Project Summary—CRII: RI: Allocating Indivisible Items in Categorized Domains

Allocating indivisible items to multiple agents without monetary transfer is one of the most important problems in multi-agent systems [32], covering a wide range of applications including allocating jobs to machines, courses to students, tasks and equipments to anti-terror troops. Designing good mechanisms to allocate *sharable* and *non-sharable* items has been a very active research topic at the interface of Computer Science and Economics.

In the past decade there has been a rapid growth in the theory and applications of allocation problems in Computer Science and Economics in parallel. These problems have been studied in many active research areas, including *multiagent resource allocation* [32], *one-sided matching* (a.k.a. *assignment problems*) [98], *voting* [80], and especially *combinatorial voting* [27, 67]. The joint consideration of computation and economics, despite being very natural and promising, has not been well-developed mainly due to three barriers: preference bottleneck, computational bottleneck, and threats of agents' strategic behavior. We are also not aware of any previous work that studies allocating sharable and non-sharable items at the same time.

In the proposed project, we plan to study a novel class of allocation problems where sharable items and non-sharable items are allocated at the same time. These problems are called *categorized domain allocation problems (CDAPs)* [106], where the items are dividied into multiple *categories*, and each agent must get at least one item from each category. For example, to organize a seminar we must allocate 10 discussion topics and 10 dates to 10 students, and we also need to allocate one classroom from a list of 5 rooms to the students for all discussions. In this example 25 items belong to 3 categories: classrooms (5 sharable items), topics (10 non-sharable items), and dates (10 non-sharable items). CDAPs extend and bring together many important areas including multiagent resource allocation, one-sided matching, and voting.

To break the three barriers encountered in previous research, we outline three interrelated fundamental directions for CDAPs: (1) Design and evaluation of new graphical languages for preference representation. (2) Analysis and prevention of agents' strategic behavior. (3) Design and evaluation of new mechanisms via extensions and compositions.

Intellectual Merit. CDAPs are a novel and general framework for allocating sharable and non-sharable items at the same time. Because of the generality of CDAPs, we can bring together methods and notions developed in many important areas, and the categorical information will greatly help us break the three barriers in previous research. Our work will directly benefit researchers working on multiagent resource allocation, one-sided matching, and voting. Moreover, since we will also bring in methodologies and ideas in algorithmic game theory, computational social choice, and knowledge representation, our work is of potential interest to researchers in those fields as well.

Broader Impact. As an interdisciplinary topic that extends and brings together many important areas, the proposed research is important in its own right. With the success of this project, we can embrace the synergy of multiagent resource allocation, one-sided matching, and combinatorial voting to design new allocation mechanisms for sharable and non-sharable items from the joint consideration of computation and economics. These new mechanisms will improve the economic efficiency, computational efficiency, and fairness in both high-stake and low-stake resource allocation problems in multiagent systems, socio-economics systems, and operations research. For example, it will push the state of the art of high-stake applications including task-equipment allocation in anti-terror activities, airport traffic management, and various scheduling problems. It will also improve many low-stake, everyday-life organizational applications, for example the allocation of topics, dates, and classroom for seminar; and the allocation of course slots under curriculum (categorical) constraints; etc.

In addition to mentoring graduate and undergraduate students, the PI also plans to develop an undergrad-level interdisciplinary course, organize workshops, give tutorials, and perform other outreach activities.

Keywords: multiagent resource allocation; one-sided matching; computational social choice; combinatorial voting.

CRII: RI: Allocating Indivisible Items in Categorized Domains

PI: Lirong Xia

1 Introduction

Allocating indivisible items to multiple agents without monetary transfer is one of the most important problems in multi-agent systems [32], covering a wide range of applications including allocating jobs to machines, courses to students, tasks and equipments to anti-terror troops.

As an illustrative running example, suppose we are organizing a seminar. We must allocate 10 discussion topics and 10 dates to 10 students, and we also need to allocate one classroom from a list of 5 rooms to the students for all discussions. Students may have different and complicated preferences over (topic, date, room) combinations, and their preferences over one component may depend on some other components they get. For example, it is quite possible that if an agent gets an easy topic, then she would prefer to have an early date to better enjoy the rest of the seminar; but if she gets a hard topic, then she would prefer to have a late date for better preparation. Agents' preferences over dates may also depend on the room allocated to them: a student may prefer a certain date because it is easier for her to come from another class on that day.

In general, there are two types of items to allocate: *sharable* items and *non-sharable* items. An item is sharable, if once it is allocated, it must be used by all agents. An item is non-sharable, if it can only be allocated to one agent. In other words, when allocating sharable items, we want to make a group decision for the agents on the items to share. On the other hand, non-sharable items are only available for personal use. In the running example, we want to allocate 25 items to 10 agents. Among the 25 items, the 5 rooms are sharable items, and the 10 topics and 10 dates are non-sharable.

The design and analysis of mechanisms to allocate sharable and non-sharable items have constituted an active research area at the interface of computer science and economics. The main challenges are to design allocation mechanisms to achieve the following two types of efficiency.

1. *Economic efficiency* concerns whether an allocation mechanism satisfies several desired economic criteria. One major criterion is how sensitive the mechanism is w.r.t. agents possibly reporting untruthful preferences.

2. *Computational efficiency* concerns communicational and computational efficiency of the allocation mechanisms, and whether it is easy to elicit agents' preferences.

A common approach is to first ask the agents to report their preferences over bundles of items to a central agent, who then decides the allocation based on all agents' reported preferences. While a lot of work has been done along this line in computer science and economics in parallel, not much has taken a joint consideration. In addition, most previous work focused on allocation problems where all items are sharable, or allocation problems where all items are non-sharable. We are not aware of a general setting that studies the allocation of both types of indivisible items at the same time.

On the computer science side, allocation problems have been studied under the active research area known as *centralized multiagent resource allocation* [32], where the items can be either all sharable or all non-sharable, and the number of items can be much larger than the number of agents. The main research agenda is to tackle computational challenges in preference representation and communication, and the computation of allocations that maximize (1) various kinds of *social welfare*, which are numerical measurements for the social satisfaction about the allocation, and (2) *fairness*, a notion formally developed in an active interdisciplinary research area called *fair division* [118, 26, 72].

On the economics side, when all items are non-sharable, the allocation problems have been studied under the research stream called *one-sided matching*, also known as *assignment problems* [98]. Most previous research focused on simple, nicely structured, yet highly practical domains, for example assigning n

houses to n families. The main objective is to design allocation mechanisms that satisfy some desired economic properties, especially *strategy-proofness*, which states that no agent has incentive to report untruthful preferences to improve her own allocation.

When all items are sharable and only one of them can be allocated to the agents, the allocation problem becomes a classical *voting* problem, where the agents must choose an item (winner) as their group decision. Previous work in voting theory focused on the design and analysis of voting mechanisms that satisfy desired economics properties [80]. Recently, computational aspects of voting have been investigated under the emerging area of computational social choice [27].

Allocation problems are also closely related to *combinatorial voting* (a.k.a. *voting in multi-issue domains*) [27, 67], which is the voting setting where the set of alternatives is extremely large and has a combinatorial structure. More precisely, in combinatorial voting, there are $p \geq 1$ issues; each issue can take a value in a *local domain*; and the alternatives are uniquely identified by their values on all issues. Combinatorial voting seeks to make a group decision for all agents on each issue based on their preferences. This can be seen as an allocation problem where all items are sharable, and we must allocate p items, one from each local domain, to the agents. Previous research in combinatorial voting studied problems considered in centralized multiagent resource allocation, one-sided matching, and classical voting: use compact languages to represent preferences, explore the computational complexity of computing a joint decision, and design and evaluate mechanisms w.r.t. economics properties including strateg-proofness.

The interplay between computer science and economics on general allocation problems is very natural and promising. However, this interplay is much more challenging than it appears, and little progress has been made in the past decade due to the following barriers.

- *Preference bottleneck:* When the number of items is not too small, it is impractical for the agents to explicitly represent their preferences over all 2^m bundles of m items, either ordinally by a full ranking over these 2^m bundles, or cardinally by giving one number (utility) for each of these 2^m bundles.
- *Computational bottleneck:* Even if the agents can report their preferences, computing an "optimal" allocation is often a hard combinatorial optimization problem [81].
- *Threats of agents' strategic behavior:* An agent may have incentive to report untruthful preferences to make her own allocation better off. This may lead to a socially undesirable allocation.

What's new in this proposal? We propose to work towards breaking the three aforementioned barriers for a class of novel and general allocation problems, called *categorized domain allocation problems (CDAPs)*, that allocate sharable and non-sharable items at the same time. CDAPs extend and bring together many important areas including resource allocation, one-sided matching, and voting.

