Andreas A. Haupt haupt@mit.edu +1 (857) 928 5787

Course of Study: Engineering-Economic Systems (coursework in Economics and Artificial Intelligence/Machine Learning)
Year: 4
Fields: Industrial Organization, Microeconomic Theory, Online Learning, Recommendation Systems.

Dissertation
Committee: Alessandro Bonatti, Dylan Hadfield-Menell, Eric Maskin, David Parkes
Working Title: Full-Stack Platform Policy
Working Abstract: Platforms that match consumers to goods gain in importance as computational abilities to understand human preferences increase. Many platforms deploy learning algorithms in their design, in particular in recommendation systems and dynamic pricing. For effective oversight over learning algorithms, diverse Economic and computational methods are needed. This thesis explores three of them: First, theoretical modeling can highlight incentives of stakeholders, learning outcomes of algorithms, and their interaction, to shed light on potential consumer harms. Second, empirical evidence helps assess which of these are economically relevant. Finally, algorithmic remedies can be proposed to mitigate harms.

Research Projects/Journal Preparations

Recommending to Strategic Users (w/ C. Podimata, D. Hadfield-Menell)
Recommendation systems are pervasive in the digital economy. An important assumption in many deployed systems is that user consumption reflects user preferences in a static sense: users consume the content they like with no other considerations in mind. However, as we document in a large-scale online survey, users do choose content strategically to influence the types of content they get recommended in the future.
We model this user behavior as a two-stage noisy signaling game between the recommendation system and users: the recommendation system initially commits to a recommendation policy, presents content to the users during a cold start phase which the users choose to strategically consume in order to affect the types of content they will be recommended in a recommendation phase. We show that in equilibrium, users engage in behaviors that accentuate their differences to users of different preference profiles. In addition, (statistical) minorities out of fear of losing their minority content exposition may not consume content that is liked by mainstream users. We next propose three interventions that may improve recommendation quality (both on average and for minorities) when taking into account strategic consumption: (1) Adopting a recommendation system policy that uses preferences from a prior, (2) Communicating to users that universally liked ("mainstream") content will not be used as basis of recommendation, and (3) Serving content that is personalized-enough yet expected to be liked in the beginning. Finally, we describe a methodology to inform applied theory modeling with survey results.

Contextually Private Mechanisms (w/ Z. Hitzig)
Consider a mechanism design environment in which a designer sequentially queries agents' private information to determine the outcome of a choice rule. The designer's social and technological environment constrains the set of access protocols that it can use. In high-tech environments, arbitrary cryptographic protocols are admissible, and so privacy concerns do not constrain the set of available choice rules. In other environments, privacy desiderata are needed to guide design. A protocol is contextually private for a choice rule if each piece of information learned about each participant is needed to determine the outcome. We characterize choice rules that can be implemented with a contextually private protocol under different assumptions about the class of admissible protocols. Under the assumption that

private information must be elicited sequentially from uniquely-identified agents, the serial dictatorship and the first-price auction have contextually private implementations. However, no other kth-price auction has a contextually private implementation, nor does any stable choice rule (in college assignment) or individually rational and efficient choice rule (in house assignment). Under a more generous assumption, which additionally allows the designer to anonymously count the number of agents whose private information has a certain property, the designer has wider scope: we describe a contextually private tâtonnement-like implementation of several choice rules that are not contextually private under the stricter assumption.

Towards Psychologically-Grounded Dynamic Preference Models (w/ M. Curmei, B. Recht, D. Hadfield-Menell)
Designing recommendation systems that serve content aligned with time varying preferences requires proper accounting of the feedback effects of recommendations on human behavior and psychological condition. We argue that modeling the influence of recommendations on people's preferences must be grounded in psychologically plausible models. We contribute a methodology for developing grounded dynamic preference models. We demonstrate this method with models that capture three classic effects from the psychology literature: Mere-Exposure, Operant Conditioning, and Hedonic Adaptation. We conduct simulation-based studies to show that the psychological models manifest distinct behaviors that can inform system design. Our study has two direct implications for dynamic user modeling in recommendation systems. First, the methodology we outline is broadly applicable for psychologically grounding dynamic preference models. It allows us to critique recent contributions based on their limited discussion of psychological foundation and their implausible predictions. Second, we discuss implications of dynamic preference models for recommendation systems evaluation and design. In an example, we show that engagement and diversity metrics may be unable to capture desirable recommendation system performance.

