

Unlabeled Debiasing in Downstream Tasks via Class-wise Low Variance Regularization

Shahed Masoudian¹, Markus Frohmann^{1,2}, Navid Rekabsaz³, Markus Schedl^{1,2}

¹ Johannes Kepler University Linz, Austria

² Linz Institute of Technology, AI Lab, Austria

³ Thomson Reuters Labs, Zug, Switzerland

shahed.masoudian@jku.at

Abstract

Language models frequently inherit societal biases from their training data. Numerous techniques have been proposed to mitigate these biases during both the pre-training and fine-tuning stages. However, fine-tuning a pre-trained debiased language model on a downstream task can reintroduce biases into the model. Additionally, existing debiasing methods for downstream tasks either (i) require labels of protected attributes (e.g., age, race, or political views) that are often not available or (ii) rely on indicators of bias, which restricts their applicability to gender debiasing since they rely on gender-specific words. To address this, we introduce a novel debiasing regularization technique based on the class-wise variance of embeddings. Crucially, our method does not require attribute labels and targets *any* attribute, thus addressing the shortcomings of existing debiasing methods. Our experiments on encoder language models and three datasets demonstrate that our method outperforms existing strong debiasing baselines that rely on target attribute labels while maintaining performance on the target task.¹

1 Introduction and Background

Language Models (LMs) based on encoders are used for a variety of purposes such as document classification (Founta et al., 2018; De-Arteaga et al., 2019), job recommendation (Kumar et al., 2023a), text generation (Eldan and Li, 2023), or as text encoder for multimodal models such as text-to-audio (Liu et al., 2023) or text-to-image (Bahani et al., 2023) models. These models often encode societal biases rooted in the corpora used for training (Mehrabi et al., 2022; Rekabsaz et al., 2021a), which causes a distributional shift of embeddings, hence affecting their outputs either with disproportionate misclassification of documents belonging

to minority groups or unfair ranking of the documents (Rekabsaz et al., 2021b; Melchiorre et al., 2021).

Several works focus on reducing the effect of these biases by improving model performance related to some specific fairness metric (*empirical fairness*) or by making the model blind to the existence of a certain attribute (*representational fairness*) (Shen et al., 2022). For instance, Shen et al. (2022) leverage contrastive learning to improve empirical fairness. Recent works focus mostly on efficiency and user flexibility when it comes to debiasing using modular approaches such as sub-networks or adapters (Houlsby et al., 2019). Hauenberger et al. (2023) introduce a modular debiasing scheme with adversarial training (Elazar and Goldberg, 2018) and mutual information reduction (Colombo et al., 2021) to control the bias in encoder LMs. Kumar et al. (2023b) use adversarial training with adapters (Pfeiffer et al., 2021) to improve representational fairness on document classification. Finally, Masoudian et al. (2024) used gated adapters to improve representational fairness while preserving task performance for classification and retrieval tasks.

Although these methods effectively reduce sensitive attribute information and enhance fairness (Zerveas et al., 2022; Shen et al., 2022) through blindness, they depend on attribute labels to align the distribution of the target attribute. Since the user input data contains numerous nuanced protected attributes, such as age, race, religion, etc., it is challenging to collect labeled data for each individual attribute across every task. Moreover, supervised debiasing methods typically require training on each attribute individually, scaling linearly with the number of attributes. This complexity highlights the need for more efficient and scalable approaches to handle multiple protected attributes in debiasing efforts.

To address this limitation, some works attempt

¹The code for the experiments is available on GitHub: <https://github.com/ShawMask/UnlabeledDebiasing>

to debias language models without using attribute labels. Zhou et al. (2022) employ contrastive learning combined with instance weighting to reduce the bias encoded in the language model. Moreover, Cheng et al. (2020) utilize post-hoc contrastive learning to enhance the fairness of pre-trained encoder language models concerning gender bias. Ghanbarzadeh et al. (2023) integrate the masking objective used during the pre-training of encoder language models with fine-tuning on gender-specific tasks to address gender bias.

