Dynamic Preference Aggregation under Preference Changes

Matthias Thimm

Institute for Web Science and Technologies
Universität Koblenz-Landau, Germany

Abstract. We consider the issue of update in settings for preference aggregation under preference changes. While the traditional problem formulation of preference aggregation assumes a fixed set of preference orders and a fixed set of domain elements we investigate how an aggregated preference order has to be updated when the input orders are dynamic. Our analysis shows that for even for some simple aggregation rules, i. e. for the plurality and the Borda rule, the dynamic setting can be handled more efficiently than recomputing the aggregated preference order from scratch when changes have to be made.

1 Introduction

Preference aggregation [1] deals with the problem of combining the preferences of multiple agents into a joint preference order that reflects the agents' preferences in a "fair" manner. This field is strongly related to voting systems and social choice theory [2] and has its main application in recommendation and decision making in groups [3, 4]. The traditional setting for preference aggregation is a static one. Given fixed preference orders on a fixed domain one asks for a preference order which jointly represents the former ones. However, in many scenarios both the input preferences and/or the domain under consideration changes. Our motivation for investigating the issue of dynamic preference aggregation stems from recommender systems for social media applications. Consider the situation of visiting a film festival with a mobile recommendation system. In this scenario, one may have preferences regarding movie genres and actors that lead to recommendations to which theater to go next. Furthermore, one may have information on the location of friends and geo-temporal information on the screening of movies which also lead to preference orders, e. g. one would prefer to go to a screening which starts in a couple of minutes or where a lot of friends are present. All these preference orders change over time and so does the aggregated preference order.

In this paper we investigate the issue of preference aggregation under changes of the input preference orders. For that we develop a framework for preference change that distinguishes between two atomic types of changes to a preference order, namely, *weakening* and *strengthening* of a specific domain element. Given a set of preference orders and an aggregation rule we investigate how the aggregated preference order changes under atomic changes of the input preference orders. We introduce the notion of a *dynamic preference aggregator* that dynamically adapts the aggregated preference order to changes and develop first dynamic approaches that implement the plurality and the Borda rule for preference aggregation. Our analysis of these approaches show that they outperform a naive re-computation of the aggregated preference order.

The issue of dynamic preference aggregation is closely related to the issue of bribery [5]. The bribery problem deals with the question of how to (minimally) adjust voters' preferences in order to establish some desired aggregated preference order. Dynamic preference aggregation is basically the reverse problem as we investigate how minimal changes of the voters' preference change the aggregated order. However, to the best of our knowledge the computational questions arising in reversing the bribery question have not been addressed before explicitly.

The remainder of this paper is structured as follows. In Section 2 we recall preliminaries on preference representation and preference aggregation. In Section 3 we develop our general framework of dynamic preference aggregation under preference change and investigate its general properties. Afterwards we have a look at specific aggregation rules and their implementation in the dynamic setting in Section 4. We review some related work in Section 5 and conclude with a summary and discussion in Section 6.

2 Preferences and Preference Aggregation

Let $\mathcal{O} = \{o_1, \ldots, o_n\}$ be a set of outcomes. In order to have a more general setting, in contrast to the majority of literature on preference aggregation we will use total preorders, instead of linear orders, to represent preferences. Therefore, a *preference order* $\preceq$ on $\mathcal{O}$ is a total preorder on $\mathcal{O}$, i. e. a relation $\preceq \subseteq \mathcal{O} \times \mathcal{O}$ that satisfies

1. if $o_1 \preceq o_2$ and $o_2 \preceq o_3$ then $o_1 \preceq o_3$ (transitivity) and
2. for all $o_1, o_2 \in \mathcal{O}$ is holds $o_1 \preceq o_2$ or $o_2 \preceq o_1$ (totality)

If $o \preceq o'$ then we say that o' is at least as preferred as o. We abbreviate $o \sim o'$ if both $o \preceq o'$ and $o' \preceq o$, and we abbreviate $o \prec o'$ if $o \preceq o'$ and $o' \not\preceq o$. Let $\mathcal{P}_{\mathcal{O}}$ be the set of all preference orders on $\mathcal{O}$. Preference orders can be used to represent a single individual's preferences on the possible outcomes (of some action) in $\mathcal{O}$.

Example 1. Let $\mathcal{O} = \{\mathit{rock}, \mathit{pop}, \mathit{country}, \mathit{electronic}\}$ be a set of outcomes that describe the choices for the music genre being played at some event. A possible preference order $\preceq_{music}$ on $\mathcal{O}$ can be given via

$$\begin{aligned} \mathit{country} &\prec_{music} \mathit{pop} \\ \mathit{country} &\prec_{music} \mathit{electronic} \\ \mathit{pop} &\prec_{music} \mathit{rock} \\ \mathit{electronic} &\prec_{music} \mathit{rock} \end{aligned}$$

In $\preceq_{music}$ the outcome *rock* is the most preferred option, both *pop* and *electronic* are equally preferred after *rock*, and *country* is the least preferred option.

