

Fairness for AUC via Feature Augmentation

Hortense Fong
Yale University
USA

Anay Mehrotra
Yale University
USA

Vineet Kumar
Yale University
USA

Nisheeth K. Vishnoi
Yale University
USA

ABSTRACT

We study fairness in the context of classification where the performance is measured by the area under the curve (AUC) of the receiver operating characteristic. AUC is commonly used when both Type I (false positive) and Type II (false negative) errors are important. However, the same classifier can have significantly varying AUCs for different protected groups and, in real-world applications, it is often desirable to reduce such cross-group differences. We address the problem of how to select additional features to most greatly improve AUC for the disadvantaged group. Our results establish that the unconditional variance of features does not inform us about AUC fairness but class-conditional variance does. Using this connection, we develop a novel approach, fairAUC, based on feature augmentation (adding features) to mitigate bias between identifiable groups. We evaluate fairAUC on synthetic and real-world (COMPAS) datasets and find that it significantly improves AUC for the disadvantaged group relative to benchmarks maximizing overall AUC and minimizing bias between groups.

CCS CONCEPTS

• **Computing methodologies** → **Supervised learning by classification.**

KEYWORDS

Classification, Area under the ROC curve (AUC), feature augmentation, data collection and curation

ACM Reference Format:

Hortense Fong, Vineet Kumar, Anay Mehrotra, and Nisheeth K. Vishnoi. 2022. Fairness for AUC via Feature Augmentation. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3531146.3533126>

ACKNOWLEDGMENTS

This project is supported in part by an NSF Award (CCF-2112665).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
FAccT '22, June 21–24, 2022, Seoul, Republic of Korea
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9352-2/22/06.
<https://doi.org/10.1145/3531146.3533126>