

Single-Peaked Opinion Updates

Robert Brederick^{1,2}, Anne-Marie George³, Jonas Israel⁴ and Leon Kellerhals⁵

¹Institut für Informatik, TU Clausthal, Germany

²Algorithm Engineering, Humboldt-Universität zu Berlin, Germany

³Analytical Solutions and Reasoning, University of Oslo, Norway

⁴Research Group Efficient Algorithms, Technische Universität Berlin, Germany

⁵Algorithmics and Computational Complexity, Technische Universität Berlin, Germany
 robert.bredereck@tu-clausthal.de, annemage@ifi.uio.no, {j.israel, leon.kellerhals}@tu-berlin.de

Abstract

We consider opinion diffusion for undirected networks with sequential updates when the opinions of the agents are single-peaked preference rankings. Our starting point is the study of *preserving single-peakedness*. We identify voting rules that, when given a single-peaked profile, output at least one ranking that is single peaked w.r.t. a single-peaked axis of the input. For such voting rules we show convergence to a stable state of the diffusion process that uses the voting rule as the agents' update rule. Further, we establish an efficient algorithm that maximises the spread of *extreme* opinions.

1 Introduction

Ahead of elections, but also for competing products on a market, finding a way to maximally spread one particular opinion through exposure of contents to targeted agents has become of increasing interest. Advances in understanding the diffusion of opinions help to grasp the extent to which opinions can be manipulated and spread in networks. We study opinion diffusion [Grandi, 2017] in a setting where agents' opinions are modelled as single-peaked rankings over a set of candidates. In each update step one agent observes the preferences of all their neighbours in the network, aggregates these by a given voting rule and changes their opinion accordingly. For issues where preferences are naturally single-peaked, it seems reasonable to assume that also the updated preferences of an agent in a diffusion process remain single-peaked. We investigate which voting rules are applicable in this sense, which lead to converging diffusion dynamics, and whether it is tractable to find update sequences that maximally spread an extreme opinion.

Research has found that computing an optimal sequence of updates to spread a specific opinion is easy for two competing opinions [Bredereck and Elkind, 2017] but it turns out to be hard in most cases involving multiple independent opinions [Auletta *et al.*, 2019; Auletta *et al.*, 2020]. However, for other scenarios, notably elections, the agents' opinions are better modeled by rankings over some candidates than by independent opinions. In this paper, we explore this additional structure on opinions which allows us to consider various known

voting rules (such as Kemeny, (weak) Dodgson, and Minimax Condorcet) as update rules. Since different opinions are not necessarily treated equally under some voting rules, our work significantly differs from work studying opinion diffusion of multiple independent opinions which often use some simple threshold function for the updating process. For example, under some voting rule one opinion (i.e. ranking) might only be adopted if and only if the majority of the agent's neighbours bares this opinion, whereas the same might not be true for other opinions under the same rule.

Originally motivated by preference aggregation in context of economic phenomena such as prices or quantities [Black, 1948], single-peakedness is probably the most prominent restricted preference domain in social choice [Brandt *et al.*, 2015; Faliszewski *et al.*, 2014; Faliszewski *et al.*, 2011]. This domain restriction solves many computational and conceptual issues of preference aggregation: among other inviting properties, the aggregation of rankings becomes tractable (e.g. for Kemeny), Condorcet cycles cannot exist, and Arrow's impossibility theorem does not apply anymore. While political elections are often not (perfectly) single-peaked, preferences in other settings often depend on some one-dimensional criterion, such as when voting on the temperature in a room, choosing the starting time of some event, fixing the voting age for an election, or considering an adequate price for a product. For Kemeny's rule (and thus the many other rules that coincide with Kemeny in the single-peaked domain), it is easy to see that, given single-peaked preferences, also at least one Kemeny outcome ranking is single-peaked. We investigate whether the same can be said for other ranking rules.

The paper first provides basic definitions in Section 2. The main part revolves around the following key questions which are discussed in Sections 3 to 5 and for each of which we briefly discuss related work here. We conclude in Section 6.

Which rules preserve single-peakedness? As we consider opinion diffusion of single-peaked preferences, we want to identify ranking rules that allow agents' preferences to remain single-peaked after updates. Under the single-peaked domain, it is known that Condorcet winners exist such that ranking adaptations of Slater's rule [Slater, 1961], Ranked Pairs/Tideman [Tideman, 1987], Beat Path/Schulze [Schulze, 2011], and Split Cycle [Holliday and Pacuit, 2020] – which repeatedly pick Condorcet winners – coincide with Ke-

Rule	CWC	CLC	SPP	EMC
Kemeny	✓	✓	✓	✓
Minimax Condorcet	✓	✗	✓	✓
Weak Dodgson	✓	✗	✓	✓
Dodgson	✓	✗	✗	✗
Copeland	✓	✓	✗	✗
Single Transferable Vote	✗	✗	✗	✗
Borda	✗	✗	✗	✗

Table 1: An overview which ranking rules are weak Condorcet winner (CWC) and loser (CLC) consistent, single-peaked preserving (SPP), and extremist majority consistent (EMC). See Sections 2 and 5.1 for definitions of the properties.

meny’s rule. Furthermore, Kemeny’s rule preserves single-peakedness [Truchon, 1998]. We show that the same is true for some rules that do not coincide with Kemeny (in particular Minimax Condorcet and weak Dodgson, as well as Borda’s and Copeland’s rule when restricted to three candidates). Furthermore, we identify that, in general, Dodgson’s rule, Copeland’s rule, Borda’s rule, and Single Transferable Vote do not preserve single-peakedness. Table 1 gives an overview over the properties of the listed rules.

