End-to-End Bias Mitigation in Candidate Recommender Systems with Fairness Gates

Adam Arafan
Eindhoven University of Technology
a.m.arafan@tue.nl

David Graus
Randstad
david.graus@randstad.com

Fernando P. Santos
University of Amsterdam
f.p.santos@uva.nl

Emma Beauxis-Aussalet
Vrije Universiteit Amsterdam
e.m.a.l.beauxisaussalet@vu.nl

Introduction

- Machine learning applications can unintentionally amplify biases, particularly in sensitive domains like recidivism prediction and facial recognition, with gender and race being common dimensions of bias.
- Our paper addresses **fairness within candidate recommender systems** (CRS), focusing on mitigating bias to prevent discrimination in employment recommendations.
- We introduce a novel approach that **integrates two fairness gates (FGs)** at both pre- and post-processing stages to improve fairness without significantly compromising utility.

Methodology

- Our work relies on a learning-to-rank powered candidate recommendation system (CRS)
- We introduce two Fairness Gates: FG1 through **synthetic data generation** to balance gender distribution in training datasets, and FG2 for **re-ranking to ensure gender balance** in candidate recommendations
- For evaluating the impact of fairness gates on **utility and fairness**, we report respectively Normalised Discounted Cumulative Gain (NDCG) and Normalized Discounted Cumulative Kullback-Leibler Divergence (NDKL)

Results

- We observe a significant **uplift in NDCG** scores **when training CRS models on synthetic data**, especially in heavily imbalanced scenarios.
- We demonstrate that **the application of both FGs leads to improved fairness**, as measured by reduced NDKL scores after re-ranking, across all levels of data imbalance.

Average *NDCG* and *NDKL* for rankings at each level of data imbalance

Ranked Lists	*NDCG*	*NDKL* before re-ranking	*NDKL* after re-ranking
Heavy imbalance			
Real data	0.384	0.366	0.200
Synthetic data	0.693	0.358	0.197
Minor imbalance			
Real data	0.403	0.217	0.126
Synthetic data	0.647	0.213	0.126
Balanced			
Real data	0.403	0.213	0.124
Synthetic data	0.633	0.206	0.124

Comparing CRSs trained with **Real** or **Synthetic data** (FG 1)

with or without re-ranking (FG 2)

