Exploring Fairness across Many Rankings

Hilson Shrestha[*] Kathleen Cachel[†] Mallak Alkhathlan[‡] Elke Rundensteiner[§] Lane Harrison[¶]

Worcester Polytechnic Institute

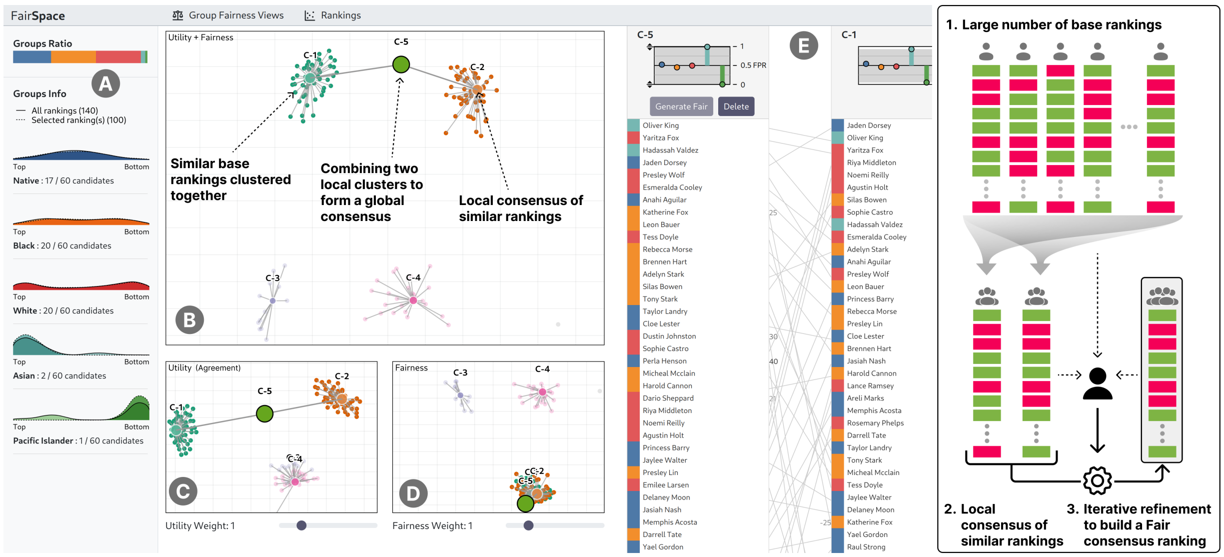

Figure 1: FairSpace is an interactive visualization system designed to explore and analyze large sets of rankings, enabling the creation of fair consensus rankings. A) Sidebar providing information of the groups in selected rankings and overall ranking distributions, B,C,D) Cluster Views displaying embedding space of rankings in a devised utility and fairness space, E) A rank comparison view. (Right) Illustrating the workflow supported by FairSpace.

Abstract

Analyzing large sets of rankings brings many unique challenges, with existing methods often obscuring key insights and overlooking outlier perspectives. This work explores these challenges through a visualization system, FairSpace, designed to effectively aggregate and visualize consensus patterns within large-scale ranking data. Addressing limitations of existing methods, FairSpace incorporates techniques for handling large numbers of individual rankings, revealing and highlighting similarities and differences across and within rankings, uncovering levels of fairness at multiple levels of aggregation, and identifying potential biases or outliers. Through interactive visualizations and data exploration tools, the system empowers users to understand the reasons behind the fair consensus, fostering transparent and informed decision-making.

Index Terms: Human-centered computing—Visualization—Visualization systems and tools

1 Introduction

In decision-making scenarios, like scholarship distribution, academic evaluation, and product recommendations, stakeholders often seek to create a consensus ranking by combining diverse perspectives. This complex process [2] requires careful consideration of different viewpoints and aggregating all preferences. Moreover, individual rankings may exhibit biases, potentially favoring certain groups over others [3]. For example, biases favoring male candidates over female ones in hiring can raise concerns about fairness in the ranking process.