Let us first briefly describe the motivation before formally introducing CDAPs. In many allocation problems, the set of indivisible items can be divided into multiple categories and each agent must get at least one item from each category. For instance, in the running example, there are three categories of items: classrooms, topics, and dates. The classroom category contains only sharable items, and the topic category and date category contain only non-sharable items. Each agent must get at least one item from each category to lead the discussion. To model such allocation scenarios, the PI introduced categorized domain allocation problems in an ongoing work [106].

Definition 1. *A* categorized domain *is a set of indivisible items with the following properties: there are* $p \geq 1$ *categories; each category is either* sharable*, meaning that it only contains sharable items, or* non-sharable*, meaning that it only contains non-sharable items; each item belongs to one of the* p *categories.*

In a categorized domain allocation problem (CDAP)*, we want to find an allocation where (1) all agents get the same set of sharable items, (2) no non-sharable item is allocated to more than one agent, (3) every agent gets at least one item from each category, and (4) no monetary transfer is allowed.*

A special yet still general case of categorized domains is *basic categorized domains*, where the number of items in each non-sharable category equals to the number of agents. The corresponding assignment problems are called *basic CDAPs*, where we further require that each agent must get exactly one item from each category. The running example illustrates a basic CDAP, since the agents must make a joint decision to choose one classroom, and in addition, each agent must get exactly one topic and one date.

Let us give a more complicated example to illustrate a non-basic CDAP. Consider a generic multi-agent setting of allocating tasks and equipments to agents, e.g. allocating anti-terror tasks and equipments to a group of SWAT troopers. There are two categories of "items": the tasks and the equipments, and it is possible that a trooper gets more than one task or more than one equipment. Moreover, both tasks and equipments may be further divided into multiple sharable or non-sharable sub-categories, so that the number of categories may be well beyond a few. For example, the troopers may want to share an armored car, and each trooper must get at least one of firearms, body armors, ballistic shields, entry tools, etc.

CDAPs are a very general class of problems that naturally combine and extend many important resource allocation problems and voting problems: when there is only one sharable category or one non-sharable category, CDAPs degenerate to centralized multi-agent resource allocation and fair division for indivisible items without monetary transfer. Even basic CDAPs generalize many important problems: when there is only one non-sharable category, basic CDAPs degenerate to one-sided matching problems; when all categories are sharable, basic CDAPs degenerate to combinatorial voting; and moreover, when there is only one sharable category, basic CDAPs become standard voting problems.

Given that CDAPs generalize these important and well-studied problems, we can bridge together notions and techniques developed from these problems. On the modeling level, when all categories are sharable or all categories are non-sharable, CDAPs can be seen as standard centralized multi-agent resource allocation without categorical information, and then one may be tempted to think that techniques in centralized multi-agent resource allocation can be directly applied to CDAPs. However, first, these techniques do not apply to general CDAPs with both sharable and non-sharable categories; and second, even for CDAPs with only sharable categories or only non-sharable categories, we still have to face the preference bottleneck and computational bottleneck after directly applying techniques in centralized multi-agent resource allocation.

Therefore, we need to develop new techniques to solve these general, exciting, and brand new CDAPs to overcome the three barriers in previous research. In doing so, we outline three interrelated fundamental directions.

1. **Design and evaluation of new graphical languages.** Compact graphical languages have been proven useful in many applications of artificial intelligence. However, not much progress has been made in using graphical languages to represent agents' preferences for allocation problems due to the lack of structures in previously-studied allocation problems [18]. We propose to utilize the categorical information to design new compact graphical languages for CDAPs with the help of ideas from recent work in knowledge representation. Then, we plan to evaluate these new languages and compare them with existing languages. These would help us overcome the preference bottleneck.

2. **Analysis and prevention of agents' strategic behavior.** To understand, alleviate, and hopefully ultimately eliminate the threats of agents' strategic behavior, we plan to focus on the following three questions: (1) Can we characterize strategy-proof mechanisms when agents' preferences are represented by compact languages? (2) For non-strategy-proof mechanisms, how can we evaluate the effect of agents' strategic behavior on social welfare? (3) If agents' strategic behavior is harmful, can we use high computational cost to prevent them? The first question mainly adopts a viewpoint in economics, the second is closely related to algorithmic game theory [79], and the third is closely related to computational social choice [27]. We note that even though some of these questions have been studied for other allocation or voting problems, the structure of agents' preferences and allocations in CDAPs make these problems highly non-trivial and challenging.

3. **Design and evaluation of new mechanisms via extensions and compositions.** Our ultimate goal is to design new mechanism for CDAPs to overcome all three barriers encountered in previous research. Given the generality of CDAPs, this is an extremely challenging task, especially because no previous research has been done for allocating sharable and non-sharable items at the same time. We believe that the categorical information can open up a world of new possibilities. In particular, the reminiscence between basic CDAPs and combinatorial voting inspired us to extend ideas that have been proven successful in combinatorial voting, to design new allocation mechanisms via **extensions** and **compositions** of existing mechanisms. After new mechanisms are designed, we plan to evaluate them w.r.t. criteria often used in the literature. This direction takes the joint consideration of computation and economics, and will benefit from the first and second proposed research directions as well as combinatorial voting.

Relationships between the three proposed research directions and knowledge representation, computational social choice, and algorithmic game theory are illustrated in Figure 1.

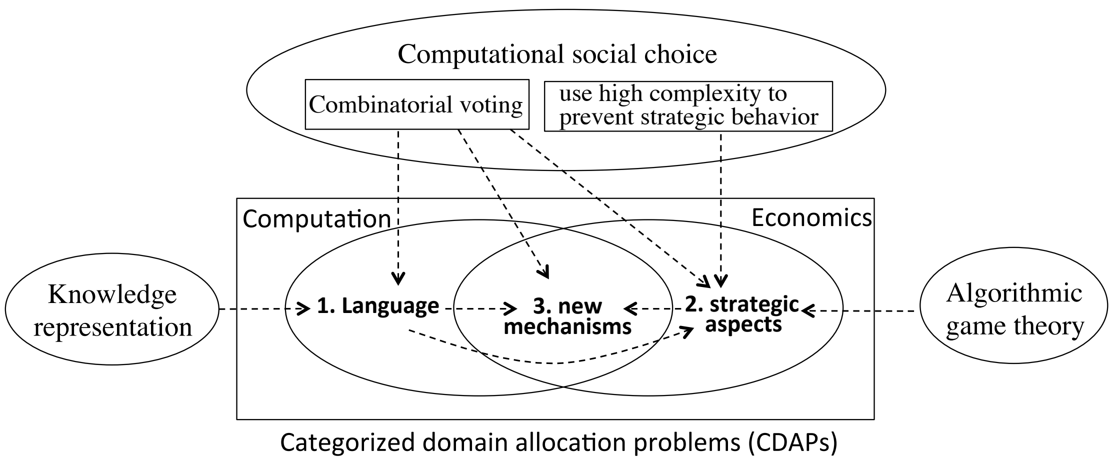

Figure 1: Relationships between the three proposed research directions and knowledge representation, computational social choice, and algorithmic game theory. An arrow from A to B means that B benefits from A.

Intellectual Merit. In the past decade there has been a rapid growth in the theory and applications of allocation problems in computer science and economics. The joint consideration of computation and economics, despite being very natural and promising, has not been well-developed and has been hindered by three barriers: preference bottleneck, computational bottleneck, and threats of agents' strategic behavior. The proposed work focuses on CDAPs, a novel and general framework for allocating sharable and non-sharable items at the same time. Because of the generality of CDAPs, our proposed work brings together methods and notions developed in many important areas, and the categorical information will greatly help us break the three barriers in previous research to advance the state of the art. Our work will directly benefit researchers working on multiagent resource allocation, one-sided matching, and voting. Moreover, since we will also bring in methodologies and ideas in algorithmic game theory, computational social choice, and knowledge representation, our work is of potential interest to researchers in those fields as well.

Broader Impact. As an interdisciplinary topic that extends and brings together many important areas, the proposed research is important in its own right. With the success of this project, we can embrace the synergy of multiagent resource allocation, one-sided matching, and combinatorial voting to design new allocation mechanisms for sharable and non-sharable items from the joint consideration of computation and economics. These new mechanisms will improve the economic efficiency, computational efficiency, and fairness in both high-stake and low-stake resource allocation problems in multiagent systems, socio-economics systems, and operations research. For example, it will push the state of the art of high-stake applications including task-equipment allocation in anti-terror activities, airport traffic management, and various scheduling problems. It will also improve many low-stake, everyday-life organizational applications, for example the allocation of topics, dates, and classroom for seminar; and the allocation of course slots under curriculum (categorical) constraints; etc.

In addition to mentoring graduate and undergraduate students, the PI also plans to develop an undergrad-level interdisciplinary course, organize workshops, give tutorials, and perform other outreach activities.