These methods address gender bias without requiring labeled data using explicit gender indicators present within the text. However, they are ineffective against other biases such as age, race, or political view of the user, as well as implicit gender bias when gender information is removed from the text, limiting their possible use cases.

In this work, we bridge this gap by introducing a new *regularization scheme based on class-wise variance* to reduce unknown (unlabeled) representational bias in the embeddings of LMs. Our regularization enforces low-variance embeddings, which results in mitigating any possible distributional shift caused by unknown attributes in the model’s embeddings. With this method we force the model to produce robust embeddings that are informative about the classification task but contain less information about the protected attributes resulting in fair representation of the protected attributes. This gives our method the advantage of not relying on any type of information on the attribute during debiasing. To the best of our knowledge, we are the first to address the debiasing of *arbitrary* attributes without having access to attribute labels.

We demonstrate the effectiveness of our method on document classification tasks using adapters (Pfeiffer et al., 2021) and two commonly used encoder LMs, BERT-BASE and ROBERTA-BASE. Furthermore, we show that our method, when compared to existing *supervised* debiasing methods, can enhance attribute removal while still showing competitive classification task performance.

2 Problem Formulation

In recent years, adapter networks (Houlsby et al., 2019; Pfeiffer et al., 2021) have emerged as an efficient way of training models on downstream tasks. In addition to their improved training efficiency,

Table 1: Result of adapter training of BERT (BERT-BASE) and a gender-debiased version of BERT (BERT-NLI) on two datasets with gender as the protected attribute. Here, *Task* corresponds to accuracy on the main classification task, and *Gender* is the balanced accuracy of the model concerning the protected attribute.

Model	BIOS		PAN16	
	Task↑	Gender↓	Task↑	Gender↓
BERT-BASE	84.3 _{0.1}	67.0 _{0.1}	92.4 _{0.1}	70.7 _{0.1}
BERT-NLI	84.1 _{0.1}	64.5 _{0.1}	88.2 _{0.1}	73.7 _{0.1}

adapters keep the backbone LM weights fixed, helping preserve information within the model.

In our initial study, we assess how much gender information can be extracted from commonly used encoder LMs. We use adapters (Pfeiffer et al., 2021, Pfeiffer variation) in combination with (Devlin et al., 2019, BERT-BASE) and, additionally, a gender-debiased version of the same model (Wu et al., 2022, BERT-NLI) debiased for empirical fairness, on two downstream classification tasks. We then train probes on the embeddings of the fine-tuned models to check how much information about gender can be extracted from both model variations and report the average balanced accuracy as indicator of gender information in the embeddings.

Table 1 shows the result of both models for occupation prediction (BIOS (De-Arteaga et al., 2019)) and mention prediction (PAN16 (Rangel et al., 2016)) datasets. We observed that task performance of the debiased LM, BERT-NLI, is consistently lower than BERT-BASE, which aligns with observations by Ghanbarzadeh et al. (2023). Moreover, training adapters contain the gender information to a great extent on BIOS; while on PAN16, BERT-NLI leaks more gender information in the embeddings, although it has already been subject to debiasing.

This provides strong motivation for using debiasing methods during fine-tuning, even when using an already debiased pre-trained LM. However, as surveyed in § 1, existing debiasing methods either rely on attribute labels or are limited to attributes with explicit indicators in the text, such as gender. Furthermore, there exists a plethora of sensitive attributes, and labeling them all is challenging across tasks. This increase in number also affects debiasing complexity as it scales with the number of attributes. Thus, a method that addresses this gap would be highly desirable. In the following, we outline how we solve this gap.

We start with our problem setting, formulated as follows: Given a set of N documents with k classes, we are interested in having robust high-dimensional embeddings ($Z \in \mathbb{R}^d$) for document classification which are (i) informative about the classes but (ii) contain as little information as possible about any arbitrary protected attribute (ρ) not directly related to the classification task. Our approach to debiasing deviates from existing ones in two crucial ways: (i) It is independent of labeled attributes, and (ii) it targets *any* protected attribute simultaneously.

