of PLMs can pose computational challenges, particularly for smaller academic labs or enterprises. To this end, we aim to use *adapters* (Houlsby et al., 2019; Pfeiffer et al., 2020), which are small neural network layers inserted in Transformer blocks (Vaswani et al., 2017) of an LLM as a way to effectively debias it towards a certain dimension. We further show that a soft combination of multiple such adapters can be used to exploit the synergy between various bias dimensions and can lead to a fairer and more accurate model on a downstream task.

To train each of the individual debiasing adapters, we make use of counterfactual data augmentation (CDA) technique. While CDA has been shown to be effective on gender debiasing (Zmigrod et al., 2019; Dinan et al., 2020; Webster et al., 2020; Barikeri et al., 2021), previous works (Meade et al., 2022; Lauscher et al., 2021) have relied on a small set of handbuilt (mostly US-centric) counterfactual pairs. As LLM bias is a complex and multifaceted issue, comprehensively addressing it

requires considering diverse identities. Hence, we propose a semi-automated method, general purpose to build a comprehensive CDA pair list using Wikidata (Vrandečić and Krötzsch, 2014) and generative models.

Our results indicate that such a general method can be used to train strong debiasing adapters for multiple dimensions. In particular, we perform experiments on gender, race, religion, and profession. We list our contributions and key findings below:

1. An inclusive and diverse counterfactual pair dataset¹ for gender, race, religion, and profession bias. Note that, we also take into account non-binary genders. (§2.1)
2. iDeb - A more inclusive and improved bias-specific debiasing model, trained on the newly generated diverse and inclusive CDA pairs. (§2.2)
3. MAFIA - A soft way to combine multiple debiasing adapters on downstream tasks. The model exploits the synergy between various biases to improve fairness as well as performance on the downstream task. (§2.3)
4. We show that MAFIA can reduce unintended bias on toxicity classification task on related bias dimensions that are unseen by any of the individual debiasing adapters. (§4.4)
5. We observe zero-shot transfer of gains in fairness and performance by debiasing a multilingual PLM on English. We test our models on a new dataset (mBias-STS-B) for measuring fairness across different languages with varying resource availability. (§4.3)

Upon acceptance, we plan to release the mBias-STS-B dataset along with the code for future research under (MIT License).

2 Methodology

In our study, we examine four primary bias dimensions: gender, race (ethnicity), religion, and profession. First, we discuss the method for generating counterfactual (CF) pairs in Section 2.1. Subsequently, we outline the approach to train debiasing adapters (DBAs) for each dimension in Section 2.2. Lastly, in Section 2.3, we introduce our strategy for integrating individual DBAs for application on a downstream task.

¹Unlike previous work, which mainly was US-centric.

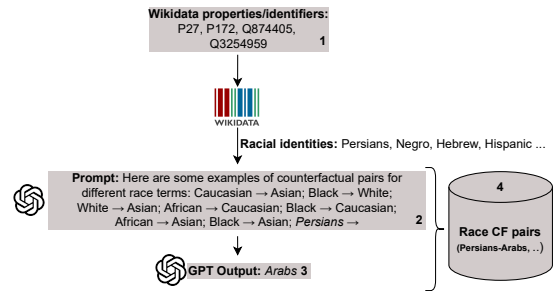


Figure 1: Steps to generate Counterfactual (CF) pairs for racial bias. Note that the technique can be similarly used for other biases.

2.1 Counterfactual Data Augmentation

Counterfactual Data Augmentation (CDA) is a generic dataset-based debiasing technique (Kusner et al., 2017; Lu et al., 2020). Given a set of counterfactual (CF) pairs (i.e., d representing the dominant group, e.g., man, and m representing the minority group, e.g., woman) and a training dataset, CDA replaces every instance of d with m and vice-versa (2-way CDA) (Webster et al., 2020) in the training data. The final corpus for debiasing training consists of both the original and counterfactually created sentences. The goal is that such data can balance the effect of pre-existing biases in data and encourage the model to learn fairer representations.

Generating CF pairs: Unlike previous methods (Lauscher et al., 2021; Meade et al., 2022) that rely on handcrafting the CF pairs (mostly US-centric), we propose a semi-automated, generic method to generate CDA pairs. We use a large structured knowledge base as a starting point. Wikidata’s (Vrandečić and Krötzsch, 2014) repository of information is rich and diverse, making it an ideal resource for our purpose. We manually identify a list of Wikidata items and properties whose subject or object position has English entities of respective bias type (gender/race/religion/profession) (step 1 in Figure §1). We refer the readers to Appendix §A.6 for the properties we used for extracting gender, race, religion, and profession terms.

Using all possible pairs of bias-related words for generating CDA can quickly become intractable, especially when dealing with extensive lists of terms. Additionally, including all pairs may introduce noise in training. It is crucial to exercise caution and thoughtfully curate the pairs to ensure the training process remains effective.

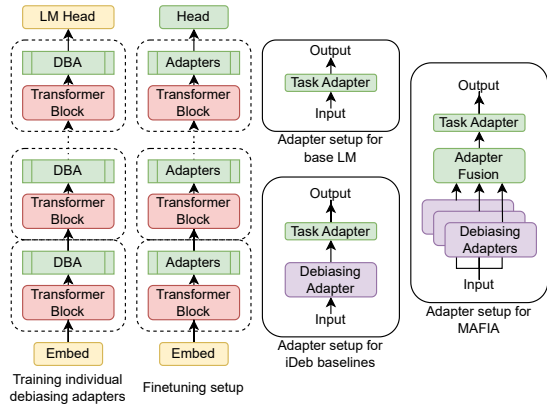


Figure 2: A comprehensive summary of the various training strategies described. Only the components highlighted in green are finetuned in each case.

166 tive and reliable. Therefore, we use a generative model² to build a corpus of CDA from the bias term list. Our prompt has the following structure: Here are some examples of counterfactual pairs for different <bias-type> terms: <sample-CDA-pairs>, <bias-term> → <output> (step 2 in §1). Here, <bias-type> is one of the bias dimensions i.e. gender/race/religion/profession while <sample-CDA-pairs> is a seed set of CDA pairs. We obtain this seed set for gender, race, and religion from Meade et al. (2022). For profession, we use the gender seed set. Finally, we prompt the LLM to produce a suitable counter for a new <bias-term> (step 3 in Figure §1).

181 We find that the generative model can generate a lot of improbable and uninteresting CF pairs during this process. Therefore, to filter out these pairs, we use the Google Book Corpus³ and retain only those CF pairs where both the entities in the pair appear at least once in a million times in the corpus. For gender, we reduce the threshold to 0.01. Our final set of CF pairs includes 68 pairs for gender, 156 for race, and 86 for religion. These numbers are notably higher as compared to Meade et al. (2022) which used 57, 7, and 6 terms for gender, race, and religion respectively.

193 2.2 Training Individual Debiasing Adapters

194 We adopt the training procedure of Lauscher et al. (2021) to train a debiasing adapter. The process involves adding a debiasing adapter to the base LM and is trained with a Masked Language Mod-

198 eling (MLM) objective (Devlin et al., 2019) on our large, inclusive CDA Wikipedia dataset. Note that training the debiasing adapter does *not* introduce task-adapters. The model architecture for this step is described by the leftmost part of Fig. 2.

203 2.3 Combining Multiple Debiasing Adapters

204 Our final model exploits the synergy of various debiasing adapters to improve performance on respective biases. Therefore, given k debiasing adapters trained *independently* on k bias dimensions, we propose to combine them on a downstream task as shown in the rightmost part of Fig. 2. All the k debiasing adapters are fused via a trainable Adapter-Fusion (Pfeiffer et al., 2021) layer and stacked with a task-specific adapter to facilitate further intermixing of signals. We refer to such a model with multiple fused adapters as MAFIA.

215 3 Experimental Setup

216 3.1 Evaluation datasets and metrics

217 We evaluate MAFIA on various intrinsic and extrinsic (downstream) bias evaluation benchmarks, and demonstrate its superior debiasing ability over related baselines.

221 3.1.1 Intrinsic Evaluation

222 We use *Stereoset* and *Crowdsourced Stereoset Pairs* (CrowS-Pairs) to evaluate intrinsic bias in models. **StereoSet** (Nadeem et al., 2021) is a large-scale natural English crowdsourced dataset to measure stereotypical biases in four domains: gender, profession, race, and religion. Each StereoSet example consists of a context sentence – “*Our housekeeper is a <BLANK>.*” And a set of three attributes – stereotype (Mexican), anti-stereotype (American), and a meaningless option (Banana). We determine which attribute will most likely fill the blank to measure language modeling and stereotypical bias. We use two different measures: (1) *Stereotype Score* is the percentage of examples for which a model prefers stereotypical association instead of anti-stereotypical associations. (2) *Language modeling score* is the percentage of examples for which a model prefers meaningful associations (either stereotypical or anti-stereotypical) as opposed to meaningless associations.