Which rules converge to a stable state? While some works consider convergence of diffusion processes under simultaneous updates of the agents [Frischknecht *et al.*, 2013; Zhuang *et al.*, 2020; Chistikov *et al.*, 2020], we focus on sequential updates. For the case of only two opinions a standard update rule is to follow the (strict) majority of the neighbours’ opinions. Here, the diffusion can be shown to always converge within a bounded number of steps [Frischknecht *et al.*, 2013]. Considering more than two opinions allows to consider thresholds and in particular majority update rules or even averaging operators in the case of continuous numbers as opinions [Noorazar, 2020; Auletta *et al.*, 2019; Auletta *et al.*, 2020]. Bredereck *et al.* [2020] show that any sequence of majority updates is finite when given $k \in \mathbb{N}$ opinions. Faliszewski *et al.* [2018] consider rankings as opinions and establish among other results that following majority updates converges. In their model however every vertex corresponds to a cluster of exactly those voters that have the same ranking preference and two such clusters are connected when the rankings are the same up to one swap. Most similarly to our work, Hassanzadeh *et al.* [2013] and Brill *et al.* [2016] also consider rankings as opinions of individual agents in a network. Both consider random update sequences and allow opinion updates in form of single swaps of adjacent candidates within a voter’s ranking. We show that these types of updates can replicate any Kemeny update sequence in the single-peaked domain when replacing one Kemeny update by repeated updates of single swaps by one user. However, Hassanzadeh *et al.* consider only complete networks and Brill *et al.* consider acyclic directed networks and simple directed cycles; hence their convergence results are not transferable to our setting. We show that for single-peaked opinions Kemeny and Minimax Condorcet updates converge, i.e., any update sequence is finite.

Is it tractable to find a sequence of updates that promotes the maximal spread of an (extreme) opinion? In line with many works on opinion diffusion we study the maximal spread of one particular opinion. While for the case of two opinions a simple greedy approach guarantees a stable state with a maximal spread of one opinion [Frischknecht *et al.*, 2013], the problem of finding such a sequence when three or more opinions are present becomes NP-hard for strict majority updates even on very restricted graph structures [Bredereck *et al.*, 2020]. In addition, Auletta *et al.* [2019] show that this problem is hard for three opinions when using weak majority updates but identify some tractable graph structures. As we consider single-peaked preference rankings as opinions, we focus on the spread of *extreme* opinions, i.e., the opinions that rank the left- or rightmost candidates on the single-peaked axis the highest. Here we show that under voting rules for which an extreme ranking can only be adopted if and only if a (weak) majority of neighbours hold this opinion, a greedy approach similar to the approach for two opinions and majority updates can guarantee a maximum number of stable agents with the extreme opinion (but the remaining agents do not necessarily become stable). This case applies to Kemeny and Minimax Condorcet updates as well as updates based on weak Dodgson (as defined in [Fishburn, 1977]). Note that these results only hold since only updates to extreme opinions must follow the weak majority of their neighbours, in contrast to the general hardness results for majority updates on more than two opinions for converging sequences established by Auletta *et al.* [2019] and Auletta *et al.* [2020].

2 Preliminaries

A *preference profile* $\mathcal{P} = (C, V)$ consists of a set C of m candidates and a set V of n voters. Each voter v is associated with a preference list (or ranking) \succ_v over the candidates. If, for instance, $C = \{a, b, c\}$, the order $c \succ_v b \succ_v a$ means that voter v prefers c the most and a the least. We omit the subscript from \succ if it is clear from context. We omit the set of rankings from the profile \mathcal{P} for convenience.

Restricted Preferences. Central to our work is the concept of single-peakedness, a domain restriction that assumes that voters’ preferences are mainly influenced by the position of the candidates on a one-dimensional axis.

Definition 1. We call a preference profile $\mathcal{P} = (C, V)$ single-peaked if there is some ordering \triangleright of the candidates C , the so called single-peaked axis, such that for all $c_i, c_j, c_k \in C$ with $c_i \triangleright c_j \triangleright c_k$ or $c_k \triangleright c_j \triangleright c_i$ each voter $v \in V$ satisfies $c_j \succ_v c_k$ if $c_i \succ_v c_j$. We call the most preferred candidate of voter v ’s preference its peak $\top(v)$.

In the remainder of the paper, when talking about a single-peaked preference profile $\mathcal{P} = (C, V)$ without further specification, we assume $|C| = m$ and that \mathcal{P} is single-peaked with respect to the axis $c_1 \triangleright \dots \triangleright c_m$. We denote by \mathcal{L} the set of all rankings over candidates C and by $\mathcal{L}_{sp}^{\triangleright}(C)$ the set of single-peaked rankings over C w.r.t. the single-peaked axis \triangleright . The *extreme opinion* $(c_1 \succ \dots \succ c_m) \in \mathcal{L}_{sp}^{\triangleright}(C)$ is denoted by r^\uparrow and the opposite extreme $(c_m \succ \dots \succ c_1) \in \mathcal{L}_{sp}^{\triangleright}(C)$ by r^\downarrow .

