Available online at www.sciencedirect.com

ScienceDirect

To appear in Procedia Computer Science (2011)

Computational Social Choice: Prospects and Challenges[1]

Ulle Endriss

Institute for Logic, Language and Computation (ILLC), University of Amsterdam
Postbus 94242, 1090 GE Amsterdam, The Netherlands

Abstract

How should we aggregate the individual views of the members of a group so as to arrive at an adequate representation of the collective view of that group? This is a fundamental question of deep philosophical, economic, and political significance that, around the middle of 20th century, has given rise to the field of Social Choice Theory. More recently, a research trend known as Computational Social Choice has emerged, which studies this question from the perspective of Computer Science. This "computational turn" is fuelled both by the fact that questions of social choice have turned out to be central to a range of application areas, notably in the domain of Information and Communication Technologies, and by the insight that many concepts and techniques originating in Computer Science can be used to solve (or provide a new angle on) problems in Social Choice Theory. In this paper, I give a brief introduction to Computational Social Choice and discuss some of the prospects and challenges for this fast growing area of research.

Keywords: Computational Social Choice, Economics and Computation, Artificial Intelligence

---

1. Introduction

Computational Social Choice is a novel area of research that recently formed at the interface of, on the one hand, a research tradition in Economic Theory known as *Social Choice Theory* and, on the other, a broad range of concepts and techniques originating in Computer Science.

Social Choice Theory [1] addresses a fundamental question: How should we aggregate the preferences of a group of individuals so as to arrive at an adequate collective preference, which can serve as the basis for making acceptable group decisions? Not only is this question relevant to the quality of everyday interactions between people, but it strikes at the very heart of our understanding of democracy. Indeed, both the importance of the question and the significance of the research tradition that has sought to address it are widely recognised, as witnessed, for instance, by the Nobel prizes awarded to social choice theorists such as Kenneth Arrow and Amartya Sen.

Arrow's doctoral thesis, published in 1951 under the title "*Social Choice and Individual Values*", is often regarded as the starting point of modern Social Choice Theory. By providing a rigorous mathematical formulation of the problem of fairly aggregating the preferences of a group of individuals into a collective preference order, Arrow made this problem amenable to systematic scientific analysis. Sixty years on, in the light of extraordinary technological advances, we now have to recast the fundamental question of social choice. It is not anymore just a matter for groups of people who have to choose between a small number of alternatives. Instead, a wide range of modern applications,

---

[1] This paper is based on an expository talk given during a session on Computational Social Choice at the *Second European Future Technologies Conference and Exhibition* (FET'11) in Budapest in May 2011, which also included talks by Péter Biró, Ioannis Caragiannis, Britta Dorn, Jérôme Lang, and Francesca Rossi on some of the challenges for the field. Their contribution, which helped shape this paper, is gratefully acknowledged.

in the domain of Information and Communication Technologies (ICT) and beyond, require agents, both human and artificial, to make collective decisions, often choosing from a huge number of alternatives and under conditions where some of the classical assumptions made in Social Choice Theory, e.g., regarding the nature of preferences or the definition of what it means to be "fair", may fail to apply. Examples for such applications include the integration of search results produced by multiple search engines, the assignment of junior doctors to hospitals in view of both their preferences and their qualifications, the amalgamation of user preferences to issue recommendations for other users in online recommender systems, the support of group decision making technologies for social networks, and a multitude of problems in *e*-commerce and *e*-governance. Indeed, for any application that requires us to aggregate information supplied by different agents it may be beneficial to refer back to the principles of Social Choice Theory to understand the fundamental dynamics of that application domain. This is also why social choice plays an important role in Multiagent Systems, the subfield of Artificial Intelligence (AI) concerned with the study of systems in which several intelligent agents interact.

But it is not just the case that social choice is relevant to ICT and Computer Science. The synergy also goes the other way round: As a surge of recent research demonstrates, many concepts and techniques originating in Computer Science have turned out to be useful in addressing problems in Social Choice Theory and have provided the impetus for a new and innovative perspective on the old problem of collective decision making. Contributions of this kind have come, amongst others, from algorithm design, computational complexity theory, knowledge representation, uncertainty in AI, constraint reasoning, machine learning, symbolic logic, and automated theorem proving. That is, this trend is complementary to the by now common use of information technology in the social sciences, namely social simulation and the collection and visualisation of large amounts of data. Together, the fact that Social Choice Theory is relevant to a range of applications with a computational component and the insight that Computer Science provides a helpful new perspective on social choice have given rise to the discipline of *Computational Social Choice* [2].

