

Preference Elicitation and Preference Learning in Social Choice

Craig Boutilier

Department of Computer Science, University of Toronto,
Toronto, ON, Canada
cebly@cs.toronto.edu
<http://www.cs.toronto.edu/~cebly>

Social choice has been the subject of intense investigation within computer science, AI, and operations research, in part because of the ease with which preference data from user populations can now be elicited, assessed, or estimated in online settings. In many domains, the preferences of a group of individuals must be aggregated to form a single consensus recommendation, placing us squarely in the realm of social choice.

I argue that the application of social choice models and voting schemes to domains like web search, product recommendation and social networks places new emphasis, in the design of preference aggregation schemes, on issues such as: articulating decision criteria suitable to the application at hand; approximate winner determination; incremental preference elicitation; learning methods for models of population preferences; and more nuanced analysis of the potential for manipulation.

In this talk, I'll provide an overview of some of these challenges and outline some of our recent work tackling of them, including methods for: learning probabilistic models of population preferences from choice data; robust optimization (winner determination) in the presence of incomplete user preferences; and incremental vote/preference elicitation for group decision making. Each of these poses interesting modeling and optimization challenges that are best tackled using a combination of techniques from AI, operations research, and statistics.

Parts of this talk describe joint work with Tyler Lu, Department of Computer Science, University of Toronto.