Welfare Maximization in Perpetual Voting (Student Abstract)

Tzeh Yuan Neoh, Nicholas Teh

Department of Computer Science, University of Oxford, UK
tzeh.neoh@stcatz.ox.ac.uk, nicholas.teh@cs.ox.ac.uk

Abstract

We study the computational problems associated with maximizing various welfare objectives—namely utilitarian welfare, egalitarian welfare, and Nash welfare—in perpetual voting, a sequential collective decision-making framework. Prior work look into notions of fairness over time and study extensions of single-round voting rules to the multi-round setting. We show that while a utilitarian-welfare maximizing outcome can be computed efficiently, an outcome that maximizes egalitarian or Nash welfare is computationally intractable, even in the case of two candidates. We complement this by showing that maximizing egalitarian welfare is fixed-parameter tractable in the number of agents, and maximizing egalitarian or Nash welfare is W[2]-hard and slicewise polynomial in the number of timesteps. We also provide an approximation algorithm for maximizing egalitarian welfare and study strategyproofness with respect to these welfare objectives. Finally, we show that a simple greedy algorithm can achieve approximate proportionality in this setting.

Introduction

Consider the scenario where a group of friends are on a two-week post-graduation trip across Europe, and faced with the task of choosing their daily meals for the duration of the trip. If there are diverse preferences for various cuisines, especially exceeding the number of daily meals (and assuming a maximum of three meals per day), it might be challenging to accommodate everyone's desires. Nevertheless, with the span of multiple days, in may be possible to satisfy everyone's preferences. An important question arises: what would be an appropriate measure of *satisfaction*, and if so, can we (efficiently) obtain an outcome that satisfies such a measure?

This problem can be studied under the *perpetual voting* framework, a model where a sequence of decisions are made and the desirability of outcomes is studied with respect to agents' temporal preferences (also refer to a survey by Elkind, Obraztsova, and Teh (2023) on temporal multiwinner voting). Perpetual voting was first studied by Lackner (2020), which looks into a formalism of the model, and studies several perpetual extensions of traditional voting rules and axioms. Bulteau et al. (2021) subsequently looked into formalizing notions of proportionate representation in this context. This framework also captures several other scenarios, including fair scheduling (Elkind, Kraiczy, and Teh 2022).

However, most (if not all) of these works focus on fairness and proportionality notions. Maximizing welfare, in particular utilitarian welfare (sum of agents' utilities), egalitarian welfare (minimum utility of any agent), and Nash welfare (product of agents' utilities), have been widely studied in problems and frameworks across computational social choice. In this paper, we study the computational problems associated with maximizing these welfare objectives in the context of perpetual voting. We first formally introduce the framework.

The Perpetual Voting Model

Let $[k] := \{1, \dots, k\}$ for any positive integer k. In the perpetual voting model, we are given a set $N = [n]$ of n agents (or voters), a set of $C = \{c_1, \dots, c_m\}$ of m candidates (or alternatives), and a set $T = [\ell]$ of ℓ timesteps. At each timestep $k \in T$, each agent $i \in N$ has a set of approved candidates $S_{ik} \subseteq C$. Denote an agent's *preference vector* by $\mathbf{S}_i = (S_{i1}, \dots, S_{i\ell})$. An *outcome* $\mathbf{o} = (o_1, \dots, o_\ell)$ is a vector of ℓ candidates, one chosen at each timestep. An agent's *utility* for an outcome is then denoted by $u_i(\mathbf{o}) = |\{t \in [\ell] : o_t \in S_{it}\}|$.

Since we focus on the computational problems associated with maximizing various welfare objectives in this work, we introduce the associated decision problem. Moreover, we consider the more general *p-mean welfare* objective, which captures all three welfare objectives that we intend to study. The decision problem is defined as follows.

MAXIMIZING p-MEAN WELFARE:
Input: A problem instance $I = (N, C, T, (\mathbf{S}_i)_{i \in N})$, and a parameter $\lambda \in \mathbb{Z}^+$.
Question: Is there an outcome $\mathbf{o}$ such that $\left(\frac{1}{n}\sum_{i \in N} u_i(\mathbf{o})^p\right)^{1/p} \geq \lambda$?

Note that letting $p = 1, -\infty, 0$ would correspond to the utilitarian, egalitarian, and Nash welfare respectively. For ease of exposition, we refer to the decision problems corresponding to the utilitarian, egalitarian, and Nash welfare as UTIL, EGAL, and NASH respectively.

Finding an outcome that maximizes utilitarian welfare can be easily obtained through a greedy algorithm: at each timestep, simply select the project that has the highest number of approvals. However, it is easy to see that if a simple majority of the population approves of one common candidate at each timestep, then the winning candidate at each timestep will be exactly the candidate approved by this same group of agents. This may be seen as unfair, as possibly close to half of the agent population may not get a single candidate approved at any timestep. As such, there is a need to consider other welfare measures as a potential balance between utilitarian welfare and fairness. In this respect, we focus on two other commonly studied welfare objectives for achieving fairness in the social choice literature—maximizing egalitarian welfare or Nash welfare.