Condorcet Winners and Losers. For a preference profile $\mathcal{P} = (C, V)$, a (weak) Condorcet winner c is a candidate that is not defeated by another candidate in a pairwise comparison, i.e., for all $c' \in C \setminus \{c\}$ we have $|\{v \in V \mid c \succ_v c'\}| \geq |\{v \in V \mid c' \succ_v c\}|$. A strict Condorcet winner is a candidate c that strictly defeats all other candidates in a pairwise comparison, i.e., for all $c' \in C \setminus \{c\}$ we have $|\{v \in V \mid c \succ_v c'\}| > |\{v \in V \mid c' \succ_v c\}|$. The notions of (weak) Condorcet loser and strict Condorcet loser are defined analogously. While this is not true for general profiles, for single-peaked profiles a weak Condorcet winner and loser always exist and the set of weak Condorcet winners can be determined as indicated by the Median Voter Theorem.

Lemma 1 ([Black, 1948]). *Let \mathcal{P} be a single-peaked preference profile. Consider the order of the voters' top-ranked candidates along the single-peaked axis. Then the set of median peaks of the voters' preferences and any candidate between median peaks is the set of (weak) Condorcet winners.*

Ranking Rules. A ranking rule is an aggregator function $\mathcal{R}: \mathcal{L}^n \rightarrow \mathcal{P}(\mathcal{L})$ (where \mathcal{P} denotes the power set) that aggregates the preference rankings of voters into an output set of preference rankings. If ranking rule \mathcal{R} always outputs rankings that rank a (weak) Condorcet winner highest whenever such a candidate exists, we call \mathcal{R} (weak) Condorcet winner consistent. Similarly, a ranking rule \mathcal{R} is called (weak) Condorcet loser consistent, if \mathcal{R} always outputs rankings that rank a (weak) Condorcet loser lowest whenever one exists. Note that while it is usually desired to aggregate the preference rankings of voters into a single output ranking, a property often called *resoluteness*, most ranking rules fail to achieve this. In this case, one often uses tie-breaking rules to resolve irresoluteness. Specific ranking rules that we consider more closely and a general tie-breaking rule that is applied throughout this paper are defined at the end of this section.

Diffusion Process. For a subset of voters $V' \subseteq V$ of a preference profile $\mathcal{P} = (C, V)$ we define the preference profile induced by V' as $\mathcal{P}[V'] = (C, V')$. A preference network is a graph $G = (V, E)$ with some profile (C, V) where the voters coincide with the vertices.

Given a preference network $G = (V, E)$ with preference profile $\mathcal{P} = (C, V)$, we consider the following opinion diffusion process. At each *update step*, a voter v (also called the *active voter*) applies a ranking rule \mathcal{R} on the preference profile induced by their neighbourhood $N(v)$ in the preference network and takes the aggregated ranking as their new opinion. Formally, we replace the preference ranking of v by a ranking in $\mathcal{R}(\mathcal{P}[N(v)])$. We speak of update rules instead of ranking rules when considering opinion diffusion processes.

If an active voter v does not change their opinion, i.e., a ranking in $\mathcal{R}(\mathcal{P}[N(v)])$ coincides with v 's current preference ranking, then we call the voter *stable*. A preference network $G = (V, E)$ with preference profile $\mathcal{P} = (C, V)$ is in a *stable state* if all the voters in V are stable. We call a sequence of voters (v_1, \dots, v_k) an *update cycle* in preference network G with preference profile \mathcal{P} and ranking rule \mathcal{R} , when updating voters' opinions along the voter sequence leads to the same preference profile \mathcal{P} for G . A ranking rule \mathcal{R} is said to *converge* if any sequence of voter updates in any preference

network G with profile \mathcal{P} is finite, i.e., results in a stable state. That is, if \mathcal{R} converges there cannot exist any update cycles.

Kemeny's rule. Kemeny's rule returns those rankings that minimise Kendall's tau distance to all voters' preference rankings. For a pair of rankings \succ, \succ' , *Kendall's tau distance* between \succ and \succ' is defined as $d_{\text{Kt}}(\succ, \succ') = \sum_{\{c, c'\} \subseteq C} d_{\succ, \succ'}(c, c')$, where $d_{\succ, \succ'}(c, c')$ is set to 0 if \succ and \succ' rank c and c' consistently, and is set to 1 otherwise.

The *Kemeny score* of a ranking r with respect to a profile $\mathcal{P} = (C, V)$ is defined as $d_{\text{Kt}}(r, \mathcal{P}) := \sum_{v \in V} d_{\text{Kt}}(r, \succ_v)$. A ranking r with a minimal Kemeny score is called a *Kemeny ranking* of \mathcal{P} and its score is the *Kemeny score* of the profile.

Brill *et al.* [2016] propose an update procedure by which voters swap the positions of two adjacent candidates in their ranking whenever the (strict) majority of their neighbours orders the two candidates in this way. When all voters' preferences are single-peaked, their update rule can replicate Kemeny's rule by making every voter execute such updates repeatedly until they reach a fixed preference. However, as the work of Brill *et al.* concerns opinion diffusion in directed graphs (that are acyclic or simple cycles), their results are not transferable to our setting. Further details on this relation are deferred to the full version of this paper.¹

Minimax Condorcet. Minimax Condorcet (MMC), also known as Simpson's rule, ranks candidates by their worst defeat in pair-wise comparisons. Let $d_{\text{pm}}(c', c) := |\{v \in V \mid c' \succ_v c\}| - |\{v \in V \mid c \succ_v c'\}|$ be the *popularity margin* between $c, c' \in C$. Then MMC ranks candidates $c \in C$ in non-decreasing order by their *MMC score* $\text{MMC}(c) := \max_{c' \in C} d_{\text{pm}}(c', c)$. MMC differs from Kemeny's rule even in the single-peaked domain, e.g., by violating Condorcet loser consistency.