3 Low Variance Regularization (LVR)

We formulate our regularization scheme based on k centers, each representing a class in the dataset with d dimension $\{C_1, C_2, \dots, C_k | C_i \in \mathbb{R}^d\}$, where d is the model’s embedding size. We aim to adjust the parameters of the network if the variance of the embeddings in a batch is high, which intuitively results in the mitigation of any undesirable distributional shift that might exist in the embeddings. Since we have k classes, class-wise variance is a good proxy for this regularization loss.

We define the regularization loss as the distance between embeddings ($Z \in \mathbb{R}^d$) of class i in a given batch from their corresponding center. For each batch, we calculate the center of embeddings that belong to the same class (C_i), which results in k centers. To account for noisy data points and empty batches, we use the weighted sum of the current batch center C_i^b and the normalized weighted sum of previous batch centers C_i^{b-1} where ω is a hyperparameter to control the influence of previous batch and found through grid search. The centers are calculated as follows:

$$C_i^b = (1 - \omega) \frac{Z_1 + Z_2 \dots Z_m}{m} + \omega C_i^{b-1}, \quad (1)$$

where m is the number of samples for the i^{th} class in a batch. In practice, if there are no samples of a class within a batch, we ignore it; and if only one sample of the class is in the batch, the center becomes the sample itself. We then define the regularization loss as the sum of distances for each specific sample belonging to class i from the estimated center of the batch:

$$\mathcal{L}_r = \sum_{i=1}^k \sum_{r=1}^m \sqrt{\sum_{j=1}^d (z_{jr}^i - c_j^i)^2}, \quad (2)$$

where c_j^i is the center value for the j^{th} dimension of the i^{th} class and z_{jr}^i is the value for the j^{th} dimension of the r^{th} embedding for the i^{th} class. This corresponds to reducing the class-wise variance of the embeddings created by the model, which in turn reduces distributional shift that might exist in the data points of the same class and results in the alignment of the embeddings. We also use the calculated centers as extra input for the classification task and calculate the loss of the centers. We show later in § 5 that this added loss term is essential to mitigate degradation in task performance.

The overall loss then becomes a linear combination:

$$\mathcal{L}_{total} = \mathcal{L}_t + \lambda \mathcal{L}_r + \mathcal{L}_c \quad (3)$$

where \mathcal{L}_t is the classification loss, \mathcal{L}_r is the regularization loss, and \mathcal{L}_c is the loss to classify the calculated centers belonging to each class.

4 Experimental Setup

For our experiments, we follow previous works and focus on transformer-based language models. We use BERT-BASE and ROBERTA-BASE, in combination with adapters (Pfeiffer et al., 2021) for each task. Trained in this way, we denote models using our debiasing method as **ADPLVR**.

We use the following document classification datasets: occupation prediction (**BIOS**; De-Arteaga et al. (2019)) with gender as protected attribute, hate speech detection (**FCDL18**; Founta et al. (2018)) with race for protected attribute, and mention detection (**PAN16**; Rangel et al. (2016)), corresponding to a multi-attribute setting with age and gender as protected attributes. For each dataset, we remove all explicit indicators of protected attributes following previous works (Hauzenberger et al., 2023; Kumar et al., 2023b; Masoudian et al., 2024) from the text.

Baselines. We choose baselines as follows: **FT**, **ADP** and **ADPNLI** as fine-tuned versions of the entire model and adapter-based training of the model and adapter-based training for BERT model trained on debiased NLI, respectively, without using any additional bias mitigation method. We also select recent in-process debiasing algorithms as strong baselines, relying either on adversarial training (Elazar and Goldberg, 2018, **FTADV**; **ADPADV**) or mutual information reduction (Colombo et al., 2021, **ADPMMD**) to reduce the bias encoded within the embeddings and increase repre-

Table 2: Results of debiasing using BERT-BASE and ROBERTA-BASE. ADPLVR has no access to labeled attribute while FTADV, ADPADV and ADPMMD have access to the attribute information in the form of attribute labels.

