A Brief Introductory Tutorial on Computational Social Choice

Vincent Conitzer

Abstract

This is a brief description of the introductory tutorial given at COMSOC 2010.

1 Focus of the Tutorial

This tutorial gives a brief introduction to computational social choice. It is directed especially at the workshop participants who are new to this community, to give them a foothold from which to appreciate the rest of the workshop. Because the workshop program is densely packed, there is too little time to give an exhaustive overview of all the exciting current research topics in computational social choice. Hence, this tutorial focuses strictly on computational aspects of common voting rules. There are two main reasons for this. First, a large fraction of the current research in computational social choice concerns such topics. Second, it gives good insight into the *type* of problem in which the computational social choice community is interested.

2 Topics

In this tutorial, after a quick review of voting rules, we consider some representative problems from computational social choice. For each voting rule, we are confronted with the following computational problems:

1. How hard is it to *execute* the voting rule, that is, to determine the winning alternative(s)?
2. How hard is it to *manipulate* the voting rule by misreporting one's preferences?
3. How hard are other types of undesirable behavior? For example, how hard is it for the chair of the election to *control* the outcome of the election, for instance by introducing additional candidates? How hard is it for an outside party to effectively *bribe* voters?
4. If we have partial information about the votes, how hard is it to determine whether a particular alternative is still a *possible winner*?
5. How can the voters effectively *communicate* their preference information to determine the winning alternative?

It should be noted that for topics 2 and 3 above, computational hardness is *desirable*, because it may prevent the undesirable behavior. This raises interesting questions about whether the worst-case nature of computational complexity theory is appropriate here.

3 Materials and Further Reading

The slides will be made available (at least) on the presenter's website, where the slides of a longer tutorial on the same topic, given jointly with Ariel Procaccia, can also be found.

There are several overview articles of research in this area (e.g., [1, 4, 3, 2, 5]), which also provide references to more focused technical papers.

References

[1] Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas 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)*, 2007.

[2] Vincent Conitzer. Comparing multiagent systems research in combinatorial auctions and voting. *Annals of Mathematics and Artificial Intelligence*. To appear.

[3] Vincent Conitzer. Making decisions based on the preferences of multiple agents. *Communications of the ACM*, 53(3):84–94, 2010.

[4] Piotr Faliszewski, Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*, chapter 14, pages 375–406. Springer, 2009.

[5] Piotr Faliszewski and Ariel D. Procaccia. AI's war on manipulation: Are we winning? Working paper.

Vincent Conitzer
Department of Computer Science
Levine Science Research Center, Box 90129
Duke University
Durham, NC 27708, USA
Email: `conitzer@cs.duke.edu`