Structure of the proposal. In Section 2 we point out a simple yet important connection between basic categorized domains and multi-issue domains in combinatorial voting. In Section 3, Section 4 and Section 5, we discuss the three proposed directions in more detail by highlighting some concrete research questions. For each of these research questions we will give a short-term plan (focusing on basic CDAPs) and a long-term plan (focusing on general CDAPs). In particular, we will illustrate how considering the category information can significantly help us break the three barriers in previous research. Afterwards, we describe a more detailed timeline and international collaborations opportunities in Section 6. We conclude with a statement of educational activities and prior NSF support.

Prior Work and Capabilities. The PI has extensive experience in the design and analysis of mechanisms without money from the joint consideration of computation and economics. Specifically, CDAPs were recently introduced by the PI in an ongoing work [106]. Some of the PI's previous work in computational social choice directly benefits the proposed work. In particular, the PI is writing a book chapter with Jérôme Lang on combinatorial voting [67], which directly benefits all three proposed directions as illustrated in Figure 1. We will describe the PI's relevant expertise for each proposed research direction in more detail.

1.1 Background and Related Work

Since no monetary transfers are allowed in CDAPs, our proposed work is from some important areas that study allocation problems *with* monetary transfer, including *mechanism design with money* [55] and *combinatorial auctions* [40], though many high-level principles in these areas have been adopted to allocation problem without monetary transfer.

In centralized multiagent resource allocation, agents often use a *compact language* to represent preferences. When agents' preferences are cardinal, most previous work allows the agents to use the following compact languages to represent their utilities over bundles of items, mostly originated from combinatorial auctions [78]: XOR [94], k-additive language [37, 33], and propositional logic formulas [62, 17, 63]. In these cases, it is natural to compute an allocation that maximizes one of the three types of social welfare: utilitarian social welfare, which is the sum of all agents' utilities of their allocated bundles; egalitarian social welfare, which is the minimum utility of the agents; and Nash social welfare, which is the product of all agents' utilities. Computational complexity of computing these welfare-maximizers for various compact languages has been studied in a series of work [33, 69, 43, 22, 89, 90]. Studying (in)approximability of welfare-maximizers is also an active research area, which we refer to a recent survey by Nguyen et al. [77]. When agents' preferences are ordinal, much less work has been done due to the lack of structure in agents' preferences [18]. We will elaborate more on ordinal languages in Section 3.

Another important criterion that has been used to evaluate the fairness of allocation mechanisms is *envy-freeness*. A mechanism satisfies envy-freeness if no agent prefers the bundle allocated to another agent over the bundle allocated to her. Computing or approximating envy-free allocations for indivisible items has been studied in a series of papers [69, 32, 20, 90, 19, 3]. There is also a rich line of research on fair division of *divisible* items, which we refer to a recent survey by Procaccia [87].

Traditional one-sided matching problems focused on two settings: house allocation [54] and housing market [96, 98]. In both problems there are n agents and equal number of indivisible items. Agents' preferences are ordinal, and are represented by full rankings over the items. Items do not have preferences over the agents, which distinguishes one-sided matching problems from the Nobel-prize-winning two-sided matching problems [50], where there are two groups of agents and each agent has preferences over the agents in the *other* group.

In house allocation problems, we want to find an allocation where each agent gets exactly one item. Previous work focused on characterizing deterministic and randomized mechanisms that satisfy some desired economic properties, the most important one being strategy-proofness [100, 46, 1, 12].

In housing market problems, it is further assumed that every agent initially owns an item. The most popular mechanism is Gale's *top-trading-cycles* mechanism [96, 2] that reallocates the items in multiple rounds: in each round, each remaining agent j points to another remaining agent who has j's most-preferred item that is still available. In the induced graph, whenever there is a cycle, we implement the cycle by giving the item of each agent in this cycle to her predecessor. Then, we remove all agents in cycle and proceed to the next round. This process is repeated until no cycle exists. Top-trading-cycles satisfies many desired economic properties [96, 92, 91, 70] and has been extended to handle indifferences between items [56].

More recently, a third type of one-sided matching problems have been introduced, where some agents initially own items and some do not. This is a natural combination of house allocation and housing market. For this setting, new mechanisms have been designed and analyzed [2, 97, 98, 45].

Designing allocation mechanisms for applications where the number of agents is not the same as the number of items has also been studied from an economic perspective: strategy-proof mechanisms have been characterized [83, 82, 84, 53, 60, 44], and practical mechanisms have been used to allocate courses to students at Harvard Business School [29]. Recently, Budish et al. studied the implementation of randomized mechanisms [30]. We will discuss some of these mechanisms in more detail in Section 5.

In addition to strategy-proofness, mechanisms developed for one-sided matching problems are often evaluated w.r.t. other normative properties, including Pareto-optimality, non-bossiness, and neutrality. A mechanism satisfies

• *Pareto-optimality*, if for any preference profile, compared to the bundles the agents obtain under the mechanism, there is no allocation where no agent is worst off and at least one agent is better off;

• *non-bossiness*, if for any preference profile, no agent can change any other agent's allocation under the mechanism without changing her own allocation by reporting differently;

• *neutrality*, if the names of the items do not matter. More precisely, if we apply a permutation on the names of the items, then the allocation under the mechanism should be permuted in the same way.

2 Basic Categorized Domains and Multi-issue Domains

As mentioned in the Introduction, combinatorial voting over multi-issue domains can be viewed as a basic CDAP where all categories are sharable. In this section, we use formal definitions to make clear a connection between multi-issue domains and basic categorized domains that may contain non-sharable categories. This connection can help us adopt techniques developed for combinatorial voting to CDAPs, thus will provide a good starting point for our proposed research.

In a CDAP instance, we let $\mathcal{D}$ denote a set of m indivisible items. Let $C_1^S, \ldots, C_{p_1}^S$ denote the sharable categories and let $C_1^N, \ldots, C_{p_2}^N$ denote the non-sharable categories. That is, $\mathcal{D} = \mathcal{D}_S \cup \mathcal{D}_N$, where $\mathcal{D}_S = C_1^S \cup \ldots \cup C_{p_1}^S$ and $\mathcal{D}_N = C_1^N \cup \ldots \cup C_{p_2}^N$. A bundle of items is *admissible* if it contains at least one item from each non-sharable category. Let $\mathcal{B}$ denote the set of all admissible bundles. In CDAPs, we only need to represent agents' preferences over admissible bundles.

For basic CDAPs, each agent must get exactly one item from each category. Therefore, the set of admissible bundles can be viewed as the Cartesian product of all (sharable and non-sharable) categories. That is, $\mathcal{B} = C_1^S \times \cdots \times C_{p_1}^S \times C_1^N \times \cdots \times C_{p_2}^N$. This bears a strong resemblance to the structure of multi-issue domains, where there are p issues, each issue i takes a value in a *local domain* D_i, and the multi-issue domain $\mathcal{X}$ consists of alternatives that are uniquely identified by their values on all issues. That is, $\mathcal{X} = D_1 \times \cdots \times D_p$. Hence, we immediately have the following observation.

Observation 1. *A basic categorized domain can be viewed as a multi-issue domain where an admissible bundle naturally corresponds to an alternative in the multi-issue domain.*

We note that Observation 1 only concerns the similarity between basic categorized domains and multi-issue domains. The allocation problems for basic categorized domains (i.e. basic CDAPs) are much more general than combinatorial voting.

3 Design and Evaluation of New Graphical Languages

Compact languages for preference representation in combinatorial space, especially for multiagent resource allocation, combinatorial auctions, and voting, are usually evaluated by the following major criteria [78].

• *Complexity.* What is the communication complexity of eliciting the preferences represented by the language? What is the computational complexity of determining basic properties of the preferences under the language, e.g. comparing two given bundles?

• *Expressive power.* What kind of preferences can be represented by the language?

• *Usability.* What is the cognitive difficulty for the agents to express their preferences?

To overcome the preference bottleneck, one promising compromise is *graphical languages*, which often have low complexity, satisfactory expressive power, and high usability. *Conditional preference networks (CP-nets)* [16] are one of the most popular and useful graphical language for representing ordinal preferences, especially in combinatorial voting [64]. A CP-net models an agent's preferences over multi-issue domains by using a directed graph to represent conditional independence relations among agents' preferences over different issues. A CP-net consists of two parts: a directed graph over all issues; and a set of *conditional preference tables*, one for each issue, containing one linear order over the issue's local domain for each combination of values of the issue's parents in G. In this way, the agents' preferences over an issue only depends on the values of its parents in G.

