

Guest editorial: special issue on computational social choice

Edith Elkind Jérôme Lang

Computational social choice is a new interdisciplinary research area at the intersection of social choice theory and computer science that studies computational aspects of social choice mechanisms as well as applications of social choice ideas in the design and analysis of computing systems. Research on computational social choice now encompasses a variety of topics, including the algorithmic aspects of voting rules, computational barriers to strategic behaviour, collective decision-making in multi-agent systems, preference elicitation and communication issues in voting, fair division, computational aspects of weighted voting games, collective decision-making in combinatorial domains, logic-based formalisms for social choice problems, belief and judgement aggregation, and social software. The aim of this special issue is to build on the spectacular success of two workshops on computational social choice (COMSOC'06, Amsterdam and COMSOC'08, Liverpool), and provide a forum for researchers in the area to present their ideas on a variety of topics.

This special issue is composed of nine research articles and a short note that span a fairly wide range of issues, including voting, mechanism design, judgment aggregation, and fairness, and a similarly wide range of techniques, such as logic, experiments, and computational complexity.

The paper by Ågotnes, van der Hoek and Wooldridge defines a new modal logic for representing and reasoning about preference and judgment aggregation. The authors show that their logic can be used to express a variety of concepts and results from the social choice domain, including majority voting, the discursive paradox, Arrow's theorem and Condorcet's paradox.

Baharad, Goldberger, Koppel and Nitzan propose a maximum likelihood estimation-based approach for aggregating voters' opinions over several issues. They argue that their method is vastly superior to simple majority voting, and use both analytical and empirical evaluation to support their argument.

Balan, Richards and Luke deal with the problem of selecting a sequence of actions when the actions have multiple beneficiaries. Their goal is to select the actions according to a lexicographic fairness criterion. They explore the complexity of this problem for both finite and infinite sequences of actions and present both algorithms and intractability results for this problem.

The paper by Caminada and Pigozzi applies argumentation theory to judgment aggregation, and proposes aggregation operators satisfying not only the usual property of collective rationality, but another highly desirable property,

namely, that the social outcome should be 'compatible' with the judgment of every individual, so that any individual is able to defend the collective decision.

De Jong and Tuyls propose a new model of fairness in multi-agent settings, which is inspired by human behavior. Their work is inspired by the observation that humans are very good at handling social dilemmas, i.e., settings where socially optimal action is not individually optimal. They provide computational models of human decision-making mechanisms and explore their performance in experimental settings.

Farfel and Conitzer develop techniques for preference aggregation and elicitation in the setting where agents have preferences over intervals of values. They generalize the theory of single-peaked preferences from points to ranges, and provide a truthful mechanism for preference elicitation under a condition on preferences. They demonstrate that their algorithm needs fewer queries than the naïve algorithm, and discuss ways to speed it up when the space of preferences has a special structure.

Kalech, Kraus, Kaminka and Goldman propose and evaluate iterative variants of many well-known voting rules. Iterative voting allows the agents to submit partial information about their preferences at each step, thus lowering their cognitive and communication load. Kalech et al. experimentally demonstrate that iterative voting leads to significant savings in communication.

Pini, Rossi, Venable and Walsh consider the issue of manipulation in stable marriage procedures. They show that, while the well-known Gale–Shapley algorithm is easy to manipulate, there exists a stable matching algorithm that is computationally hard to manipulate. They also show how to make stable marriage procedures gender-neutral, and demonstrate that gender neutrality and resistance to manipulation can be achieved simultaneously.

The paper by Auletta, Penna, Persiano, and Ventre studies computational aspects of mechanism design with partial verification, where the set of types that the agent can report depends on his true type. In this setting, the revelation principle no longer holds, so some choice rules may only be implementable in a non-truthful manner. The paper shows that given a social choice rule, it is computationally hard to decide whether it has a non-truthful implementation, either with or without payments. Moreover, it can be hard to decide if a given rule can be implemented without payments even if it is straightforwardly implementable when payments are allowed.

After the paper by Auletta et. al. appeared online, Yu observed that one of the results presented there (namely, a characterization of the domains with partial verification for which the revelation principle holds, in settings with money) is incorrect. She was then able to derive a correct characterization. Her proof, together with the response by Auletta et al., also appears in this special issue.

We would like to thank the Journal of Autonomous Agents and Multiagent Systems editors-in-chief for giving us the opportunity to put together this special issue. We are also grateful to the many anonymous referees for the hard work they put into helping us select the best papers and providing the authors with detailed and thoughtful comments.