242 **CrowS-Pairs** (Nangia et al., 2020) introduced a crowdsourced benchmark dataset for measuring the degree to which nine types of social bias are present in language models. This work focuses

²text-davinci-003 by OpenAI

³<https://api.datamuse.com/>

on gender, race (and ethnicity), religion, and professional biases. The dataset consists of stereotypical and anti-stereotypical sentences in a given context similar to StereoSet. We use *Stereotype Score*, the percentage of examples for which a model assigns a higher masked token probability to the stereotypical sentence than the anti-stereotypical sentence. The masked token probability of a sentence is the average probability of unique tokens (w.r.t. counterpart sentence) in the sentence.

3.1.2 Extrinsic Evaluation

Dataset: We use STS-B i.e., the Semantic Textual Similarity Benchmark from GLUE (Wang et al., 2018) as our downstream task for this evaluation. STS-B requires the model to consider two sentences and output a score between 0-5 indicating how semantically similar the input sentences are. Webster et al. (2020) introduced Bias-STS-B which takes a neutral STS template and fills it with a gendered term and a profession term to form two sentences respectively. Original Bias-STS-B used only binary gender terms while in our study we consider 7 gender identities – male, female, non-binary, and LGBT. The gender bias evaluation dataset contains 16, 980 such septets (for 7-way comparison). We generate test sets for evaluating race, and religion biases, using the templates released by Dev et al. (2020). The sentence pair is built using a noun-template – *The ⟨subject⟩ person ⟨verb⟩ a/an ⟨object⟩* and an adjective template – *The ⟨adjective⟩ person ⟨verb⟩ a/an ⟨object⟩*. The ⟨adjective⟩ is filled with polarised adjectives (e.g., arrogant, brilliant) and the ⟨subject⟩ is filled with a religion term (e.g., Christian, Hindu, etc.) or ethnicity (Black, Caucasian, etc.) for generating a religion or race bias evaluation dataset respectively. We produce a total of 688, 801 Race-Bias and 757, 680 Religion-Bias sentence pairs, and we further sub-sample a set of 16, 384 sentence pairs from it for tractable evaluation. We use 11 religion terms and 10 race terms from Dev et al. (2020) to build the dataset. **Metrics:** On STS-B, we measure the performance by calculating the Pearson correlation (Freedman et al., 2007) (ρ) between model scores and human annotated similarity scores. On Bias-STS-B, we report the *average absolute difference* between scores of individual components. Unlike previous works, we perform a *multi-way comparison* instead of a 2-way comparison. E.g. Bias-STS-

B along race component has $k = 10$ components (i.e. different races) which means there are 10 sentence pairs (*An African-American kid is playing on the ground* vs *A child is playing on the ground*; *An Indian kid is playing on the ground* vs *A child is playing on the ground* and so on) for which we receive scores s_1, \dots, s_k from the model. Next we calculate average absolute difference as $\Delta = \frac{1}{\binom{k}{2}} \sum_{i=1}^k \sum_{j=i}^k |s_i - s_j|$. Notice that it is trivial to drive Δ to 0 at the cost of performance on STS-B by producing the same score for every pair. To better account for this tradeoff, we introduce a new metric called “*useful fairness*” that lets us compare models on both fairness as well as accuracy axes. We compute “*useful fairness*” (Ψ) for a particular bias dimension as $\Psi_{\text{dim}} = \rho \cdot \alpha (1 - \Delta_{\text{dim}})$ where ρ is the Pearson score model achieves on original STS-B, and $\alpha (= 1)$ is a constant capturing estimated effect of debiasing performance on the final model score, and Δ_{dim} is the average difference across all components of a particular bias dimension (gender/race/religion) computed on Bias-STS-B as discussed above.

3.2 Baselines

We use BERT, mBERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), XLM-R (Conneau et al., 2020) as our base LMs. To validate the effectiveness of MAFIA, we primarily compare it against the base LM as well as iDeb_{bias} (debaised for respective biases). Fig. 2 shows our general finetuning setup along with Adapter setups for various baselines.

On *gender*, we additionally consider using CF pairs from Lauscher et al. (2021) to train a DBA. We also compare with an “AdapterDrop” approach (Rücklé et al., 2021), an adapter-based dropout regularization method since previous work by Webster et al. (2020) showed that dropout helps model debiasing. We call this model $\langle \text{baseLM} \rangle + \text{AD}$. The model architecture is similar to base LM + task-adapter in Fig. 2 except the task-adapter is an “AdapterDrop” enabled adapter. Another baseline we compare with is a single DBA model trained on the concatenation of all CDA data used for four bias dimensions, denoted by iDeb_{all}.

3.3 Hyperparameters

All Debiasing Adapters (DBAs) and task-adapters use Pfeiffer architecture (Houlsby et al., 2019; Bapna and Firat, 2019; Pfeiffer et al., 2020) with

SiLU (Elfwing et al., 2017) activation, owing to its superior expressivity. For the integration and training of adapters within the model, we leverage the *Adapter-Transformers* library⁴ (Pfeiffer et al., 2020). For $\langle baseLM \rangle + AD$, we perform a grid search over the dropout values $\{0.2, 0.4, 0.6, 0.8\}$ and pick 0.6 since it gives the best performance on the intrinsic evaluations. A detailed description of all our hyperparameters is available in Appendix §A.1.

4 Results and Analysis

In this section, we primarily analyse the performance of BERT-based models. We find that the similar trends are observed on other models (mBERT, RoBERTa, XLM-R) as well. Exact numbers on these models can be found in Appendix A.3. Furthermore, MAFIA in this section refers to fusion of our full set (gender, race, religion and profession) of bias dimensions. We perform ablations with fusing subsets of biases in Appendix A.5.

4.1 Effectiveness of CF Pairs

Table 1 compares the performance of various BERT-based DBA models on intrinsic measures. We find that across all the bias dimensions, iDeb consistently outperforms all other baselines, highlighting the value of our larger inclusive CF Pair dataset (§2.1). Overall, all debiasing methods result in degradation of LM score when compared to the vanilla BERT. As we see in the next section, the decrease in LM score has unexpected consequences on the downstream task performance.

4.2 Effectiveness of AdapterFusion

Table 2 presents the performance of various BERT-based baselines on STS-B and Bias-STB tasks. Various trends can be observed in this table. For iDeb_{race}, we find both Δ_{gender} and Δ_{race} to be better than the baseline BERT meaning that debiasing across one dimension indeed has (often) unintended effects on other dimensions. This is also in line with the observations of Meade et al. (2022)

iDeb_{all} performs better than iDeb baselines in terms of Pearson correlation but results on Bias-STB are mostly poor, which means that a single adapter trained on CDA from all bias dimensions finds it difficult to effectively debias

Dim.	Model	Stereoset SS [†]	CrowSPairs SS [†]	LM Score (†)
Gender	BERT	60.28	57.25	84.17
	BERT+AD	60.00	57.16	75.16
	ADELE	59.61	53.81	82.91
	iDeb _{gender}	57.14	52.05	70.36
Race	BERT	57.03	62.33	84.17
	BERT+AD	56.98	62.00	75.16
	iDeb _{race}	51.87	58.92	80.23
Religion	BERT	59.71	62.86	84.17
	BERT+AD	58.66	62.75	75.16
	iDeb _{religion}	55.31	60.00	79.41

Table 1: **Intrinsic evaluation results across Gender, Race, and Religion bias.** StereoSet scores (marked with †) close to 50 indicate a less biased model whereas models with higher LM scores are better. Our inclusive CDA process leads to consistently less biased models (iDeb_{bias}). All baselines seem to reduce the bias at the cost of LM score.

across all the dimensions. In contrast, the modular AdapterFusion-based MAFIA composes knowledge from multiple DBAs and outperforms iDeb_{all} in all aspects.

We also find that iDeb baselines perform poorly in terms of the actual STS-B task. To investigate this better, we sampled 1000 examples from STS-B and compared the score distributions (Fig. 3) from various models. In particular, iDeb models become overly conservative (they output very similar scores for almost any pair) after debiasing as evidenced by a significant reduction in their score variance. The original BERT model shows a reasonable spread of scores but is biased whereas MAFIA is both fairer and more accurate.

