

Exploring Fusion Techniques in Multimodal AI-Based Recruitment: Insights from FairCVdb

Swati Swati^{1,*}, Arjun Roy^{1,2} and Eirini Ntoutsi¹

¹Research Institute CODE, University of the Bundeswehr Munich, Germany

²Institute of Computer Science, Free University Berlin, Germany

Abstract

Despite the large body of work on fairness-aware learning for individual modalities like tabular data, images, and text, less work has been done on multimodal data, which fuses various modalities for a comprehensive analysis. In this work, we investigate the fairness and bias implications of multimodal fusion techniques in the context of multimodal AI-based recruitment systems using the FairCVdb dataset. Our results show that *early-fusion* closely matches the ground truth for both demographics, achieving the lowest MAEs by integrating each modality's unique characteristics. In contrast, *late-fusion* leads to highly generalized mean scores and higher MAEs. Our findings emphasise the significant potential of *early-fusion* for accurate and fair applications, even in the presence of demographic biases, compared to *late-fusion*. Future research could explore alternative fusion strategies and incorporate modality-related fairness constraints to improve fairness. For code and additional insights, visit: <https://github.com/Swati17293/Multimodal-AI-Based-Recruitment-FairCVdb>

Keywords

Multimodal bias, Multimodal fairness, Algorithmic Fairness, Fairness, Early Fusion, Late Fusion

1. Introduction

The increasing popularity of decision-making algorithms has raised concerns about bias in decision-making, especially towards specific social groups defined by protected attributes such as gender and ethnicity [1]. Research on fairness-aware learning primarily focuses on individual modalities, such as tabular data [2], text [3], images [4], and graphs [5]. However, there has been less focus on bias in multimodal systems [6], which can result from integration complexity, unbalanced representation, alignment, and the compounding effect of biases present in each modality. To this end, in this work, we investigate the bias and fairness implications of multimodal AI in automated recruitment systems using the FairCVdb [7] dataset. We use it as a testbed, as it offers diverse data, including images, text, and structured data with intentionally designed gender and ethnicity biases. We focus on fusion techniques for integrating information from different modalities [8], specifically analysing *early-* and *late-* fusion techniques known for their straightforward interpretability and widespread usage in multimodal AI systems [9, 10].

EWAF'24: European Workshop on Algorithmic Fairness, July 01–03, 2024, Mainz, Germany

*Corresponding author.

✉ swati.swati@unibw.de (S. Swati); arjun.roy@unibw.de (A. Roy); eirini.ntoutsi@unibw.de (E. Ntoutsi)

🆔 0000-0002-7637-6640 (S. Swati); 0000-0002-4279-9442 (A. Roy); 0000-0001-5729-1003 (E. Ntoutsi)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

Early-fusion typically concatenates features from different modalities early on, creating a unified representation of the data [10], which simplifies training and effectively captures interactions between modalities [11]. *Late-fusion*, on the other hand, processes each modality individually before combining their outputs at a later stage, offering flexibility by allowing different processing pathways for individual modalities [12]. While *late-fusion* captures modality-specific patterns more accurately, it may overlook lower-level interactions between modalities [13]. By investigating these two fusion strategies, we aim to gain insight into how they impact bias and fairness in automated recruitment processes.

2. Experimental Setup

Dataset: The FairCVdb dataset [7] comprises of 24, 000 synthetic resume profiles, each featuring demographic characteristics (gender and ethnicity), *textual* data (a short biography), *visual* data (a facial image), and *tabular* data (seven common resume attributes). The resume attributes include occupation, suitability, education, previous experience, recommendation, availability, and language proficiency. Each profile has been generated based on two gender categories and three ethnic categories. The profiles in the dataset are scored based on the likelihood of a candidate being invited to an interview, yielding a numerical score. These scores are assigned either blindly (i.e., without any bias), leading to bias-neutral scores, or with a penalty factor applied to specific individuals within a demographic group, resulting in biased scores. See [14] for more details. This setup simulates scenarios where cognitive biases, introduced by humans, protocols, or automated systems, influence the decision-making process.

Evaluation Metrics: Following [14], we use Mean Absolute Error (MAE) to measure prediction error and Kullback-Leibler divergence (KL) to assess demographic bias. For gender, we compare score distributions for males and females; for ethnicity, we perform pairwise comparisons and report the average divergence.

