

I am a microeconomic theorist whose work is at the intersection of economic theory and machine learning (ML). My research uses ML techniques to improve economic models, and conversely uses economic frameworks to better design the acquisition and use of novel data. Specifically, my research focuses on three areas:

Machine Learning for Evaluating and Improving Economic Models. ML algorithms can be very effective at prediction, but they are often opaque and uninformative about underlying economic forces. Can these *black box* methods help us to build better interpretable models and predictions of behavior? [1], [4], [5], and [8] demonstrate various complementarities between ML techniques and traditional economic modeling.

In [4], we use ML algorithms to uncover the limits of prediction in different economic problems. In many problems, perfect prediction is not achievable, so we need a domain-appropriate benchmark against which to assess our models' errors. We use ML algorithms to construct such a benchmark. When an economic model substantially underperforms this benchmark, then there is important structure in that domain that the economic model does not capture. In [1], we further use ML algorithms to uncover some of that missing structure. By identifying new regularities, we are able to extend the existing economic models, producing new *interpretable* models that approximate the predictive performance of the black box.

Is performance at the level of an ML black box always desirable? In [5] we observe that one explanation for high predictive performance is that the economic model is in fact as flexible as the black box (at an extreme, the model may not be falsifiable). We propose a new computational method for quantifying the “restrictiveness” of an economic model based on how well the model fits randomly generated data. We apply the proposed method to evaluate models from three economic applications.

One important goal of economic models is to identify common structure across disparate (but related) economic domains. A model that falls short of the predictive limit may still be valuable for its generality. In new work [8], we propose a novel measure for the “transferability” of economic models.

Economic Implications of Big Data. Algorithms are increasingly used to guide consequential decisions (e.g., loans, bail, insurance rates). These algorithms are often optimized for the accuracy of their predictions—what unintended consequences may result for other economic objectives, such as incentives or redistribution?

In [6], we study the economic implications of big data in settings with moral hazard, and show that whether new data exacerbates or reduces moral hazard depends crucially on whether the new data is informative about a persistent quality of the agent versus a transient shock.

In work in progress [7], we consider the “fairness” of algorithmic decision-making. Use of new data inputs reduces *overall* error in the population, but it may shift the burden of errors disproportionately to members of a specific social group. We characterize the Pareto frontier between (aggregate) accuracy and group fairness, and show how this tradeoff is moderated by underlying statistical properties of the data.

(Dynamic) Information Acquisition. As the economy becomes more “information-rich,” firms and individuals must decide how to allocate limited attention and resources across an increasing variety of informational sources. In [2] and [3] we develop tractable models of this dynamic information acquisition problem.

In [2] we focus on the normative question of how an overabundance of informational sources affects the efficiency of social learning. We demonstrate an information acquisition externality: whom an agent chooses to listen to today will impact whom agents consider “most informative” tomorrow. This can lead to a *learning trap* where society persistently acquires information from lower quality sources and thus learns slowly. We provide precise conditions for when a learning trap can and cannot emerge, and relate these conditions to the underlying complementarities across the informational sources.

In [3] we consider the positive question of how to optimally allocate limited resources across informational sources. We show that when sources are not too strongly correlated, then the optimal information acquisition strategy takes a simple form, which we completely characterize. We then apply this solution to obtain new results in three economic applications: binary choice problems, attention manipulation, and dynamic news provision in a game between competing news sources.

REFERENCES

- [1] Fudenberg, Drew and Annie Liang (2019), “Predicting and Understanding Initial Play,” *American Economic Review*.
- [2] Liang, Annie and Xiaosheng Mu (2020), “Complementary Information and Learning Traps,” *Quarterly Journal of Economics*.
- [3] Liang, Annie, Xiaosheng Mu and Vasilis Syrgkanis (2021), “Dynamically Aggregating Diverse Information,” forthcoming at *Econometrica*.
- [4] Fudenberg, Drew, Jon Kleinberg, Annie Liang and Sendhil Mullainathan (2021), “Measuring the Completeness of Economic Models,” forthcoming at *Journal of Political Economy*.
- [5] Fudenberg, Drew, Wayne Gao and Annie Liang (2021), “How Flexible is that Functional Form? Measuring the Restrictiveness of Theories,” Working Paper.
- [6] Liang, Annie and Erik Madsen (2021), “Data and Incentives,” Working Paper.

Works in Progress:

- [7] Liang, Annie, Jay Lu and Xiaosheng Mu, “Information Design for Algorithms: Fairness and Accuracy.”
- [8] Andrews, Isaiah, Drew Fudenberg, Annie Liang and Chaofeng Wu, “Measuring Theory Transfer.”