Adila
Fairness-Informed Neural Team Recommendation

Hamed Loghmani — University of Windsor — ghasrlo@uwindsor.ca
Mahdis Saeedi — University of Windsor — msaeedi@uwindsor.ca
Gabriel Rueda — University of Windsor — ruedag@uwindsor.ca
Edwin Paul — University of Windsor — paul43@uwindsor.ca
Hossein Fani — University of Windsor — hossein.fani@uwindsor.ca

01 A Story to Tell

Team recommendation aims at forming a collaborative group of experts to accomplish complex tasks, which is a recognized objective in the industry. While state-of-the-art neural team recommenders can efficiently analyze massive sets of candidate experts to form effective collaborative teams, they overlook fairness. Due to this critical ethical issue in AI-based decision making, in this work, we adopt a various greedy reranking algorithms to achieve fairness with respect to (1) popularity or (2) gender in neural models in view of two notions of fairness, demographic parity and equality of opportunity.

02 Research Questions

RQ1) If team recommendation models, when recommending teams of experts, perpetuate biases, particularly concerning popularity and gender as protected attributes.

RQ2) If state-of-the-art greedy reranking algorithms are capable of enhancing the fairness of neural team recommendation models without compromising on their utility.

RQ3) How effective post-processing methods are in mitigating severe pre-existing biases within training datasets, and under what conditions these methods uphold the integrity and utility of the generated models across various application domains.