For example, consider the CP-net in Figure 2 with three issues where all local domains are $\{0, 1\}$. Issue 1 has no parent, so its conditional preference table only contains one pairwise comparison (linear order over two alternatives) between 0 and 1; issue 2 has two parents, so its conditional preference table contains four comparisons, one for each combination of values of issue 1 and issue 3. Each entry in conditional preference tables has the form $\vec{d} : a \succ b$, encoding the following semantics: "given that the values of parents are $\vec{d}$, I would prefer the current issue taking a than taking b, regardless of the values of other issues."

Under this interpretation, the $0 \succ 1$ entry for issue 1 is decoded to four pairwise comparisons: $000 \succ 100$, $001 \succ 101$, $010 \succ 110$, $011 \succ 111$, where ijk represents the alternative whose values on issue 1, 2, 3 are i, j, k, respectively. We note that in each of these four pairs, the two alternatives in comparison only differ on their values for issue 1.

Similarly, the $0 : 0 \succ 1$ entry for issue 3 is decoded to two pairwise comparisons: $000 \succ 001$ and $100 \succ 101$. Putting all decoded pairwise comparisons together gives us a (possibly cyclic) partial order over the multi-issue domain, which is the preferences represented by the CP-net.

The idea behind CP-nets has been extended in *UCP-nets* [15] to model cardinal preferences in multi-issue domains. UCP-nets model conditional utility functions in a way similar to conditional preferences tables in CP-nets. While both CP-nets and UCP-nets have been proven very useful in many applications in artificial intelligence, none of them can be easily used in general allocation problems, as Bouveret et al. commented [18]:

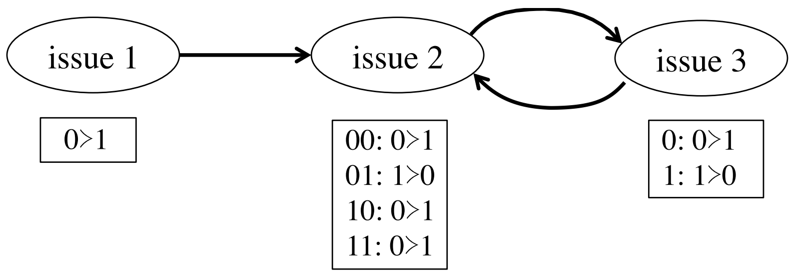

Figure 2: A CP-net for three issues.

And after all, the set of all bundles has a combinatorial structure, so we might wonder whether these languages [CP-nets] *would not be good candidates for our concern. It turns out that they are not, because they cannot easily express statements such as everything else being equal, I prefer to have the bundle* $\{a, b\}$ *rather than the bundle* $\{a, c, d\}$.

The above comment is based on the fact that in general, agents need to express preferences over *all* bundles of items, but graphical languages like CP-nets only provide localized comparisons between two bundles that differ on one issue/item. To address this, Bouveret et al. proposed a graphical language called *CI-nets* [18], allowing the agent to express arbitrary conditional independence relationships among *bundles*. This language, due to its generality, is much more complicated than CP-nets, and it is not clear whether such generality is necessary in practice. Hence, designing good compact graphical languages for allocation problems is still an open question.

Fortunately, for categorized domains, the categorical information provides a natural way to model conditional preferences as in CP-nets. Due to the relationship between basic categorized domains and multi-issue domains revealed in Observation 1, we can directly use CP-nets in basic categorized domains. By doing so, the comment by Bouveret et al. may be alleviated or even completely avoided!

There are still a lot of challenges behind this encouraging first step. For example, even though CP-nets and UCP-nets have been proven superior in many applications, it is not clear whether they are still superior for basic categorized domains compared to other popular (non-graphical) compact languages. More importantly, designing new graphical languages for general categorized domains is still highly challenging. Therefore, we propose to work on the following research question.

- **Research Question 1. Explore the extensions of CP-nets and other graphical languages to categorized domains.**

Short term plan: We believe that compact graphical languages including CP-nets and UCP-nets are advantageous compared to existing ordinal and cardinal languages. Therefore, we plan to conduct a thorough comparison between CP-nets (respectively, UCP-nets) and existing languages w.r.t. the criteria described in the beginning of this section. Then, we plan to investigate the extensions of other popular graphical languages to basic categorized domains including but is not limited to TCP-nets [23] and LP-trees [13].

Long term plan: Once the properties of these new graphical languages for basic categorized domains are well-understood, we plan to move on to develop novel graphical languages for general categorized domains. This is highly challenging since none of CP-nets and their extensions can be directly applied, as we no longer have a similar relationship as in Observation 1.

We plan to start with a graphical language that is similar to CP-nets in spirit, but instead of using a linear order to represent an agent's conditional local preferences over all items in a category, we propose to use an existing compact language to represent an agent's conditional local preferences over *bundles* of items. This gives us a whole new class of composite languages. Then, we plan to conduct formal comparisons between these new languages and the existing ones w.r.t. the criteria described in the beginning of this section.

4 Analysis and Prevention of Agents' Strategic Behavior

In one-sided matching or more generally mechanism design, strategy-proofness is one of the most desired criterion for mechanisms for various reasons including simplicity (agents can just report their true preferences without putting efforts into computing a false report) and fairness (agents do not benefit from being more sophisticated and computationally powerful) [86].

While strategy-proof mechanisms with other desired economic properties exist when monetary transfer is allowed, for example the Vickrey-Clarke-Groves (VCG) mechanism in auctions [102, 34, 52], in many settings without monetary transfer, strategy-proofness was shown to be incompatible with many other desired economic properties. For example, in social choice theory, the Gibbard-Satterthwaite theorem [51, 95] shows that dictatorships are the only strategy-proof mechanisms that satisfy a few natural normative properties. A dictatorship always selects the top choice of a special agent (called the dictator), regardless of other agents' preferences.

For house allocation problems, a similar negative result was proved by Svensson [100], which states that the only strategy-proof allocation mechanism that also satisfies non-bossiness and neutrality (definitions can be found in the end of Introduction) are *serial dictatorships*, which are defined as follows.

Definition 2 (Serial dictatorships). *A serial dictatorship for n agents is characterized by an order $\mathcal{O}$ over the agents, and it allocates the n items in n rounds: in the j-th round, the j-th agent in $\mathcal{O}$ chooses her most-preferred item that is still available.*

Similar characterizations have been shown for other more general allocation problems [83, 82, 84, 53, 60, 44]. As we have mentioned in the Introduction, CDAPs generalize one-sided matching and voting, so one may wonder if Svensson's characterization or the Gibbard-Satterthwaite theorem would immediately work for general CDAPs with multiple categories.

The answer is no. Without touching technical details, the best way to understand the reason might be the following: in Svensson's characterization and Gibbard-Satterthwaite theorem, agents' preferences are represented by linear orders over the items (alternatives). For CDAPs with more than one category, agents' preferences are much more complicated and structured, so that the allocation problem cannot be simply viewed as a classical one-sided matching or voting problem. Moreover, it is well-known in social choice that when agents' preferences are represented by a compact language, it is possible that other natural mechanisms are strategy-proof [71, 14, 8, 9, 10, 68, 57, 76, 109]. For the same reason, other similar characterizations mentioned above cannot be directly applied to CDAPs.

Therefore, our first research question in this section is the following.

- **Research Question 2. Characterization of strategy-proof allocation mechanisms in general categorized domains.** In an ongoing work [106], the PI proved a non-trivial extension of Svensson's characterization of serial dictatorships to any basic categorized domains with only non-sharable categories, assuming that each agent reports a full ranking over all admissible bundles.

  It would be very interesting to know whether other types of strategy-proof mechanisms exist when a categorized domain contains sharable categories. When the categorized domain contains only one sharable category, as we mentioned, this degenerates to the classical voting setting, where we can immediately apply the Gibbard-Satterwaite theorem. But for the more general cases the answer is not clear. The same question can also be asked when agents' preferences are represented by a compact language. Even for basic categorized domains these questions are highly non-trivial. The PI's experience in studying a similar problem for combinatorial voting [109] will be helpful. Investigating these questions for general categorized domains is a natural yet challenging next step.

So far most previous work on characterizing strategy-proof mechanisms in social choice and house allocation problems are negative, suggesting that if we insist on using a strategy-proof mechanism, then we are only left with mechanisms that do not satisfy many other desired economic properties. But then why should we insist on using a strategy-proof mechanism? What would be the consequence of using a non-strategy-proof mechanism?

To understand the effect of agents' strategic behavior, a common approach is to apply game theory [49] and then study the *price of anarchy* of the game, as being often done in the active interdisciplinary field called algorithmic game theory, initialed by Koutsoupias and Papadimitriou [61]. The price of anarchy of a game is the *optimal* social welfare in all possible outcomes divided by the *worst* social welfare in the outcomes of the game, which are usually the *Nash Equilibria* [75]. Price of anarchy provides a quantifiable way to evaluate the effects of agents' strategic behavior: a high price of anarchy indicates a highly negative worst-case effect of agents' strategic behavior on the social welfare.