Observation 1 (★¹). *MMC is not Condorcet loser consistent even in the single-peaked domain.*

Tie-Breaking. Since the active voter in the opinion diffusion process can change only to a single new opinion we introduce a tie-breaking rule for irresolute ranking rules. First, since the focus of this paper is single-peaked preferences, we always restrict the set of winning rankings to those that are single-peaked w.r.t. the given single-peak axis, if existent. Further, we envision the active voter to tend to the single-peaked ranking closest to their own, i.e., we break possible further ties by minimising Kendall's tau distance to the current opinion of the active voter. This assumption models a reluctance of agents to change their opinions and is needed in the proofs of the presented results. Any remaining ties are broken arbitrarily (deterministic). Note that a random tie-breaking rule would hinder our analysis of convergence, as stability cannot be guaranteed.

3 Preserving Single-Peakedness

In this section, we consider the question of whether a ranking rule preserves single-peakedness, i.e., outputs at least one single-peaked ranking when given a single peaked profile.

¹ Available at <https://arxiv.org/abs/2204.14094>. Proofs of results marked with ★ are deferred to the full version.

This is not only relevant for our later analysis of diffusion of single-peaked opinions, but also extends the existing research on this domain restriction.

Definition 2. A ranking rule \mathcal{R} preserves single-peakedness if, given a single-peaked preference profile $\mathcal{P} = (C, V) \in \mathcal{L}_{sp}^{\triangleright}(C)$ with axis \triangleright , it always outputs at least one single-peaked ranking w.r.t. \triangleright , i.e., $\mathcal{R}(\mathcal{P}) \cap \mathcal{L}_{sp}^{\triangleright}(C) \neq \emptyset$.

We next investigate which ranking rules preserve single-peakedness. Intriguingly, our results show that preserving single-peakedness is independent of being (weak) Condorcet winner or loser consistent: Copeland’s rule is weak Condorcet winner and loser consistent but does not preserve single-peakedness. Conversely, Minimax Condorcet preserves single-peakedness but is not Condorcet loser consistent. It is easy to construct rules that preserve single-peakedness but are not Condorcet winner consistent. One such example is the rule that always returns the ranking r^{\downarrow} .

3.1 Kemeny’s Rule and its Equivalents on the Single-Peaked Domain

Kemeny’s rule preserves single-peakedness [Cornaz *et al.*, 2013; Truchon, 1998; Betzler *et al.*, 2014]. For self containment we provide an induction-based proof in the full version.

Observation 2 (★). *Kemeny’s rule preserves single-peakedness.*

Many social choice functions can be defined by specifying their selection process on (weighted) majority graphs. By the Median Voter Theorem (Lemma 1), a majority graph under single-peaked preferences is acyclic. Thus, if two social choice functions that are defined over such graphs only differ in the way they handle cycles, then they output the same winning set under single-peaked preferences. Further, we can adapt any social choice function into a ranking rule by repeatedly deleting one of the (possibly multiple) winners from the profile and appending it to the ranking. It is easy to see that, in the single-peaked domain, the ranking adaptations of Slater’s rule [Slater, 1961], Ranked Pairs/Tideman [Tideman, 1987], Beat Path/Schulze [Schulze, 2011] and Split Cycle [Holliday and Pacuit, 2020] coincide with Kemeny’s rule.

3.2 Minimax Condorcet

To prove that MMC preserves single-peakedness we use the following two lemmas. The first shows that when computing the MMC score of a candidate we only need to consider the candidate’s direct neighbours on the single-peaked axis. The second provides us with a structure on the set of all rankings that are single-peaked with respect to the same axis.

Lemma 2 (★). *Let \mathcal{P} be a single-peaked profile on the axis $c_1 \triangleright \dots \triangleright c_m$. Then any c_i , $i \in [m]$, experiences a worst defeat in \mathcal{P} against c_{i-1} or c_{i+1} (if they exist), i.e., $d_{pm}(c_\ell, c_i) \geq d_{pm}(c_j, c_i)$ for all $j \in [m] \setminus \{i\}$ and some $\ell \in \{i-1, i+1\}$.*

Lemma 3 (★). *For an axis $c_1 \triangleright \dots \triangleright c_m$ there exists an ordering of all single-peaked rankings $r_1, \dots, r_{2^m-1} \in \mathcal{L}_{sp}^{\triangleright}(C)$ such that for all $c_i, c_{i+1} \in C$, $i \in [m-1]$, there exists a threshold $h \in [2^m-1]$ with*

$$\begin{aligned} c_i \succ_r c_{i+1} & \text{ for all rankings } r \in \{r_1, \dots, r_h\} \text{ and} \\ c_{i+1} \succ_r c_i & \text{ for all rankings } r \in \{r_{h+1}, \dots, r_{2^m-1}\}. \end{aligned}$$

Further, r_1, \dots, r_h are exactly those single-peaked rankings where the candidates c_1, \dots, c_i are the peaks.