A preference order $\preceq$ on $\mathcal{O}$ can also be concisely characterized by its *ranking function* $rank_{\preceq} : \mathcal{O} \to \mathbb{N}$ defined via

$$rank_{\preceq}(o) = |\{o' | o' \prec o\}|$$

Note that we have $rank_{\preceq}(o) \leq rank_{\preceq}(o')$ if and only if $o \preceq o'$. In the following, we use both the preference order $\preceq$ and its ranking function $rank_{\preceq}$ interchangeably. For $\preceq_{\mathrm{music}}$ of Example 1 we have

$$\begin{aligned} rank_{\preceq_{music}}(country) &= 0 \\ rank_{\preceq_{music}}(pop) &= 1 \\ rank_{\preceq_{music}}(electronic) &= 1 \\ rank_{\preceq_{music}}(rock) &= 3 \end{aligned}$$

We define the sets of *direct predecessors* $pre_{\preceq}(o)$ and *direct successors* $suc_{\preceq}(o)$ of an outcome o as follows:

$$\begin{aligned} pre_{\preceq}(o) &= \{o' \in \mathcal{O} \mid o' \prec o \wedge \; \neg\exists\, o'' : o' \prec o'' \prec o\} \\ suc_{\preceq}(o) &= \{o' \in \mathcal{O} \mid o \prec o' \wedge \; \neg\exists\, o'' : o \prec o'' \prec o'\} \end{aligned}$$

We further define the *set of most preferred outcomes* $\mathsf{top}(\preceq)$ and the *set of least preferred outcomes* $\mathsf{bot}(\preceq)$ via

$$\begin{aligned} \mathsf{top}(\preceq) &= \{o \in \mathcal{O} \mid suc_{\preceq}(o) = \emptyset\} \\ \mathsf{bot}(\preceq) &= \{o \in \mathcal{O} \mid pre_{\preceq}(o) = \emptyset\} \end{aligned}$$

In Example 1 we have e. g.

$$\begin{aligned} suc_{\preceq_{music}}(country) &= \{pop, electronic\} \\ \mathsf{top}(\preceq_{music}) &= \{rock\} \end{aligned}$$

The field of preference aggregation [1] is concerned with how potentially conflicting preferences of different individuals are to be aggregated in order to come up with a single preference order that reflects the group's joint preference in a "fair" sense. Formally, a *preference aggregator* is defined as follows.

Definition 1. *A* preference aggregator Θ *(of order* $m \in \mathbb{N}^+$*) is a function* $\Theta : \mathcal{P}_{\mathcal{O}}^m \to \mathcal{P}_{\mathcal{O}}$.

A vector of preference orders $\underline{\preceq} \in \mathcal{P}_{\mathcal{O}}^m$ is also called a *preference profile.*

In the past 60 years a series of different implementations for preference aggregators have been proposed. The goal of this research is to define a preference aggregator that is, in some sense, "fair", i. e. reflects the individual preferences in an unbiased and intuitively correct way. Some properties that formalize the notion of "fairness" are as follows.

Pareto-efficiency (PE) For every $o, o' \in \mathcal{O}$, if $o \preceq_i o'$ for all $i = 1, \ldots, m$ then $o \preceq o'$ (for preference orders $\preceq_1, \ldots, \preceq_m$ and $\preceq = \Theta(\preceq_1, \ldots, \preceq_m)$).

Non-dictatorship (ND) There is no $i \in \{1, \ldots, m\}$ such that $\preceq_i = \Theta(\preceq_1, \ldots, \preceq_m)$ for *every* $\langle \preceq_1, \ldots, \preceq_m \rangle \in \mathcal{P}_{\mathcal{O}}^m$.

Independence of irrelevant alternatives (IIA) If for two profiles $\langle \preceq_1, \ldots, \preceq_m \rangle$ and $\langle \preceq'_1, \ldots, \preceq'_m \rangle$ and every $i = 1, \ldots, m$ it holds $o \preceq_i o'$ whenever $o \preceq'_i o'$ then $o \preceq o'$ whenever $o \preceq' o'$ (with $\preceq = \Theta(\preceq_1, \ldots, \preceq_m)$ and $\preceq' = \Theta(\preceq'_1, \ldots, \preceq'_m)$).

Monotonicity (Mon) If for two profiles $\langle \preceq_1, \ldots, \preceq_m \rangle$ and $\langle \preceq'_1, \ldots, \preceq'_m \rangle$ we have that $o \preceq_i o'$ implies $o \preceq'_i o'$ then $o \preceq o'$ implies $o \preceq o'$ (with $\preceq = \Theta(\preceq_1, \ldots, \preceq_m)$ and $\preceq' = \Theta(\preceq'_1, \ldots, \preceq'_m)$).

Pareto-efficiency says that if all preference orders agree that o' is at least as preferred as o then o' should also be at least as preferred as o in the joint preference order. The property *non-dictatorship* says that there is no single individual (the dictator) whose preference order is always adopted as the joint preference order, regardless of what the preference orders actually look like. The property *independence of irrelevant alternatives* can be best explained in the context of voting. There, this property states that given two elections A and B with voters $1, \ldots, m$, if every voter i prefers candidate 1 to candidate 2 in A whenever he prefers candidate 1 to candidate 2 in B, then candidate 1 is preferred to candidate 2 in the outcome of the election A if and only if candidate 1 is preferred to candidate 2 in the outcome of the election B. This means that preferences regarding other combinations of candidates are irrelevant for the decision regarding just candidates 1 and 2. Monotonicity basically demands that strengthening the position of an outcome o in one preference order cannot result in a worse situation for o in the aggregated preference order. Arrow's famous impossibility result [6] states that there is no preference aggregator Θ which satisfies (PE), (ND), and (IIA).[1]

