Trimmed Ranking Welfare Rules

Roy Fairstein
Ben-Gurion University
royfa@post.bgu.ac.il

Georgios Papasotiropoulos
Athens University of Economics & Business
gpapasotiropoulos@aueb.gr

Bin Sun
RWTH Aachen University
bin.sun1@rwth-aachen.de

Consider an election setting with n voters and m candidates, where each voter provides a ranking over all candidates. Our proposal primarily involves the use of `welfare rules`, a family of voting methods designed to incorporate voters' preferences, resulting in a ranking of the candidates that maximizes the electorates' welfare. By adopting this family, the focus shifts to capturing the nuanced welfare of candidates, providing a more refined and context-specific evaluation of their desirability. Within the family of `welfare rules`, we introduce the `top-x` rule, for any $x \leq m$, which draws inspiration from a modified version of the well-established Borda rule. While Borda assigns utility values to all but the last candidate in each voter's ranking, the `top-x` rule is restricted to the candidates ranked in the top x positions according to the voters' preferences. We suggest the term "Trimmed Ranking Welfare Rule" to signify the process of curtailing voters' rankings to their top positions.

To calculate the welfare of a candidate c in relation to a particular voter v, denoted as $sc_v(c)$, we employ a formula that considers c's position within the top x positions of v's ranking. Let $A_x(v)$ be the set of candidates that belong to the top x positions of v's ranking and say that $r_v(c)$ is a number in $[0, m-1]$ that denotes the position of candidate c in v's ranking, where $r_v(c) = 0$ if c is ranked first for voter v and $r_v(c) = m-1$ if it is ranked last. We suggest a scoring function that assigns a score of 1 to v's most favored candidate, while gradually decreasing the score for other candidates within v's top x choices. Specifically,

$$sc_v(c) = \begin{cases} \frac{m - r_v(c)}{m}, \text{if } c \in A_x(v) \\ 0, \text{otherwise.} \end{cases}$$

Example. *Let's illustrate the concept with an example involving 5 candidates. In the case of the* `top-3` *rule, each voter's top three choices would be assigned scores of 1, 0.8, and 0.6, respectively. This reflects the emphasis on the welfare of candidates within the top positions of voters' rankings and contrasts with the Borda scoring system, where scores would be assigned as (4, 3, 2, 1, 0). Similarly, for the* `top-5` *rule, scores would be assigned to every candidate based on their ranking in voters' preferences, following the pattern (1, 0.8, 0.6, 0.4, 0.2).*

By adopting the `top-x` rule, which accounts for welfare scores, we aim to construct a ranking that captures the preferences of the voters, emphasizing on the candidates favored within the top x positions of their rankings. This approach provides a nuanced evaluation of candidates and seeks to generate a socially desirable ranking that reflects the overall satisfaction and preferences of the electorate. The essence of this rule lies in determining the top welfare through an analysis of the welfare scores derived from a truncated ranking comprising the highest-ranked candidates for each voter. By integrating this welfare-based approach and refining the Borda ranking based on the top choices, `welfare rules` offer a more comprehensive evaluation of candidates' desirability.

Having defined the family of `welfare rules` one could also consider the average of multiple voting rules within the family. Let `top-i`(c) be the score assigned to candidate c, under `top-i` voting rule. For any $x \leq m$, we introduce a rule, called `avg-x`, that assigns a score to every candidate c, according to the relation

$$sc(c) = \frac{1}{x} \sum_{i=1}^{x} \texttt{top-i}(c).$$

We highlight that the `top-x` approach is effective in terms of welfare when a specific value of x is selected. By considering an average, a rule that may not excel for any particular x value, but instead achieves a satisfactory balance across several values, is being established. Furthermore, the approach becomes less sensitive to distortions in rankings, resulting in a more robust and reliable rule.

For the purposes of the competition, we have opted to adopt the `avg-x` rule, where the specific value of x is being determined as a function of the total number of candidates of the given instance. More precisely we have used $\texttt{avg-(}min(\frac{m}{2}\texttt{, } 2\log(m)\texttt{))}$.

Submitted at The 1st Computational Social Choice Competition (COMPSOC'23, in conjunction with IJCAI'23).