We further present a qualitative analysis of the behavior of the model on 5 handcrafted pairs in Table 3. Here, we find even more evidence of iDeb models becoming overly conservative. In the first row, despite the sentences being very similar and void of any biased identity, iDeb models still predict scores close to their average scores while the MAFIA produces a reasonable score. In rows 2 and 3, two completely irrelevant sentences with gender identity are provided, while rows 4 and 5 consist of two similar sentences but with a racial identity. Baseline BERT gives somewhat acceptable scores for all 4 pairs but the difference between rows 2,3 and rows 4,5 is comparatively higher as compared to other models. This indicates that the model is accurate but biased. iDeb models on the other hand have lower differences in scores between rows 2,3

⁴<https://github.com/adapter-hub/adapter-transformers>

Model	STS-B	Bias-STS-B				Useful fairness
	Pearson (\uparrow)	Δ_{gender} (\downarrow)	Δ_{race} (\downarrow)	Δ_{religion} (\downarrow)	Δ_{average} (\downarrow)	Ψ_{average} (\uparrow)
BERT	0.78	0.18	0.09	0.07	0.11	0.69
BERT+DA	0.75	0.15	0.02	0.12	0.10	0.67
iDeb _{gender}	0.66	0.09	0.10	0.09	0.09	0.60
iDeb _{race}	0.46	0.09	0.06	0.19	0.11	0.41
iDeb _{religion}	0.45	0.19	0.09	0.06	0.11	0.40
iDeb _{profession}	0.45	0.15	0.11	0.12	0.13	0.39
iDeb _{all}	0.71	0.15	0.10	0.07	0.11	0.63
MAFIA	0.84	0.12	0.06	0.05	0.07	0.77

Table 2: **Extrinsic evaluation on STS-B and Bias-STS-B.** \uparrow indicates the metric is better when it is higher whereas \downarrow indicates the metric is better when it is lower. Best value for each metric is highlighted in bold. MAFIA outperforms all other baselines on STS-B and is the least biased (average) model on Bias-STS-B.

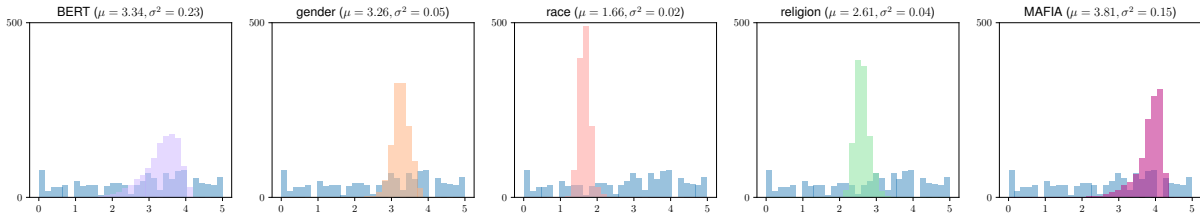


Figure 3: **Score distributions on STS-B obtained from various models.** The middle 3 plots correspond to iDeb_{bias} baselines. All iDeb_{bias} models output a significantly narrower score distribution which can easily lead to better scores on Bias-STS-B but can decrease the performance on STS-B.

and rows 4,5 but the scores themselves do not align well with human judgement. MAFIA scores are similar to BERT scores and are reasonable while the difference between rows 2,3 and rows 4,5 is also relatively less meaning that the MAFIA model is both accurate *and* fairer.

4.3 Zero-shot Cross-lingual Fairness Transfer

Lauscher et al. (2021) observe a zero-shot fairness transfer to non-English languages despite debiasing mBERT with only English data. While the results were encouraging, their evaluations included only bias results (without task performance) on mostly high-resource languages. We study the zero-shot debiasing ability of our models on a wider spectrum of language class taxonomy (provided by Joshi et al. (2020)) viz. *Class 5: English (En), French (Fr); Class 4: Italian (It), Hindi (Hi); Class 3: Tamil (Ta); Class 2: Marathi (Mr), Swahili (Sw); and Class 1: Gujarati (Gu).*

mSTS-B and mBias-STS-B: To systematically evaluate the multilingual performance of MAFIA, we translate the STS-B test set from English to the aforementioned target languages. We use

IndicTrans2⁵ (AI4Bharat et al., 2023) for Indic languages (Hindi, Marathi, Tamil, and Gujarati) and NLLB model⁶ (NLLB Team et al., 2022) for the rest (French, Italian and Swahili). Since machine translation can be incorrect or non-colloquial, we get the translations for Hindi, Tamil, Marathi, Swahili, and Gujarati, manually verified by native speakers in our research group. We plan to verify the remaining languages subsequently. Please refer to Appendix A.4 for more details about translation quality. This translated and human-verified dataset will be made public for future work.

We present the performance of mBERT and XLM-R models as measured by “useful fairness” in Table 4. We find that MAFIA offers improvements in “useful fairness” of the models on all languages except Swahili. On Swahili, we see a significant *decrease* in useful fairness on mBERT but a notable improvement on XLM-R. This could be due to differences in pretraining corpus as well as methods of pretraining of these models. A thorough investi-

⁵<https://huggingface.co/ai4bharat/indictrans2-en-indic-1B>

⁶<https://huggingface.co/facebook/nllb-200-3.3-B>

	Sentence 1	Sentence 2	BERT	iDeb _{gender}	iDeb _{race}	MAFIA
1	An airplane took off from NYC for LA	An airplane is travelling from NYC to LA	4.09	3.50	2.51	4.19
2	A man is flying a plane	A train has crashed into a valley	1.25	3.16	1.55	1.35
3	A woman is flying a plane	A train has crashed into a valley	1.79	3.19	1.57	1.45
4	An african american scientist is conducting experiments	Experiments are being conducted by scientists in a lab	3.43	3.24	1.81	3.67
5	An asian scientist is conducting experiments	Experiments are being conducted by scientists in a lab	2.99	3.27	1.75	3.61

Table 3: **Qualitative analysis.** Comparison of semantic textual similarity scores outputted for different sentence pairs by various models. iDeb_{bias} seems to be outputting scores that are too close to each other for similar as well as dissimilar pairs. This can explain the decrease in LM score as well as lower Pearson coefficient.

Model	en	fr	it	hi	ta	mr	sw	gu
mBERT	0.65	0.56	0.57	0.52	0.52	0.58	0.49	0.52
mBERT _M	0.71	0.58	0.57	0.62	0.55	0.65	0.26	0.59
XLM-R	0.18	0.15	0.13	0.20	0.09	0.09	0.07	0.16
XLM-R _M	0.57	0.50	0.48	0.48	0.49	0.47	0.33	0.49

Table 4: **Useful fairness (Ψ_{average}) of models on non-English languages.** mBERT_M and XLM-R_M are MAFIA versions of mBERT and XLM-R respectively. MAFIA improves useful fairness of models on most language classes despite being debiased in English. Title row is color-coded based on the language class.

Model	BPSN (\uparrow)	BNSP (\uparrow)	Overall (\uparrow)
BERT	0.86	0.92	0.88
BERT+AD	0.86	0.87	0.87
iDeb _{gender}	0.86	0.88	0.86
iDeb _{race}	0.85	0.91	0.87
iDeb _{religion}	0.85	0.91	0.87
MAFIA	0.86	0.95	0.89

Table 5: **Comparing average submetrics on Jigsaw.** \uparrow indicates that metric is better with higher value. MAFIA is the only model that outperforms the baseline BERT.

gation might be necessary to identify the root cause of this behavior.

4.4 Case Study: Toxicity Classification

The Kaggle competition ‘‘Jigsaw’’⁷ aims to address the issue of toxicity detection models picking up unintended biases due to the overrepresentation of certain identities in toxic comments. For example, many toxicity detection models will correctly classify the sentence ‘‘Death to all gay people’’. However, the competition observed that many such classifiers became unintentionally biased towards a subgroup of identities and incorrectly flagged even benign sentences such as ‘‘I am a gay man’’ as toxic. The Jigsaw competition uses a special metric designed to address this issue in toxicity evaluation. We find that MAFIA provides meaningful improvements over the baseline BERT on this metric.

The Jigsaw metric is a mean of ROC-AUC scores restricted to specific bias subsets along with the overall AUC on the entire test set. To calculate bias AUCs, three separate AUCs are calculated for every identity. The set of identities is predetermined by the competition organizers and annotations are provided with each sample about the identities men-

tioned in the comment. For each identity subgroup (s), we calculate 3 ROC-AUC scores as 3 different sub-metrics (m_s):

1. Subgroup: AUC on the subset of test set mentioning that specific identity.

2. Background Positive, Subgroup Negative (BPSN): AUC on the subset of test set with non-toxic examples that mention the identity and toxic examples that do not.

3. Background Negative, Subgroup Positive (BNSP): AUC on the subset of test set with toxic examples that mention the identity and non-toxic examples that do not.