Research in preference aggregation did come up with several preference aggregators that are meaningful under certain circumstances. In this paper, we consider two of the most simple preference aggregators as suitable examples, namely *plurality preference aggregation* and *Borda preference aggregation* which are both special cases of scoring rules, cf. [1]. Note that the original definitions for these preference aggregators are given for total orders instead of total preorders as we use here. We adapt those definitions appropriately for total preorders as follows, cf. [7] for some general discussions regarding partially ordered preferences. Let $\underline{\preceq} = \langle \preceq_1, \ldots, \preceq_m \rangle$ be a preference profile. The *plurality preference aggregator* $\Theta_p : \mathcal{P}_{\mathcal{O}}^m \to \mathcal{P}_{\mathcal{O}}$ and the *Borda preference aggregator* $\Theta_b : \mathcal{P}_{\mathcal{O}}^m \to \mathcal{P}_{\mathcal{O}}$

[1] Note that the original impossibility result refers to total orders instead of total preorders but a similar result holds for total preorders as well.

are defined via $\Theta_p(\preceq) = \preceq_p$ and $\Theta_b(\preceq) = \preceq_b$, respectively, and

$$o \preceq_p o' \quad \text{iff} \quad |\{i \mid o \in \mathsf{top}(\preceq_i)\}| \leq |\{i \mid o' \in \mathsf{top}(\preceq_i)\}|$$

$$o \preceq_b o' \quad \text{iff} \quad \sum_{i=1}^{m}(rank_{\preceq_i}(o)) \leq \sum_{i=1}^{m}(rank_{\preceq_i}(o')) \quad .$$

Plurality preference aggregation generalizes majority voting and says, that o' is at least as preferred as o if and only if o' appears at least as often as top element in the preference orders $\preceq_1, \ldots, \preceq_m$ as o. The intuition behind Borda preference aggregation is that o' is at least as preferred as o, if the sum of all ranks over all considered preference orders for o' is smaller or equal to the sum for o. Note, that we generalized the standard definition of the Borda rule by assigning equal rank to equally classified outcomes.

3 Dynamic Preference Aggregation

In the traditional setting for preference aggregation, preference orders are static entities that reflect the preferences of some individual at some given point in time. Furthermore, a preference aggregator is a simple function that takes this static view and returns a coherent joint preference order for the very same point in time. However, in reality preferences are rarely static but change frequently. In the following, we consider the problem of how the joint preference order changes given that one of the input preference orders changes.

In order to describe "change" of preference orders we need some further notation. We focus one specific atomic change operations for preference orders where we only change the preference of a single outcome. More precisely, for a preference order $\preceq$, an outcome $o \in \mathcal{O}$, the *strengthening* of $\preceq$ by o, denoted by $\preceq + o$, is the preference order $\preceq' = \preceq + o$ characterized via

$$rank_{\preceq'}(\alpha) = \begin{cases} rank_{\preceq}(\alpha) + |\{o' | o' \sim \alpha\} \setminus \{\alpha\}| \\ \qquad\qquad\qquad \text{(if } \alpha = o) \\ rank_{\preceq}(\alpha) - 1 \\ \text{(if } \alpha \in suc_{\preceq}(o) \text{ and } |\{o' \mid o' \sim o\}| = 1) \\ rank_{\preceq}(\alpha) \qquad \text{(otherwise)} \end{cases}$$

for every $o' \in \mathcal{O}$. Similarly, the *weakening* of $\preceq$ by o, denoted by $\preceq - o$, is the preference order $\preceq' = \preceq - o$ characterized via

$$rank_{\preceq'}(\alpha) = \begin{cases} rank_{\preceq}(\alpha) - |\{o' | o' \in pre(\alpha)\}| \\ \qquad \text{(if } \alpha = o \text{ and } |\{o' \mid o' \sim o\}| = 1) \\ rank_{\preceq}(\alpha) + 1 \qquad \text{(if } \alpha \neq o \text{ and } \alpha \sim o) \\ rank_{\preceq}(\alpha) \qquad \text{(otherwise)} \end{cases}$$

for every $o' \in \mathcal{O}$.

Example 2. We continue Example 1 and assume the music genre *electronic* has to be weakened in $\preceq_{music}$. Then we have $\preceq'_{music} = \preceq_{music} -electronic$ with

$$country \prec'_{music} electronic \prec'_{music} pop \prec'_{music} rock$$

with

$$rank_{\preceq'_{music}}(country) = 0 \qquad rank_{\preceq'_{music}}(electronic) = 1$$
$$rank_{\preceq'_{music}}(pop) = 2 \qquad rank_{\preceq'_{music}}(rock) = 3$$

Furthermore, if we weaken *electronic* even further we obtain $\preceq''_{music} = \preceq'_{music} -electronic$ with

$$country, electronic \prec''_{music} pop \prec''_{music} rock$$

and

$$rank_{\preceq''_{music}}(country) = 0 \qquad rank_{\preceq''_{music}}(electronic) = 0$$
$$rank_{\preceq''_{music}}(pop) = 2 \qquad rank_{\preceq''_{music}}(rock) = 3$$

Strengthening and weakening of outcomes can be regarded as the most basic change operation for a preference order. Some simple observations on strengthening and weakening are summarized in the following proposition. Proofs of technical results are omitted due to space restrictions.

