Artificial Intelligence 227 (2015) 190–213

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Optimal social choice functions: A utilitarian view ☆

Craig Boutilier [a,1], Ioannis Caragiannis [b], Simi Haber [c], Tyler Lu [a,2], Ariel D. Procaccia [d,*], Or Sheffet [e]

[a] Dept. of Computer Science, University of Toronto, Canada
[b] Computer Technology Institute "Diophantus" and Dept. of Computer Engineering and Informatics, University of Patras, Greece
[c] Dept. of Mathematics, Bar-Ilan University, Israel
[d] Computer Science Dept., Carnegie Mellon University, United States
[e] Center for Research on Computation and Society, Harvard SEAS, United States

A R T I C L E   I N F O

Article history:
Received 12 April 2014
Received in revised form 11 June 2015
Accepted 14 June 2015
Available online 25 June 2015

Keywords:
Computational social choice

A B S T R A C T

We adopt a utilitarian perspective on social choice, assuming that agents have (possibly latent) utility functions over some space of alternatives. For many reasons one might consider mechanisms, or *social choice functions*, that only have access to the ordinal rankings of alternatives by the individual agents rather than their utility functions. In this context, one possible objective for a social choice function is the maximization of (expected) social welfare relative to the information contained in these rankings. We study such *optimal* social choice functions under three different models, and underscore the important role played by *scoring functions*. In our worst-case model, no assumptions are made about the underlying distribution and we analyze the worst-case *distortion*—or degree to which the selected alternative does not maximize social welfare—of optimal (randomized) social choice functions. In our average-case model, we derive optimal functions under neutral (or impartial culture) probabilistic models. Finally, a very general learning-theoretic model allows for the computation of optimal social choice functions (i.e., ones that maximize expected social welfare) under arbitrary, sampleable distributions. In the latter case, we provide both algorithms and sample complexity results for the class of scoring functions, and further validate the approach empirically.

1. Introduction

Classic models in social choice theory assume that the preferences of a set of *agents* over a set of *alternatives* are represented as linear orders; a *social choice function*, given these preferences as input, outputs a single socially desirable alternative. A host of clever social choice functions have been designed to satisfy various *normative* criteria. Most work in *computational social choice* studies computational aspects of these models, addressing questions such as the complexity of computing social choice functions [5,17] or manipulating them (see the survey by Faliszewski and Procaccia [13]).

☆ A preliminary version of this paper appeared in the proceedings of EC'12.

* Corresponding author.
E-mail addresses: cebly@cs.toronto.edu (C. Boutilier), caragian@ceid.upatras.gr (I. Caragiannis), simi@math.biu.ac.il (S. Haber), tl@cs.toronto.edu (T. Lu), arielpro@cs.cmu.edu (A.D. Procaccia), osheffet@seas.harvard.edu (O. Sheffet).

[1] Currently on leave at Google, Inc., Mountain View, CA.
[2] Currently at Google, Inc., Mountain View, CA.

http://dx.doi.org/10.1016/j.artint.2015.06.003
0004-3702/

Under ordinal preferences, an axiomatic approach to obtaining a socially desirable outcome seems—on the face of it—necessary, absent concrete measures of the quality of an alternative. In contrast, some work in economics assumes *cardinal* preferences and takes a *utilitarian* approach. This viewpoint dates to the work of Bentham at the end of the 18th century, who argued that *"it is the greatest happiness of the greatest number that is the measure of right and wrong."* This axiom suggests that happiness can be quantified, and indeed, having coined the term *utility*, Bentham proposed that the goal of government is to maximize the sum of individual utilities—the *social welfare* (defying contemporary wisdom that the goal of government is to enrich the coffers of the ruler). The utilitarian approach is prevalent, for example, in mechanism design, and perhaps even more so in *algorithmic* mechanism design [25].