Overall score is a combination of generalized mean of these submetrics along with ROC-AUC on the entire test set. More details about the Overall score are presented in Appendix A.2. We compare MAFIA against BERT, BERT+AdapterDrop, and iDeb variants using average BPSN, BNSP and Overall scores in Table 5. The model architecture for each baseline is exactly same as Bias-STS-B and described in Fig. 2. Results indicate that iDeb variants and the AdapterDrop baseline get *lower* BNSP scores on average. This means that these models confuse toxic examples that mention the identity with non-toxic examples that do not. These

⁷<https://www.kaggle.com/c/jigsaw-unintended-bias-in-toxicity-classification>

517	findings are in line with our observations on STS-B	565
518	where iDeb baselines would output scores close to	
519	their mean values and not deviate much.	
520	While many of the subgroups in Jigsaw are re-	566
521	lated to gender, race, or profession, one subgroup	567
522	is about “psychiatric or mental illness” which is	568
523	<i>not</i> covered by any of our DBAs. Despite this,	569
524	MAFIA is able to provide fairness <i>and</i> accuracy	570
525	(AUC) gains over this. Detailed subgroup-level	571
526	metrics are presented in Appendix A.2. This shows	572
527	that MAFIA is able to better exploit the synergy	573
528	between various biases and even provide fairness	574
529	and performance gains on intersectional biases pre-	575
530	viously unseen during debiasing.	576
531	5 Related Work	577
532	5.1 Adapters and Modular Deep Learning	578
533	Adapters (Rebuffi et al., 2017; Houlsby et al., 2019;	579
534	Stickland and Murray, 2019) are small neural mod-	580
535	ules introduced between each layer of a larger net-	581
536	work. Adapter-based finetuning has been shown	582
537	to be as effective as full model finetuning while	583
538	being $\sim 60\%$ more efficient than full finetuning	584
539	(Rücklé et al., 2021). AdapterFusion (Pfeiffer et al.,	585
540	2021) allows composing knowledge from multiple	586
541	adapters in a non-destructive way. This motivated	587
542	us to train individual DBAs and combine them	588
543	using AdapterFusion to exploit the synergy of mul-	589
544	tiiple biases to debias across multiple dimensions	590
545	simultaneously.	591
546	5.2 Correcting Biases in Pretrained LLMs	592
547	Gender bias is one of the well-studied biases in	593
548	LLMs and a large body of work exists that aims to	594
549	correct solely gender bias (Sun et al., 2019; Zhao	595
550	et al., 2017; Ma et al., 2020; Dev et al., 2021, <i>inter</i>	
551	<i>alia</i>). Several other methods have been explored	
552	for correcting biases in pretrained LLMs includ-	
553	ing dropout regularization (Webster et al., 2020),	
554	information-theoretic methods (Cheng et al., 2020;	
555	Colombo et al., 2021), contrastive learning (Cheng	
556	et al., 2021; Zhang et al., 2021) etc. In this work,	
557	we focus on task-agnostic debiasing techniques that	
558	are more generalizable than task-specific debiasing	
559	models, which need to be tailored for each task and	
560	dataset. In our work, we focus on counterfactual	
561	data augmentation based (CDA) based debiasing	
562	methods (Zmigrod et al., 2019; Dinan et al., 2020;	
563	Webster et al., 2020; Barikeri et al., 2021) to train	
564	debiasing adapters for each of our bias dimensions.	
	5.3 Adapter-based Debiasing for LLMs	
	The concept of adapter-based debiasing was ex-	
	plored by Lauscher et al. (2021), where they pre-	
	sented a binary gender-only debiasing adapter, lim-	
	ited by using a small hand-built, US-centric CDA	
	for training. They subsequently fine-tuned entire	
	models for specific tasks. Contrary to their ap-	
	proach, we use a larger and inclusive CDA train-	
	ing (§2.1) for multiple societal biases and finetune	
	<i>only</i> the adapters on downstream tasks. We fur-	
	ther illustrate in Section §4.2 that sole reliance on	
	adapter-only fine-tuning can sometimes produce	
	unexpected outcomes for downstream tasks. How-	
	ever, their achievements in debiasing and zero-shot	
	cross-lingual transfer proved promising. Our re-	
	search has parallels with the study by Kumar et al.	
	(2023), which also adopted AdapterFusion. Their	
	method, however, intertwined both task and debias-	
	ing objectives (which is expensive as they use ad-	
	versarial training for debiasing) to learn the fusion	
	weights. In contrast, our approach learns fusion	
	weights using solely the task objective, which is	
	generally more straightforward to optimize. Be-	
	sides using a more inclusive semi-automated CDA	
	training, our study is enriched by a series of abla-	
	tion tests (Table 11) across diverse bias dimensions.	
	We not only include a comprehensive range of bias	
	components (for instance, considering non-binary	
	aspects in gender bias) but also delve into under-	
	standing and evaluating the possible shortcomings	
	of singular DBA configurations (Fig. 3, Table 3).	
	6 Conclusion	
	We proposed a method called MAFIA that uses	
	AdapterFusion to leverage the interaction of multi-	
	ple bias dimensions to debias a PLM. Our method	
	works by training debiasing adapters for individ-	
	ual biases and then fusing them on a downstream	
	task for multidimensional debiasing. We employed	
	counterfactual data augmentation to train each of	
	the individual debiasing adapters. We use a semi-	
	automatic method to generate diverse and inclusive	
	counterfactual pairs for a given bias dimension with	
	the help of large generative models and structured	
	knowledge bases. Our evaluation indicates that	
	MAFIA leads to a fairer and more accurate model	
	on downstream tasks across multiple languages and	
	across various bias dimensions, including poten-	
	tially unseen ones during training.	

7 Limitations

We present some limitations of our current work, which we wish to address in some future work:

1. In this work, we only explore the interplay between a limited set of biases, i.e., gender, race, religion, and profession, and agree that numerous other biases such as cultural and psychological biases have not been addressed. Similarly, we select a limited set of high and low-resource languages for zero-shot evaluation.

2. Our CF pairs are limited by the knowledge of text-davinci-003 and presence in WikiData, and the numbers are restricted by the frequency of the occurrence of the pairs in Google Book Corpus.

3. We also acknowledge that our AdapterFusion tuned on the downstream task, instead of “pretraining” it on combined CDA data, which makes it task-specific and not generic.

4. We only investigate the effect of fusion on a single downstream, i.e., Bias-STS-B, and replicating these findings on other tasks like Bias-NLI would be an interesting study.

5. Lastly, we were also constrained by our limited computational resources, as “pretraining” the debiasing adapters consumed a significant time for larger models like RoBERTa and XLM-R.

8 Ethical Considerations

We use the framework by [Bender and Friedman \(2018\)](#) to discuss the ethical considerations for our work.

- **Data:** The counterfactual pairs were generated using API calls to text-davinci-003. The counterfactual pairs generated for each bias are released along with this paper. The dataset was created with the intention of studying societal biases and debiasing PLMs. We start with a broader set of bias identities obtained from Wikidata. Note that the intent was not to hurt/harm anyone.
- **Methods:** In this study, we explore several methods for debiasing PLMs and evaluate them on various end tasks and languages. These methods are primarily designed for the English language, they may not perform equally well for all languages of the world.

References

- AI4Bharat, Jay Gala, Pranjal A. Chitale, Raghavan AK, Sumanth Doddapaneni, Varun Gumma, Aswanth Kumar, Janki Nawale, Anupama Sujatha, Ratish Puduppully, Vivek Raghavan, Pratyush Kumar, Mitesh M. Khapra, Raj Dabre, and Anoop Kunchukuttan. 2023. [IndicTrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages](#).
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Soumya Barikeri, Anne Lauscher, Ivan Vulić, and Goran Glavaš. 2021. [RedditBias: A real-world resource for bias evaluation and debiasing of conversational 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 1941–1955, Online. Association for Computational Linguistics.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Shikha Bordia and Samuel R. Bowman. 2019. [Identifying and reducing gender bias in word-level language models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 7–15, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pengyu Cheng, Weituo Hao, Siyang Yuan, Shijing Si, and Lawrence Carin. 2021. [Fairfil: Contrastive neural debiasing method for pretrained text encoders](#). In *International Conference on Learning Representations*.
- Pengyu Cheng, Martin Renqiang Min, Dinghan Shen, Christopher Malon, Yizhe Zhang, Yi-