The next section details our results. Due to space constraints, all proofs are omitted.

Our Results

As mentioned above, while computing an outcome that maximizes utilitarian welfare is computationally tractable, the same cannot be said for the other two objectives. The following result shows that, perhaps surprisingly, both EGAL and NASH are NP-complete even when guaranteeing each agent a utility of 1, and when there are only two candidates.

Theorem 1. EGAL *and* NASH *are both* NP-*Complete, even for* $\lambda = 1$ *and* $m = 2$.

The above negative result effectively rules out the possibility of maximizing the egalitarian or Nash welfare objectives even in simple settings.

However, when the number of agents or timesteps is constant, we are able to efficiently find a solution. This follows from the parameterized complexity results we state below. We first show that EGAL is fixed-parameter tractable (FPT) with respect to the number of agents, but the same observation cannot be extended to NASH.

Theorem 2. EGAL *is* FPT *with respect to* n.

The proof of the above relies on the construction of an integer linear program, but such an approach would fail when the objective function contains a non-linear logarithmic function (which is present in NASH). Thus, an analogous FPT (in n) result for NASH remains open.

Next, we show that when the number of timesteps is constant, both EGAL and NASH can be solved in polynomial time, i.e., slicewise polynomial (XP) with respect to the number of timesteps.

Theorem 3. EGAL *and* NASH *are both* XP *with respect to* ℓ.

Furthermore, we complement the above result by showing that EGAL and NASH is W[2]-hard with respect to ℓ, essentially indicating that an FPT (in ℓ) algorithm does not exist unless FPT = W[2], and hence the XP result is tight.

Theorem 4. EGAL *and* NASH *are both* W[2]-*Hard with respect to* ℓ.

Now, by Theorem 1, EGAL and NASH is NP-complete even when $\lambda = 1$. Hence, no approximation algorithm can exist. However, if we augment each agent's utility function by defining it to be $u'_i(\mathbf{o}) = 1 + u_i(\mathbf{o})$, then, with respect to u', we can obtain an $\frac{1}{4 \log n}$-approximation to the optimal egalitarian welfare.

Theorem 5. *With respect to the augmented utility functions, there exists a* $\frac{1}{4 \log n}$*-approximation algorithm for* EGAL.

Next, we study *strategyproofness* with respect to the three welfare objectives.

Strategyproofness

Another important consideration in the social choice literature is strategyproofness, which states that no agent should be able to misreport their preferences so as to obtain strictly better utility. We show that the greedy algorithm used in obtaining a utilitarian welfare-maximizing outcome—which we term GREEDYUTIL—satisfies this property.

Theorem 6. GREEDYUTIL *is strategyproof.*

However, this positive result does not extend to the other two welfare objectives.

Theorem 7. *Any deterministic mechanism maximizing egalitarian or Nash welfare cannot be strategyproof.*

Additionally, we include a discussion on an adjacent property which may be of interest in this setting.

Proportionality

Another common property studied in related settings is the concept of *proportionality* (PROP). However, as it is typically defined, PROP is not always achievable. Consider the setting with two agents and a single timestep, with preference vectors $\mathbf{S}_1 = (c_1)$ and $\mathbf{S}_2 = (c_2)$. No matter which candidate is selected, one agent will have utility of 0. Hence, we extend the notion of proportionality appropriately by adding a floor function.

Definition 1 (Proportionality)**.** An outcome $\mathbf{o}$ is *proportional (PROP)* if for all $i \in N$ it holds that $u_i(\mathbf{o}) \geq \lfloor \frac{\mu_i}{n} \rfloor$, where $\mu_i = |\{k \in [\ell] : S_{ik} \neq \emptyset\}|$.

Then, we get the following positive result.

Theorem 8. *A* PROP *outcome always exists and can be computed by a polynomial-time greedy algorithm.*

Conclusion and Future Work

In this work, we studied various computational questions related to three main welfare objectives in the context of perpetual voting, and looked at additional properties such as strategyproofness and proportionality. Future directions include exploring the computational hardness (or easiness) for more general p-mean welfare objectives, or considering a different definition of EGAL or NASH where utilities are augmented, and studying if the results would differ.

References

Bulteau, L.; Hazon, N.; Page, R.; Rosenfeld, A.; and Talmon, N. 2021. Justified Representation for Perpetual Voting. *IEEE Access*, 9: 96598–96612.

Elkind, E.; Kraiczy, S.; and Teh, N. 2022. Fairness in Temporal Slot Assignment. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory (SAGT)*, 490–507.

Elkind, E.; Obraztsova, S.; and Teh, N. 2023. Temporal Fairness in Multiwinner Voting. *arXiv preprint arXiv:2312.04417*.

Lackner, M. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2103–2110.