To address these limitations, we designed FairSpace, an interactive visualization system for exploring large datasets of rankings and iteratively constructing fair consensus rankings. We use dimensional reduction techniques to visualize similar rankings in terms of utility and fairness and hierarchically construct a fair consensus ranking.

We begin with the data model used in FairFuse [4], which defines a list of candidates ranked by a set of rankers. Each candidate has a protected attribute such as race or gender, defining groups such as White, Black, Asian, etc. For demonstrating this system, we incorporate the fairness metrics Favored Pair Representation (FPR) and Attribute Pair Representation (ARP) to quantify the fairness score of rankings [1]. Similarly, we use the Kendall Tau distance to measure utility. Finally, we apply the Fair-Copeland Algorithm [1] for the auto-generation of fair consensus rankings by applying a certain fairness threshold. However, these metrics and algorithms can be replaced with others depending on the scenarios and requirements.

In addition to the existing tasks from FairFuse [4], we add new tasks considering the possibility of an increased number of rankings, such as: A) Identifying similar rankings and forming local clusters to simplify comparison, B) Supporting comparisons between clusters in terms of fairness and agreement, C) Supporting comparisons of individual rankings with their local cluster, D) Constructing a global consensus through a hierarchical approach, E) Constructing a global fair consensus ranking and analyzing its agreement with local

*e-mail:hshrestha@wpi.edu
†e-mail:kcachel@wpi.edu
‡e-mail:malkhathlan@wpi.edu
§e-mail:rundenst@wpi.edu
¶e-mail:ltharrison@wpi.edu

consensuses as well as individual rankings.

2 FairSpace Overview

FairSpace is designed to analyze and establish a fair consensus when dealing with a large number of rankings. The system comprises several views that offer a holistic perspective by presenting all the rankings collectively, as well as individual rankings. Additionally, we integrate visualizations to illustrate fairness metrics [4] and fair consensus ranking generation process.

2.1 Distribution Plots: Fairness across Groups

To provide a holistic view of the data with respect to protected groups, FairSpace uses distribution plots as shown in the sidebar of Fig. 1A. When a ranking or a cluster of rankings is selected, this view is updated with an overlay of groups' distribution of the selected ranking(s). Alongside the distribution plot, we display the ratio of candidates by the groups in the protected attribute to provide the idea of majority and minority groups in the candidates pool.

2.2 Cluster Views: Fairness+Utility Embeddings

FairSpace uses a dimensional reduction technique, specifically Multidimensional Scaling (MDS), to visualize the large number of rankings. However, this can be switched with other techniques such as t-SNE if required. These views represent the similarity of rankings in utility space, fairness space, and combination of both. This enables users to form cluster of similar rankings and analyze the clusters as a whole instead of going through individual rankings.

Using the Cluster Views (see Fig. 1), the decision maker can identify similar rankings in terms of both the fairness and utility metrics. These similar rankings can be selected by drawing a lasso around them, eventually allowing comparison of these rankings and generating their local consensus. The generated local consensus ranking is projected back to the cluster views enabling comparison between multiple local consensuses. Additionally, we incorporate the Group Fairness View from FairFuse [4] directly into the Cluster Views. Having the local consensus and the Group Fairness View within the Cluster Views allows the decision-maker to quickly grasp the underlying rankings without having to review each individual ranking. Any rankings in this view, when selected, are displayed in the Rank Comparison View Fig. 1E.

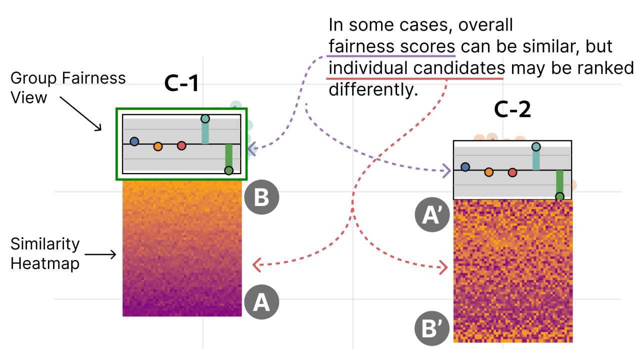

Figure 2: We design a Fair Divergence View by combining a heat map and Group Fairness View. The heatmap shows the (dis)similarity between rankings, in this case cluster C-1 is selected and its divergence from C-1 is displayed on C-2. Some of the bottom ranked candidates from C-1 (A) are seen to be ranked higher in C-2 (A'), and top ranked candidates from C-1 (B) are ranked at the bottom in C-2 (B').