Proposition 1. *Let $\preceq \in \mathcal{P}_{\mathcal{O}}$ be a preference order.*

1. $(\preceq - o) + o = \preceq$ *for all* $o \notin \mathsf{bot}(\preceq)$,
2. $(\preceq + o) - o = \preceq$ *for all* $o \notin \mathsf{top}(\preceq)$,
3. *for every other preference order* $\preceq' \in \mathcal{P}_{\mathcal{O}}$ *there are sequences* $\langle \pm_1, \ldots, \pm_l \rangle$, $\langle o_1, \ldots, o_l \rangle$ *with* $\pm_i \in \{+, -\}$, $o_i \in \mathcal{O}$ *for* $i = 1, \ldots, l$ *for some* $m \in \mathbb{N}$ *such that* $\preceq = \preceq' \pm_1 o_1 \ldots \pm_l o_l$.
4. *computing* $\preceq \pm o$ *has time complexity* $O(n)$.

Items 1.) and 2.) state that strengthening and weakening are inverse operations as long as the element is not a most or least preferred outcome, respectively. Item 3.) states that strengthening and weakening operations are generating operations of the set of preference orders, i. e. every preference order can be represented as applying a series of strengthening/weakening operations to every other preference relation. Finally, item 4.) states that determining $\preceq \pm o$ from $\preceq$ has time complexity $O(n)$.

If $\underline{\preceq} = \langle \preceq_1, \ldots, \preceq_m \rangle \in \mathcal{P}_{\mathcal{O}}^m$ is a preference profile we abbreviate

$$\underline{\preceq} \pm_i o = \langle \preceq_1, \ldots, \preceq_i \pm o, \ldots, \preceq_m \rangle$$

with $\pm \in \{+, -\}$, $o \in \mathcal{O}$, and $i = 1, \ldots, , m$. The single-step operators $+$ and $-$ can be extended to multi-step operators $+^h$ and $-^h$ for $h \in \mathbb{N}$ via

$$\preceq +^h o = \preceq + \underbrace{o + \ldots + o}_{h \text{ times}}$$
$$\preceq -^h o = \preceq - \underbrace{o - \ldots - o}_{h \text{ times}}$$

Some observations relating change operations and preference aggregation are as follows.

Proposition 2. *Let* $\underline{\preceq}$ *be a preference profile,* $i \in \{1, \ldots, m\}$*,* $\pm \in \{+, -\}$*,* $o \in \mathcal{O}$*,* Θ *a preference aggregator,* $\preceq = \Theta(\underline{\preceq})$*, and* $\preceq' = \Theta(\underline{\preceq} \pm_i o)$*.*

1. *If* Θ *satisfies (IIA) then* $o_1 \preceq o_2$ *iff* $o_1 \preceq' o_2$ *for all* $o_1, o_2 \in \mathcal{O} \setminus \{o\}$*.*
2. *If* $\pm = +$ *and* Θ *satisfies (Mon) then* $o_1 \preceq o_2$ *implies* $o_1 \preceq' o_2$ *for all* $o_1 \in \mathcal{O} \setminus \{o\}, o_2 \in \mathcal{O}$*.*
3. *If* $\pm = -$ *and* Θ *satisfies (Mon) then* $o_1 \preceq o_2$ *implies* $o_1 \preceq' o_2$ *for all* $o_1 \in \mathcal{O}, o_2 \in \mathcal{O} \setminus \{o\}$*.*

The above proposition states that there are many cases in which the aggregated preference order only changes minimally on atomic changes of an input preference order. For that, it seems beneficial to investigate the computational issue of dynamic preference aggregation. The core concept of our formalism is the dynamic preference aggregator which can naively be defined as follows.

Definition 2. *A* dynamic preference aggregator Λ *is a function* $\Lambda : \mathcal{P}_{\mathcal{O}} \times \mathbb{N} \times \{+, -\} \times \mathcal{O} \to \mathcal{P}_{\mathcal{O}}$*.*

Let $\preceq \in \mathcal{P}_{\mathcal{O}}$ be some preference order that is the result of applying some aggregation rule on the profile $\underline{\preceq} = \langle \preceq_1, \ldots, \preceq_m \rangle$. Furthermore, let $i \in \{1, \ldots, m\}$, $\pm \in \{+, -\}$, and $o \in \mathcal{O}$. Then the idea behind a dynamic preference aggregator Λ is that $\Lambda(\preceq, i, \pm, o) = \preceq'$ is the result of first updating the preference profile $\underline{\preceq}$ via $\underline{\preceq} \pm_i o$ and then aggregating again to $\preceq'$. However, this update operation should be performed on $\preceq$ directly. Formally, the intuition of dynamic preference aggregation can be phrased as follows.

Definition 3. *Let* Θ *be a preference aggregator and* Λ *be a dynamic preference aggregator. We say that* Λ *is a* faithful representation *of* Θ *if and only if for all* $\underline{\preceq} = \langle \preceq_1, \ldots, \preceq_n \rangle \in \mathcal{P}^m_{\mathcal{O}}$*, all* $i = 1, \ldots, n$*,* $\pm \in \{+, -\}$*, and all* $o \in \mathcal{O}$ *it holds*

$$\Theta(\underline{\preceq} \pm_i o) = \Lambda(\Theta(\underline{\preceq}), i, \pm, o)$$

In other words, a faithful representation of a preference aggregator is a dynamic preference aggregator that exhibits the same behavior as the former but adapts to changes of the input preference orders. Given a preference profile $\underline{\preceq}$ we expect updating $\Theta(\underline{\preceq})$ using Λ to be same as first updating $\underline{\preceq}$ and then applying Θ again. Figure 1 illustrates this relationship between a preference aggregator Θ and a faithful representation Λ using a commuting diagram. However, the definition of a faithful representation is very restricting as it requires a functional dependency between the old result of aggregation and the new one. In general, this demand is not satisfiable for most interesting aggregation operators.