When agents only have ordinal preferences, the question is much harder to answer, since social welfare is not well-defined. One may still assume that the agents have certain kinds of utility functions (for example, Borda utility function), and then conduct usual price of anarchy analysis [28]. However, this type of results may not be very informative when agents' utility functions are different from the assumption.

Another approach is to develop ordinal versions of the price of anarchy, by using ordinal measurements to describe the quality of the equilibrium outcomes [42, 108, 112]. For example, if we can show that when agents are strategic, *everyone* will be allocated a bundle that is ranked below 99% of all bundles w.r.t. her true preferences, while in the optimal allocation everyone gets an almost-in-top bundle, then we can argue that strategic behavior is highly undesirable. This type of phenomena were sometimes called *paradoxes* in voting [24].

These considerations naturally lead to our second research question in this section.

- **Research Question 3. Characterization of equilibria outcomes and (ordinal) price of anarchy for allocation mechanisms in categorized domains.** For classical price of anarchy, the PI plans to collaborate with Prof. Elliot Anshelevich at RPI, who is an expert in price of anarchy. Study of the ordinal price of anarchy is much more challenging. The PI's experience in studying ordinal price of anarchy for voting games [108, 112], especially for combinatorial voting [112], will be helpful.

Even if we observe high price of anarchy of a non-truthful mechanism, we may still want to use it because it may satisfy many desired properties that are not satisfied by other mechanisms. In such cases we want to develop practical ways to prevent agents' strategic behavior for a given mechanism. One promising idea is to explore the possibility of using high computational complexity to prevent agents' strategic behavior. This idea was first explored by Bartholdi et al. in the voting setting [11], and has become a major research direction in computational social choice. See [47, 48, 93, 27] for recent surveys. The idea is, if we can show that it is computationally too hard for an agent to compute a beneficial false report, then she may give up and report truthfully instead.

To the best of our knowledge, using high computational complexity to prevent agents' strategic behavior has not been investigated for multiagent resource allocation or one-sided matching problems. Therefore, our final research question in this section is the following.

- **Research Question 4. Explore the possibility of using high computational complexity to prevent agents' strategic behavior in categorized domains.** The PI has extensive experience in developing this idea in the voting setting, including characterizing the worst-case complexity [117, 35, 113, 73, 103, 85, 41] and typical-case complexity [107, 113, 73, 41, 105] of agents' strategic behavior, and using information constraints to make the computation of strategic behavior harder [38].

Short-term plan. We plan to investigate the three research questions for existing mechanisms and compact languages in basic categorized domains. More precisely, for Question 2, we plan to characterize strategy-proof mechanisms for basic categorized domains that contain both sharable and non-sharable items under existing compact languages, including XOR, k-additive, and CI-nets. For Question 3, we plan to start with the price of anarchy for existing mechanisms. For Question 4, we can immediately work on characterizing the computational complexity of strategic behavior for non-strategy-proof mechanisms including the three types of welfare-maximizers w.r.t. existing compact languages.

Long-term plan. In the long run, we plan to answer the three questions for general categorized domains under the new languages proposed in Section 3. Specifically, for the new mechanisms proposed in the next section, we can compare them with existing mechanisms w.r.t. their answers to the three questions proposed in this section.

5 Design and Evaluation of New Mechanisms

In this section, we first describe criteria that are often used to evaluate allocation mechanisms in the literature. Then, we propose to explore two new types of mechanisms in categorized domains that are **extensions** and **compositions** of existing mechanisms. In the literature, the following criteria are often used to evaluate allocation mechanisms.

• *Welfare properties.* For any allocation mechanism, we can evaluate it by the social welfare of the allocation *assuming that the agents are truthful*, for worst case or for average case. Here the worst case is taken over all possible preferences profiles and the average case is computed w.r.t. a distribution over preference profiles.

• *Strategic aspects of agents.* These include the three research questions we proposed in Section 4, i.e., is the mechanism strategy-proof? If the mechanism is not strategy-proof, then what is the price of anarchy? What is the computational complexity of agents' strategic behavior?

• *Fairness.* Lipton et al. [69] defined two types of degree of envy. Both first define the degree of envy for a single agent, then take the maximum degree of envy across the agents. For the first type, a single agent's degree of envy is the maximum difference in utility between her current bundle and bundle of any agent (including herself). For the second type, a single agent's degree of envy is the maximum ratio of her utility of the bundle of any agent (including herself) over her current utility.

• *Other normative properties, including Pareto-optimality, non-bossiness, and neutrality.* These criteria are briefly described in the end of Introduction. We note that unlike welfare properties, price of anarchy, and fairness, these normative properties are 0-1 valued, meaning that for any of them, a mechanism either satisfies it or does not satisfy it.

To study welfare properties, price of anarchy, and fairness, it is required that the agents provide utilities for bundles. When ordinal languages are used, we may use Borda utility, where the social welfare has a close relationship to the average rank as studied in e.g. [21, 58]. Alternatively, we may study these properties in a setting where the agents do have utilities over bundles in their mind, but are required to report preferences in an ordinal language, as has been recently studied in the voting setting [88, 31].

Due to the generality of CDAPs, designing new mechanisms for CDAPs is extremely challenging. Instead of working completely from scratch, we propose to utilize the categorical information to explore the following two possibilities: design by extensions and design by compositions of existing mechanisms.

5.1 Extensions of existing mechanisms

Let us first describe *sequential allocation protocols* [21]. A sequential allocation protocol for m non-sharable items and n agents is defined by an order $\mathcal{O}$ over the agents, where some agents may appear more than once. Items are allocated in m rounds: in round i, the i-th agent in $\mathcal{O}$ chooses an item that is still available. A special sequential allocation protocol called *draft* has been used in Harvard Business School to allocate course slots to students, and its practical economic efficiency and students' strategic behavior have been studied recently [29]. Draft uses a balanced order to ensure fairness: $(1, 2, \ldots, n, n, n-1, \ldots, 1, 1, 2, \ldots)$.

For CDAPs, the PI has proposed a natural family of allocation mechanisms inspired by the sequential allocation protocols [106]. We call these new mechanisms *categorized sequential mechanisms*. Take basic categorized domains for example, a categorized sequential mechanism is defined by an order over all (agent, category) pairs, so that in each round the designated agent picks an item from the designated category.

To understand the properties of categorized sequential mechanisms and compare them to existing mechanisms, we propose to investigate the following question.

- **Research Question 5. Multi-criteria evaluation of categorized sequential mechanisms.** We aim at achieving two levels of goals: in the first level, we want to understand which categorized sequential mechanism is best among all categorized sequential mechanisms w.r.t. criteria described in the beginning of this section. In the second level, we want to compare the best categorized sequential mechanisms to existing mechanisms including welfare-maximizers, serial dictatorships, and sequential allocation protocols, w.r.t. the criteria described in the beginning of this section. In addition, it would also be very interesting to study the probabilistic combination of categorized sequential mechanisms as in random serial dictatorships [99].

5.2 Compositions of existing mechanisms

Inspired by seat-by-seat voting rules [25] and sequential voting rules [64, 66] in multi-issue domains, we plan to design and evaluate a framework called *composite allocation mechanisms* that combine sub-allocation-mechanisms, one for each category, to allocate items in a sequential way. Similar ideas were also explored recently in a mechanism design setting [101].

More precisely, fix an order over the categories, w.l.o.g. $[1 \succ 2 \succ \cdots \succ p]$. Given p sub-mechanisms $M_1, \ldots, M_p$, one for each category, the composite allocation mechanism $\text{Comp}(M_1, \ldots, M_p)$ allocates the items in the following p steps: in step i, we ask the agents to report their preferences over category i conditioned on the allocation of previous steps, then we apply M_i on the reported preferences to allocate items in category i, and finally announce the allocation to the agents and move on to the next category. The architecture of composite allocation mechanisms is illustrated in Figure 3.

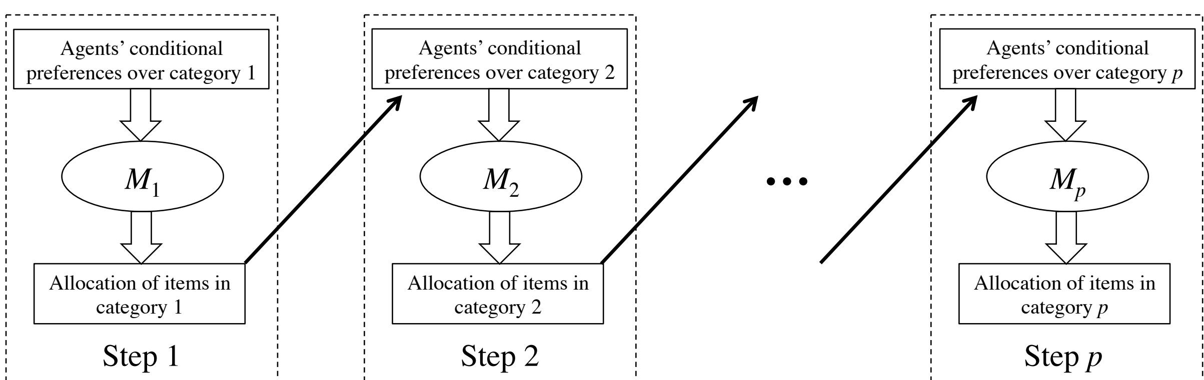

Figure 3: Composite allocation mechanism $\text{Comp}(M_1, \ldots, M_p)$.