Model	Type	BIOS		FDCL18		PAN16-Gender		PAN16-Age	
		Task↑	Probe↓	Task↑	Probe↓	Task↑	Probe↓	Task↑	Probe↓
BERT-BASE	FT	84.6 _{0.4}	67.3 _{0.8}	81.0 _{1.0}	92.9 _{1.8}	93.6 _{1.8}	69.6 _{0.8}	93.6 _{1.8}	42.3 _{0.9}
	ADP	84.3 _{0.1}	67.0 _{0.1}	80.0 _{0.1}	93.3 _{0.4}	92.4 _{0.1}	70.7 _{0.1}	92.4 _{0.1}	42.4 _{0.}
	ADPNLI	84.1 _{0.1}	64.5 _{0.1}	81.2 _{0.6}	93.5 _{0.6}	88.2 _{0.1}	73.7 _{0.1}	88.2 _{0.1}	42.5 _{0.1}
	FTADV	84.0 _{0.3}	60.8 _{0.2}	81.0 _{1.0}	84.4 _{4.0}	92.4 _{0.8}	59.8 _{0.7}	92.4 _{0.8}	31.3 _{1.1}
	ADPADV	84.2 _{0.1}	61.9 _{0.5}	79.8 _{0.3}	75.6 _{0.5}	92.2 _{0.1}	54.2 _{0.4}	92.1 _{0.1}	21.7 _{0.1}
	ADPMMD	84.4 _{0.2}	65.3 _{0.3}	80.1 _{0.2}	81.0 _{0.3}	91.4 _{0.4}	67.4 _{0.3}	92.0 _{0.8}	36.8 _{0.7}
	ADPLVR	84.0 _{0.2}	59.2 _{0.3}	81.7 _{0.1}	66.7 _{0.9}	91.3 _{0.1}	54.4 _{0.1}	91.3 _{0.1}	21.9 _{0.2}
ROBERTA-BASE	FT	84.5 _{0.4}	66.2 _{0.7}	80.6 _{0.4}	93.2 _{1.2}	98.5 _{0.1}	63.6 _{0.4}	98.5 _{0.1}	22.7 _{0.8}
	ADP	84.3 _{0.1}	67.3 _{0.7}	80.0 _{0.6}	94.0 _{0.6}	98.2 _{0.1}	62.8 _{0.4}	98.1 _{0.1}	31.9 _{0.1}
	FTADV	84.1 _{0.3}	61.6 _{0.3}	80.5 _{1.0}	83.6 _{1.9}	98.2 _{0.1}	52.0 _{0.9}	98.2 _{0.1}	24.1 _{1.4}
	ADPADV	84.0 _{0.1}	62.9 _{0.1}	80.0 _{0.5}	79.7 _{0.3}	98.1 _{0.1}	53.7 _{0.7}	98.0 _{0.1}	22.3 _{1.0}
	ADPMMD	84.3 _{0.2}	64.2 _{0.3}	80.0 _{0.1}	80.0 _{0.5}	97.8 _{0.1}	60.4 _{0.3}	98.0 _{0.4}	27.1 _{0.3}
	ADPLVR	83.8 _{0.1}	55.6 _{0.3}	81.5 _{0.2}	77.3 _{0.1}	97.7 _{0.1}	51.1 _{0.4}	97.7 _{0.1}	20.6 _{0.8}

Table 3: Results for adapter based training of BERT-BASE with different memory (ω) and center loss (\mathcal{L}_c) combinations.

