Application for
NBER Digital Economics and AI Tutorial

Zijun Ding
Machine Learning Department and Heinz College, Carnegie Mellon University

A. Contact Info

Zijun Ding
5000 Forbes Ave HBH 3041
Pittsburgh PA 15213
Email: dingzj@cmu.edu
Phone: 412-320-5260

B. Courses Pursued

I am a fourth-year PhD in Machine Learning and Public Policy at Carnegie Mellon University. The unique PhD program, jointly offered by CMU's Machine Learning Department at School of Computer Science and Heinz College of Information Systems and Public Policy, allows me to take a combination of rigorous courses in machine learning/AI and economics. I got trained in a holistic series of economic courses: Microeconomics, Econometrics series, Behavioral Economics and Information Economics. ML courses include but are not limited to Statistics and ML Theory, ML Algorithm, Advanced NLP, and Reinforcement Learning and Decision Making. Furthermore, I got the chance to explore frontier research in the intersection of economics, information technology, and AI. Courses I have taken include Advanced Machine Learning and Game Theory, ML Ethics and Society, and Privacy in GenAI. My GPA is 4.04/4.33. Besides being a student, I have TA'ed master level microeconomics course (CMU 95-710) and machine learning course (CMU 95-828, 10-605).

C. Research Plan

The past decade has witnessed an increasing amount of online data and the development of computational technologies that utilize personal data to achieve better performance. Foundation models, which have been used in various applications, also build on massive public and private data. The concerns on digital privacy lead to rapid development of privacy regulations and privacy preserving algorithms. I have been researching both aspects, with a recent focus on generative AI.

Economics is my main lens to examine the societal impact of online tracking and privacy regulations. Specifically, I am interested in how people and companies will be affected by

algorithms that use personal data for targeted advertising or technologies that protect online users from being tracked. In one of my ongoing projects, we aim to analyze a full picture of privacy regulation, by incorporating different stakeholders in the online advertising industry into our analysis. We find a potential that privacy regulations may benefit both advertising intermediaries (that collect data from users and provide targeting ad technology) and publishers (websites that sell spaces for advertisement) by alleviating competition between advertising intermediaries on the over-investment of targeting ad technology. A first-authored paper was accepted at the Workshop on Information Systems and Economics. For the consumer side, I am running a large-scale field experiment that investigates the effect of ad-blocking and anti-tracking technologies on people's online browsing and shopping. We monitor people's online browsing and shopping behavior and compare their satisfaction with the products they purchase. The results would illuminate the effect of data protection technologies on online users.

My second part of research starts with examining pitfalls of existing technologies and algorithms, and moves on to remedies with algorithms themselves or economic mechanisms. Generative algorithms based on foundation models have generated high quality contents as well as privacy and copyright concerns. From an algorithmic perspective, I am developing privacy preserving algorithms, such as unlearning and watermarking, to mitigate potential risks in generative AI. From users' perspective, I am curious about how privacy will be preserved when people (unconsciously) fail to protect their data under some scenarios. From some preliminary analysis, we have found a spillover effect of ad opting-out that people will be less targeted if they opt out of targeted ads from some other advertisers. The results are accepted by the Annual Computer Security Applications Conference, a top conference in computer security. The remaining of my PhD research will focus on developing a more responsible pipeline of (generative) AI. It ranges from designing economic mechanisms to collect training data from different sources, to building robust evaluation methods on AI applications. Economic incentives and downstreaming impact will be incorporated into the pipeline.

In summary, my research goal is to combine the power of privacy protection algorithms and regulations to guide the development of computational technologies and AI to maximize the benefits they bring to society.

While I have not decided on my dissertation title, I have been working closely with faculty members with a diverse expertise:
Ramayya Krishnan, Dean and William W. And Ruth F. Cooper Professor of Management Science And Information Systems, Carnegie Mellon University
Alessandro Acquisti, Trustees Professor of Information Technology and Public Policy, Carnegie Mellon University
Steven Zhiwei Wu, Assistant Professor of Computer Science, Carnegie Mellon

University
Lowell Taylor, H. John Heinz III University Professor of Economics, Carnegie Mellon University and Research Associate, NBER
Yue Wu, James Allen Associate Professor of Marketing and Business Economics, University of Pittsburgh