

Winston Chou

wchou@alumni.princeton.edu · New York, NY

EMPLOYMENT

Netflix

Staff Data Scientist

New York, NY

3/2022 –

- Designed north star metric used to decide all personalization A/B tests at Netflix
- Conducted experiments to optimize \$15bn content budget; presented results to C-level execs
- Developed statistical methodology used in Netflix-wide observational causal inference platform
- Promoted from Senior (L5) to Staff (L6) Data Scientist in 2024

Apple

Staff Data Scientist

New York, NY

4/2020 – 3/2022

- Designed, executed, and analyzed the first A/B tests on the Siri voice assistant
- Wrote statistical software for analyzing experiments on 250m+ devices using Python and Spark
- Promoted from Senior (ICT4) to Staff (ICT5) Data Scientist in 2021

Facebook

Research Scientist

New York, NY

12/2018 – 4/2020

- Drove statistical models of missing data in advertising measurement from R&D to launch
- Partnered with globally distributed engineering teams to implement statistical models in production

EDUCATION

Princeton University

PhD, Political Science

Princeton, NJ

2019

SKILLS

- *Causal inference and experimentation.* Experiment design and analysis, A/B testing platforms, causal machine learning, observational causal inference
- *Statistics and machine learning.* Regression analysis, hierarchical models, Bayesian statistics, maximum likelihood estimation, missing data, variance reduction, data visualization
- *Statistical computing.* Python, R, Spark, SQL, Stan, PyMC

SELECTED RESEARCH

- [1] “[The value of personalized recommendations: Evidence from Netflix](#),” with Kevin Zielnicki, Guy Aridor, Aurélien Bibaut, Allen Tran, and Nathan Kallus. Under review.
- [2] “[Evaluating decision rules across many weak experiments](#),” with Colin Gray, Nathan Kallus, Aurélien Bibaut, and Simon Ejdemyr. *KDD '25: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. **Best Paper Award (Applied Data Science)**. [Netflix TechBlog post](#).
- [3] “[Learning the covariance of treatment effects across many weak experiments](#),” with Aurélien Bibaut, Simon Ejdemyr, and Nathan Kallus. *KDD '24: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. [Netflix TechBlog post](#).
- [4] “[Does residuals-on-residuals regression produce representative estimates of causal effects?](#)” with Apoorva Lal. *3rd Workshop on Causal Inference and Machine Learning in Practice, KDD '25*.