Models: We extend the testbed [14] to facilitate multimodal recruitment learning by including *early-fusion* and *late-fusion* techniques for all three modalities (*textual*, *visual*, *tabular*).

Simulated setups: We investigated both i) unbiased ideal world setup and ii) real-world setups (gender- and ethnicity- biased).

3. Evaluation Results

Figure 1 depicts the KL-divergence (KL), Mean Absolute Error (MAE), and score distributions for gender and ethnicity across different modalities and bias setups. A smaller KL-divergence indicates better alignment between distributions, implying less bias, while a lower MAE indicates a smaller margin of error. Analysis of the score distributions offers additional insights into the models' predictive performance.

In unbiased ideal-world (Neutral): We note that the *ground-truth* distributions are closely aligned for both demographics (c.f., Figure 1a). W.r.t. individual modalities (c.f., Figure 1b - 1d), we can see that *tabular* modality exhibits a lower score distribution centred around a mean of




Figure 1: KL-divergence between score distributions across Gender and Ethnicity demographics for different modalities and bias setups. Lower KL and MAE scores are better.

0.4 with a negatively-skewed distribution, indicating that it tends to underestimate the *ground-truth*. The presence of a bimodal distribution in the *textual* modality is specially intriguing, demonstrating its ability to differentiate between instances with high and low scores. The *Visual*

modality, on the other hand, exhibits extreme behaviour by concentrating the distribution of nearly the entire population within a very narrow range [0.39–0.44] (c.f., Figure 1d), pointing an over-generalization of the mean score to all instances. Interestingly, *late-fusion* produces the least biased results for both demographics. However, while aggregating the decisions from different modalities, its average decision gets affected by the extremity of the *visual* modality, leading to over-generalization of the mean score, consequently resulting in higher MAEs (c.f., Figure 1f). In contrast, *early-fusion* delivers the most accurate predictions with the lowest MAEs (c.f., Figure 1e) by effectively learning and resolving the unique peculiarities of each modality, such as underestimation, over-generalization, and bimodal distribution, resulting in a shape that resembles the *ground-truth* (c.f., Figure 1a, 1e).

In biased real-world setups (Gender/Ethnicity-Biased): We observe that the *ground-truth* distributions are not aligned for both demographics (c.f., Figure 1a). W.r.t. individual modalities (c.f., Figure 1b - 1d), we see that the *tabular* modality continues to exhibit underestimation across all demographics, which leads to close alignment of the demographic specific distributions (c.f., Figure 1b(2) and b(4)). With *textual* modality we notice a misalignment of distribution w.r.t. gender demographics with a favourable skewness for males. However, no such bias is observed w.r.t. ethnicity, indicating a possibility of gender-skewness being much higher than the ethnicity-skewness for the job-related words in the embedded space. Conversely, the *visual* modality demonstrates the most extreme bias for both demographics. Regarding gender, it shows a positive bias towards males, while for ethnicity, it overgeneralizes Asians, discriminates against Blacks, and favours Caucasians. Continuing the trend established in the neutral setup, *Early* fusion consistently mimics the *ground-truth* for both demographics, yielding the lowest MAEs while maintaining fairness. *Late-fusion*, while also following its trend, tends to over-generalize the mean score, resulting in higher MAEs but also higher KL scores.

In general, leveraging multimodal data can enhance performance and mitigate bias compared to relying on a single modality. However, blindly fusing all modalities may not always yield the best results. For instance, the *tabular* in *gender-biased* setup (c.f., Figure 1b(2)) and the *textual* in *ethnicity-biased* setup (c.f., Figure 1c(4)) outperformed both fusion strategies. We hypothesise that *late-fusion* exacerbates biases by independently learning biased models for each modality, cumulatively impacting decision fairness, while *early-fusion* offers greater flexibility and generally yields fairer outcomes with lower prediction error. Dataset diversity and biases may have influenced these findings, highlighting the need to assess robustness across multiple datasets, domains, and fusion strategies. We contemplate that in the future, exploring mid-fusion strategies could enhance fairness and accuracy in decision-making through strategic selection and a combination of modalities.