Get It in Writing: Formal Contracts Mitigate Social Dilemmas in Multi-Agent RL (w/ P. Christoffersen, D. Hadfield-Menell)
Multi-agent reinforcement learning (MARL) is a powerful tool for training automated systems acting independently in a common environment. However, it can lead to sub-optimal behavior when individual incentives and group incentives diverge. Humans are remarkably capable at solving these social dilemmas. It is an open problem in MARL to replicate such cooperative behaviors in selfish agents. In this work, we draw upon the idea of formal contracting from economics to overcome diverging incentives between agents in MARL. We propose an augmentation to a Markov game where agents voluntarily agree to binding state-dependent transfers of reward, under pre-specified conditions. Our contributions are theoretical and empirical. First, we show that this augmentation makes all subgame-perfect equilibria of all fully observed Markov games exhibit socially optimal behavior, given a sufficiently rich space of contracts. Next, we complement our game-theoretic analysis by showing that state-of-the-art RL algorithms learn socially optimal policies given our augmentation. Our experiments include classic static dilemmas like Stag Hunt, Prisoner's Dilemma and a public goods game, as well as dynamic interactions that simulate traffic, pollution management and common pool resource management.

Prior-Independent Auctions for the Demand Side of Federated Learning (w/ V. Mugunthan)
Federated learning (FL) is a paradigm that allows distributed clients to learn a shared machine learning model without sharing their sensitive training data. While largely decentralized, FL requires resources to fund a central orchestrator or to reimburse contributors of datasets to incentivize participation. Inspired by insights from prior-independent auction design, we propose a mechanism, FIPIA (Federated Incentive Payments via Prior-Independent Auctions), to collect monetary contributions from self-interested clients. The mechanism operates in the

semi-honest trust model and works even if clients have a heterogeneous interest in receiving high-quality models, and the server does not know the clients' level of interest. We run experiments on the MNIST, FashionMNIST, and CIFAR-10 datasets to test clients' model quality under FIPIA and FIPIA's incentive properties.

Risk aversion in learning algorithms and recommendation systems (w/ A. Narayanan)
Consider online learning algorithms that simultaneously make decisions and learn from feedback. Such algorithms are widely deployed in recommendation systems for products and digital content. This article exhibits a bias of online learning algorithms towards less risky alternatives, and how it shapes demand on recommendation systems. First, we consider k-armed bandits. We prove that ε-Greedy chooses a riskless arm over a risky arm of equal expected reward with probability arbitrarily close to one. This is a consequence of undersampling of arms with bad reward estimates. Through experiments, we show that other online learning algorithms exhibit risk aversion as well. In a recommendation system environment we show that content that yields less noisy reward from users is favored by the algorithm. Combined with equilibrium forces driving strategic content creators towards content of similar expected quality, the advantage for content that is not necessarily better, just less volatile, is exaggerated.

Opaque Contracts (w/ Z. Hitzig)
Firms have access to abundant data on market participants. They use these data to target contracts to agents with specific characteristics, and describe these contracts in opaque terms. In response to such practices, recent proposed regulations aim to increase transparency, especially in digital markets. In order to understand when opacity arises in contracting and the potential effects of proposed regulations, we study a moral hazard model in which a risk-neutral principal faces a continuum of weakly risk-averse agents. The agents differ in an observable characteristic that affects the payoff of the principal. In a described contract, the principal sorts the agents into groups, and to each group communicates a distribution of output-contingent payments. Within each group, the realized distribution of payments must be consistent with the communicated contract. A described contract is transparent if the principal communicates the realized contract to the agent ex-ante, and otherwise it is opaque. We provide a geometric characterization of the principal's optimal described contract as well as conditions under which the optimal described mechanism is transparent and opaque. We apply our results to the design and description of driver payment schemes on ride-hailing platforms.