Example 1 (*composite allocation mechanism for basic categorized domain*)**.** Consider again the running example of topic-date-room allocation. We can allocate the bundles in three steps: first, we ask the students to report their preferences about the classrooms, and use a commonly studied voting rule, for example the Borda rule [80], to decide the room. Given this, we ask the agents to report preferences over the topics, and then use a serial dictatorship to allocate the topics. Then, we decide an initial allocation of the dates, which may correspond to a good order to present the topics. Finally, we use the top-trading-cycles mechanism to reallocate the dates.

This is an example of Comp(Borda, serial dictatorship, top-trading-cycles). In fact, it was the mechanism the PI designed in the first place to organize the seminar in the course he taught in 2013 fall, except that the classroom was decided by university registrar rather than by student voting (see Section 7 for more information on this course). □

Example 2 (*composite allocation mechanisms for non-basic categorized domain)***.** The actual allocation mechanism used by the PI was slightly different: 6 out of 10 students formed 3 groups, each has two members, and the remaining 4 students decided to lead the discussions alone. To ensure fairness, each group must get 2 topics and 2 time slots. This shows a non-basic CDAP.

Let us assume again that the students use the Borda voting rule to decide the room in the first step. In the second step, the PI used a variant of the draft mechanism to allocate topics to groups and individuals. The mechanism for reallocating time slots was still top-trading-cycle. Another feature was that the order used in draft was determined by the chronological order the PI received students' reported preferences, to encourage early reports. This is an example of Comp(Borda, draft, top-trading-cycles). □

Composite allocation mechanisms have a number of potential advantages over existing mechanisms. (1) They are much more flexible. The designer is free to choose any combination of sub-mechanisms. (2) They may break the computational bottleneck since we only need to compute an allocation of items in the same category in each step, instead of solving a complicated global optimization problem as in most classical multiagent resource allocation approaches. (3) The localized structure of composite allocation mechanisms copes very well with the new languages we proposed in Section 3. (4) Since the agents may not need to specify their preferences in full, composite allocation mechanisms preserve more privacy as well.

We are still facing a lot of challenges. For example, the properties of composite allocation mechanisms w.r.t. the criteria mentioned in the beginning of this section are not very clear, and they can be hard to analyze due to the combinatorial structure of composite allocation mechanisms. Also it is not clear how to determine the order to execute sub-mechanisms. These lead to our last research question in this section.

- **Research Question 6. Design and evaluation of composite allocation mechanisms for categorized domains.**

5.3 Research agenda

Short-term plan. In the short term, we plan to work on Research Questions 5 and 6 for basic categorized domains. Even for this case the work is highly nontrivial. For welfare properties, strategic aspects, and fairness, we plan to conduct worst-case and average-case evaluations by theoretical analysis as well as computer simulations.

Given the reminiscence of basic categorized domains and multi-issue domains described in Observation 1, the PI's extensive experience in the design and analysis of mechanism for combinatorial voting in multi-issue domains [115, 116, 110, 66, 114, 35, 111, 109, 112, 36, 39, 65] put him in a good position to address these questions. Specifically, two of the PI's recent work are directly relevant to the questions proposed in this section: one is a published work on the strategic behavior and the computational complexity of manipulation under sequential allocation protocols [59], and the other is an ongoing work that contains some preliminary results for categorized sequential mechanisms [106].

More concretely, for Research Question 5, we plan to extend the techniques used by Kalinowski et al. [59] to understand the welfare properties of categorized sequential mechanisms. For Research Question 6, we plan to explore the relationship between the properties of composite allocation mechanisms and the properties of its sub-mechanisms. For example, we are interested in the following type questions "if all sub-mechanisms satisfy a normative property X, does their composition also satisfy X?", where X can be any properties mentioned in the beginning of this section. Similar questions have been studied for sequential voting by the PI [66, 114].

Long-term plan. We plan to work on Research Questions 5 and 6 for general categorized domains. Again, we plan to conduct both theoretical analysis and computer simulations for the evaluation w.r.t. welfare properties, strategic aspects, and fairness. With the success in the design of the new languages in Section 3, we hope to design composite mechanisms that satisfy many desired properties.

6 Timeline and International Collaborations

We propose to conduct the proposed work according to the timeline in Figure 4.

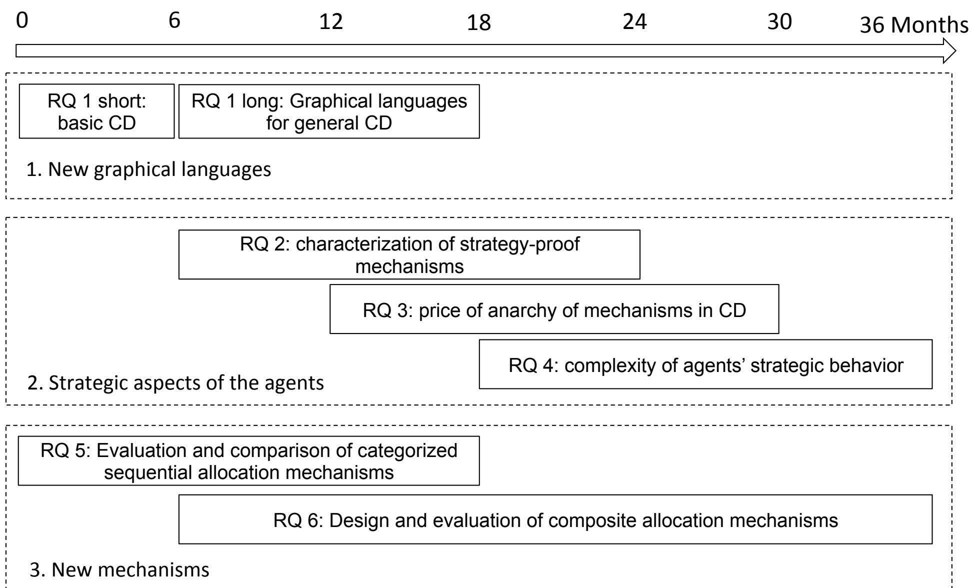

Figure 4: Timeline of the proposed research. RQ="Research Question". CD="categorized domains".

Direction 1: New graphical languages (Section 3). We want to give this direction the highest priority since many other proposed research directions depend on it. In the first 6 months we plan to work intensively on the short-term plan to extend existing graphical languages to basic categorized domains and evaluate them. In M6–M18 we plan to design and evalute new graphical languages for general categorized domains.

Direction 2: Strategic aspects of the agents (Section 4). We plan to build a pipeline to study the short-term and long-term plans for RQ 2 (M6–M24), RQ 3 (M12–M30), and RQ 4 (M18–M36). This order is because, as explained in Section 4, RQ 3 only makes sense for non-strategy-proof mechanisms, and RQ 4 only makes sense for mechanisms that have high price of anarchy. Within the 18 months in each of RQ 2, RQ 3, and RQ 4, the first 6 months are devoted to the short-term plan to work on basic categorized domains and the rest 12 months are devoted to the long-term plan for general categorized domains.

Direction 3: New mechanisms (Section 5). We plan to immediately start to work on RQ 5 (M1–M18), where the time for short-term plan and long-term plan is divided half-half. RQ 6 is the most involved research question, for which we plan to spend 30 months from M6 to M36, and the time for short-term plan and long-term plan is divided half-half.

International collaborations. In the past few years the PI has collaborated with Jérôme Lang (Paris

Dauphine University) on a wide range of topics in combinatorial voting [115, 116, 110, 114, 35, 111, 112, 36, 65, 67], including a book chapter on this topic [67]. The PI plans to continue this collaboration especially on RQ 1, RQ 5 and RQ 6.

The PI has also collaborated with Toby Walsh (UNSW and NICTA) on a wide range of related topics, including using computational complexity to protect elections [38, 73, 103, 74] and strategic behavior of sequential allocation protocols [59]. The PI plans to continue this collaboration especially on RQ 4 and the analysis of welfare properties of the new mechanisms proposed in RQ 6.