ω	\mathcal{L}_c	BIOS		FDCL18		PAN16-Gender		PAN16-Age	
		Task↑	Probe↓	Task↑	Probe↓	Task↑	Probe↓	Task↑	Probe↓
-	-	80.2 _{0.8}	60.9 _{0.3}	81.1 _{0.1}	72.8 _{0.2}	91.2 _{0.1}	54.4 _{0.3}	91.2 _{0.1}	24.3 _{0.1}
✓	-	83.5 _{0.2}	59.6 _{0.1}	81.1 _{0.2}	72.0 _{0.3}	91.1 _{0.1}	54.4 _{0.1}	91.1 _{0.1}	22.3 _{0.1}
-	✓	83.8 _{0.4}	59.8 _{0.5}	81.5 _{0.2}	70.9 _{0.8}	91.4 _{0.2}	55.5 _{0.3}	91.4 _{0.2}	23.1 _{0.1}
ADPLVR		84.0 _{0.3}	59.2 _{0.2}	81.7 _{0.1}	66.7 _{0.9}	91.3 _{0.1}	54.4 _{0.1}	91.3 _{0.1}	21.9 _{0.2}

sentational fairness. Note that all supervised methods use labels of the target attribute to align the embeddings, while ADPLVR does *not* have access to any attribute label throughout the training.

Training. We follow the setup of previous works using the same datasets (Hauzenberger et al., 2023; Kumar et al., 2023b; Masoudian et al., 2024). Specifically, we use a maximum of 120 tokens for the BIOS dataset and 40 tokens for FDCL18 and PAN16 since they comprise comparatively short tweets. We train each model for 15 epochs with a learning rate of 2×10^{-5} . We select reduction factors for adapters on BIOS, PAN16, and FCDL18 as 2, 1, and 2, respectively, as they led to the best task performance. Since each loss term affects each model differently, we train baselines with a fixed $\lambda = 1$ for supervised debiasing and our unsupervised ADPLVR with $\lambda = 0.1$. We also select $\omega = 0.3$ as it performed best across all datasets in our grid search.

We train five probes consisting of two-layer fully connected networks and tanh activation function for 40 epochs and a learning rate of 1×10^{-4} to predict protected attributes from embeddings (Hauzenberger et al., 2023).

To evaluate task performance, we use accuracy as the evaluation metric. In order to evaluate the performance of bias mitigation methods we use balanced accuracy of the probes. We choose balanced accuracy to account for any unbalanced dataset with regard to the distribution of the protected attributes. For gender and dialect attributes, if balanced accuracy is around 50% it shows that the model is confusing the protected attribute; for age this value should be close to 20% because there are five classes for age. Furthermore, we run our experiments three times for each model and report the mean and standard deviation of three runs to account for variations in the training process.

5 Results and Discussion

Table 2 shows the task and probe performance of the baselines and ADPLVR.

In our single attribute experiments on FCDL18 and BIOS, using both BERT-BASE and ROBERTA-BASE, ADPLVR is able to remove information about protected attributes considerably better than *all* the baselines on BIOS and FCDL18. As for task performance, we observe a decrease in accuracy with ADPLVR compared to the baselines on BIOS. Remarkably, our regularization method

even shows an improvement in task performance on FCDDL18, demonstrating the robustness of its embeddings.

In our multi-attribute experiment on PAN16 (2 protected attributes), we observe that ADPLVR performs slightly worse than the best-performing model, ADPADV, on the main task. However, unlike ADPADV, which has access to the protected attribute label during training, ADPLVR crucially does not rely on attribute labels for bias mitigation; yet, it still outperforms the baselines in protected attribute information removal. Overall, we observe that, with BERT-BASE, ADPLVR shows slightly higher balanced accuracy compared to ADPADV for both protected attributes, while with ROBERTA-BASE, ADPLVR similarly shows improved mitigation performance.

Notably, other debiasing methods show similar decreases in task performance. Still, on FCDDL18, ADPLVR clearly outperforms all *supervised* baselines on the main task and information removal.