Example 3. Let $\preceq_1, \preceq_2, \preceq_3$ be preference orders on $\mathcal{O} = \{o_1, o_2, o_3\}$ defined via

$$o_3 \prec_1 o_1 \prec_1 o_2$$
$$o_2 \prec_2 o_3 \prec_2 o_1$$
$$o_3 \prec_3 o_2 \prec_3 o_1$$

$$\begin{tikzcd}
\preceq_1, \ldots, \preceq_n \arrow[r, "\Theta"'] \arrow[d, "\preceq'_i = \preceq_i \pm o"] & \preceq \arrow[d, "\Lambda"'] \\
\preceq_1, \ldots, \preceq'_i, \ldots, \preceq_n \arrow[r, "\Theta"] & \preceq'
\end{tikzcd}$$

Fig. 1. Commuting diagram for dynamic preference aggregation

By using plurality aggregation we obtain $\Theta_p(\preceq_1, \preceq_2) = \Theta_p(\preceq_1, \preceq_3) = \preceq$ with

$$o_1 \sim o_2, o_3 \prec o_1, o_3 \prec o_2 \quad .$$

That is, we obtain that profiles $\langle \preceq_1, \preceq_2 \rangle$ and $\langle \preceq_1, \preceq_3 \rangle$ are aggregated yielding the same result. Let $\underline{\preceq} = \langle \preceq_1, \preceq_2 \rangle$, $\underline{\preceq}' = \langle \preceq_1, \preceq_3 \rangle$ and consider $\hat{\preceq} = \Theta_p(\underline{\preceq} +_2 o_2)$ resp. $\hat{\preceq}' = \Theta_p(\underline{\preceq}' +_2 o_2)$ with

$$o_1 \hat{\sim} o_2, o_3 \hat{\prec} o_1, o_3 \hat{\prec} o_2$$
$$o_1 \hat{\prec}' o_2, o_3 \hat{\prec}' o_2, o_3 \hat{\prec}' o_1$$

In summary, we get $\Theta_p(\underline{\preceq}) = \Theta_p(\underline{\preceq}')$ but $\Theta_p(\underline{\preceq} +_2 o_2) \neq \Theta_p(\underline{\preceq}' +_2 o_2)$. Therefore, there can be no dynamic preference aggregator Λ_p that is a faithful representation of Θ_p.

The above example (unsurprisingly) shows that there is no direct functional dependency between $\preceq$ and $\preceq'$. In order to get to actual approaches for addressing the dynamic preference problem we need to have some more information that is carried from one update to the other.

Definition 4. *A* state-based dynamic preference aggregator Δ *is a pair* $\Delta = \langle \iota, \Lambda \rangle$ *such that*

1. ι *is function* $\iota : \mathcal{P}_{\mathcal{O}}^m \to \mathcal{S} \times \mathcal{P}_{\mathcal{O}}$ *and*
2. Λ *is a function* $\Lambda : \mathcal{S} \times \mathbb{N} \times \{+, -\} \times \mathcal{O} \to \mathcal{S} \times \mathcal{P}_{\mathcal{O}}$

where $\mathcal{S}$ *is some set of states.*

The intuition behind a state-based dynamic preference aggregator $\Delta = \langle \iota, \Lambda \rangle$ is as follows. Given a preference profile $\underline{\preceq} = \langle \preceq_1, \ldots, \preceq_m \rangle$ the function ι is the *initialization function* that delivers $\iota(\underline{\preceq}) = (S, \preceq)$ where $\preceq$ is the aggregated preference order of $\underline{\preceq}$ and S some state (defined in some way suitable for the preference aggregator) that is carried over to the next update. Given a state S, $i \in \{1, \ldots, m\}$, $\pm \in \{+, -\}$, and $o \in \mathcal{O}$ the value $\Lambda(S, i, \pm, o) = (S', \preceq')$ then updates the state S to S' and the aggregated order $\preceq$ to $\preceq'$, given the change operation $\pm_i$ on o. Formally, this intuition extends our notion of faithfulness as follows.

Definition 5. *Let Θ be a preference aggregator and $\Delta = \langle \iota, \Lambda \rangle$ be a state-based dynamic preference aggregator. We say that Δ is a* state-based faithful representation *of Θ if and only if for all $\underline{\preceq} = \langle \preceq_1, \ldots, \preceq_m \rangle \in \mathcal{P}_{\mathcal{O}}^m$, all sequences $\langle \langle \pm_1, i_1, o_1 \rangle, \ldots, \langle \pm_k, i_k, o_k \rangle \rangle \in (\{+,-\} \times \mathbb{N} \times \mathcal{O})^k$ there are states $S_1, \ldots, S_{k-1} \in \mathcal{S}$ such that*

$$\langle S_1, \Theta(\underline{\preceq}) \rangle = \iota(\underline{\preceq})$$
$$\langle S_j, \Theta(\underline{\preceq}(\pm_1)_{i_1} o_1 \ldots (\pm_j)_{i_j} o_j) \rangle = \Lambda(S_{j-1}, i_j, \pm_j, o_j)$$

for $j = 2, \ldots, k$.

In the next section we investigate actual approaches for state-based dynamic preference aggregators that effectively implement the preference aggregators discussed before.

4 Approaches and Analysis

We first discuss a naive implementation which just applies the original preference aggregator on every change operation.