The *Unified Cluster View* (Fig. 1B) allows decision-makers to visualize differences between clusters or between a cluster and a ranking using the *Fair Divergence View* (Fig. 2). Utilizing its heatmap, users can identify candidates ranked differently, such as how Cluster 2 compares to Cluster 1.

Finally, Cluster Views enable decision-makers to select local consensuses and construct a global consensus ranking. When building a global consensus, we consider the base rankings of the local consensuses. To emphasize the hierarchy, we use lines and opacity—lower levels in the tree have lower opacity, and vice versa.

2.3 Rank Comparison View and Group Fairness View

The *Rank Comparison View* (Fig. 1E) displays any selected ranking(s) from the *Cluster Views* enabling detailed comparison. Each candidate's name along with the color associated with the protected attribute is displayed, which was found to be effective in FairFuse [5]. This view utilizes Parallel Coordinates Plot design for comparison of candidates between the rankings.

At the top of each ranking, we have the *Group Fairness View* which displays the fairness scores (FPR and ARP) [1] of the ranking. The colored dots represent a group, like White, Asian, Black, etc. The dots above 0.5 FPR line represent advantaged groups and vice-versa. This view is also displayed on the *Cluster Views* for quick comparison. Additionally, the *Group Fairness View* is overlayed with a fairness slider. This slider can used to modify the ranking or a local consensus ranking to build a fairer ranking, essentially controlling the fairness threshold for the resulting fair ranking. This *modified ranking* is then projected back to the *Cluster Views* to provide feedback on how close/far it is from other rankings. This feedback can drive an iterative approach to build a global consensus that is fair and representative of all rankings.

3 Conclusion

The process of constructing a consensus ranking from a large set of individual rankings is complex, especially to mitigate biases. To address this challenge, we introduced FairSpace, a system to visualize and interact with many rankings. By facilitating the comparison of group distributions, fairness scores, and similarity scores, FairSpace empowers exploratory analysis of rankings, clusters of rankings, and consensus rankings. Through the integration of algorithms for generating fair consensus rankings and enabling manual ranking adjustments, along with a hierarchical approach for constructing a fair global consensus, FairSpace offers an iterative refinement approach to building a fair consensus from vast sets of rankings.

In the future, we plan to conduct a user study to validate the system and understand how people might use the system with large datasets. Additionally, we plan to expand this system to accommodate multiple protected attributes, provide support for partial rankings, and explore potential ways of integrating fairness in more realistic, uncertain contexts such as incomplete rankings.

Acknowledgments

This work was supported in part by NSF IIS #2007932.

References

[1] K. Cachel, E. Rundensteiner, and L. Harrison. Mani-rank: Multiple attribute and intersectional group fairness for consensus ranking. In *2022 IEEE 38th Intl. Conf. on Data Engineering (ICDE)*. IEEE, 2022.

[2] C. Kuhlman and E. Rundensteiner. Rank aggregation algorithms for fair consensus. *Proceedings of the VLDB Endowment*, 13(12), 2020.

[3] D. Pedreshi, S. Ruggieri, and F. Turini. Discrimination-aware data mining. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 560–568, 2008.

[4] H. Shrestha, K. Cachel, M. Alkhathlan, E. Rundensteiner, and L. Harrison. Fairfuse: Interactive visual support for fair consensus ranking. In *2022 IEEE Visualization and Visual Analytics (VIS)*, pp. 65–69. IEEE, 2022.

[5] H. Shrestha, K. Cachel, M. Alkhathlan, E. Rundensteiner, and L. Harrison. Help or hinder? evaluating the impact of fairness metrics and algorithms in visualizations for consensus ranking. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1685–1698, 2023.