Ablation Study. To ensure all parts of our method are necessary to achieve its performance, we conduct an ablation study where we remove (i) the memory of the previous batch, ω , and (ii) the center loss \mathcal{L}_c introduced in section 3. Table 3 shows the result of this ablation study. By removing ω , the balanced accuracy of the probe considerably increases, meaning that the robustness of the embeddings toward protected attributes is reduced. Thus, more information about unknown, unrelated attributes influences the final output of the model to a larger extent.

Moreover, we observe that task performance clearly degrades when removing \mathcal{L}_c . Overall, the best-performing model, both in terms of task performance and probe balanced accuracy, is the one that has both ω acting as memory of previous batches for the model and \mathcal{L}_c , corresponding to our class-center-based loss.

6 Conclusion

In this work, we focus on representational fairness and introduce a novel regularization and optimization scheme to debias encoder LMs without accessing protected attribute labels. We show the effectiveness of our method using two encoder LMs across three datasets and multiple protected attributes. We demonstrate that our method enhances debiasing while maintaining task performance compared to strong baselines. To the best

of our knowledge, our method is the first that can mitigate bias of any *arbitrary* target attribute by generating robust embeddings best suited for the classification task. Since our method does not rely on attribute labels, we hope it paves the future for more accessible, effective, and efficient debiasing of encoder-based transformer models.

Limitations

One limitation of this work is the definition of gender used in all datasets, which is limited to binary female/male, lacking an inclusive and nuanced definition of gender. Moreover, although our method proved independent of attribute labels, a thorough evaluation would require more datasets with a variety of defined attributes. Another limitation of this work is the task in which we narrowed our study to classification tasks. We acknowledge that the findings of this paper might not be applicable to other tasks such as retrieval or recommendation. Furthermore, our study is focused on transformer-based language models which put an additional limitation on the generalization of the work to other models such as CNNs or LSTM-based language models. Due to the lack of suitable datasets, we relied on datasets commonly used in the debiasing literature. In FCDDL18, race is restricted to *African American* and *White American*, which does not reflect real-life scenarios. Furthermore, we follow previous works (Sap et al., 2019; Ravfogel et al., 2020; Zhang et al., 2021) and use labels of protected attributes assigned using another model, making them not fully representational of the real data distribution. A final limitation of this work is the lack of suitable datasets for multi-attribute settings, in which we could demonstrate that our approach can handle even more attributes than demonstrated with PAN16 simultaneously.

Acknowledgements

This research was funded in whole or in part by the Austrian Science Fund (FWF): <https://doi.org/10.55776/P33526>, <https://doi.org/10.55776/DFH23>, <https://doi.org/10.55776/COE12>, <https://doi.org/10.55776/P36413>.

References

Mourad Bahani, Aziza El-Ouaazizi, and Khalil Maalmi. 2023. The effectiveness of t5, gpt-2, and BERT on text-to-image generation task. *Pattern Recognit. Lett.*, 173:57–63.