In this paper I give a brief sketch of Computational Social Choice as a field of scientific enquiry. I start by revisiting the fundamental question of how to aggregate preferences, as studied in classical Social Choice Theory (Section 2); I then describe some of the modern applications that, arguably, rely on principles of social choice (Section 3); and I discuss examples of recent research demonstrating how the computational approach can contribute to our understanding of the problem of collective decision making (Section 4). I conclude with a discussion of some of the prospects and challenges for the field (Section 5). For the sake of brevity, references are restricted to a small number of expository works; for extensive references to original research please consult the bibliographies of the cited works.

2. Classical Social Choice Theory: From a Curious Observation to a Fully Fledged Mathematical Theory

Social Choice Theory is the formal study of mechanisms for collective decision making. Such mechanisms include voting rules for conduction elections, protocols for the fair division of a set of goods, procedures for aggregating the opinions of a court of several judges into a single judgment, and algorithms for matching the elements of two sets (e.g., pairing up men and women). As an example, let us consider the domain of preference aggregation. Suppose there are five experts who are asked to rank three alternative proposals, imaginatively referred to as A, B, and C. We model each expert's preferences as a linear order over the set $\{A, B, C\}$. They might have these preferences:

Expert 1: $A \succ B \succ C$
Expert 2: $B \succ C \succ A$
Expert 3: $C \succ A \succ B$
Expert 4: $C \succ A \succ B$
Expert 5: $B \succ C \succ A$

What would be a good mechanism to aggregate this into a single collective preference order? A promising approach is to use the *majority rule:* rank X above Y if and only if a majority of the individual experts do. If we do this, we must rank A above B (as experts 1, 3, and 4 do), B above C, and C above A. Thus, the collective preference fails to be a linear order but instead we obtain a cycle! This surprising outcome is an instance of the so-called *Condorcet Paradox*, deriving its name from the 18th century mathematician and political scientist who first discussed it at length.

Are there better preference aggregation mechanisms that do not suffer from this problem? The seminal contribution of Arrow, in his 1951 doctoral thesis, has been to cast this question in a mathematically rigorous way by formulating desirable properties of such mechanisms as "axioms". Specifically, he proposed two such axioms:

- *Unanimity:* If every individual agrees that $X \succ Y$, then the collective preference order should also rank $X \succ Y$.
- *Independence of Irrelevant Alternatives (IIA):* If the collective preference order ranks $X \succ Y$ and one of the individuals changes her ranking of Z (different from both X and Y), then the collective should still rank $X \succ Y$.

Arrow then proved a truly astonishing theorem: If there are at least three alternatives, then *the only* kind of mechanism that will respect both unanimity and IIA and that will return a collective preference that is a linear order for any combination of individual preferences is a *dictatorship*, i.e., a function that simply copies the preferences of a fixed individual and returns it as the collective preference! That is, we now have a mathematically precise way of saying that the curious problem observed by Condorcet cannot be circumvented. This is not only a seminal, but also a typical result in Social Choice Theory: It uses a mathematical approach, and specifically the *axiomatic method*, to arrive at deep insights into the dynamics of collective decision making.

For a broad overview of Social Choice Theory refer to the *Handbook of Social Choice and Welfare* [1]. The most recent developments in the field typically get published in journals in Economic Theory, particularly *Social Choice and Welfare*, and are often presented at the biannual congress of the *Society for Social Choice and Welfare*.

3. Modern Applications of Social Choice Theory in ICT and Beyond

The original domain of application for Social Choice Theory has been that of political elections: We are given a group of voters who each express preferences over a set of candidates and we are asked to determine the "best" candidate. But the fundamental concepts of social choice are considerably more pervasive than that. They are relevant to a wide range of current and future applications, many of them in the ICT domain. I briefly outline some of them here:

- *Internet search:* Social choice plays a role in more than one way here. First, if we want to aggregate the output of several search engines that rank webpages by relevance, then we need to perform an operation similar to preference aggregation. Second, if we think of a link from one page to another as a vote of endorsement, then the problem of determining the importance of a page becomes similar to an exercise in voting. Indeed, Google's PageRank algorithm uses this idea and has been analysed in terms of the axiomatic method.
- *Decision support systems:* Any kind of system designed to support groups of users, such as the members of a local community or the shareholders of a company, to make collective decisions, such as whether to fund a particular project or which job candidate to hire, is trying to solve a problem of social choice.
- *Social networks, e-commerce, e-governance:* In all of these areas, groups of users need to make decisions together. Social Choice Theory can contribute the foundations for how such processes should be set up so as to avoid counter-intuitive outcomes and to guarantee that appropriate fairness conditions are being respected.

Importantly, none of these problems are perfect instances of the classical problems studied in Social Choice Theory. That is, we need to adapt the classical models and methods to these new challenges.