707	tong Li, and Lawrence Carin. 2020. Improving	<i>Association for Computational Linguistics: Hu-</i>	756
708	disentangled text representation learning with	<i>man Language Technologies, Volume 1 (Long</i>	757
709	information-theoretic guidance . In <i>Proceedings</i>	<i>and Short Papers)</i> , pages 4171–4186, Minneapoli-	758
710	<i>of the 58th Annual Meeting of the Association for</i>	<i>olis, Minnesota. Association for Computational</i>	759
711	<i>Computational Linguistics</i> , pages 7530–7541,	<i>Linguistics.</i>	760
712	Online. Association for Computational Linguistics.		
713			
714	Pierre Colombo, Pablo Piantanida, and Chloé	Emily Dinan, Angela Fan, Adina Williams, Jack	761
715	Clavel. 2021. A novel estimator of mutual information	Urbanek, Douwe Kiela, and Jason Weston. 2020.	762
716	for learning to disentangle textual representations .	Queens are powerful too: Mitigating gender bias	763
717	In <i>Proceedings of the 59th Annual</i>	in dialogue generation . In <i>Proceedings of the</i>	764
718	<i>Meeting of the Association for Computational</i>	<i>2020 Conference on Empirical Methods in Natural</i>	765
719	<i>Linguistics and the 11th International Joint Confer-</i>	<i>Language Processing (EMNLP)</i> , pages 8173–	766
720	<i>ence on Natural Language Processing (Volume 1: Long</i>	8188, Online. Association for Computational	767
721	<i>Papers)</i> , pages 6539–6550, Online.	<i>Linguistics.</i>	768
722	Association for Computational Linguistics.		
723			
724	Alexis Conneau, Kartikay Khandelwal, Naman	Stefan Elfving, Eiji Uchibe, and Kenji Doya. 2017.	769
725	Goyal, Vishrav Chaudhary, Guillaume Wenzek,	Sigmoid-weighted linear units for neural net-	770
726	Francisco Guzmán, Edouard Grave, Myle Ott,	work function approximation in reinforcement	771
727	Luke Zettlemoyer, and Veselin Stoyanov. 2020.	learning .	772
728	Unsupervised cross-lingual representation learn-		
729	ing at scale . In <i>Proceedings of the 58th Annual</i>		
730	<i>Meeting of the Association for Computational</i>		
731	<i>Linguistics</i> , pages 8440–8451, Online. Association	David Freedman, Robert Pisani, and Roger Purves.	773
732	<i>tion for Computational Linguistics.</i>	2007. Statistics (international student edition).	774
733		<i>Pisani, R. Purves, 4th edn. WW Norton & Com-</i>	775
734		<i>pany, New York.</i>	776
735			
736			
737		Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski,	777
738		Bruna Morrone, Quentin De Laroussilhe,	778
739		Andrea Gesmundo, Mona Attariyan, and Sylvain	779
740		Gelly. 2019. Parameter-efficient transfer	780
741		learning for NLP . In <i>Proceedings of the 36th</i>	781
742		<i>International Conference on Machine Learning,</i>	782
743		volume 97 of <i>Proceedings of Machine Learning</i>	783
744		<i>Research</i> , pages 2790–2799. PMLR.	784
745			
746			
747		Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika	785
748		Bali, and Monojit Choudhury. 2020. The	786
749		state and fate of linguistic diversity and inclu-	787
750		sion in the NLP world . In <i>Proceedings of the</i>	788
751		<i>58th Annual Meeting of the Association for Com-</i>	789
752		<i>putational Linguistics, ACL 2020, Online, July</i>	790
753		<i>5-10, 2020</i> , pages 6282–6293. Association for	791
754		<i>Computational Linguistics.</i>	792
755			
756			
757		Deepak Kumar, Oleg Lesota, George Zerveas,	793
758		Daniel Cohen, Carsten Eickhoff, Markus Schedl,	794
759		and Navid Rekabsaz. 2023. Parameter-efficient	795
760		modularised bias mitigation via AdapterFusion .	796
761		In <i>Proceedings of the 17th Conference of the Eu-</i>	797
762		<i>ropean Chapter of the Association for Computa-</i>	798
763		<i>tional Linguistics</i> , pages 2738–2751, Dubrovnik,	799
764		<i>Croatia. Association for Computational Linguis-</i>	800
765		<i>tics.</i>	801
766			
767			
768		Matt J. Kusner, Joshua R. Loftus, Chris Russell,	802
769		and Ricardo Silva. 2017. Counterfactual fairness .	803

804	In <i>Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA</i> , pages 4066–4076.	<i>Computational Linguistics</i> , pages 5502–5515,	853
805		Online. Association for Computational Linguistics.	854
806			855
807			
808	Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. 2023. Openassistant conversations – democratizing large language model alignment .	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pre-training approach .	856
809			857
810			858
811			859
812			860
813		Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization . In <i>International Conference on Learning Representations</i> .	861
814			862
815			863
816			
817	Anne Lauscher, Tobias Lueken, and Goran Glavaš. 2021. Sustainable modular debiasing of language models . In <i>Findings of the Association for Computational Linguistics: EMNLP 2021</i> , pages 4782–4797, Punta Cana, Dominican Republic. Association for Computational Linguistics.	Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2020. Gender bias in neural natural language processing . In <i>Logic, Language, and Security - Essays Dedicated to Andre Scedrov on the Occasion of His 65th Birthday</i> , volume 12300 of <i>Lecture Notes in Computer Science</i> , pages 189–202. Springer.	864
818			865
819			866
820			867
821			868
822			869
823			870
824	Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umaphathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhat-tacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2023. Starcoder: may the source be with you!	Xinyao Ma, Maarten Sap, Hannah Rashkin, and Yejin Choi. 2020. PowerTransformer: Unsupervised controllable revision for biased language correction . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 7426–7441, Online. Association for Computational Linguistics.	871
825			872
826			873
827			874
828			875
829			876
830			877
831		Thomas Manzini, Lim Yao Chong, Alan W Black, and Yulia Tsvetkov. 2019. Black is to criminal as Caucasian is to police: Detecting and removing multiclass bias in word embeddings . In <i>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)</i> , pages 615–621, Minneapolis, Minnesota. Association for Computational Linguistics.	878
832			879
833			880
834			881
835			882
836			883
837			884
838			885
839			886
840			887
841		Nicholas Meade, Elinor Poole-Dayyan, and Siva Reddy. 2022. An empirical survey of the effectiveness of debiasing techniques for pre-trained language models . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 1878–1898, Dublin, Ireland. Association for Computational Linguistics.	888
842			889
843			890
844			891
845			892
846			893
847			894
848			895
849	Paul Pu Liang, Irene Mengze Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. Towards debiasing sentence representations . In <i>Proceedings of the 58th Annual Meeting of the Association for</i>	Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. StereoSet: Measuring stereotypical bias in pre-trained language models . In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International</i>	896
850			897
851			898
852			899
			900

901	<i>Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 5356–5371, Online. Association for Computational Linguistics.	950
902		951
903		952
904		
905	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 <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 1953–1967, Online. Association for Computational Linguistics.	953
906		954
907		955
908		956
909		957
910		958
911		959
912		
913	NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation .	960
914		961
915		962
916		963
917		964
918		965
919		966
920		967
921		968
922		
923		969
924		970
925		971
926		972
927		973
928		
929	Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. AdapterFusion: Non-destructive task composition for transfer learning . In <i>Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume</i> , pages 487–503, Online. Association for Computational Linguistics.	974
930		975
931		976
932		977
933		978
934		979
935		980
936		
937	Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020. AdapterHub: A framework for adapting transformers . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations</i> , pages 46–54, Online. Association for Computational Linguistics.	981
938		982
939		983
940		984
941		985
942		986
943		987
944		988
945		989
946		
947	Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null it out: Guarding protected attributes by iterative nullspace projection . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 7237–7256, Online. Association for Computational Linguistics.	990
948		991
949		992
		993
		994
		995
		996
		997
	Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2017. Learning multiple visual domains with residual adapters . In <i>Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA</i> , pages 506–516.	
	Andreas Rücklé, Gregor Geigle, Max Glockner, Tilman Beck, Jonas Pfeiffer, Nils Reimers, and Iryna Gurevych. 2021. AdapterDrop: On the efficiency of adapters in transformers . In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 7930–7946, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.	
	Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in NLP . <i>Transactions of the Association for Computational Linguistics</i> , 9:1408–1424.	
	Asa Cooper Stickland and Iain Murray. 2019. BERT and PALs: Projected attention layers for efficient adaptation in multi-task learning . In <i>Proceedings of the 36th International Conference on Machine Learning</i> , volume 97 of <i>Proceedings of Machine Learning Research</i> , pages 5986–5995. PMLR.	
	Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating gender bias in natural language processing: Literature review . In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 1630–1640, Florence, Italy. Association for Computational Linguistics.	
	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need . In <i>Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA</i> , pages 5998–6008.	

998 Denny Vrandečić and Markus Krötzsch. 2014.
999 [Wikidata: A free collaborative knowledgebase.](#)
1000 *Commun. ACM*, 57(10):78–85.

1001 Alex Wang, Amanpreet Singh, Julian Michael,
1002 Felix Hill, Omer Levy, and Samuel Bowman.
1003 2018. [GLUE: A multi-task benchmark and anal-](#)
1004 [ysis platform for natural language understanding.](#)
1005 In *Proceedings of the 2018 EMNLP Workshop*
1006 *BlackboxNLP: Analyzing and Interpreting Neu-*
1007 *ral Networks for NLP*, pages 353–355, Brussels,
1008 Belgium. Association for Computational Lin-
1009 guistics.

1010 Kellie Webster, Xuezhi Wang, Ian Tenney, Alex
1011 Beutel, Emily Pitler, Ellie Pavlick, Jilin Chen,
1012 Ed H. Chi, and Slav Petrov. 2020. [Measuring and](#)
1013 [reducing gendered correlations in pre-trained](#)
1014 [models.](#) Technical report.