- Pengyu Cheng, Weituo Hao, Siyang Yuan, Shijing Si, and Lawrence Carin. 2020. Fairfil: Contrastive neural debiasing method for pretrained text encoders. In *International Conference on Learning Representations*.
- Pierre Colombo, Pablo Piantanida, and Chloé Clavel. 2021. A novel estimator of mutual information for learning to disentangle textual representations. 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)*.
- Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 120–128.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. ACL.
- Yanai Elazar and Yoav Goldberg. 2018. [Adversarial removal of demographic attributes from text data](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 11–21. Association for Computational Linguistics.
- Ronen Eldan and Yuanzhi Li. 2023. [Tinystories: How small can language models be and still speak coherent english?](#) *CoRR*, abs/2305.07759.
- Antigoni Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of twitter abusive behavior. In *Twelfth International AAAI Conference on Web and Social Media*.
- Somayeh Ghanbarzadeh, Yan Huang, Hamid Palangi, Radames Cruz Moreno, and Hamed Khanpour. 2023. [Gender-tuning: Empowering fine-tuning for debiasing pre-trained language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 5448–5458. ACL.
- Lukas Hauenberger, Shahed Masoudian, Deepak Kumar, Markus Schedl, and Navid Rekabsaz. 2023. Modular and On-demand Bias Mitigation with Attribute-Removal Subnetworks. In *Findings of the Association for Computational Linguistics: ACL (Findings of ACL)*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *International Conference on Machine Learning*, volume 97, pages 2790–2799. Proceedings of Machine Learning Research.
- Deepak Kumar, Tessa Grosz, Navid Rekabsaz, Elisabeth Greif, and Markus Schedl. 2023a. Fairness of recommender systems in the recruitment domain: an analysis from technical and legal perspectives. *Frontiers in big Data*, 6.
- Deepak Kumar, Oleg Lesota, George Zerveas, Daniel Cohen, Carsten Eickhoff, Markus Schedl, and Navid Rekabsaz. 2023b. [Parameter-efficient modularised bias mitigation via AdapterFusion](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2738–2751, Dubrovnik, Croatia. Association for Computational Linguistics.
- Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo P. Mandic, Wenwu Wang, and Mark D. Plumbley. 2023. [Audioldm: Text-to-audio generation with latent diffusion models](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 21450–21474. PMLR.
- Shahed Masoudian, Cornelia Volaucnik, Markus Schedl, and Navid Rekabsaz. 2024. [Effective controllable bias mitigation for classification and retrieval using gate adapters](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2024 - Volume 1: Long Papers, St. Julian's, Malta, March 17-22, 2024*, pages 2434–2453. ACL.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2022. [A survey on bias and fairness in machine learning](#). *ACM Comput. Surv.*, 54(6):115:1–115:35.
- Alessandro B. Melchiorre, Navid Rekabsaz, Emilia Parada-Cabaleiro, Stefan Brandl, Oleg Lesota, and Markus Schedl. 2021. [Investigating gender fairness of recommendation algorithms in the music domain](#). *Information Processing Management (IP&M)*, 58(5):102666.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. Adapterfusion: Non-destructive task composition for transfer learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503.
- Francisco Rangel, Paolo Rosso, Ben Verhoeven, Walter Daelemans, Martin Potthast, and Benno Stein. 2016. Overview of the 4th author profiling task at pan 2016: cross-genre evaluations. In *Working Notes Papers of the CLEF 2016 Evaluation Labs. CEUR Workshop Proceedings/Balog, Krisztian [edit.]; et al.*, pages 750–784.

- Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null it out: Guarding protected attributes by iterative nullspace projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256.
- Navid Rekabsaz, Simone Kopeinik, and Markus Schedl. 2021a. Societal biases in retrieved contents: Measurement framework and adversarial mitigation of bert rankers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 306–316.
- Navid Rekabsaz, Simone Kopeinik, and Markus Schedl. 2021b. Societal biases in retrieved contents: Measurement framework and adversarial mitigation of BERT rankers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 306–316.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The risk of racial bias in hate speech detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. ACL.
- Aili Shen, Xudong Han, Trevor Cohn, Timothy Baldwin, and Lea Frermann. 2022. [Does representational fairness imply empirical fairness?](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 81–95, Online only. ACL.
- Yuxiang Wu, Matt Gardner, Pontus Stenetorp, and Pradeep Dasigi. 2022. Generating data to mitigate spurious correlations in natural language inference datasets. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. ACL.
- George Zerveas, Navid Rekabsaz, Daniel Cohen, and Carsten Eickhoff. 2022. [Mitigating bias in search results through contextual document reranking and neutrality regularization](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, page 2532–2538, New York, NY, USA. Association for Computing Machinery.
- Xiongyi Zhang, Jan-Willem van de Meent, and Byron Wallace. 2021. [Disentangling representations of text by masking transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 778–791, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kun Zhou, Beichen Zhang, Wayne Xin Zhao, and Ji-Rong Wen. 2022. [Debiased contrastive learning of unsupervised sentence representations](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2022, Dublin, Ireland, May 22-27, 2022, pages 6120–6130. ACL.