

Achieving Fairness in AI with Synthetic Data

Artificial intelligence and machine learning increasingly inform decisions in hiring, lending, healthcare, and justice. Yet real-world datasets often encode historical bias, and models trained on them can reproduce or amplify inequities. Pre-processing via fair synthetic data is a promising: if we can generate data that mitigates bias at the source while preserving signal, downstream models can be both fair and useful. This talk introduces FDA (Fair synthetic data via Data Augmentation), a statistically principled framework that makes the fairness–faithfulness trade-off explicit and controllable. FDA jointly models a fair submodel and a faithful submodel, coupled by a single parameter $\alpha \in [0,1]$ that quantifies the fraction of bias removed. We prove clear operating points: $\alpha=0$ yields maximal fairness (with larger deviation from the original distribution), $\alpha=1$ recovers the original data in probability (hence in distribution), and intermediate α values guarantee calibrated compromises with interpretable bounds. Practically, FDA samples directly from simple predictive distributions, avoiding heavy black-box training. We further provide theory connecting FDA's α to fairness of downstream models. Together, these results deliver a transparent, efficient, and deployable path to generating fair synthetic data without sacrificing essential statistical structure.

Dr. Bei Jiang is an Associate Professor in the Department of Mathematical and Statistical Sciences at the University of Alberta, a Fellow of the Alberta Machine Intelligence Institute (Amii), and a Canada CIFAR AI Chair. She received her PhD in Biostatistics from the University of Michigan in 2014, followed by a postdoctoral appointment in the Department of Biostatistics at Columbia University (2014–2015), before joining the University of Alberta as an Assistant Professor in 2015. Dr. Jiang has authored more than 50 journal articles—including in the *Annals of Statistics*, *Journal of the American Statistical Association* and the *Journal of Machine Learning Research* and over 20 peer-reviewed conference papers at venues such as NeurIPS, ICML, ICLR, and AAAI. Her research focuses on Bayesian hierarchical modeling, statistical learning methods that advance privacy and fairness, and federated statistical inference. Dr. Jiang has an extensive record of service to the statistical community. She is currently serving on the SSC Equity, Diversity, and Inclusion Committee, the CANSSI Showcase Organizing Committee, the Committee of the COPSS Presidents' Award, and the JSM 2026 Program Committee. She is an Associate Editor for the *Journal of the American Statistical Association*. Dr. Jiang is the 2025 recipient of the COPSS Emerging Leaders Award, recognizing early-career statistical scientists whose leadership and scholarship are shaping the field.