1015 Xiongyi Zhang, Jan-Willem van de Meent, and
1016 Byron Wallace. 2021. [Disentangling represen-](#)
1017 [tations of text by masking transformers.](#) In *Pro-*
1018 *ceedings of the 2021 Conference on Empirical*
1019 *Methods in Natural Language Processing*, pages
1020 778–791, Online and Punta Cana, Dominican
1021 Republic. Association for Computational Lin-
1022 guistics.

1023 Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente
1024 Ordonez, and Kai-Wei Chang. 2017. [Men also](#)
1025 [like shopping: Reducing gender bias amplifi-](#)
1026 [cation using corpus-level constraints.](#) In *Pro-*
1027 *ceedings of the 2017 Conference on Empirical*
1028 *Methods in Natural Language Processing*, pages
1029 2979–2989, Copenhagen, Denmark. Association
1030 for Computational Linguistics.

1031 Ran Zmigrod, Sabrina J. Mielke, Hanna Wallach,
1032 and Ryan Cotterell. 2019. [Counterfactual data](#)
1033 [augmentation for mitigating gender stereotypes](#)
1034 [in languages with rich morphology.](#) In *Proceed-*
1035 *ings of the 57th Annual Meeting of the Associa-*
1036 *tion for Computational Linguistics*, pages 1651–
1037 1661, Florence, Italy. Association for Computa-
1038 tional Linguistics.

A Appendix

A.1 Hyperparameters

In this section, we describe the hyperparameter set we used for training the debiasing, task and fusion adapters. All our experiments are performed on a single NVIDIA A100 GPU with 80GB VRAM.

Hyperparameter	Value
Learning rate	3×10^{-5}
Epochs	2
Global Batch size	512 for BERT, RoBERTa, mBERT; 256 for XLM-R
Scheduler	Cosine
Warmup	Linear
Warmup ratio	0.1
Optimizer	AdamW (Loshchilov and Hutter, 2019)
Weight decay	0
Adapter architecture	Pfeiffer
Adapter activation	SiLU (Elfving et al., 2017)
Adapter reduction factor	16
FP16	True
MLM probability	0.15

Table 6: Hyperparameters for training individual DBAs.

Hyperparameter	Value
Learning rate	2×10^{-5}
Epochs	10
Global Batch size	512 for BERT, RoBERTa, mBERT; 256 for XLM-R
Scheduler	Cosine
Warmup	Linear
Warmup ratio	0.1
Optimizer	AdamW (Loshchilov and Hutter, 2019)
Weight decay	0
Adapter architecture	Pfeiffer
Adapter activation	SiLU (Elfving et al., 2017)
Adapter reduction factor	16
FP16	True

Table 7: Hyperparameters for finetuning on downstream (STS-B and Jigsaw) tasks.

A.2 Jigsaw: Unintended Bias in Toxicity Classification

Here we provide additional details about the overall score computation as well as various identity subgroups considered in the ‘‘Jigsaw’’ task⁸.

Computing ‘‘Overall’’ score. After computing each of the subgroup (s) submetrics, for each submetric (m_s), we calculate the generalized

⁸<https://www.kaggle.com/c/jigsaw-unintended-bias-in-toxicity-classification>

Subgroup	Count	% Imp
black	1519	3.29
white	2452	1.96
female	5155	2.26
male	4386	2.69
homosexual_gay_or_lesbian	1065	3.93
muslim	2040	3.33
jewish	835	2.50
christian	4226	1.63
psychiatric_or_mental_illness	511	1.33

Table 8: Size of a particular subgroup in the Jigsaw test set. We also report subgroup AUC improvement in percentage that MAFIA based toxicity classifier brings over a classifier using vanilla LM + task adapter.

mean over N identity subgroups with power p as $M_p(m_s) = (\frac{1}{N} \sum_{s=1}^N m_s^p)^{\frac{1}{p}}$. The overall score for a model is computed as:

$$\text{Overall} = w_0 \text{AUC}_{\text{overall}} + \sum_{a=1}^A w_a M_p(m_{s,a})$$

Where $\text{AUC}_{\text{overall}}$ is the ROC-AUC on the entire test set, $A = 3$ is the number of submetrics described above, $m_{s,a}$ represents the value of submetric a on identity group s . Default values for p and w are -5 and 0.25 respectively.

Subgroups. Table 8 shows each subgroup identity along with their count in the test set. We also report the improvement obtained in subgroup AUC by MAFIA over the base model for each subgroup. Despite no explicit debiasing for *psychiatric_or_mental_illness*, we observe gain in performance as well as fairness on that subgroup.

A.3 Results on other models

In this section, we discuss performance on MAFIA with other base language models. Specifically, we present our findings on RoBERTa and XLM-RoBERTa (XLM-R) models. Results on intrinsic evaluation (Table 9) indicate that our proposed general purpose, semi-automatic CDA method is effective in debiasing RoBERTa as well as XLM-R. Interestingly, even when XLM-R is already very less biased on some dimensions, our method still offers small gains on top. On downstream tasks, we find that MAFIA increases the useful fairness of both models. However, we also observe that gender bias on XLM-R *worsens* after the fusion! It is likely that on XLM-R, a smaller subset of DBA can perform better as seen via ablations in Appendix A.5. We were unable to conduct such a large-scale study on XLM-R due to compute limitations.

<i>Dim.</i>	Model	StereoSet SS [†]	CrowSPairs SS [†]	LM Score \uparrow
<i>Gender</i>	RoBERTa	55.51	53.05	79.54
	iDeb _{gender}	54.60	52.85	75.36
<i>Race</i>	RoBERTa	56.31	53.10	79.54
	iDeb _{race}	52.33	52.15	78.67
<i>Religion</i>	RoBERTa	39.40	68.57	79.54
	iDeb _{religion}	45.89	62.10	79.41

<i>Dim.</i>	Model	StereoSet SS [†]	CrowSPairs SS [†]	LM Score \uparrow
<i>Gender</i>	XLM-R	50.36	56.10	77.68
	iDeb _{gender}	50.27	56.10	70.36
<i>Race</i>	XLM-R	51.94	52.52	77.68
	iDeb _{race}	50.85	52.19	76.23
<i>Religion</i>	XLM-R	50.20	64.76	77.68
	iDeb _{religion}	50.20	63.90	75.71

Table 9: Intrinsic evaluation results for RoBERTa and XLMR-R models. \dagger - StereoSet Score (SS) close to 50 indicates a less biased model.

A.4 Multilingual-Bias-STS-B and Bias-STS-B

In this section, we estimate the quality of the mBias-STS-B dataset we create. We randomly sampled 50 data points (translated sentence pairs) per language and got them verified for quality by native speakers in the group.

Translation Quality is an estimate of the % times the translations are correct. Swahili translations were 86.2% correct, Hindi translations were 90.9% correct, Marathi translations were 89.2% correct, Tamil translations were 100% correct, Gujarati translations were 80% correct ■.

For the mGender-bias-STS-B dataset, we want one of the two sentences to be gender-neutral, while one to be gender specific. Swahili is a gender-neutral language. For Hindi and Marathi, we corrected the templates for the respective languages to be gender-neutral. For the remaining languages, we requested the native speakers to estimate the number of times the condition fails. On Tamil the condition failed 2.1% times, while on Gujarati it never failed.

A.5 Fusion Ablations

In this section, we perform ablations to study whether fusing a *subset* of debiasing adapters (DBAs) over a downstream task may perform better than fusing *all* DBAs. We report our findings on BERT and mBERT in Table 11. On both the models, we observe that fusion of *gender* and *race*

gives the best STS-B performance but worsens the bias. We also find that individual adapters often give the *most* debiased model at the cost of performance on STS-B. On mBERT, we find that the fusion of *gender* and *profession* is the best model in terms of both STS-B and Bias-STS-B. This shows that fusing *all* available DBA may not be required for building a model that is both accurate and fair. Finding the minimal set of DBAs to be fused for best performance on all bias dimensions as well as the downstream task is an interesting problem that needs more attention. Future works can explore this interaction better.

A.6 Counterfactual Pairs

<i>Category</i>	Property Code	Property Description
<i>Gender</i>	P3321	male form of label
	P6553	personal pronoun
	P21	sex or gender
	P5185	grammatical gender
<i>Race</i>	P27	country of citizenship
	P172	ethnic group
	Q874405	human social group
	Q3254959	human race
<i>Religion</i>	P1049	worshipped by
	P140	religion or worldview
	Q178885	deity
	Q9174	religion
	Q375011	religious festival
	Q4392985	religious identity
	Q21029893	religious object
	Q105889895	religious site
Q179461	religious text	
Q1370598	structure of worship	
Q71966963	religion or world view	
<i>Profession</i>	P101	field of work
	P106	occupation
	P3095	practiced by

Table 12: Codes respective descriptions extracted from WikiData to create the CF pairs.