Definition 6. *Let Θ be a preference aggregator and let $\mathcal{S}_{can} = \mathcal{P}_{\mathcal{O}}^m$. The* canonical state-based dynamic preference aggregator *Δ_Θ^{can} for Θ is the pair $\Delta_\Theta^{can} = \langle \iota_\Theta^{can}, \Lambda_\Theta^{can} \rangle$ with $\iota_\Theta^{can} : \mathcal{P}_{\mathcal{O}}^m \to \mathcal{S}_{can} \times \mathcal{P}_{\mathcal{O}}$ and $\Lambda_\Theta^{can} : \mathcal{S}_{can} \times \mathbb{N} \times \{+,-\} \times \mathcal{O} \to \mathcal{S}_{can} \times \mathcal{P}_{\mathcal{O}}$ defined via*

$$\iota_\Theta^{can}(\underline{\preceq}) = \langle \underline{\preceq}, \Theta(\underline{\preceq}) \rangle$$
$$\Lambda_\Theta^{can}(\underline{\preceq}, i, \pm, o) = \langle \underline{\preceq} \pm_i o, \Theta(\underline{\preceq} \pm_i o) \rangle$$

The canonical state-based dynamic preference aggregator simply carries over the whole preference profile from one iteration to the next and applies the preference aggregator in a direct way. This is the baseline approach for solving the dynamic preference aggregation problem but, obviously, the effort required at each iteration should not be necessary. The following proposition follows by construction of the canonical state-based dynamic preference aggregator.

Proposition 3. *Let Θ be a preference aggregator. Then Δ_Θ^{can} is a state-based faithful representation of Θ.*

Example 4. We continue Example 3 and define $\underline{\preceq} = \langle \preceq_1, \preceq_2, \preceq_3 \rangle$. For the plurality preference aggregator Θ_p we obtain

$$\iota_{\Theta_p}^{can}(\underline{\preceq}) = \langle \underline{\preceq}, \preceq \rangle$$

where $\preceq$ is defined via $o_3 \prec o_2 \prec o_1$. Therefore, $\preceq$ is the initial result of aggregating $\underline{\preceq}$. Assume now that in the preference order $\preceq_3$ the outcome o_2 is to be strengthened, i. e., we have $\preceq_3' = \preceq_3 + o_2$ with $o_3 \preceq_3' o_1, o_2$. Then we have

$$\Lambda_{\Theta_p}^{can}(\underline{\preceq}, 3, +, o_2) = \langle\langle \preceq_1, \preceq_2, \preceq_3' \rangle, \preceq' \rangle$$

where $\preceq' = \Theta_p(\preceq_1, \preceq_2, \preceq_3')$ is defined via $o_3 \prec' o_1, o_2$. In this example, a (simple) strengthening of o_2 in one of the input preference orders caused a strengthening of o_2 in the aggregated order as well.

Proposition 4. *Let Θ be a preference aggregator of time complexity $O(f(n,m))$. Then ι_Θ^{can} has time complexity $O(f(n,m))$ and Λ_Θ^{can} has time complexity $O(n + f(n,m))$. The space complexity for storing a state in $\mathcal{S}_{can}$ is $O(nm)$.*

The question we are addressing in the following is whether the time complexity $O(n + f(n,m))$ of updating the aggregated preference order can be improved. Proposition 2 already suggested that there are many cases under which the aggregated preference order only changes slightly. We now have a look at some concrete preference aggregators and, first, give a straightforward approach for the plurality aggregator.

Definition 7. *Let $\underline{\preceq} = \langle \preceq_1, \ldots, \preceq_m \rangle$ and let $\mathcal{S}_p = \mathcal{P}_\mathcal{O}^m \times \mathbb{N}^\mathcal{O}$. Define*

$$\iota_p(\underline{\preceq}) = \langle \langle \underline{\preceq}, g \rangle, \Theta_p(\underline{\preceq}) \rangle$$

with $g : \mathcal{O} \to \mathbb{N}$ and $g(o) = |\{i \mid o \in \mathsf{top}(\preceq_i)\}|$. Define

$$\Lambda_p(\langle \underline{\preceq}, g \rangle, i, \pm, o) = \langle \langle \underline{\preceq} \pm_i o, g' \rangle, \preceq' \rangle$$

with $g' : \mathcal{O} \to \mathbb{N}$ defined via[2]

$$\begin{aligned} g'(o') = g(o') &- [o' \in \mathsf{top}(\preceq_i) \wedge o \notin \mathsf{top}(\preceq_i \pm o)] \\ &+ [o' \notin \mathsf{top}(\preceq_i) \wedge o \in \mathsf{top}(\preceq_i \pm o)] \end{aligned}$$

and

$$o_1 \preceq' o_2 \qquad \textit{iff} \qquad g'(o_1) \leq g'(o_2)$$

Then $\Delta_p = \langle \iota_p, \Lambda_p \rangle$ is called the dynamic plurality preference aggregator*.*

The above definition avoids re-computing the whole aggregated preference order and only considers changes in the top elements of the preference order that is being changed. More precisely, the value $g(o')$—which stores the number of times each element appears as a top element in the preference orders—is updated by subtracting 1 if the element has been top-ranked in $\preceq_i$ and is no longer top-ranked in $\preceq_i \pm o$ or by adding 1 if the element has not been top-ranked in $\preceq_i$ but is top-ranked in $\preceq_i \pm o$.

[2] $[P]$ is the Iverson bracket defined via $[P] = 1$ if P is true and $[P] = 0$ otherwise.