7 Educational Activities

To promote more interdisciplinary collaborations between CS and Economics in preference handling, the PI is currently co-organizing the 8th multidisciplinary workshop on preference handling (M-PREF-14). The PI plans to continue playing an important role in this area, including submitting a proposal for M-PREF-15. The PI also plans to help with the organization of computational social choice (COMSOC) workshop in 2016, which is one of the most important venues in the computational social choice community.

The primary goal of writing this grant is to support one Ph.D. student to work on the interdisciplinary area of allocation problems. The PI will keep a high priority on inspiring and encouraging female students to pursue research careers in academia or industry. In the last year, the PI developed and taught a multidisciplinary course called "Computational Social Choice", which provides lectures, seminars, and hands-on research experiences, where allocation problems are an important topic. The course attracted both graduate and senior undergraduate students from various departments including CS, Economics, ISE, and Math, and the PI was very delighted and encouraged by a high average overall score of 4.75/5.0. The running example stems from the real-world situation in this course. The PI plans to continue teaching it in the next 3 years, and develop an undergrad-level introductory course on economics and computation, which will be a project-oriented course that not only teaches students principles in economics and computation, but also helps them develop and deploy practical preference handling systems at RPI.

The PI is currently supervising Zach Jablons, a very talented undergraduate student who did very well in the PI's course last year and is now working on the continuation of his coursework project. The PI has a plan to actively attract undergraduate students, especially women and minority, to participate in research in the interdisciplinary area of economics and computation, which is jointly supported by RPI's undergraduate research program.

The PI has given tutorials on computational social choice at three conferences: ACM EC-12, IJCAI-13, and WINE-13 conferences, and a winter school in Singapore in 2013. The PI has also been invited, and will give talks soon at two active and well-organized local student groups: Rensselaer Center for Open Source Software and a newly launched Computing Club. These outreach activities provide opportunities to inspire and encourage undergrad students and new researchers to the field.

Prior NSF Support. From 2011 to 2013 the PI was supported by NSF under Grant #1136996 to the Computing Research Association for the CIFellows Project, for conducting postdoctoral research at Harvard University. While at Harvard, the PI and colleagues have made significant progress in computational social choice and machine learning. The work was published at top-tier conferences in AI and machine learning, including AAAI, AAMAS, ACM EC, CP, ICML, KR, NIPS, UAI. Among other things, the PI took an important step forward towards better fraud detections in elections by studying the computation of the margin of victory [104]; we obtained new results in using computational complexity to protect elections [85, 103]; we provided approximation and constraint-satisfaction viewpoints of combinatorial voting [39, 65]; and we designed several new computational methods to tackle statistical inferences in a well-known statistical model in economics and cognitive science that was not previously known to be computationally tractable [5, 6, 4, 7].

References

[1] Atila Abdulkadiroğlu and Tayfun Sönmez. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica*, 66(3):689–702, 1998.

[2] Atila Abdulkadiroğlu and Tayfun Sönmez. House Allocation with Existing Tenants. *Journal of Economic Theory*, 88:233–260, 1999.

[3] Arash Asadpour and Amin Saberi. An Approximation Algorithm for Max-Min Fair Allocation of Indivisible Goods. *SIAM Journal on Computing*, 39(7):2970–2989, 2010.

[4] Hossein Azari Soufiani, William Chen, David C. Parkes, and Lirong Xia. Generalized method-of-moments for rank aggregation. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, Lake Tahoe, NV, USA, 2013.

[5] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Random utility theory for social choice. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, pages 126–134, Lake Tahoe, NV, USA, 2012.

[6] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Preference Elicitation For General Random Utility Models. In *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, Bellevue, Washington, USA, 2013.

[7] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Computing Parametric Ranking Models via Rank-Breaking. In *Proceedings of the 31st International Conference on Machine Learning*, Beijing. China, 2014.

[8] Salvador Barbera, Faruk Gul, and Ennio Stacchetti. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61(2):262–289, 1993.

[9] Salvador Barbera, Jordi Masso, and Alejandro Neme. Voting under constraints. *Journal of Economic Theory*, 76(2):298–321, 1997.

[10] Salvador Barbera, Hugo Sonnenschein, and Lin Zhou. Voting by committees. *Econometrica*, 59(3):595–609, 1991.

[11] John Bartholdi, III, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[12] Anna Bogomolnaia and Hervé Moulin. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

[13] Richard Booth, Yann Chevaleyre, Jérôme Lang, Jérôme Mengin, and Chattrakul Sombattheera. Learning conditionally lexicographic preference relations. In *Proceeding of the 2010 conference on ECAI 2010: 19th European Conference on Artificial Intelligence*, pages 269–274, Amsterdam, The Netherlands, The Netherlands, 2010.

[14] Kim C. Border and J. S. Jordan. Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1):153–170, January 1983.

[15] Craig Boutilier, Fahiem Bacchus, and Ronen I. Brafman. UCP-Networks: A Directed Graphical Representation of Conditional Utilities. *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pages 56–64, 2001.

[16] Craig Boutilier, Ronen Brafman, Carmel Domshlak, Holger Hoos, and David Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[17] Craig Boutilier and Holger H. Hoos. Bidding Languages for Combinatorial Auctions. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 1211–1217, 2001.

[18] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In *Proceedings of the 21st international jont conference on Artifical intelligence*, IJCAI'09, pages 67–72, Pasadena, California, USA, 2009.

[19] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-2010)*, 2010.

[20] Sylvain Bouveret and Jérôme Lang. Efficiency and Envy-freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity. *Journal of Artificial Intelligence Research*, 32:525–564, 2008.

[21] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 73–78, Barcelona, Catalonia, Spain, 2011.

[22] Sylvain Bouveret and Michel Lemaître. Computing leximin-optimal solutions in constraint networks. *Artificial Intelligence*, 173(343—364), 2009.

[23] Ronen I. Brafman, Carmel Domshlak, and Solomon E. Shimony. On Graphical Modeling of Preference and Importance. *Journal of Artificial Intelligence Research*, 25(1):389–424, 2006.

[24] Steven J. Brams, D Marc Kilgour, and William S. Zwicker. Voting on referenda: the separability problem and possible solutions. *Electoral Studies*, 16(3):359–377, 1997.

[25] Steven J. Brams, D. Marc Kilgour, and William S. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

[26] Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

[27] Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. In G. Weiss, editor, *Multiagent Systems*. MIT Press, 2013.

[28] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 138–144, 2013.

[29] Eric Budish and Estelle Cantillon. The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. *American Economic Review*, 102(5):2237–71, 2012.

[30] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing Random Allocation Mechanisms: Theory and Applications. *American Economic Review*, 103(2):585–623, 2013.

[31] Ioannis Caragiannis and Ariel D. Procaccia. Voting Almost Maximizes Social Welfare Despite Limited Communication. *Artificial Intelligence*, 175(9–10):1655–1671, 2011.

[32] Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[33] Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Multiagent resource allocation in k-additive domains: Preference representation and complexity. *Annals of Operations Research*, 163(1):49–62, 2008.

[34] Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[35] Vincent Conitzer, Jérôme Lang, and Lirong Xia. How hard is it to control sequential elections via the agenda? In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 103–108, Pasadena, CA, USA, 2009.

[36] Vincent Conitzer, Jérôme Lang, and Lirong Xia. Hypercubewise preference aggregation in multi-issue domains. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, Barcelona, Catalonia, Spain, 2011.

[37] Vincent Conitzer, Tuomas Sandholm, and Paolo Santi. Combinatorial auctions with k-wise dependent valuations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 248–254, Pittsburgh, PA, USA, 2005. Draft in Oct., 2003.

[38] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 638–643, San Francisco, CA, USA, 2011.

[39] Vincent Conitzer and Lirong Xia. Paradoxes of Multiple Elections: An Approximation Approach. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 179–187, Rome, Italy, 2012.

[40] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

[41] Jessica Davies, George Katsirelos, Nina Narodytska, Toby Walsh, and Lirong Xia. Complexity of and Algorithms for the Manipulation of Borda, Nanson and Baldwin's Voting Rules. *Artificial Intelligence, to appear*, 2012.

[42] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 347–356, Cambridge, MA, USA, 2010.

[43] Dmitri A. Dolgov and Edmund H. Durfee. Resource Allocation Among Agents with MDP-Induced Preferences. *Journal of Artificial Intelligence Research*, 27:505–549, 2006.

[44] Lars Ehlers and Bettina Klaus. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Social Choice Welfare*, 21:265—280, 2003.

[45] Özgün Ekici. Reclaim-proof allocation of indivisible objects. *Games and Economic Behavior*, 81:1–10, 2013.

[46] Haluk Ergin. Consistency in House Allocation Problems. *Journal of Mathematical Economics*, 34:77–97, 2000.

[47] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53:74–82, 2010.

[48] Piotr Faliszewski and Ariel D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[49] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, October 1991.

