Optimal Multi-Attribute Decision Making in Social Choice Problems

Sujoy Sikdar
Department of Computer Science
Rensselaer Polytechnic Institute
sikdas@rpi.edu

Abstract

My thesis solves problems of *decision making when alternatives are characterized by multiple attributes*, under natural restrictions on agents' preferences that are motivated by practical and cognitive considerations. Computing optimal decisions in these settings is often hard in general. Fortunately, agents' preferences often have some natural structure, which have been studied in cognitive psychology literature. This makes several important problems tractable. I identify cases where such structure accurately models preferences in real world data, and provide efficient mechanisms to compute optimal outcomes for important social choice problems with theoretical guarantees.

Introduction

Several social choice problems involve making decisions over alternatives that are characterized by multiple attributes. For example, students may want to exchange research papers and time slots for presentation in a seminar class, in cloud computing, agents may want to allocate multiple types of resources, including CPU, memory, and storage, or agents may vote on multiple referenda addressing different issues. The problem of finding the best allocation, or the best decision on all issues can become challenging due to: (i) the number of alternatives grows exponentially with the number of attributes, and (ii) agents' preferences over the alternatives may have a complex combinatorial structure. My thesis addresses the following challenges that arise in multi-attribute decision making problems:

Preference Representation. I consider representations that balance the cognitive load of forming preferences over all alternatives, and practical considerations of elicitation.

Computation. Computing optimal decisions becomes challenging in combinatorial domains. I provide efficient algorithms for several problems, or complexity results otherwise.

Strategic Behavior. Strategic manipulation is a serious concern in several settings. Wherever possible, I address this by providing strategyproof mechanisms.

Modeling and Learning Preferences. I develop models and methods to learn users' preferences from real world data of their implicit preferences and opinions, which often do not capture full preferences in a structured form.

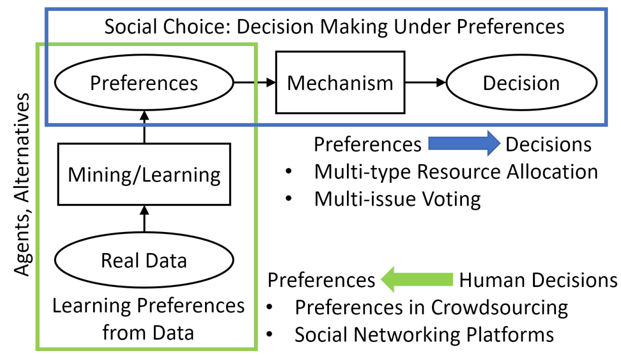

Figure 1: The structure of my thesis.

My Contributions

I approach multi-attribute decision making from two directions, as summarized in Figure 1. Several positive results in decision making disappear when alternatives are characterized by multiple attributes. The main theme of my work is the identification of reasonable restrictions on the problem domain under which the positive results may be recovered. My thesis delivers the following positive message: *We can design efficient mechanisms with desirable properties for multi-attribute decision making under natural assumptions on agents' preferences.*

Direction 1: Decision Making under Preferences

Multi-type Resource Allocation. Here the attributes correspond to different types of items. In (Sikdar, Adali, and Xia 2017a), we provide the first positive results for multi-type housing markets (Moulin 1995), when agents' preferences are lexicographic with possibly different importance orders over the types. Here, a collection of agents each endowed with a set of indivisible items of different types, have preferences over *bundles* consisting of subsets of all items. The goal is to find a redistribution of items that best satisfies agents' preferences without exchanging money. The notion of the *core* (Shapley and Scarf 1974) of the market, the set of allocations where no group of agents has incentive to deviate by exchanging their initial endowments within the group, is the most widely accepted and studied notion of what the best such redistribution may be, as it is intuitively stable, and implies individual rationality and Pareto optimality.

When each agent owns a single item, and has strict preferences over items, a core allocation always exists, and can be computed in polynomial time by Gale's celebrated Top-Trading-Cycles (TTC) algorithm (Shapley and Scarf 1974), which is also strategyproof. Despite the vast literature on housing markets, multi-type housing markets have received very little attention because the core may be empty (Konishi, Quint, and Wako 2001). Indeed, as Sönmez and Ünver (2011) noted: *"Positive results of this section [on housing markets] no longer hold in an economy in which one agent can consume multiple houses or multiple types of houses."* This is the problem we address.