Proposition 5. *Δ_p is a state-based faithful representation of Θ_p.*

Proposition 6. *ι_p has time complexity $O(nm)$ and Λ_p has time complexity $O(n)$. The space complexity for storing a state in $\mathcal{S}_p$ is $O(nm)$.*

Note that the time complexity of computing $\Theta_p(\preceq)$ is $O(nm)$. Thus Λ_p has better time complexity ($O(n)$) than the canonical solution $\Lambda^{can}_{\Theta_p}$ with $O(nm)$, cf. Proposition 4.

Example 5. We continue Example 4. For $\preceq = \langle \preceq_1, \preceq_2, \preceq_3 \rangle$ we obtain

$$\iota_p(\preceq) = \langle\langle \preceq, g\rangle, \preceq\rangle\rangle$$

with $\preceq$ as in Example 4 and $g : \mathcal{O} \to \mathbb{N}$ given via $g(o_1) = 2$, $g(o_2) = 1$, and $g(o_3) = 0$. Given that o_2 is to be strengthened in $\preceq_3$, i. e. $\preceq'_3 = \preceq_3 +o_2$, we obtain

$$\Lambda_p((\preceq, g), 3, +, o_2) = \langle\langle\langle \preceq_1, \preceq_2, \preceq'_3\rangle, g'\rangle, \preceq'\rangle$$

with

$$\begin{aligned} g'(o_1) &= g(o_1) - [o_1 \in \mathsf{top}(\preceq_3) \wedge o_1 \notin \mathsf{top}(\preceq'_3)] \\ &\quad + [o_1 \notin \mathsf{top}(\preceq_3) \wedge o_1 \in \mathsf{top}(\preceq'_3)] = 2 - 0 + 0 = 2 \\ g'(o_2) &= 1 - 0 + 1 = 2 \\ g'(o_3) &= 0 - 0 + 0 = 0 \end{aligned}$$

and therefore $o_3 \prec' o_1, o_2$.

We now turn to the Borda preference aggregator.

Definition 8. *Let $\preceq = \langle \preceq_1, \ldots, \preceq_m \rangle$ and let $\mathcal{S}_b = \mathcal{P}^m_{\mathcal{O}} \times \mathbb{N}^{\mathcal{O}}$. Define*

$$\iota_b(\preceq) = \langle (\preceq, g), \Theta_b(\preceq)\rangle$$

with $g : \mathcal{O} \to \mathbb{N}$ and $g(o) = \sum_{i=1}^{m} rank_{\preceq_i}(o)$. Define

$$\Lambda_b((\preceq, g), i, \pm, o) = \langle (\preceq \pm_i o, g'), \preceq' \rangle$$

with $g' : \mathcal{O} \to \mathbb{N}$ defined via

$$g'(o') = g(o') - rank_{\preceq_i}(o') + rank_{\preceq_i \pm o}(o')$$

and

$$o_1 \preceq' o_2 \qquad \text{iff} \qquad g'(o_1) \leq g'(o_2)$$

Then $\Delta_b = \langle \iota_b, \Lambda_b \rangle$ is called the dynamic Borda preference aggregator.

The dynamic Borda preference aggregator is defined in a similar way as the dynamic plurality preference aggregator (in fact, both definitions can easily be generalized to obtain a scheme for dynamic preference aggregators for preference aggregators based on scoring rules). For each $o' \in \mathcal{O}$ only the rank changes of o' in $\preceq_i$ are taken into account when updating the value $g(o')$, which stores the sum of all ranks for all $\preceq_j$, $j = 1, \ldots, m$.

Proposition 7. *Δ_b is a state-based faithful representation of Θ_b.*

Proposition 8. *ι_b has time complexity $O(nm)$ and Λ_b has time complexity $O(n)$. The space complexity for storing a state in $\mathcal{S}_b$ is $O(nm)$.*

Note that the time complexity of computing $\Theta_b(\preceq)$ is $O(nm)$. Thus Λ_b has better time complexity ($O(n)$) than the canonical solution $\Lambda^{can}_{\Theta_b}$ with $O(nm)$, cf. Proposition 4.

Example 6. We continue Example 4. For $\preceq = \langle \preceq_1, \preceq_2, \preceq_3 \rangle$ we obtain

$$\iota_b(\preceq) = \langle\langle \preceq, g\rangle, \preceq\rangle\rangle$$

with $\preceq = \Theta_b(\preceq)$ given via $o_3 \prec o_2 \prec o_1$ and $g : \mathcal{O} \to \mathbb{N}$ given via $g(o_1) = 5$, $g(o_2) = 3$, and $g(o_3) = 1$. Given that o_2 is to be strengthened $\preceq_3$, i. e. $\preceq_3' = \preceq_3 +o_2$, we obtain

$$\Lambda_b((\preceq, g), 3, +, o_2) = \langle\langle\langle \preceq_1, \preceq_2, \preceq_3' \rangle, g'\rangle, \preceq'\rangle$$

with

$$\begin{aligned} g'(o_1) &= g(o_1) - rank_{\preceq_3}(o_1) + rank_{\preceq_3'}(o_1) \\ &= 5 - 2 + 1 = 4 \\ g'(o_2) &= 3 - 1 + 1 = 3 \\ g'(o_3) &= 1 - 0 + 0 = 1 \end{aligned}$$

and therefore $\preceq' = \preceq$.

The computational approaches discussed so far illustrate that in dynamic settings, preference aggregation can be implemented more effectively. This work is the first step towards investigating the issue of dynamic preference aggregation. Current work deals with investigating more complex preference aggregation mechanisms such as the Dodgson and Kemeny rules, cf. [8, 9].