[50] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[51] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.

[52] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[53] John William Hatfield. Strategy-proof, efficient, and nonbossy quota allocations. *Social Choice and Welfare*, 33(3):505–515, 2009.

[54] Aanund Hylland and Richard Zeckhauser. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy*, 87(2):293–314, 1979.

[55] Matthew O. Jackson. *Mechanism Theory*. Encyclopedia of Life Support Systems. EOLSS Publishers, 2000.

[56] Paula Jaramilloa and Vikram Manjunath. The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory*, 147(5):1913–1946, 2012.

[57] Biung-Ghi Ju. A characterization of strategy-proof voting rules for separable weak orderings. *Social Choice and Welfare*, 21(3):469–499, 2003.

[58] Thomas Kalinowski, Nina Narodytska, and Toby Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of IJCAI-2013*, 2013.

[59] Thomas Kalinowski, Nina Narodytska, Toby Walsh, and Lirong Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of AAAI-2013*, 2013.

[60] Bettina Klaus and Eiichi Miyagawa. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, 30(3):421–435, 2002.

[61] Elias Koutsoupias and Christos Papadimitriou. Worst-case Equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*, pages 404–413, Trier, Germany, 1999.

[62] Céline Lafage and Jérôme Lang. Logical representation of preference for group decision making. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning*, pages 457–468, 2000.

[63] Jérôme Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1):37–71, 2004.

[64] Jérôme Lang. Vote and aggregation in combinatorial domains with structured preferences. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1366–1371, Hyderabad, India, 2007.

[65] Jérôme Lang, Jérôme Mengin, and Lirong Xia. Aggregating conditionally lexicographic preferences on multi-issue domains. In *Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming (CP)*, pages 973–987, 2012.

[66] Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

[67] Jérôme Lang and Lirong Xia. Voting in Combinatorial Domains. Handbook of Computational Social Choice, chapter 9. https://sites.google.com/site/handbookcomsoc/chapters, 2014.

[68] Michel Le Breton and Arunava Sen. Separable preferences, strategyproofness, and decomposability. *Econometrica*, 67(3):605–628, 1999.

[69] Richard Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 125–131, New York, NY, USA, 2004.

[70] Jinpeng Ma. Strategy-proofness and the Strict Core in a Market with Indivisibilities. *International Journal of Game Theory*, 23:75–83, 1994.

[71] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

[72] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2003.

[73] Nina Narodytska, Toby Walsh, and Lirong Xia. Manipulation of Nanson's and Baldwin's rule. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, San Francisco, CA, USA, 2011.

[74] Nina Narodytska, Toby Walsh, and Lirong Xia. Combining voting rules together. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, pages 612–617, 2012.

[75] John Nash. Non-cooperative games. *Annals of Mathematics*, 54:289–295, 1951.

[76] Klaus Nehring and Clemens Puppe. Efficient and strategy-proof voting rules: A characterization. *Games and Economic Behavior*, 59(1):132–153, 2007.

[77] Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence*, 68(1–3):65–90, 2013.

[78] Noam Nisan. Bidding languages for combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 9, pages 215–231. MIT Press, 2006.

[79] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[80] Hannu Nurmi. *Comparing voting systems*. Springer, 1987.

[81] Christos H. Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, 1998.

[82] Szilvia Pápai. Strategyproof assignment by hierarchical exchange. *Econometrica*, 68(6):1403–1433, 2000.

[83] Szilvia Pápai. Strategyproof multiple assignment using quotas. *Review of Economic Design*, 5:91–105, 2000.

[84] Szilvia Pápai. Strategyproof and nonbossy multiple assignments. *Journal of Public Economic Theory*, 3(3):257–71, 2001.

[85] David C. Parkes and Lirong Xia. A Complexity-of-Strategic-Behavior Comparison between Schulze's Rule and Ranked Pairs. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1429–1435, Toronto, Canada, 2012.

[86] Parag A. Pathak and Tayfun Sönmez. Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism. *American Economic Review*, 98(4):1636–52, 2008.

[87] Ariel D. Procaccia. Cake Cutting: Not Just Child's Play. *Communications of the ACM*, 56(7):78–87, 2013.

[88] Ariel D. Procaccia and Jeffrey S. Rosenschein. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents*, volume 4149 of *LNAI*, pages 317–331. 2006.

[89] Sara Ramezani and Ulle Endriss. Nash social welfare in multiagent resource allocation. In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, volume 59 of *Lecture Notes in Business Information Processing*, pages 117–131. Springer Berlin Heidelberg, 2010.

[90] Magnus Roos and Jörg Rothe. Complexity of social welfare optimization in multiagent resource allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 641–648, 2010.

[91] Alvin E. Roth. Incentive Compatibility in a Market with Indivisibilities. *Economics Letters*, 9:127–132, 1982.

[92] Alvin E. Roth and Andrew Postlewaite. Weak versus Strong Domination in a Market with Indivisible Goods. *Journal of Mathematical Economics*, 4:131–137, 1977.

[93] Jörg Rothe and Lena Schend. Challenges to Complexity Shields That Are Supposed to Protect Elections Against Manipulation and Control: A Survey. *Annals of Mathematics and Artificial Intelligence*, 68(1–3):161–193, 2013.

[94] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, January 2002. Early version in IJCAI-99.

[95] Mark Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[96] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.

[97] Tayfun Sönmez and M. Utku Ünver. House Allocation with Existing Tenants: A Characterization. *Games and Economic Behavior*, 69(2):425–445, 2010.

[98] Tayfun Sönmez and M. Utku Ünver. Matching, Allocation, and Exchange of Discrete Resources. In Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, editors, *Handbook of Social Economics*, chapter 17, pages 781–852. North-Holland, 2011.

[99] Lars-Gunnar Svensson. Queue allocation of indivisible goods. *Social Choice and Welfare*, 11:323–330, 1994.

[100] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.

[101] Vasilis Syrgkanis and Eva Tardos. *Composable and Efficient Mechanisms*. arXiv:1211.1325, 2012.

[102] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[103] Toby Walsh and Lirong Xia. Lot-based voting rules. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 603–610, Valencia, Spain, 2012.

[104] Lirong Xia. Computing the margin of victory for various voting rules. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 982–999, Valencia, Spain, 2012.

[105] Lirong Xia. How many vote operations are needed to manipulate a voting system? In *Proc. COMSOC-12*, 2012.

[106] Lirong Xia. Assigning indivisible and categorized items. To be presented at ISAIM-14 special session on Computational Social Choice, 2014.

[107] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 109–118, Chicago, IL, USA, 2008.

[108] Lirong Xia and Vincent Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 921–926, Atlanta, GA, USA, 2010.

[109] Lirong Xia and Vincent Conitzer. Strategy-proof voting rules over multi-issue domains with restricted preferences. In *Proceedings of the Sixth Workshop on Internet and Network Economics (WINE)*, pages 402–414, Stanford, CA, USA, 2010.

[110] Lirong Xia, Vincent Conitzer, and Jérôme Lang. Voting on multiattribute domains with cyclic preferential dependencies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 202–207, Chicago, IL, USA, 2008.

[111] Lirong Xia, Vincent Conitzer, and Jérôme Lang. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 399–406, Toronto, Canada, 2010.

[112] Lirong Xia, Vincent Conitzer, and Jérôme Lang. Strategic sequential voting in multi-issue domains and multiple-election paradoxes. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 179–188, San Jose, CA, USA, 2011.

[113] Lirong Xia, Vincent Conitzer, and Ariel D. Procaccia. A scheduling approach to coalitional manipulation. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 275–284, Cambridge, MA, USA, 2010.

[114] Lirong Xia and Jérôme Lang. A dichotomy theorem on the existence of efficient or neutral sequential voting correspondences. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 342–347, Pasadena, CA, USA, 2009.

[115] Lirong Xia, Jérôme Lang, and Mingsheng Ying. Sequential voting rules and multiple elections paradoxes. In *Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 279–288, Brussels, Belgium, 2007.

[116] Lirong Xia, Jérôme Lang, and Mingsheng Ying. Strongly decomposable voting rules on multiattribute domains. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 776–781, Vancouver, BC, Canada, 2007.

[117] Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 348–353, Pasadena, CA, USA, 2009.

[118] H.Peyton Young. *Equity: In theory and practice*. Princeton University Press, 1994.

Facilities, Equipment, and Other Resources

Since most of the proposed work is theoretical, we do not expect to use computational clusters available at RPI.

RPI provides facilities that support regular everyday work, including high-speed wifi, library space, shared PCs. RPI library has access to articles in many major journals.

RPI computer science department offers major software packages including various scientific computing softwares. The PI has his own office space. Graduate students share offices.

Data Management Plan

Most of the proposed work is theoretical. We will generate synthetic datasets for our simulations, and will make them available.