03 Gender Distribution

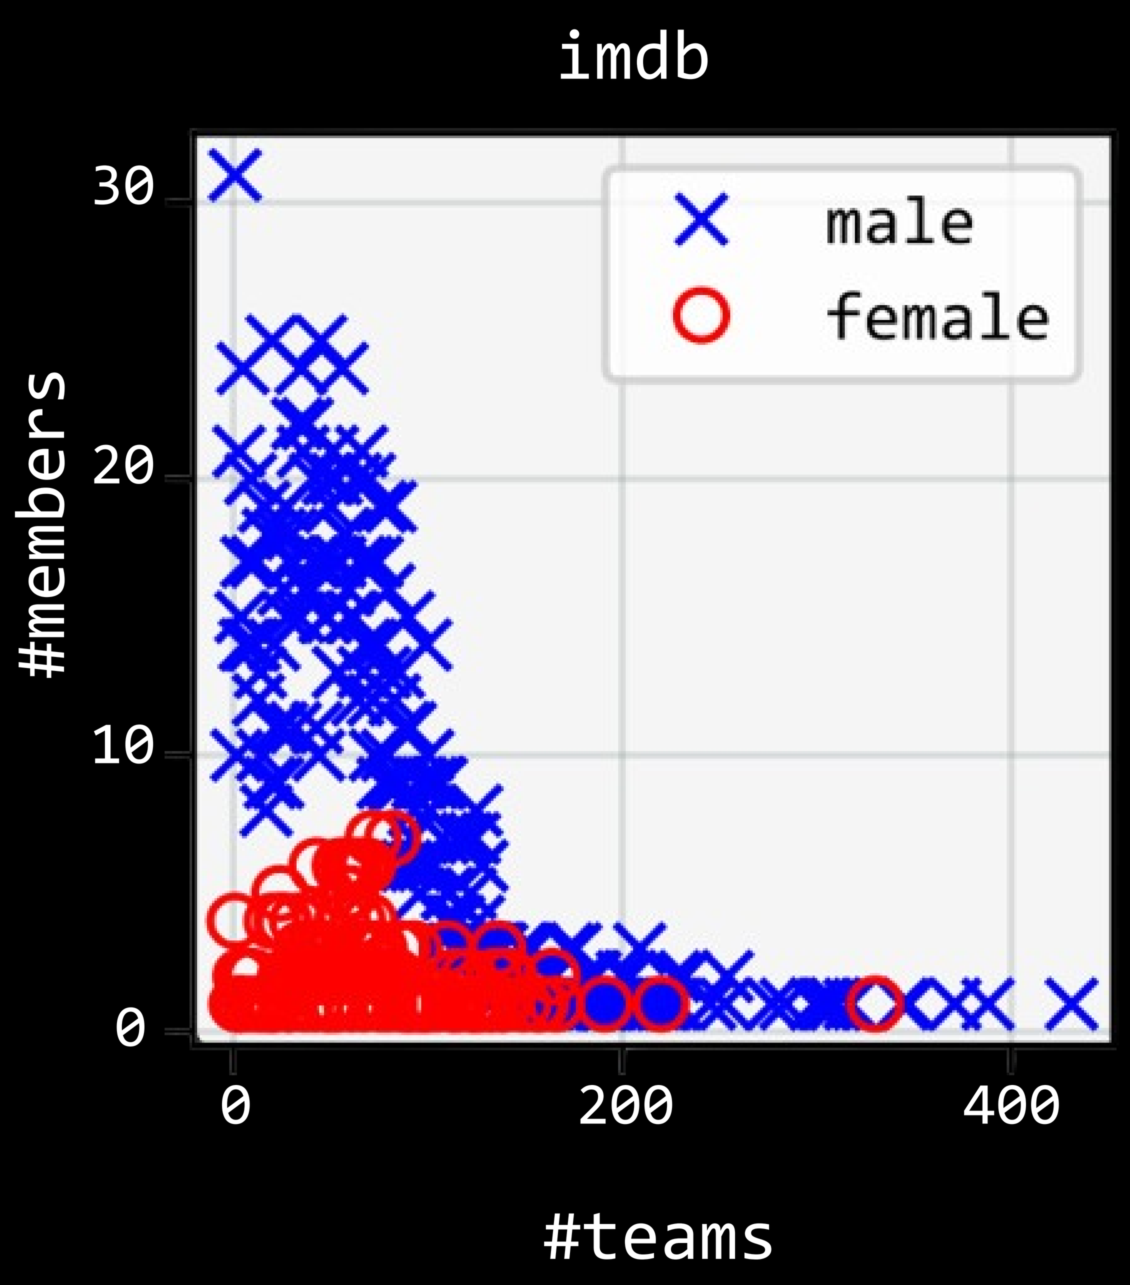

04 In Search of Light

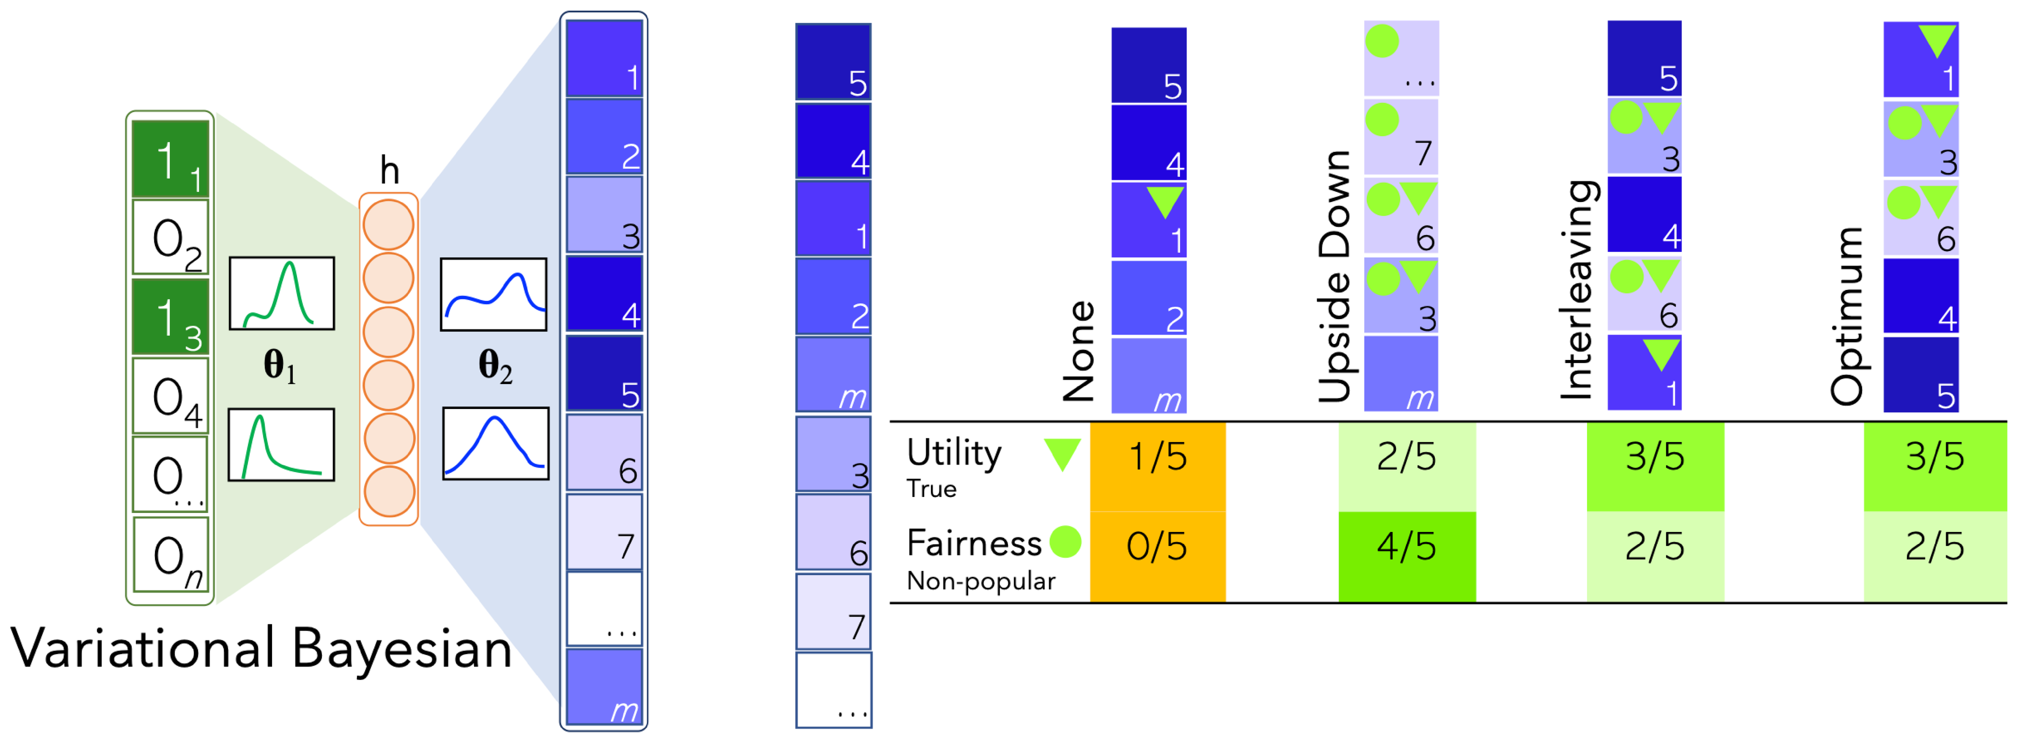

05 Popularity Labeling

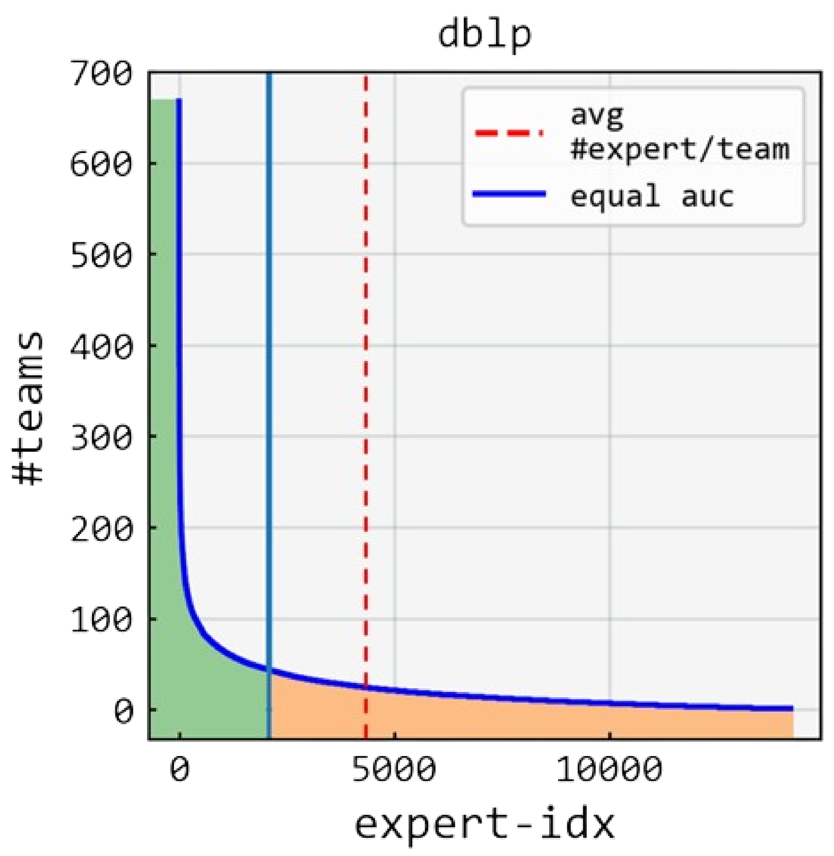

06 Let There Be Light

1. Upon a comprehensive fairness evaluation, it has been determined that the output of team recommendation methods can exhibit bias towards female/nonpopular experts.
2. When considering popularity as the protected attribute, our findings confirm its influence.
3. We determined that while reranking methods can be notably effective in addressing biases, their efficacy diminishes when they are employed single-handedly. Specifically, when confronting extreme biases in data, these methods struggle to rectify them without a consequential loss in utility.