Proof sketch. We construct the ordering iteratively. For $m = 3$ choose $(c_1 \succ c_2 \succ c_3, c_2 \succ c_1 \succ c_3, c_2 \succ c_3 \succ c_1, c_3 \succ c_2 \succ c_1)$. We then iteratively include candidates at the end of the single-peaked axis. For this we take the rankings of the previous ordering (over $m - 1$ candidates) and replace them with an ordering of rankings where the new candidate c_m is included in every possible position (not violating single-peakedness, starting with the highest position possible. Lastly we append the sequence with $c_m \succ c_{m-1} \succ \dots \succ c_1$. \square

We can now prove that MMC always outputs at least one single-peaked ranking if the input profile is single-peaked. Assuming that this is not the case, we reach a contradiction by analysing the MMC scores of three neighbouring candidates that are creating a dip with the help of the order of single-peaked rankings constructed in Lemma 3.

Theorem 1 (★). *MMC preserves single-peakedness.*

3.3 Other Rules

We show that many other ranking rules that do not coincide with the above-mentioned rules fail to preserve single-peakedness; see the full version for details. We identify counterexamples for Dodgson’s and Copeland’s rule (with 5 candidates), Single Transferable Vote (with 3 candidates), and Borda’s rule (with 4 candidates), which all can be easily extended to examples with more candidates.

Proposition 1 (★). *Dodgson’s, Copeland’s and Borda’s rule and Single Transferable Vote do not preserve single-peakedness.*

However, we can show that on 3-candidate profiles Borda’s rule preserves single-peakedness. Copeland’s rule with 3 candidates coincides with Kemeny’s rule due to weak Condorcet loser and winner consistency. It thus preserves single-peakedness on these profiles. Note that any ranking on 2 candidates is single-peaked; thus under the here discussed rules the only open cases are for Copeland’s rule on profiles with 4 candidates and for Dodgson’s rule with 3 or 4 candidates.

We remark that *weak* Dodgson [Fishburn, 1977], in which one ranks candidates by the number of pairwise swaps needed to make a candidate a *weak* Condorcet winner, *does* preserve single-peakedness; see the full version for details.

4 Convergence of Single-Peaked Updates

We identified in Section 3 that Kemeny’s rule (or its equivalents on the single-peaked domain) and MMC preserve single-peakedness. In this section, we show first for Kemeny updates and then for MMC updates (with tie-breaking as described in Section 2) that any sequence of updates is finite, i.e., ends in a stable state after finitely many updates. Note that convergence of weak Dodgson remains an open problem.

4.1 Convergence of Kemeny Updates

In the following, we outline how a stable state can be found in any network with single-peaked voter preferences that use Kemeny’s rule as an update rule (while maintaining single-peakedness and using the above mentioned tie-breaking rule).

Theorem 2. *Let $G = (V, E)$ be a network with single-peaked preference profile $\mathcal{P} = (C, V)$. Then any update sequence using Kemeny’s rule is of length at most $|E| \cdot \binom{|C|}{2}$. That is, Kemeny’s rule converges.*

Proof. We show that any update to G decreases the sum $\sum_{\{u,v\} \in E} d_{\text{Kt}}(\succ_u, \succ_v)$ of the distances by at least 1. Consider an update by Kemeny’s rule for a voter $v \in V$ with opinion \succ_v and let r be the newly obtained opinion of v , i.e., the Kemeny ranking of the profile $\mathcal{P}[N(v)]$ of v ’s neighbours in G . Let $\kappa := \sum_{u \in N(v)} d_{\text{Kt}}(\succ_u, \succ_v)$ and let $\kappa^* := \sum_{u \in N(v)} d_{\text{Kt}}(\succ_u, r)$. As r is an opinion minimising Kendall’s tau distance to the opinions of $N(v)$, we have $\kappa^* \leq \kappa$. Further, only v changes its opinion, so the sum of Kendall’s tau distances cannot increase. But if it does not change, then the opinion \succ_v is used as a tie breaker, and we have $r = \succ_v$, that is, v does not update its opinion. Lastly, no update sequence under Kemeny can have more than $|E| \cdot \binom{|C|}{2}$ updates as for any two rankings r, r' we have $d_{\text{Kt}}(r, r') \leq \binom{|C|}{2}$. \square

For a general complexity result, we show that we can efficiently compute an update based on Kemeny’s rule. For this, note that one of the ends of the single-peaked axis is always among the Condorcet losers.

Lemma 4 (★). *Computing an update step with Kemeny’s rule takes $O(|V| \cdot |C|^2)$ time.*

4.2 Convergence of MMC Updates

Recall that $\top(v)$ is the peak in voter v ’s single-peaked preference ranking. Further, denote by $d_{\text{sp}}(c_i, c_j) = |i - j|$ the distance between candidates c_i and c_j according to the single-peaked axis $c_1 \triangleright \dots \triangleright c_m$. We show that if voters update their opinion by applying MMC, the total distance between the peaks of the opinions of neighbours can never increase.

Lemma 5 (★). *Let $G = (V, E)$ be a network with single-peaked preference profile \mathcal{P} . Under MMC updates the sum of distances $\sum_{\{v,v'\} \in E} d_{\text{sp}}(\top(v), \top(v'))$ cannot increase.*

Note that while a single MMC update can be computed in polynomial time, Lemma 5 does not give us a distance measure that constantly decreases. Thus we cannot estimate how long an update sequence can be. Nevertheless, we now show the convergence of MMC as an update rule.

