

I am a microeconomic theorist whose work is at the intersection of information economics and machine learning (ML). I describe some of my papers below, organized into three branches.

(Dynamic) Information Acquisition. Firms and individuals must decide how to allocate limited resources across an increasing variety of informational sources. In [2] we study how an overabundance of informational sources affects the efficiency of social learning. We identify an information-acquisition externality: whom an agent chooses to acquire information from today impacts whom other agents consider “most informative” tomorrow. This can lead to a *learning trap* where society persistently acquires information from lower-quality sources and 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 characterize the dynamically optimal strategy for information acquisition. We show that when sources are Gaussian and not too strongly correlated, then the optimal information strategy can be exactly characterized. This tractable solution further allows us to generalize recent results about dynamic information acquisition, and to derive new results in applications such as dynamic competition among news sources.

Social Implications of Algorithmic Predictions. ML algorithms have traditionally been optimized to maximize the accuracy of predictions. Now that these algorithms also guide consequential predictions about people—such as who is creditworthy or who needs a medical procedure—objectives beyond accuracy have become relevant for their design and regulation. One emerging concern is the possibility that algorithms are very accurate for individuals in one social group while highly error-prone for another. In [6] we introduce a fairness-accuracy frontier that traces out the tradeoff between making predictions that are more accurate (i.e., achieving lower error *on average* in a population), versus predictions that are more group-fair (i.e., leading to similar error rates across groups), and show how this tradeoff is shaped by the statistical properties of the data. In [7], we study the disparate impact of algorithms through another lens: how use of new data for forecasting affects economic incentives in settings with moral hazard. We show that a simple property of the data (related to the persistence of the data’s predictive power) pins down its average effect on incentives. But this average effect masks a redistributive implication, with new data potentially overpowering incentives for some agents and underpowering incentives for others. This heterogeneity is first-order for welfare (even when inequality does not explicitly appear in the welfare function) and we describe how it shapes optimal data regulation.

Black Boxes and Economic Models. ML algorithms can be very effective at prediction, but 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? In [4], we propose using ML algorithms to construct a predictive benchmark for economic models. This is useful for understanding whether a model predicts imperfectly because the economic observables are of limited informativeness about the outcome, or because the model misses important regularities. In [1], we show that when ML algorithms substantially outpredict economic models, we can sometimes use those algorithms to uncover useful directions to extend our models. In a particular setting, we are able to identify a one-parameter extension of a relevant economic model, which is nearly as predictive as the black boxes.

Is predictive performance at the level of ML algorithms always desirable? In [5] we observe that one explanation for why a model fits many data sets well is if it is unrestrictive (potentially unfalsifiable). We propose a new computational method for quantifying the *restrictiveness* of an economic model based on how well the model fits randomly generated data, and show that some economic models with a small number of parameters in fact impose very few restrictions on data. In [8] we note that an important goal of economic models is to identify common structure across distinct (but related) economic domains. An economic model that performs worse than a black box when the two are evaluated in a specific economic domain may nevertheless generalize better to new domains where we have no data. We propose a novel measure of the *transferability* of economic models and black boxes, and demonstrate in an application that more structured economic models indeed transfer better than black-box methods do.

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 (2022), “Dynamically Aggregating Diverse Information,” *Econometrica*.
- [4] Fudenberg, Drew, Jon Kleinberg, Annie Liang and Sendhil Mullainathan (2022), “Measuring the Completeness of Economic Models,” *Journal of Political Economy*.
- [5] Fudenberg, Drew, Wayne Gao and Annie Liang (2022), “How Flexible is that Functional Form? Quantifying the Restrictiveness of Theories,” Working Paper.
- [6] Liang, Annie, Jay Lu and Xiaosheng Mu (2022), “Algorithm Design: Fairness and Accuracy,” Working Paper.
- [7] Liang, Annie and Erik Madsen (2022), “Data and Incentives,” *R&R at Theoretical Economics*.
- [8] Andrews, Isaiah, Drew Fudenberg, Annie Liang, and Chaofeng Wu (2022), “The Transfer Performance of Economic Models,” Working Paper.