07 Partial Results

dblp-popularity

		%ndkl before↓	%ndkl after↓	skew before→ 0 nonprotected	skew before→ 0 protected	skew after→ 0 nonprotected	skew after→ 0 protected	%map10 Δ↑	%ncdg10 Δ↑
demographic parity									
bnn	det-cons	109.56	14.64	1.1343	-19.9704	0.6484	-0.5462	-0.28	-0.58
	det-greedy		14.64			0.6484	-0.5462	-0.28	-0.58
	det-relaxed		18.31			0.6413	-0.5360	-0.28	-0.58
	fa*ir		19.71			0.2639	-0.1524	0.00	00.00
bnn-emb	det-cons	110.31	14.09	1.1415	-20.7584	0.6262	-0.5161	-0.28	-0.58
	det-greedy		14.09			0.6262	-0.5161	-0.28	-0.58
	det-relaxed		17.65			0.6189	-0.5063	-0.28	-0.58
	fa*ir		19.61			0.2686	-0.1531	0.00	0.00
equality of opportunity									
bnn	det-cons	102.01	13.12	1.0560	-19.9253	0.5773	-0.5113	-0.28	-0.58
	det-greedy		13.16			0.5773	-0.5113	-0.28	-0.58
	det-relaxed		16.15			0.5729	-0.5050	-0.28	-0.58
	fa*ir		18.96			0.2499	-0.1631	0.00	0.00
bnn-emb	det-cons	102.85	12.65	1.0641	-20.6268	0.5555	-0.4813	-0.28	-0.58
	det-greedy		12.67			0.5555	-0.4813	-0.28	-0.58
	det-relaxed		15.63			0.5512	-0.4752	-0.28	-0.58
	fa*ir		18.39			0.2526	-0.1645	0.00	0.00

08 Future Remarks

- Experiment on different cutting-edge fairness-informed reranking algorithms.
- Include new datasets and domains
- Experiment on pre-processing fairness methods.

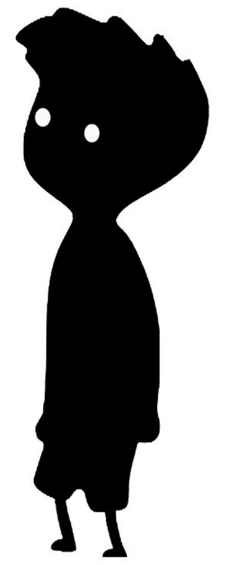

FANI-LAB/ADILA