Theorem 3 (★). *Let $G = (V, E)$ be a network with single-peaked preference profile \mathcal{P} . Then any update sequence using MMC is finite. That is, MMC converges.*

Proof sketch. We prove the statement by induction over the number of candidates m . By Lemma 5, the sum of distances $\sum_{\{v,v'\} \in E} d_{\text{sp}}(\top(v), \top(v'))$ cannot increase with any update. In particular, in an update cycle the distances must stay equal for every update. The top choice of a voter after that voter updated must always be a weak Condorcet winner and, because of the tie breaking rule, must remain the same Condorcet winner throughout all updates. Furthermore, because of the tie breaking rule, and because the only single-peaked rankings with c_1 , respectively c_m , on the top are the

extreme opinions, Lemma 5 implies that there can be no update cycle for which opinions are switched from or to an extreme opinion. The base case of the induction, $m = 2$, is thus settled. For the induction step assume there exist no cycles of voter updates in any network with single-peaked opinions over m candidates, but that there exists an update cycle K in G with single-peaked preference profile \mathcal{P} over $m + 1$ candidates. We show that by eliminating the last candidate of the single-peaked axis of \mathcal{P} , the relative ranks of the candidates remain the same for every voter update in K and thus K must be an update cycle on G with a single-peaked preference profile over m candidates – a contradiction. \square

5 Maximally Spreading Extreme Opinions

Especially in the political context, studying extreme opinions and their spread on social networks is often seen as more important than the spread of more centralist views. Platforms such as Facebook, Youtube or Instagram invest vast resources to combat the spread of hateful and extreme opinions.

Recall that for single-peaked axis $c_1 \triangleright \dots \triangleright c_m$ the rankings $r^\uparrow = c_1 \succ \dots \succ c_m$ and $r^\downarrow = c_m \succ \dots \succ c_1$ are called the *extreme opinions*, as they rank the extreme candidates on the single-peaked axis highest. In the following, we consider under which conditions an extreme opinion can be maximally spread in a single-peaked preference network. For this purpose we define the notion of *extremist majority consistency* for ranking rules that preserve single-peakedness – a property that holds for both Kemeny’s rule and MMC. We then formulate a general algorithm for extremist majority consistent rules that maximally spreads extreme opinions.

5.1 Extremist Majority Consistency

Informally, extremist majority consistency states that an extreme opinion can only be adopted by an agent when a majority of the agent’s neighbours hold this opinion.

Definition 3. *Let $\mathcal{P} = (C, V) \in \mathcal{L}_{\text{sp}}^\triangleright(C)$. We call a ranking rule \mathcal{R} that preserves single-peakedness extremist majority consistent if for any extreme opinion $r^* \in \{r^\uparrow, r^\downarrow\}$ we have*

$$r^* \in \mathcal{R}(\mathcal{P}) \iff a \text{ (weak) majority in } V \text{ has opinion } r^*.$$

One might think that for extremist majority consistency it is sufficient that a ranking \mathcal{R} preserves single-peakedness and is weak Condorcet winner consistent. However, this is not the case since these two properties alone do not enforce that there exists an output ranking that is single-peaked and ranks the desired weak Condorcet winner (c_1 or c_m) highest.

We proceed by showing that Kemeny’s update rule and MMC are extremist majority consistent.

Theorem 4. *Kemeny’s rule is extremist majority consistent.*

Proof. Due to symmetry we need to prove the statement only for the extreme opinion r^\uparrow . Assume that at least half of the voters in $\mathcal{P} = (C, V)$ have opinion r^\uparrow . By definition of Kemeny’s rule, r^\uparrow is among the output of the rule. For the other direction assume that r^\uparrow is a Kemeny ranking and consider the ranking $r' = c_2 \succ c_1 \succ c_3 \succ \dots \succ c_m$, which is the only single-peaked ranking with $d_{\text{Kt}}(r^\uparrow, r') = 1$. Then

$d_{\text{Kt}}(r^\uparrow, r) \geq d_{\text{Kt}}(r', r) + 1$ for all single-peaked rankings $r \in \mathcal{L}_{sp}^\triangleright(C) \setminus \{r^\uparrow\}$. Also:

$$\begin{aligned} d_{\text{Kt}}(r', \mathcal{P}) &= \sum_{v \in V} d_{\text{Kt}}(r', \succ_v) \\ &= |\{v \in V \mid \succ_v = r^\uparrow\}| + \sum_{v \in V: \succ_v \neq r^\uparrow} d_{\text{Kt}}(r', \succ_v) \\ &\leq |\{v \in V \mid \succ_v = r^\uparrow\}| + \sum_{v \in V: \succ_v \neq r^\uparrow} d_{\text{Kt}}(r^\uparrow, \succ_v) - 1 \\ &\leq |\{v \in V \mid \succ_v = r^\uparrow\}| + d_{\text{Kt}}(r^\uparrow, \mathcal{P}) - |\{v \in V \mid \succ_v \neq r^\uparrow\}| \\ &\leq |\{v \in V \mid \succ_v = r^\uparrow\}| + d_{\text{Kt}}(r', \mathcal{P}) - |\{v \in V \mid \succ_v \neq r^\uparrow\}| \end{aligned}$$

and hence, $|\{v \in V \mid \succ_v = r^\uparrow\}| \geq |\{v \in V \mid \succ_v \neq r^\uparrow\}|$. \square

Note that for non-extreme opinions this equivalence need not hold, i.e., a Kemeny ranking need not be supported by a majority of the voters. However, if the majority of voters has opinion r , then r is a Kemeny ranking; see the full version of this paper for details.