5 Related Work

Preference reasoning and preference aggregation is a very active area within artificial intelligence research, economics, and other fields. However, there are only few works that deal with dynamics in preference aggregation settings. Related to our work is e. g. [10] that deals with influence of one agent's preferences to other agents' preferences. This setting also entails some dynamics as preference relations might be changed through influence. Maudet et al. investigate this setting within the framework of CP-nets [3], a specific approach to reason with preferences. They are mainly interested in computational properties in this framework but do not investigate preference change in our more general setting. The work [11] explicitly considers updates to input preference orders and their influence to aggregation. They introduce the property of *update monotonicity* for (static) preference aggregators as a novel property to assess the quality of an aggregator. An aggregator satisfies *update monotonicity* if under changes of one input preference order towards the aggregated order, the aggregation does not change. Therefore, [11] do not consider general updates and computational approaches.

Furthermore, there are approaches in belief revision that deal with dynamics of epistemic states given in form of preorders. For example, the work [12] considers iterated belief revision based on enriched preference states. There, a preference state is basically a preference order on possible worlds that is revised upon newly received evidence. The work [13] deals with revising a given preference relation with another (partial) one such that the former is modified in a minimal way to incorporate the latter. Although these works also deal with issues related to temporal evolution of (preference) orders they do not address the evolution of the aggregated orders.

Top-k querying [14] deals with effective approaches to determine the best k answers to a (relational) query, given possibly multiple preference orders. In many approaches, the final aggregated order is constructed by incrementally taking more result items into account. Similar to the setting of dynamic preference aggregation the aggregated order is also dynamically updated. However, top-k querying mechanisms consider fixed preference orders and dynamics is only implicitly present by taking more result items into account. The work [15] considers a similar setting of taking more result items with fixed preferences into account but also uses methods from preference aggregation. However, the issue of dynamics is only briefly discussed and not elaborated.

6 Summary and Conclusion

We discussed the issue of dynamics in general settings of preference aggregation under preference change. We introduced the concept of dynamic preference aggregators and investigated the consequences of atomic changes in the input preference orders. Besides some general results on the relationships between the original aggregated order and the updated aggregated order we also established a framework for computational approaches to dynamic preference aggregation. We developed dynamic preference aggregators for two simple aggregation mechanisms, namely plurality and Borda, and discussed their properties.

As mentioned in the introduction the work developed here is applied in the field of social web recommendation systems and a first prototype of a working application is currently under development. Another application of dynamic preference aggregators can also been seen in preference elicitation [16]. Preference elicitation describes the process of iteratively updating an initially empty preference order in order to determine an agent's preference order. When preference orders of multiple agents have to be elicited and their orders have to be aggregated methods for dynamic preference aggregation can be utilized.

References

1. Walsh, T.: Representing and Reasoning with Preferences. AI Magazine **28**(4) (2007) 59–70
2. Kelly, J.: Social Choice Theory: An Introduction. Springer-Verlag (1988)
3. Boutilier, C., Brafman, R.I., Domshlak, C., Hoos, H.H., Poole, D.: CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. Journal of Artificial Intelligence Research **21** (2004) 135–191
4. Rossi, F., Venable, K.B., Walsh, T.: mCP nets: Representing and Reasoning with Preferences of Multiple Agents. In: Proceedings of the 19th National Conference on Artifical Intelligence (AAAI-2004). AAAI'04 (2004)
5. Faliszewski, P., Hemaspaandra, L., Hemaspaandra, E.: How Hard Is Bribery in Elections? Journal of Artificial Intelligence Research **35** (2009) 485–532
6. Arrow, K.J.: A Difficulty in the Concept of Social Welfare. Journal of Political Economy **58**(4) (1950) 328–346
7. Pini, M.S., Rossi, F., Venable, K.B., Walsh, T.: Aggregating Partially Ordered Preferences: Impossibility and Possibility Results. In: Proceedings of the 10th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2005), Singapore, Singapore (2005)
8. Dodgson, C.L.: A Method for Taking Votes on More than Two Issues. Clarendon Press (1876)
9. Kemeny, J.: Mathematics without numbers. Daedalus **88** (1959) 577–591
10. Maudet, N., Pini, M.S., Rossi, F., Venable, K.B.: Influence and Aggregation of Preferences over Combinatorial Domains. In: Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012). (2012)
11. Burak, C., Ton, S.: Update Monotone Preference Rules. Technical Report 048, Maastricht Research School of Economics of Technology and Organization (2011)
12. Booth, R., Meyer, T., Wong, K.S.: A Bad Day Surfing is Better than a Good Day Working: How to Revise a Total Preorder. In: Proceedings of the Tenth International Conference on the Principles of Knowledge Representation and Reasoning (KR-2006). (2006)
13. Chomicki, J., Song, J.: Monotonic and Nonmonotonic Preference Revision. The Computing Research Repository **abs/cs/0503092** (2005)
14. Ilyas, I.F., Beskales, G., Soliman, M.A.: A Survey of Top-k Query Processing Techniques in Relational Database Systems. ACM Computing Surveys **40**(4) (2008)
15. Dividino, R., Gröner, G., Scheglmann, S., Thimm, M.: Ranking RDF with Provenance via Preference Aggregation. In: Proceedings of the 18th International Conference on Knowledge Engineering and Knowledge Management (EKAW-2012). (2012)
16. Sandholm, T., Boutilier, C.: Preference Elicitation in Combinatorial Auctions. In: Combinatorial Auctions. MIT Press (2006)