The main contributions are: (i) An extension of the TTC mechanism to multi-type housing markets which is strict core selecting, non-bossy, strong group strategyproof when agents cannot lie about importance orders, and runs in polynomial time. (ii) We show that no mechanism can satisfy both strict core selection and strategyproofness, when agents are allowed to lie about importance orders. (iii) Computational complexity results for checking if a given allocation is in the strict core.

Multi-issue voting. CP-nets (Boutilier et al. 2004) are a popular and natural preference representation language to compactly represent agents' preferences over multiple issues. Previous positive results assume acyclic CP-nets, or a common structure for all agents (see (Lang and Xia 2016)), leaving the following open questions: How to aggregate CP-nets without assuming a common structure? How to define the optimal outcome for cyclic CP-nets?

In (Sikdar, Adali, and Xia 2017b), I provide a loss-minimization framework that allows reasoning about CP-nets, and their probabilistic extension PCP-nets (Bigot et al. 2013; Cornelio et al. 2013) with full generality, allowing for the aggregation of possibly cyclic CP-nets with different structures. I define natural notions of the loss of an alternative as a function of other alternatives that are preferred over it. The optimal outcome minimizes the loss. The main contributions are: (i) Computational complexity results, and identifying cases when computing an optimal outcome is computationally tractable. (ii) A new class of voting rules for aggregating CP-nets that satisfy desirable properties.

Direction 2: Learning Preferences from Data

I model preferences using representation schema inspired by work on lexicographic heuristic decision making involving multiple factors from psychology literature (Gigerenzer and Goldstein 1996), as well as machine learning techniques. My contributions are: (i) Novel natural language and semantic features relevant to human decision making. (ii) An analysis of why users give some responses with higher frequency in a question answering setting using "fast and frugal heuristic (FFH)" based decision models borrowed from psychology literature that require very little training data. (iii) In (Horne, Adali, and Sikdar 2017), we model factors that drive discussions and voting behavior on reddit by learning to rank comments on noisy data, and analyze the learned models to understand factors affecting users' behavior.

Future Work

I plan to converge on a unified approach by developing the theory on mechanism design for preference representations that accurately model agents' preferences in practice.

Direction 1. For housing markets, I plan to consider more general problem settings, more expressive preference representations, and develop mechanisms that satisfy desirable properties in these settings.

In multi-type resource allocation (Mackin and Xia 2016), often, different organizations are responsible for the allocation of different types of items. I am interested in the following research questions: What properties can we expect from the resulting allocations? What properties must local mechanisms satisfy to obtain desirable allocations of all types?

Direction 2. I plan to build problem specific models for a better understanding of online communities discussing topics ranging from news and politics to question answering.

References

Bigot, D.; Zanuttini, B.; Fargier, H.; and Mengin, J. 2013. Probabilistic Conditional Preference Networks. In *UAI-13*.

Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. Preference-based constrained optimization with CP-nets. *Computational Intelligence* 20(2):137–157.

Cornelio, C.; Goldsmith, J.; Mattei, N.; Rossi, F.; and Venable, K. B. 2013. Updates and uncertainty in CP-nets. In *Advances in Artificial Intelligence*. Springer. 301–312.

Gigerenzer, G., and Goldstein, D. G. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103(4):650–669.

Horne, B. D.; Adali, S.; and Sikdar, S. 2017. Identifying the social signals that drive online discussions: A case study of reddit communities. 1–10.

Konishi, H.; Quint, T.; and Wako, J. 2001. On the Shapley–Scarf economy: the case of multiple types of indivisible goods. *Journal of Mathematical Economics* 35(1):1—15.

Lang, J., and Xia, L. 2016. Voting in Combinatorial Domains. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 9.

Mackin, E., and Xia, L. 2016. Allocating Indivisible Items in Categorized Domains. In *IJCAI-16*, 359–365.

Moulin, H. 1995. *Cooperative Microeconomics: A Game-Theoretic Introduction*. Prentice Hall.

Shapley, L., and Scarf, H. 1974. On cores and indivisibility. *Journal of Mathematical Economics* 1(1):23–37.

Sikdar, S.; Adali, S.; and Xia, L. 2017a. Mechanism Design for Multi-Type Housing Markets. In *AAAI-17*.

Sikdar, S.; Adali, S.; and Xia, L. 2017b. Optimal decision making with cp-nets and pcp-nets. In *AAMAS-17*.

Sönmez, T., and Ünver, M. U. 2011. Matching, Allocation, and Exchange of Discrete Resources. In Benhabib, J.; Bisin, A.; and Jackson, M. O., eds., *Handbook of Social Economics*. North-Holland. chapter 17, 781–852.