An area of Social Choice Theory that is particularly advanced in terms of turning theory into applications is *matching theory*. Suppose we are given two sets of, say, "men" and "women", and each man has a preferences over the women and each woman has a preference over the men. What would be a good way of pairing up men and women? For instance, we may be interested in a *stable* matching, where no pair of a man and a woman have an incentive to leave their assigned partners and run off with each other. The mathematical theory developed for the analysis of this model has led to the optimisation of a number of real-world matching problems. Examples include matching junior doctors looking for internships to hospitals, matching students to universities, and matching patients requiring a kidney transplant to compatible donors (who themselves are linked to another patient they are not compatible with).

4. The Computational Turn: From Social Choice to Computational Social Choice

Taking the perspective of Computer Science when analysing problems of social choice has proven tremendously successful. Here is a small selection of the research directions that are currently being pursued:

- *Complexity Theory:* Imagine you have information suggesting that you favourite candidate has no chance of winning the election. You might then want to misrepresent your true preferences and vote for, say, your second most preferred candidate instead. A classical theorem in Social Choice Theory shows that *any* conceivable voting rule (that is not a dictatorship) suffers from this problem and will sometimes incentivise voters to lie. But maybe we can find a voting rule that will make it very difficult (maybe NP-hard) to cheat in this way? Using complexity as a means of protecting elections against manipulation is often cited as the archetypal example for work in Computational Social Choice. Faliszewski et al. [3] give a high-level review of this research direction.

- *Algorithmics:* Complex social choice mechanisms can be hard to implement and thus require sophisticated techniques from algorithm design. This might be the most obvious use of Computer Science in the field.
- *Knowledge Representation:* In many scenarios the alternatives to be decided upon will have a combinatorial structure. For instance, we might have to choose for each of a set of binary variables whether to assign them value 0 or 1, i.e., the number of alternatives is exponential in the number of those variables. This requires advanced techniques of the kind studied in Knowledge Representation to allow for a compact representation of preferences. Chevaleyre et al. [4] give a high-level introduction to social choice in combinatorial domains.
- *Logic:* Just as computer scientists have long been using logic to formally specify the behaviour of computer systems, so as to allow for the automatic verification of certain desirable properties of such systems, logic can also be used to formally specify and analyse social choice mechanisms.

Chevaleyre et al. [2] propose a classification of research in Computational Social Choice in terms of, first, the type of social choice problem addressed and, second, the type of formal or computational technique applied. Work in the field gets published in a broad range of venues, notably the archival proceedings of the major international conferences in AI. The best source for finding out about the latest developments are the proceedings of the *International Workshop on Computational Social Choice*, held biannually since 2006.

5. Prospects and Challenges

Despite its early successes, Computational Social Choice is only at the beginning of providing us with a comprehensive understanding of the formal, algorithmic, and computational properties of systems for collective decision making that would allow us to take full advantage of the insight that principles of social choice play a role in a wide range of modern applications. Here are four broad areas that all provide formidable challenges for future research:

- *Multivariate algorithmics:* We require a better understanding of the algorithmic nature of problems such as computing the range of *possible* outcomes when only partial preferences have been elicited. The tools of multivariate algorithmics promise to deliver a much more fine-grained picture than is available today.
- *Nonstandard models of preference:* We need to relax the classical assumptions made in Economic Theory, which are not always appropriate for the scenarios studied in Computational Social Choice. For instance, we might not just want to aggregate preferences that provide a relative ranking of any two alternatives, but a decision maker may also be uncertain about some of these rankings. More generally, there is room for extending the methods of Social Choice Theory to a range of other types of information aggregation.
- *Automated reasoning:* We need to adapt methods studied in AI and elsewhere to be able to automatically reason about problems of social choice. At the far end of the spectrum this includes the automatic generation of appropriate mechanisms for a given set of desiderata.
- *Novel applications:* We need to develop new application domains for Computational Social Choice, including for instance the *Semantic Web* (cf. the aggregation of information present in ontologies stemming from different sources) and *Grid Computing* (cf. the fair scheduling of jobs originating with different users).

Tackling these challenges will require a concerted effort across a number of different scientific disciplines. The recent work carried out within the nascent Computational Social Choice research community, with its remarkable achievements to date and its interdisciplinary orientation, provides a perfect position from which to start.

6. References

[1] K. J. Arrow, A. K. Sen, K. Suzumura (Eds.), Handbook of Social Choice and Welfare, North-Holland, 2002.

[2] Y. Chevaleyre, U. Endriss, J. Lang, N. Maudet, A short introduction to computational social choice, in: Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM-2007), Vol. 4362 of LNCS, Springer-Verlag, 2007, pp. 51–69.

[3] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, Using complexity to protect elections, Comm. ACM 53 (11) (2010) 74–82.

[4] Y. Chevaleyre, U. Endriss, J. Lang, N. Maudet, Preference handling in combinatorial domains: From AI to social choice, AI Magazine 29 (4) (2008) 37–46.