MMC, just like Kemeny's rule, outputs an extreme opinion if and only if a (weak) majority of voters has this opinion.

Theorem 5 (★). *MMC is extremist majority consistent.*

5.2 Maximally Spreading Extreme Opinions for Extremist Majority Consistent Update Rules

We will next prove that for an extremist majority consistent ranking rule \mathcal{R} and an extreme opinion r^* the following greedy update sequence σ^* reaches a stable state that maximises the number of voters with opinion r^* . Note that due to step (3) below this sequence is only well defined if there exists a finite sequence of updates w.r.t. \mathcal{R} that leads to a stable state for any network.

GREEDY SEQUENCE σ^* FOR EXTREME OPINION r^*

- (1) Update every non-stable voter with opinion $r' \neq r^*$ to opinion r^* if possible.
- (2) Update every non-stable voter with opinion r^* .
- (3) Stabilize network: update non-stable voters with opinions $r' \neq r^*$.

A similar algorithm was used before in the setting with two competing opinions in [Auletta *et al.*, 2015; Bredereck and Elkind, 2017] and in the setting of *discrete preference games* in [Chierichetti *et al.*, 2013]. In particular, Bredereck and Elkind [2017, Proposition 1] show that in the binary model with only two possible opinions, first updating every non-stable voter with the second opinion to the first opinion, and then updating every non-stable voter with the first opinion to the second opinion yields a stable outcome with a maximum number of voters having the first opinion. Similarly, we derive the following result.

Lemma 6 (★). *Let $G = (V, E)$ be a preference network with single-peaked preference profile \mathcal{P} and let \mathcal{R} be an extremist majority consistent ranking rule. The subsequence σ of updates in σ^* that only runs steps (1) and (2) satisfies:*

- (i) *The length of σ is at most $2|V|$.*
- (ii) *Each voter changes their opinion at most twice on σ .*
- (iii) *Computing σ requires $O(|V|^2)$ executions of \mathcal{R} .*

(iv) *After σ , the set V^* of voters that still have opinion r^* is stable.*

(v) *For every sequence σ' that maximises the number of voters with opinion r^* such that every such voter is stable, the set of voters with opinion r^* is equal to V^* .*

If the ranking rule in use converges, then we are guaranteed that step (3) has a finite number of updates. What remains to show is that, in this case, running step (3) does not affect the stability of the voters with opinion r^* and that the set V^* of voters with opinion r^* after completing σ^* is the same as for any other sequence that reaches a stable state and maximises the spread of opinion r^* .

Proposition 2 (★). *Let $G = (V, E)$ be a network with single-peaked preference profile \mathcal{P} and an extremist majority consistent and converging ranking rule \mathcal{R} , and let r^* be an extreme opinion. Let V^* be the set of voters with opinion r^* after completing steps (1) and (2) of σ^* . The sequence σ^* satisfies:*

- (i) *After completing σ^* , every voter is stable.*
- (ii) *During step (3) of σ^* , V^* remains unchanged.*
- (iii) *For every sequence σ' that maximises the number of voters with opinion r^* such that every voter (independent of its opinion) becomes stable, the set of voters with opinion r^* is equal to V^* .*

Corollary 1. *Let $G = (V, E)$ be a network with preference profile $\mathcal{P} = (C, V) \in \mathcal{L}_{sp}^\triangleright(C)$. Then one can compute a sequence of Kemeny updates that maximises the number of voters with opinion $r^* \in \{r^\uparrow, r^\downarrow\}$ in $O(|V|^3 \cdot |C|^4)$ time.*

Corollary 2. *Let $G = (V, E)$ be a network with preference profile $\mathcal{P} = (C, V) \in \mathcal{L}_{sp}^\triangleright(C)$. Then there exists a finite sequence of MMC updates that maximises the number of voters with opinion $r^* \in \{r^\uparrow, r^\downarrow\}$ and ends in a stable state.*

6 Conclusion

Motivated by the question of how single-peaked opinions propagate in social networks we studied opinion diffusion processes in this domain. We investigated which rules admit a single-peaked outcome given single-peaked preferences. For these rules we then established convergence for arbitrary update sequences and provided an algorithm that outputs an update sequence to optimally spread an extreme opinion in any preference network.

We conclude by suggesting future research directions. While single-peakedness is a well-established domain restriction, the study of opinion diffusion under single-crossing or single-caved domains is also well motivated. Further, our initial study on opinion diffusion may be extended: can we efficiently compute update sequences to optimally spread non-extreme opinions? Lastly, while many ranking rules coincide with Kemeny's rule in the single-peaked domain, this is not necessarily true for preference profiles with bounded single-peaked width [Cornaz *et al.*, 2012]. In terms of opinion diffusion, while a Kemeny ranking can be computed efficiently whenever the single-peaked width is small [Cornaz *et al.*, 2013], it is not clear whether, e.g., our greedy sequence is applicable in this scenario.

Acknowledgements

Anne-Marie George was supported under the NRC Grant No 302203 “Algorithms and Models for Socially Beneficial Artificial Intelligence”. Jonas Israel was supported by the DFG grant BR 4744/2-1.