4. Conclusions

In our study, we used the FairCVdb dataset to investigate the bias implications of *early-* and *late-* fusion strategies in multimodal AI-based recruitment. We assessed biases in gender and ethnicity demographics across both unbiased (neutral) and real-world (gender/ethnicity-biased) setups. Our findings reveal that *early-fusion* closely mimics the ground truth for both demographics, achieving the lowest MAEs by effectively incorporating the unique characteristics of

each modality. In contrast, *late-fusion* leads to highly over-generalized mean scores, resulting in higher MAEs. Our evaluation underscores the significant potential of *early-fusion* for applications requiring both accuracy and fairness, providing robust solutions even in the presence of demographic biases. Based on the results, we speculate that mid-fusion strategies may enhance fairness and accuracy by strategically selecting and combining modalities. Exploring these findings across diverse datasets and domains beyond hiring could further broaden the study's impact and relevance. **Ethics statement:** Understanding the risks of using simulated or synthetic data is crucial for fairness, transparency, and effectiveness in automated hiring processes.

Acknowledgments

This research work is funded by the European Union under the Horizon Europe MAMMOth project, Grant Agreement ID: 101070285. UK participant in Horizon Europe Project MAMMOth is supported by UKRI grant number 10041914 (Trilateral Research LTD). The research is also supported by the EU Horizon Europe project STELAR, Grant Agreement ID: 101070122.

References

- [1] M. Raghavan, S. Barocas, J. Kleinberg, K. Levy, Mitigating bias in algorithmic hiring: Evaluating claims and practices, in: Proceedings of the 2020 conference on fairness, accountability, and transparency (ACM FAT*), 2020, pp. 469–481.
- [2] T. Le Quy, A. Roy, V. Iosifidis, W. Zhang, E. Ntoutsi, A survey on datasets for fairness-aware machine learning, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 12 (2022) e1452.
- [3] Y. Cai, A. Zimek, G. Wunder, E. Ntoutsi, Power of explanations: Towards automatic debiasing in hate speech detection, in: 2022 IEEE 9th International Conference on Data Science and Advanced Analytics (DSAA), IEEE, 2022, pp. 1–10.
- [4] S. Fabbrizzi, S. Papadopoulos, E. Ntoutsi, I. Kompatsiaris, A survey on bias in visual datasets, Computer Vision and Image Understanding 223 (2022) 103552.
- [5] S. Ghodsi, S. A. Seyedi, E. Ntoutsi, Towards cohesion-fairness harmony: Contrastive regularization in individual fair graph clustering, in: Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), Springer, 2024, pp. 284–296.
- [6] B. M. Booth, L. Hickman, S. K. Subburaj, L. Tay, S. E. Woo, S. K. D'Mello, Bias and fairness in multimodal machine learning: A case study of automated video interviews, in: Proceedings of the 2021 International Conference on Multimodal Interaction (ICMI), 2021, pp. 268–277.
- [7] A. Pena, I. Serna, A. Morales, J. Fierrez, Bias in multimodal ai: Testbed for fair automatic recruitment, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (Workshop@CVPR), 2020, pp. 28–29.
- [8] Z. Xue, R. Marculescu, Dynamic multimodal fusion, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023, pp. 2574–2583.
- [9] K. Gadzicki, R. Khamsehashari, C. Zetzsche, Early vs late fusion in multimodal convolu-

tional neural networks, in: 2020 IEEE 23rd international conference on information fusion (FUSION), IEEE, 2020, pp. 1–6.

- [10] L. M. Pereira, A. Salazar, L. Vergara, A comparative analysis of early and late fusion for the multimodal two-class problem, *IEEE Access* (2023).
- [11] G. Barnum, S. Talukder, Y. Yue, On the benefits of early fusion in multimodal representation learning, *arXiv preprint arXiv:2011.07191* (2020).
- [12] L. M. Pereira, A. Salazar, L. Vergara, On comparing early and late fusion methods, in: *International Work-Conference on Artificial Neural Networks (IWANN)*, Springer, 2023, pp. 365–378.
- [13] K. Bayoudh, R. Knani, F. Hamdaoui, A. Mtibaa, A survey on deep multimodal learning for computer vision: advances, trends, applications, and datasets, *The Visual Computer* 38 (2022) 2939–2970.
- [14] A. Peña, I. Serna, A. Morales, J. Fierrez, A. Ortega, A. Herrarte, M. Alcantara, J. Ortega-Garcia, Human-centric multimodal machine learning: Recent advances and testbed on ai-based recruitment, *SN Computer Science* 4 (2023) 434.