In this paper we view the social choice problem through this utilitarian lens. Our premise is that agents have (possibly implicit) utility functions, and the goal of a social choice function is to maximize the (utilitarian) social welfare[3]—i.e., (possibly weighted) sum of agent utilities—of the selected alternative. The utilitarian perspective is not appropriate for all social choice problems (a point we discuss further below). However, the methods of social choice—especially voting systems—are finding increasing application in recommender systems, web search, product design, and many more practical domains, in which the primary aim is often, as in much of mechanism design, to aggregate preferences so that utility or efficiency is maximized. Indeed, one motivation for our work is the development of group recommendation systems for a variety of domains, including low-stakes consumer applications and higher profile public policy and corporate decisions. Our work can be viewed as a step toward supporting groups of users making decisions using social choice functions that are automatically optimized for their needs. In these settings, a utilitarian perspective is often called for.

If we could directly access the utilities of agents, the socially desirable alternative could be easily identified. However, such access is often not feasible for a variety of reasons. As a result, we use agent preference orders as a *proxy* for their utility functions; and the social choice function, taking preference orders as input, should perform well with respect to the underlying utilities. From this point of view, a social choice function is *optimal* if it maximizes social welfare given the available information. Using a preference order as proxy for utility in this fashion serves several purposes. First, behavioral economists have argued that people find it difficult to construct utilities for alternatives. Second, the cognitive and communication burden of articulating precise utilities has long been recognized within decision analysis, behavioral economics, and psychology. By contrast, simply comparing and ordering alternatives is considerably easier for most people, which makes soliciting preference orders more practical than eliciting utilities. Furthermore, choice behavior among alternatives can often be interpreted as revealing ordinal (rather than cardinal) preference information, providing ready access to (sometimes incomplete) orders in many of the domains described above. Hence we content ourselves with orders as inputs.

1.1. Our results

Our study of optimal social choice functions incorporates three distinct but related models, each with its own assumptions regarding available information and therefore its own notion of optimality. One common thread is that the family of *scoring functions*—social choice functions that score alternatives based only on their position in each agent's preference order—plays a key role in optimizing social welfare.

In Section 3 we study a model where no information about agents' utility functions is available when constructing the social choice function. A *worst-case* analysis is thus called for. We believe that the study of this model is of theoretical interest, but it is certainly the least practical of our three models. Specifically, given a collection of agents' preferences—a *preference profile*—there are many *consistent* collections of utility functions—*utility profiles*—that induce this preference profile in the natural way (by ranking alternatives with higher utility closer to the top). The *distortion* of a social choice function on a preference profile is the worst-case ratio (over feasible utility profiles) of the social welfare of the best alternative to the social welfare of the alternative that is selected by the function. A *worst-case optimal* social choice function minimizes the distortion on every preference profile.

We first derive upper and lower bounds on the least distortion that one can hope for, focusing on *randomized* social choice functions. We show that there exists a preference profile where every randomized social choice function must have distortion at least $\Omega(\sqrt{m})$, where m is the number of alternatives. We complement this result with a randomized social choice function whose distortion on *every* preference profile is $\mathcal{O}(\sqrt{m}\log^* m)$. A slightly weaker upper bound is obtained via a randomized variation of a natural scoring function that we call the *harmonic scoring function* (a new canonical scoring function that may be of independent interest). Finally, we establish that the worst-case optimal social choice function (which achieves minimum distortion on every profile) is polynomial-time computable. The proof is based on linear programming, and (roughly speaking) relies on embedding the dual of a sub-problem within a carefully constructed larger LP, in order to avoid quadratic constraints.

In Section 4 we study an *average-case model*, assuming a known distribution D over utility functions. We assume that the utility function of each agent is drawn independently from D. Given reported agent preferences, one can compute the expected utility any agent has for an alternative with respect to D. An *average-case optimal* social choice function selects an alternative that maximizes expected social welfare given the reported profile. We show that when D is *neutral*, i.e., symmetric with respect to alternatives, the average-case optimal social choice function must be a scoring function. The

[3] Hereinafter, we simply write "social welfare" to refer to "utilitarian social welfare".