References

- [Auletta *et al.*, 2015] Vincenzo Auletta, Ioannis Caragiannis, Diodato Ferraioli, Clemente Galdi, and Giuseppe Persiano. Minority becomes majority in social networks. In *Proc. 11th WINE*, pages 74–88, 2015.
- [Auletta *et al.*, 2019] Vincenzo Auletta, Diodato Ferraioli, Valeria Fionda, and Gianluigi Greco. Maximizing the spread of an opinion when tertium datur est. In *Proc. 18th AAMAS*, pages 1207–1215, 2019.
- [Auletta *et al.*, 2020] Vincenzo Auletta, Diodato Ferraioli, and Gianluigi Greco. On the effectiveness of social proof recommendations in markets with multiple products. In *Proc. 24th ECAI*, pages 19–26, 2020.
- [Betzler *et al.*, 2014] Nadja Betzler, Robert Bredereck, and Rolf Niedermeier. Theoretical and empirical evaluation of data reduction for exact kemeny rank aggregation. *Auton. Agents Multi Agent Syst.*, 28(5):721–748, 2014.
- [Black, 1948] Duncan Black. On the rationale of group decision making. *J. Polit. Econ.*, 56(1):23–34, 1948.
- [Brandt *et al.*, 2015] Felix Brandt, Markus Brill, Edith Hemaspaandra, and Lane A. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *J. Artif. Intell. Res.*, 53:439–496, 2015.
- [Bredereck and Elkind, 2017] Robert Bredereck and Edith Elkind. Manipulating opinion diffusion in social networks. In *Proc. 26th IJCAI*, pages 894–900, 2017.
- [Bredereck *et al.*, 2020] Robert Bredereck, Lilian Jacobs, and Leon Kellerhals. Maximizing the spread of an opinion in few steps: Opinion diffusion in non-binary networks. In *Proc. 29th IJCAI*, pages 1622–1628, 2020.
- [Brill *et al.*, 2016] Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. Pairwise diffusion of preference rankings in social networks. In *Proc. 25th IJCAI*, pages 130–136, 2016.
- [Chierichetti *et al.*, 2013] Flavio Chierichetti, Jon Kleinberg, and Sigal Oren. On discrete preferences and coordination. In *Proc. 14th EC*, pages 233–250, 2013.
- [Chistikov *et al.*, 2020] Dmitry Chistikov, Grzegorz Lisowski, Mike Paterson, and Paolo Turrini. Convergence of opinion diffusion is pspace-complete. In *Proc. 34th AAI*, pages 7103–7110, 2020.
- [Cornaz *et al.*, 2012] Denis Cornaz, Lucie Galand, and Olivier Spanjaard. Bounded single-peaked width and proportional representation. In *Proc. 20th ECAI*, pages 270–275, 2012.
- [Cornaz *et al.*, 2013] Denis Cornaz, Lucie Galand, and Olivier Spanjaard. Kemeny elections with bounded single-peaked or single-crossing width. In Francesca Rossi, editor, *Proc. 23rd IJCAI*, pages 76–82, 2013.
- [Faliszewski *et al.*, 2011] Piotr Faliszewski, Edith Hemaspaandra, Lane A. Hemaspaandra, and Jörg Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Inf. Comput.*, 209(2):89–107, 2011.
- [Faliszewski *et al.*, 2014] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. *J. Artif. Intell.*, 207:69–99, 2014.
- [Faliszewski *et al.*, 2018] Piotr Faliszewski, Rica Gonen, Martin Koutecký, and Nimrod Talmon. Opinion diffusion and campaigning on society graphs. In *Proc. 27th IJCAI*, pages 219–225, 7 2018.
- [Fishburn, 1977] Peter C Fishburn. Condorcet social choice functions. *SIAM J. Appl. Math.*, 33(3):469–489, 1977.
- [Frischknecht *et al.*, 2013] Silvio Frischknecht, Barbara Keller, and Roger Wattenhofer. Convergence in (social) influence networks. In *Proc. 27th DISC*, pages 433–446, 2013.
- [Grandi, 2017] Umberto Grandi. Social choice and social networks. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
- [Hassanzadeh *et al.*, 2013] Farzad Farnoud Hassanzadeh, Eitan Yaakobi, Behrouz Touri, Olgica Milenkovic, and Jehoshua Bruck. Building consensus via iterative voting. In *Proc. ISIT*, pages 1082–1086. IEEE, 2013.
- [Holliday and Pacuit, 2020] Wesley H Holliday and Eric Pacuit. Split cycle: A new condorcet consistent voting method independent of clones and immune to spoilers, 2020. arXiv:2004.02350 [cs.GT].
- [Noorazar, 2020] Hossein Noorazar. Recent advances in opinion propagation dynamics: a 2020 survey. *Eur. Phys. J. Plus*, 135(6):1–20, 2020.
- [Schulze, 2011] Markus Schulze. A new monotonic, clone-independent, reversal symmetric, and Condorcet-consistent single-winner election method. *Soc. Choice Welf.*, 36(2):267–303, 2011.
- [Slater, 1961] Patrick Slater. Inconsistencies in a schedule of paired comparisons. *Biometrika*, 48(3/4):303–312, 1961.
- [Tideman, 1987] T Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Soc. Choice Welf.*, 4(3):185–206, 1987.
- [Truchon, 1998] Michel Truchon. An extension of the Condorcet criterion and Kemeny orders. Technical report, Université Laval, Québec, Canada, 1998.
- [Zhuang *et al.*, 2020] Zhiqiang Zhuang, Kewen Wang, Junhu Wang, Heng Zhang, Zhe Wang, and Zhiguo Gong. Lifting majority to unanimity in opinion diffusion. In *Proc. 24th ECAI*, pages 259–266, 2020.