CSCE 475/875 Multiagent Systems

Handout 16: Application: A Social Ranking System

February 20, 2020

(Based on Shoham and Leyton-Brown 2011)

Introduction

We now turn to a specialization of the social choice problem that has a computational flavor, and in which some interesting progress can be made. Specifically, consider a setting in which the set of agents is *the same* as the set of outcomes—**agents are asked to vote to express their opinions about each other, with the goal of determining a social ranking**.

Such settings have great practical importance. For example, **search engines rank Web pages by considering hyperlinks** from one page to another to be votes about the importance of the destination pages. Similarly, online auction sites employ ***reputation*** *systems* to provide assessments of agents' trustworthiness based on ratings from past transactions.

Our setting is characterized by two assumptions.

- First, $N = O$: the set of agents is the same as the set of outcomes.
- Second, agents' preferences are such that each agent divides the other agents into a set that it likes *equally*, and a set that it dislikes *equally* (or, equivalently, has no opinion about). Formally, for each $i \in N$ the outcome set O (equivalent to N) is partitioned into two sets $O_{i,1}$ and $O_{i,2}$, with $\forall o_1 \in O_{i,1}, \forall o_2 \in O_{i,2}, o_1 \succ_i o_2$, and with $\forall o, o' \in O_{i,k}$ for $k \in \{1, 2\}, o \sim_i o'$.

We call this the ***ranking systems setting***, and call **a social welfare function in this setting a *ranking rule***.

Now, consider an example in which Alice votes only for Bob, Will votes only for Liam, and Liam votes only for Mary. Who should be ranked highest? Three of the five kids have received votes (Bob, Liam, and Mary); these three should presumably rank higher than the remaining two. But of the three, Mary is special: she is the only one whose voter (Liam) himself received a vote. Thus, intuitively, Mary should receive the highest rank.

This intuition is captured by the idea of transitivity.

> **Definition 9.5.2 (Strong transitivity)** *Consider a preference profile in which outcome o_2 receives at least as many votes as o_1, and it is possible to pair up all the voters for o_1 with voters from o_2 so that each voter for o_2 is weakly preferred by the ranking rule to the corresponding voter for o_1. Further assume that o_2 receives more votes than o_1 and/or that there is at least one pair of voters where the ranking rule strictly prefers the voter for o_2 to the voter for o_1. Then the ranking rule satisfies* strong transitivity *if it always strictly prefers o_2 to o_1.*

Further, consider an example in which Mary votes for almost all the kids, whereas Ray votes only for one. If Mary and Ray are ranked the same by the ranking rule, strong transitivity requires that their votes must count equally. However, we might feel that Ray has been more decisive, and therefore feel that his vote should be counted more strongly than Mary's.

Definition 9.5.3 (Weak transitivity) *Consider a preference profile in which outcome* o_2 *receives at least as many votes as* o_1*, and it is possible to pair up all the voters for* o_1 *with voters for* o_2 ***who have both voted for exactly the same number of outcomes*** *so that each voter for* o_2 *is weakly preferred by the ranking rule to the corresponding voter for* o_1*. Further assume that* o_2 *receives more votes than* o_1 *and/or that there is at least one pair of voters where the ranking rule strictly prefers the voter for* o_2 *to the voter for* o_1*. Then the ranking rule satisfies* weak transitivity *if it always strictly prefers* o_2 *to* o_1.

Definition 9.5.4 (RIIA, informal) *A ranking rule satisfies* ***ranked independence of irrelevant alternatives*** *(RIIA) if the relative rank between pairs of outcomes is always determined according to the same rule, and this rule depends* ***only*** *on* (1) *the number of votes each outcome received; and* (2) *the relative ranks of these voters.*

Definition 9.5.7 (Strong quasi-transitivity) *Consider a preference profile in which outcome* o_2 *receives at least as many votes as* o_1*, and it is possible to pair up all the voters for* o_1 *with voters from* o_2 *so that each voter for* o_2 *is weakly preferred by the ranking rule to the corresponding voter for* o_1*. Then the ranking rule satisfies* strong quasi-transitivity *if it weakly prefers* o_2 *to* o_1*, and strictly* strong *prefers* o_2 *to* o_1 *if either* o_1 *received no votes or each paired voter for* o_2 *is strictly preferred by the ranking rule to the corresponding voter for* o_1.

forall $i \in N$ **do** $rank(i) \leftarrow 0$
repeat
  forall $i \in N$ **do**
  if $|voters\_for(i)| > 0$ **then**
    $rank(i) \leftarrow (\frac{1}{n+1})[|voters\_for(i)| + \max_{j \in voters_{for(i)}} rank(j)$
  else
    $rank(i) \leftarrow 0$
until $rank$ *converges*

Figure 9.3: A ranking algorithm that satisfies strong quasi-transitivity and RIIA.