Summary of Interests in AI

My primary research interests are in market design and fundamental aspects of machine learning. In addition to topics within each field, I believe there are fruitful avenues for work at their intersection. I have recently been working on the incentives of platforms when agents make decisions with the assistance of machine learning algorithms. On Amazon, for example, a recommender system is a key input into buyers' decisions. There have been some stories about sellers attempting to manipulate the recommender system to the detriment of buyers, which raises an interesting technical question of how one should design recommender systems to be resistant to such manipulation attempts. I have been using techniques from learning theory to analyze this problem and devise and prove properties of robust procedures. There is also an interesting economic question of whether it is in Amazon's interest to deploy a manipulation-proof algorithm. This work is in the early stages, but I hope to compare the equilibrium level of manipulability with the social planner's benchmark.

I have also become interested in more direct connections between game theory and machine learning. An influential paper by Freund and Schapire shows that their boosting algorithm for online learning can be interpreted as learning to play an approximate mixed strategy Nash equilibrium against an adversary that "plays" distributions of data. Since their algorithm achieves desirable properties (low regret), it seems that it can be used for normative decision making in real world settings. In collaboration with a private equity firm, I have done some work on analyzing the performance of these types of strategies in making sequentially arriving and irreversible investment decisions.

Advances in AI have created many promising new research directions and exciting applications. I would be grateful for the opportunity to learn more about this topic at the workshop and conference in Toronto this fall.