Approach

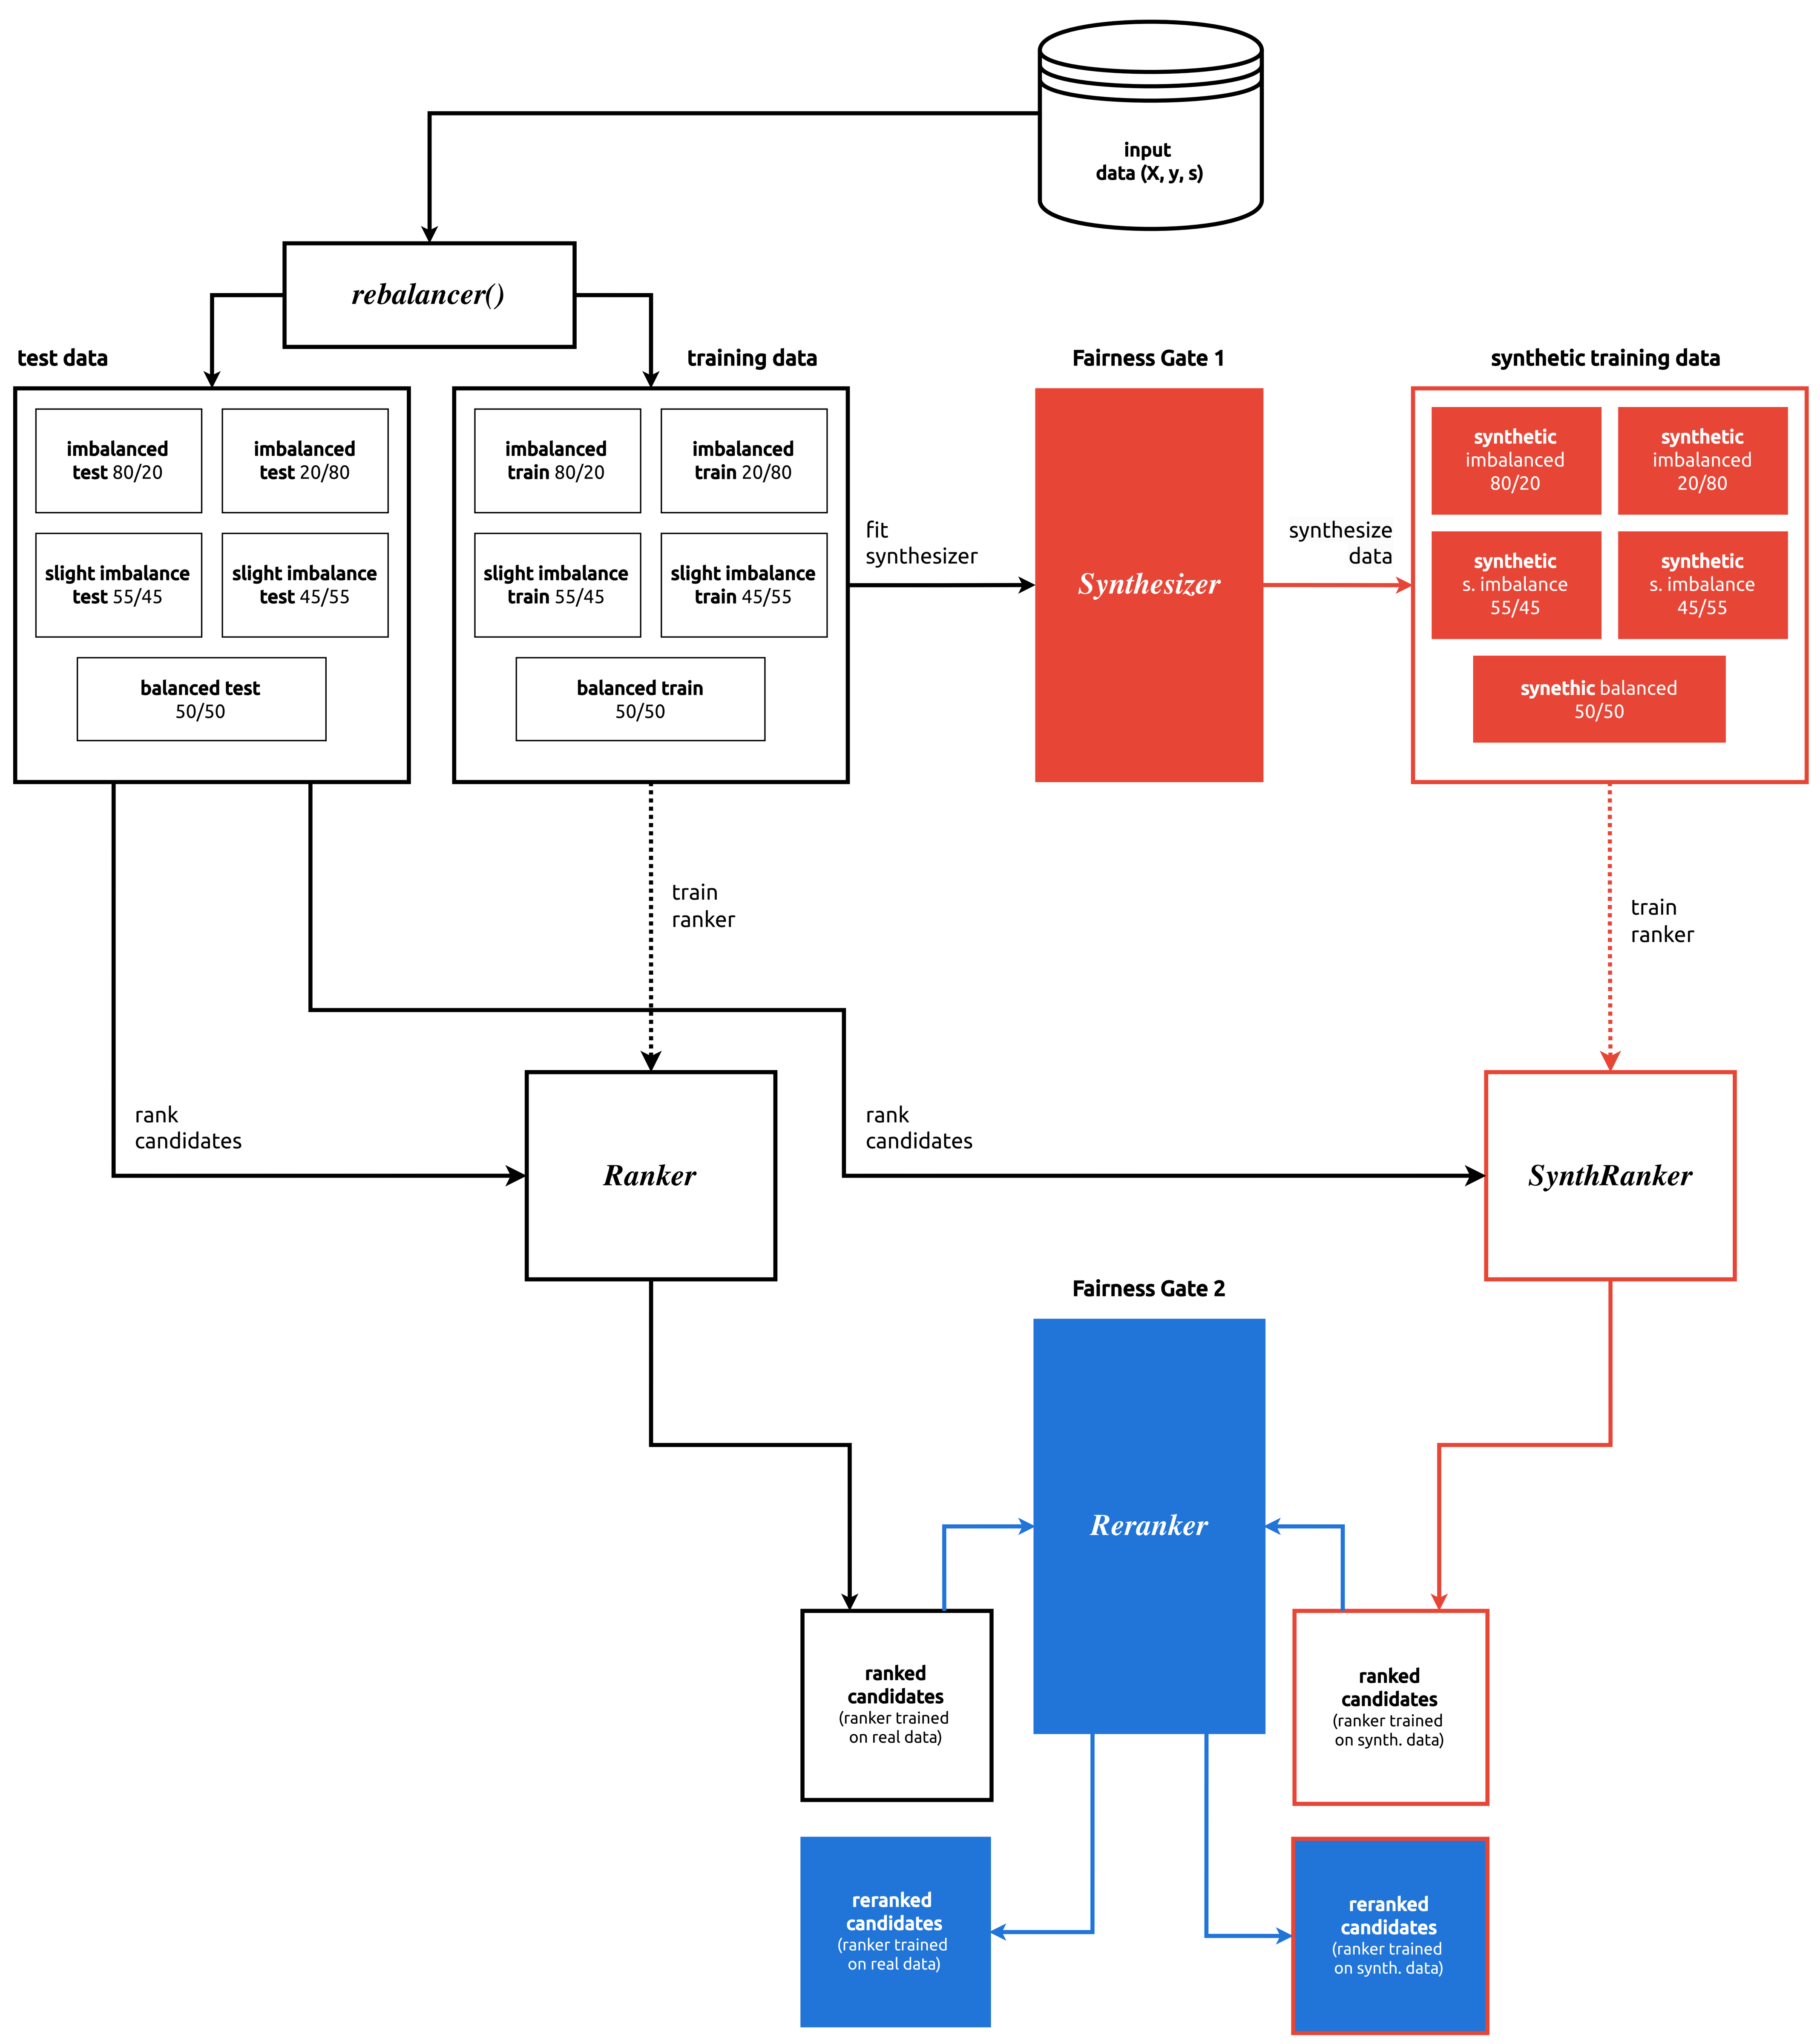

Conclusion

- We affirm the success of the CRS pipeline in **mitigating bias through the integration of synthetic data and fairness gates**.
- Our approach provides a viable pathway to **fair and useful candidate recommendations** within the CRS pipeline.

Future Work

- We recommend exploring additional evaluation methods, addressing data scarcity, and refining fairness rules to **account for real-world job domain skewness**.
- We encourage further research into less data-intensive techniques and the potential of synthetic data-driven augmentation to enhance the CRS pipeline.

Read the paper

Arafan, A., Graus, D., Santos, F., & Beauxis-Aussalet, E. (2022). End-to-End Bias Mitigation in Candidate Recommender Systems with Fairness Gates. In RecSys in HR'22: The 2nd Workshop on Recommender Systems for Human Resources.