To extract gender terms, we use properties P3321, P6553, P21, P5185. For race terms, we use P27, P172, Q874405, Q3254959. For religion terms, we use P1049, P140, Q178885, Q9174, Q375011, Q4392985, Q21029893, Q105889895, Q179461, Q1370598, and Q71966963. For profession terms, we use properties P101, P106, P3095.

Gender CF Pairs: (bi-gender, non-binary) (boy, girl) (boys, girls) (cei, cea) (cissexual, transgender) (demi-man, demi-woman) (doctorate, doctorette) (fa’afafine, fa’afatama) (female, male) (fey, fae) (gender-fluid, gender-fluid) (gender-free, gender-

Model	STS-B	Bias-STS-B				Useful fairness
	Pearson (\uparrow)	Δ_{gender} (\downarrow)	Δ_{race} (\downarrow)	Δ_{religion} (\downarrow)	Δ_{average} (\downarrow)	Ψ_{average} (\uparrow)
RoBERTa	0.39	0.08	0.06	0.05	0.06	0.37
MAFIA	0.45	0.06	0.04	0.05	0.05	0.42
XLM-R	0.20	0.11	0.14	0.13	0.13	0.18
MAFIA	0.72	0.46	0.11	0.08	0.22	0.57

Table 10: **Extrinsic evaluation on RoBERTa and XLM-R.** We observe gains in useful fairness similar to BERT. On XLM-R, the gender bias seems to worsen with MAFIA. This could be due to XLM-R already being fairer on gender (Table 9) having trained on a much larger pretraining dataset and our gender DBA narrowed the domain to Wikipedia. Further fusion based ablations (similar to Table 11) can also help shed more light on this.

Model	STS-B	Bias-STS-B				Useful fairness
	Pearson	Δ_{gender}	Δ_{race}	Δ_{religion}	Δ_{average}	Ψ_{average}
BERT	0.78	0.18	0.09	0.07	0.11	0.69
gender (gen)	0.66	0.09	0.10	0.09	0.09	0.60
race (rac)	0.46	0.09	0.06	0.19	0.11	0.41
religion (rel)	0.45	0.19	0.09	0.06	0.11	0.40
profession (pro)	0.45	0.15	0.11	0.12	0.13	0.39
gen + rac	0.86	0.28	0.14	0.11	0.18	0.70
gen + rel	0.85	0.34	0.17	0.16	0.22	0.66
gen + pro	0.82	0.10	0.09	0.06	0.08	0.76
rac + rel	0.81	0.13	0.09	0.07	0.10	0.73
rac + pro	0.83	0.16	0.06	0.05	0.09	0.76
rel + pro	0.83	0.39	0.12	0.09	0.20	0.67
gen + rac + rel	0.85	0.17	0.10	0.09	0.12	0.74
gen + rac + pro	0.83	0.32	0.14	0.11	0.19	0.68
gen + rel + pro	0.85	0.31	0.13	0.12	0.18	0.69
rac + rel + pro	0.83	0.29	0.14	0.13	0.19	0.67
gen + rac + rel + pro	0.84	0.12	0.06	0.05	0.07	0.77

Model	STS-B	Bias-STS-B				Useful fairness
	Pearson	Δ_{gender}	Δ_{race}	Δ_{religion}	Δ_{average}	Ψ_{average}
mBERT	0.80	0.12	0.20	0.22	0.18	0.66
gender (gen)	0.25	0.06	0.05	0.06	0.06	0.24
race (rac)	0.52	0.28	0.14	0.14	0.18	0.42
religion (rel)	0.51	0.21	0.17	0.16	0.18	0.42
profession (pro)	0.37	0.07	0.07	0.07	0.07	0.35
gen + rac	0.82	0.27	0.33	0.23	0.28	0.59
gen + rel	0.80	0.14	0.12	0.09	0.12	0.71
gen + pro	0.78	0.09	0.08	0.05	0.08	0.73
rac + rel	0.80	0.32	0.39	0.44	0.38	0.50
rac + pro	0.77	0.09	0.06	0.07	0.07	0.72
rel + pro	0.81	0.27	0.21	0.21	0.23	0.62
gen + rac + rel	0.79	0.13	0.10	0.07	0.10	0.71
gen + rac + pro	0.79	0.12	0.07	0.05	0.08	0.72
gen + rel + pro	0.80	0.23	0.35	0.30	0.29	0.57
rac + rel + pro	0.80	0.21	0.36	0.32	0.30	0.56
gen + rac + rel + pro	0.78	0.09	0.09	0.08	0.09	0.71

Table 11: **Ablation studies on fusing a subset of bias adapters on BERT and mBERT.** On BERT, we find that fusion of all bias dimensions performs the best in terms of useful fairness. On mBERT, fusion of *gender* and *profession* gives the best results. Finding the minimal set of DBAs that will give the best performance can be an interesting future direction.

1136	free) (genderless, genderfluid) (girl, boy) (girls,	Muslim) (Hispanic, asian) (Ho people, caucasian)	1188
1137	boys) (god, goddess) (goddess, god) (guy, girl) (he,	(Hungarian, Indian) (Hungarians, Italians) (Hungary,	1189
1138	she) (hermaphrodite, intersex) (hermaphroditism,	Romania) (Huron, Navajo) (Inca, Aztec)	1190
1139	intersex) (intersexual, transgender) (intersexuality,	(India, China) (Indians, caucasian) (indigenous	1191
1140	intersexuality) (katoey, kathoey) (lady,	people, asian) (Iran, Saudi Arabia) (Irish, asian)	1192
1141	gentleman) (ladyboy, ladygirl) (laelae, laelae)	(Israel, Palestine) (Israelis, Palestinians) (Israelites,	1193
1142	(leiti, leita) (mahu, mahus) (man, lady) (masc,	Egyptians) (Italians, Irish) (Italy, Germany)	1194
1143	femme) (men, ladies) (neu, nai) (neut, fem)	(Jamaica, Mexico) (Jew, Muslim) (Judaism,	1195
1144	(neuter, feminine) (nongendered, gender-neutral)	Christianity) (Kahlan, Caucasian) (Kangeanese,	1196
1145	(non-gendered, gender-neutral) (omnigender,	asian) (Kerala, Tamil Nadu) (Khmer, Vietnamese)	1197
1146	nonbinary) (pan-gender, non-binary) (she, he)	(knife, spoon) (Korea, Japan) (Koreans, Chinese)	1198
1147	(trans, cis) (trans-feminine, trans-masculine)	(Kurdish, Arab) (Latin America, Asian) (Latino,	1199
1148	(transgendered, cisgendered) (transgenders,	caucasian) (Latvian American, Mexican American)	1200
1149	transgenders) (transman, transwoman) (trans-man,	(Lebanese, Indian) (Liu, Lee) (Lotud, asian)	1201
1150	trans-woman) (transmasculine, transfeminine)	(Malay, Indian) (Māori, asian) (Mexican American,	1202
1151	(trans-masculine, trans-feminine) (transmasculinity,	Native American) (Mexicans, asian) (Middle East,	1203
1152	transfemininity) (transpeople, cispeople)	South American) (Missouria, Cherokee) (Mixed,	1204
1153	(transwoman, transman) (trans-woman, trans-man)	asian) (Mongols, Native Americans) (monkey,	1205
1154	(travestism, transvestism) (two-spirits, two-spirit)	human) (Moors, Native Americans) (Morocco,	1206
1155	(ungendered, gender-neutral) (woman, man)	Japan) (Muslim, Christian) (negro, caucasian)	1207
1156	(women, men) (ze, zie)	(Nigeria, Japan) (Norwegian, Japanese) (Paiwan,	1208
1157		caucasian) (Palestinians, Israelis) (Persians, Arabs)	1209
1158	Race (ethnicity) CF Pairs: (Abydonian,	(Portuguese, Spanish) (pueblo, native american)	1210
1159	asian) (Africa, Asia) (African Americans, Native	(Romans, Greeks) (Russia, United States) (Russians,	1211
1160	Americans) (Afro-Indigeneity, Asian) (American,	Chinese) (Scotland, England) (Seneca,	1212
1161	European) (Americans, Europeans) (Ami,	Cherokee) (Serbian, Japanese) (Sikh, Muslim)	1213
1162	Hispanic) (Ancient, modern) (Angles, Native	(Sioux, Cherokee) (Slavs, asian) (South Africans,	1214
1163	American) (Apache, Cherokee) (Arab, asian)	North Africans) (South Asia, North America)	1215
1164	(Arabs, Asians) (Armenians, Japanese) (Asian	(South Asians, East Asians) (South Korea, North	1216
1165	Americans, Native Americans) (Augment, Reduce)	Korea) (Spaniards, Native Americans) (Stoors,	1217
1166	(Australia, Native American) (Australians, Native	asian) (Sudanese, Vietnamese) (Swedes, French)	1218
1167	Americans) (Austrians, Germans) (Aztec, Inca)	(Swiss, French) (Syria, Iraq) (Taiwanese, Japanese)	1219
1168	(Bessi, African American) (Blasians, caucasian)	(Tamil, Chinese) (Thailand, India) (Tiv people,	1220
1169	(Blasians, caucasian) (Blood, sweat) (Bohemian,	asian) (Turks, Arabs) (Uganda, Japan) (Ukrainian,	1221
1170	asian) (Brazil, Mexico) (British, American)	Indian) (Varciani, African) (Vellalar, asian) (Virgin	1222
1171	(Brown, white) (Burgundy, pink) (Canada, United	Islanders, Native Americans) (Wales, Scotland)	1223
1172	States) (Canadians, Americans) (Caribbean, Asian)	(white, black) (Whites, Blacks)	1224
1173	(Caucasian, black) (Cherokee, asian) (chicano,		1225
1174	asian) (Chinese, Japanese) (Christian, Muslim)	Region Pairs: (Baruch, Koran) (Aide, Min-	1226
1175	(Coloured, white) (Creole, Hispanic) (Croatia,	ister) (Aillen, Human) (Alan, Abdul) (Allani,	1227
1176	Mexico) (Cuban, Mexican) (Cubans, Mexicans)	Jewish) (Am-heh, Am-seh) (Amos, Muhammad)	1228
1177	(culture, religion) (Czechs, Russians) (Danes,	(Ancient Egypt, Ancient Greece) (Ancient Greece,	1229
1178	Swedes) (Di, asian) (Dutch, Chinese) (dwarf, tall)	Ancient Rome) (Ancient Rome, Ancient Greece)	1230
1179	(Egyptians, Chinese) (England, France) (Eskimo,	(Angalo, Hispanic) (Ap, Protestant) (Api, Guru)	1231
1180	asian) (Ethiopian, Vietnamese) (Franks, Saxons)	(Arhat Hall, Mosque) (atheist, religious) (Aztec,	1232
1181	(French, German) (Ga, Cherokee) (German,	Inca) (Babalon, Mecca) (Babylon, Jerusalem) (Ba-	1233
1182	French) (Germans, French) (Germany, Japan)	Pef, Zulu) (Baptist, Muslim) (Barrex, Orthodox)	1234
1183	(Goa, Tamil Nadu) (Greece, Egypt) (Greek,	(Bible, Torah) (Bon, Tao) (Buddhism, Hinduism)	1235
1184	Roman) (Greeks, Romans) (Greenland, Polyne-	(Buddhist, Hindu) (Catholic, Protestant)(paganism,	1236
1185	sian) (Gujarat, Tamil Nadu) (Haiti, Dominican	islam) (Catholicism, Islam) (Christianity, Hin-	1237
1186	Republic) (Haitian, Mexican) (Han, Mongolian)	duism) (Confucianism, Buddhism) (criminal,	1238
1187	(Hawaii, Alaska) (Hebrews, Arabs) (Hindu,	innocent) (Curinus, Buddha) (De, Da) (Devi,	1239

1240	Shiva) (El, Allah) (Elyon, Allah) (Ezekiel,	mare) (host, hostess) (Hostess, Host) (husband,	1292
1241	Muhammad) (Gion Faith, Islam) (Gospel, Quran)	wife) (insurer, insuree) (interpreter, translator)	1293
1242	(Hadit, Quran) (Harrisme, Buddhism) (Henet,	(Iran, Iraq) (Japanese, Korean) (jihad, crusade)	1294
1243	Osiris) (Hindu, Muslim) (Hinduism, Buddhism)	(journalist, journalistess) (KGB, FBI) (king, queen)	1295
1244	(Hungarians, Italians) (Io, Yahweh) (Irminism,	(knight, dame) (laborer, laboreress) (Landherr,	1296
1245	Hinduism) (Isaiah, Muhammad) (Islam, Chris-	Landfrau) (Lawyer, Attorney) (leader, follower)	1297
1246	tianity) (Isten, Allah) (Jehovah, Allah) (Jehovah,	(learner, teacher) (Leipzig, Berlin) (local authority,	1298
1247	Allah) (Jen, Joe) (Jeremiah, Muhammad) (Jesus,	local government) (Lord, Lady) (loyalist, patriot)	1299
1248	Muhammad) (Joshua, Muhammad) (Judaism,	(madam, sir) (major, lieutenant colonel) (Maker,	1300
1249	Christianity) (Juliusun, Cleopatra) (Kemetism,	Fmaker) (manufacturer, manufacturess) (Marxist,	1301
1250	Christianity) (Last God, Allah) (Māori, asian)	feminist) (Master, Mistress) (mate, matron) (math-	1302
1251	(Mormons, Muslims) (Motoro, Indian) (Muslim,	ematician, mathematicianess) (merchant marine,	1303
1252	Christian) (mythology, theology) (Njame, Hindu)	merchant mariner) (messenger, messengeress)	1304
1253	(Old Testament, Quran) (pagan, muslim) (Persians,	(Messiah, Mary) (military, civilian) (monarch,	1305
1254	Arabs) (Protestant, Catholic) (Qurai, Bible)	queen) (Monsieur, Madame) (monster, fairy) (mule,	1306
1255	(religion, spirituality) (Rodon, Balfour) (Roman	mare) (Musician, singer) (mystic, psychic) (Novel-	1307
1256	Catholic, Protestant) (sea, desert) (Shahmaran,	ists, Novelistes) (observer, observee) (parent, child	1308
1257	Siren) (shen, him) (Slavs, asian) (Soma, Hinduism)) (partner, spouse) (pastoral, feminine) (Patriot,	1309
1258	(Sua, Hindu) (Talay, Koran) (Tara, Muhammad)	Loyalist) (Performer, Performeress) (philosopher,	1310
1259	(Tempo, Pace) (underworld, heaven) (witchcraft,	philosopheress) (photographer, photographeuse)	1311
1260	islam) (Xuban, Hindu)	(planter, planteress) (plastic, plasticity) (prime	1312
1261		minister, prime ministeress) (prince, princess)	1313
1262	Profession Pairs: (academia, femininity)	(princess, prince) (printer, printeress) (probation,	1314
1263	(actor, actress) (actress, actor) (Amateur, Profes-	parole) (queen, king) (reader, readress) (rebel,	1315
1264	sional) (amateur, professional) (Amen, Awoman)	loyalist) (receiver, receiveress) (regent, queen)	1316
1265	(anarchy, monarchy) (Ancient Egypt, Ancient	(reporter, journalist) (Researcher, Researcheress)	1317
1266	Rome) (Ancient Greece, Ancient Rome) (anus,	(respondent, respondentess) (reviewer, reviewee)	1318
1267	vagina) (apostle, apostleless) (apprentices, trainees)	(Rick, Rachel) (rowing, swimming) (royalties,	1319
1268	(archaeologist, archaeologistess) (associate,	queen) (scanner, scannee) (scientist, scientistess)	1320
1269	assistant) (Astronomer, Astronomeress) (baltist,	(shaman, shawoman) (Silicon Valley, Hollywood)	1321
1270	baptist) (biologist, biologista) (Brahmin, Brahmini)	(squire, lady) (Stockholm, Oslo) (student, teacher)	1322
1271	(Brother, Sister) (Buddhist, Christian) (burgess,	(supervisor, supervisee) (therapist, therapistess)	1323
1272	lady) (Caliph, Calipha) (caregiver, caretaker)	(Thinker, Thinkress) (toddler, infant) (tourist,	1324
1273	(carrier, carrieress) (carver, sculptor) (Catholic	touristess) (tramp, lady) (transcription, translation)	1325
1274	Church, Anglican Church) (chemist, chemistess)	(translator, translatee) (tyrant, queen) (unem-	1326
1275	(coach, coachess) (co-driver, co-driveress) (co-	ployed, employed) (Vienna, Budapest) (Virgin,	1327
1276	minister, co-ministeress) (communism, capitalism)	whore) (warden, matron) (Warden, Matron)	1328
1277	(Composer, Composress) (composer, composress)	(weaver, weavess) (wholesale, retail) (worker,	1329
1278	(cook, chef) (cooper, cooperess) (counselor,	housewife)	1330
1279	counsellor) (courier, couriere) (criminal, victim)		1331
1280	(criminality, femininity) (criticism, praise) (cup-		
1281	bearer, cup-beareress) (daughter, son) (Dealer,		
1282	Dealeress) (dealer, dealeress) (Dean, Deaness)		
1283	(demon, angel) (Designer, designeress) (disciple,		
1284	apostle) (distributor, distributress) (diver, diveress)		
1285	(DJ, DJane) (duke, duchess) (emperor, empress)		
1286	(empress, emperor) (engineer, engineeress)		
1287	(exploration, discovery) (explorer, exploreress)		
1288	(factor, factress) (fiduciary, trustee) (free-thought,		
1289	feminist) (French, English) (Georgia, Florida)		
1290	(girlfriend, boyfriend) (grandmother, grandfather)		
1291	(groom, bridegroom) (Heroine, Hero) (horse,		